

Chapman & Hall/CRC
Handbooks of Modern
Statistical Methods

Handbook of Design and Analysis of Experiments

Edited by

Angela Dean
Max Morris
John Stufken
Derek Bingham

 CRC Press
Taylor & Francis Group

A CHAPMAN & HALL BOOK

Handbook of Design and Analysis of Experiments

Chapman & Hall/CRC

Handbooks of Modern Statistical Methods

Series Editor

Garrett Fitzmaurice

*Department of Biostatistics
Harvard School of Public Health
Boston, MA, U.S.A.*

Aims and Scope

The objective of the series is to provide high-quality volumes covering the state-of-the-art in the theory and applications of statistical methodology. The books in the series are thoroughly edited and present comprehensive, coherent, and unified summaries of specific methodological topics from statistics. The chapters are written by the leading researchers in the field, and present a good balance of theory and application through a synthesis of the key methodological developments and examples and case studies using real data.

The scope of the series is wide, covering topics of statistical methodology that are well developed and find application in a range of scientific disciplines. The volumes are primarily of interest to researchers and graduate students from statistics and biostatistics, but also appeal to scientists from fields where the methodology is applied to real problems, including medical research, epidemiology and public health, engineering, biological science, environmental science, and the social sciences.

Published Titles

Handbook of Mixed Membership Models and Their Applications

*Edited by Edoardo M. Airoldi, David M. Blei,
Elena A. Erosheva, and Stephen E. Fienberg*

Handbook of Markov Chain Monte Carlo

*Edited by Steve Brooks, Andrew Gelman,
Galin L. Jones, and Xiao-Li Meng*

Handbook of Design and Analysis of Experiments

*Edited by Angela Dean, Max Morris,
John Stufken, and Derek Bingham*

Longitudinal Data Analysis

*Edited by Garrett Fitzmaurice, Marie Davidian,
Geert Verbeke, and Geert Molenberghs*

Handbook of Spatial Statistics

*Edited by Alan E. Gelfand, Peter J. Diggle,
Montserrat Fuentes, and Peter Guttorp*

Handbook of Survival Analysis

*Edited by John P. Klein, Hans C. van Houwelingen,
Joseph G. Ibrahim, and Thomas H. Scheike*

Handbook of Missing Data Methodology

*Edited by Geert Molenberghs, Garrett Fitzmaurice,
Michael G. Kenward, Anastasios Tsiatis, and Geert Verbeke*

Chapman & Hall/CRC
Handbooks of Modern
Statistical Methods

Handbook of Design and Analysis of Experiments

Edited by

Angela Dean

The Ohio State University
Columbus, Ohio, USA

Max Morris

Iowa State University
Ames, Iowa, USA

John Stufken

Arizona State University
Tempe, Arizona, USA

Derek Bingham

Simon Fraser University
Burnaby, British Columbia, Canada



CRC Press

Taylor & Francis Group
Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an **informa** business

A CHAPMAN & HALL BOOK

MATLAB® is a trademark of The MathWorks, Inc. and is used with permission. The MathWorks does not warrant the accuracy of the text or exercises in this book. This book's use or discussion of MATLAB® software or related products does not constitute endorsement or sponsorship by The MathWorks of a particular pedagogical approach or particular use of the MATLAB® software.

CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2015 by Taylor & Francis Group, LLC
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works
Version Date: 20150514

International Standard Book Number-13: 978-1-4665-0434-9 (eBook - PDF)

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>

Contents

Preface.....	xi
Editors.....	xv
Contributors.....	xvii

Section I General Principles

1. History and Overview of Design and Analysis of Experiments.....	3
<i>Klaus Hinkelmann</i>	
2. Introduction to Linear Models.....	63
<i>Linda M. Haines</i>	

Section II Designs for Linear Models

3. Blocking with Independent Responses.....	99
<i>John P. Morgan</i>	
4. Crossover Designs.....	159
<i>Mausumi Bose and Aloke Dey</i>	
5. Response Surface Experiments and Designs.....	197
<i>André I. Khuri and Siuli Mukhopadhyay</i>	
6. Design for Linear Regression Models with Correlated Errors.....	237
<i>Holger Dette, Andrey Pepelyshev, and Anatoly Zhigljavsky</i>	

Section III Designs Accommodating Multiple Factors

7. Regular Fractional Factorial Designs.....	279
<i>Robert Mee and Angela Dean</i>	
8. Multistratum Fractional Factorial Designs.....	321
<i>Derek Bingham</i>	
9. Nonregular Factorial and Supersaturated Designs.....	339
<i>Hongquan Xu</i>	

10. **Structures Defined by Factors**..... 371
R.A. Bailey
11. **Algebraic Method in Experimental Design** 415
Hugo Maruri-Aguilar and Henry P. Wynn

Section IV Optimal Design for Nonlinear and Spatial Models

12. **Optimal Design for Nonlinear and Spatial Models: Introduction and Historical Overview** 457
Douglas P. Wiens
13. **Designs for Generalized Linear Models** 471
Anthony C. Atkinson and David C. Woods
14. **Designs for Selected Nonlinear Models** 515
Stefanie Biedermann and Min Yang
15. **Optimal Design for Spatial Models** 549
Zhengyuan Zhu and Evangelos Evangelou

Section V Computer Experiments

16. **Design of Computer Experiments: Introduction and Background**..... 577
Max D. Morris and Leslie M. Moore
17. **Latin Hypercubes and Space-Filling Designs**..... 593
C. Devon Lin and Boxin Tang
18. **Design for Sensitivity Analysis** 627
William Becker and Andrea Saltelli
19. **Expected Improvement Designs**..... 675
William I. Notz

Section VI Cross-Cutting Issues

20. **Robustness of Design** 719
Douglas P. Wiens

21. **Algorithmic Searches for Optimal Designs** 755
Abhyuday Mandal, Weng Kee Wong, and Yaming Yu

Section VII Design for Contemporary Applications

22. **Design for Discrete Choice Experiments** 787
Heiko Grossmann and Rainer Schwabe
23. **Plate Designs in High-Throughput Screening Experiments for Drug Discovery** 833
Xianggui Qu (Harvey) and Stanley Young
24. **Up-and-Down Designs for Dose-Finding** 857
Nancy Flournoy and Assaf P. Oron
25. **Optimal Design for Event-Related fMRI Studies** 895
Jason Ming-Hung Kao and John Stufken

Preface

Experiments are run for many purposes and are prevalent in a wide range of application areas. Examples of experimental contexts include engineering for the design and improvement of complex systems, pharmaceutical industry efforts to discover and develop new drugs, and the improvement of crop yield and sustainability in agriculture. Recent improvements in computer speed and memory have led to important developments in the computational aspects of the statistical challenges associated with the various components of experimentation. This handbook provides a description of many of the recent advances in the field, as well as a detailed, comprehensive overview of the basic tools required for the optimal design of experiments and their analyses.

Although there are textbooks available that deal with the traditional areas of design of experiments such as block designs and factorial experiments and a few textbooks that discuss individual specialized areas, this handbook is a comprehensive book that gives a careful treatment of the latest developments of a wide range of topics in a unified whole. It consists of a carefully edited collection of 25 chapters that synthesize the state of the art in the theory and applications of designed experiments and their analyses. The chapters present a good blend of methodology and application and are written by leading researchers in the field.

The chapters have been arranged into “sections” to bring focus to the work. However, they can all be read independently, and cross-references to other chapters are given where helpful. Section I begins with a historical look at experimental design (Chapter 1), from its beginnings in agriculture for controlling the effects of environmental and soil variation to more modern uses in industry, medicine, marketing, and beyond. In particular, the chapter serves as an introduction to the material in Sections II and III. The history is followed in Chapter 2 with a detailed look at the theory of estimation of parameters in linear models and the specification of optimality criteria by which to judge the efficiency of a design. This fundamental theory is referenced by many other chapters in their discussion of optimal design.

Section II deals with a number of different settings in which the response is modeled by a linear model. An account is given of the current state of knowledge in block designs (Chapter 3), crossover and repeated measurement designs (Chapter 4), designs for estimating response surfaces (Chapter 5), and a look at optimal designs for linear models in which the error terms are correlated (Chapter 6).

In Section III, more attention is given to designs with multiple factors (both treatment and blocking factors). Factorial experiments are considered in several of these chapters. In particular, the construction of fractional factorial designs is discussed in Chapter 7, and the art of arranging factorial experiments in blocks, split plots, and split lots is covered in Chapter 8. A recent area of growth in the design of experiments lies in the need to handle large numbers of factors within a physical experiment. This has led to the development of a new methodology for screening, which is the process of searching for the most active variables. Small “nonregular” designs for screening are discussed in Chapter 9. Section III ends with two chapters (Chapters 10 and 11) that discuss the structure of multifactor designs, models, and analysis.

Another area of research emphasis that has become important in recent years is that of optimal designs for generalized linear models, other nonlinear models, and spatial models; these form the topic of Section IV. Such models are used extensively in many fields, including medical sciences, social sciences, marketing, and agriculture. There is an extensive machinery for analyzing data under such models, but while there are results that go back more than two decades, research on designs for nonlinear models, in particular, is still very much in its developmental stage. Chapters 12 through 15 give details of the current knowledge in these important areas.

Another area that has seen enormous growth in the past 20 years is that of computer experiments, which is discussed in Section V. Introduced in Chapter 16, computer experiments are performed on large-scale simulators of physical systems. Frequently, a satisfactory mathematical description of the system is computationally intensive, and thus, a careful selection of design points (simulator trials) is critical. Section V addresses general classes of experimental designs that have proven to be useful in computer experiments (Chapter 17), input sensitivity analysis (Chapter 18), and experiments designed for the more specific goals of function optimization (Chapter 19).

Two chapters (Chapters 20 and 21) deal with some “cross-cutting” issues that are relevant to all of experimental design (Section VI). Chapter 20 focuses on how experimental design can be made *robust* to errors in modeling assumptions. Chapter 21 reviews some standard computer algorithms that are associated with designs for linear models and other, more recent algorithms that are applicable to a wider variety of design problems.

Section VII explores the application of experimental design in areas that have developed more recently, and which have unique and interesting characteristics. Included are discrete choice experiments that are used in many fields of study for understanding consumer decisions (Chapter 22); high-throughput screening experiments that search among huge numbers of chemical compounds for those molecules that can be used in the next generation of drugs (Chapter 23); dose-finding designs for finding the median effective dose of a drug or, for example, determining a failure threshold in engineering (Chapter 24); fMRI experiments that are widely used in research fields such as cognitive neuroscience, medical science, and psychology for studying the functions of the brain (Chapter 25). Each of these applications presents new challenges and opens up entirely new directions of research to which not many researchers have yet been exposed in a substantial way.

Not every topic can be covered in great detail in a book such as this one. Indeed, much more could be said on the important topics of screening designs, Bayesian designs, and clinical trials, for example, as well as specialized topics in the design and analysis of computer experiments, such as the calibration of models to physical data. And of course, Section VII contains the discussion of only a small sampling of the unique design issues that arise in specific application areas, and interesting modern topics such as designs for web-based experiments and networks are not included. Nevertheless, this handbook gives a taste of the broad range of uses of experimental design, the current knowledge in these areas, and some indications for further reading on related topics. It provides a comprehensive overview of many techniques and applications that a new researcher in the area of design and analysis of experiments needs to know. In addition, the book provides a valuable reference tool for research statisticians working in engineering and manufacturing, the basic sciences, and any discipline that depends on controlled experimental investigation.

All of the chapters in this handbook are written by experts in the field and contain not only considerable detail on the state of the art in the respective areas but also numerous

examples for illustration. We thank each and every author for their willingness to share their expertise and knowledge, for their patience with our editing, and for their enthusiasm in joining this project. We also thank Rob Calver, senior acquisitions editor at Taylor & Francis Group, for inviting us to prepare this book and for his guidance throughout the process.

**Angela Dean
Max Morris
John Stufken
Derek Bingham**

MATLAB[®] is a registered trademark of The MathWorks, Inc. For product information, please contact:

The MathWorks, Inc.
3 Apple Hill Drive
Natick, MA 01760-2098 USA
Tel: 508-647-7000
Fax: 508-647-7001
E-mail: info@mathworks.com
Web: www.mathworks.com

Editors

Derek Bingham is professor in the Department of Statistics and Actuarial Science at Simon Fraser University, Burnaby, British Columbia, Canada. After receiving his PhD from the Department of Mathematics and Statistics from the same university in 1999, he joined the Department of Statistics at the University of Michigan as an assistant professor. He returned to Simon Fraser University in 2003, joining the Department of Statistics and Actuarial Science as the Canada Research Chair in Industrial Statistics. He has also held a faculty affiliate position at Los Alamos National Laboratory. His primary research interests lie in the design and analysis of physical and computer experiments. He has coauthored research papers that have won the Youden and Wilcoxon prizes in Technometrics. He was the 2013 recipient of the Centre de Recherche Mathématique—Statistical Society of Canada prize awarded in recognition of “a statistical scientist’s professional accomplishments in research during the first fifteen years after earning a doctorate.”

Angela Dean is professor emeritus in the Statistics Department and a member of the Emeritus Academy at The Ohio State University, Columbus, Ohio. She is a fellow of the American Statistical Association and the Institute of Mathematical Statistics and an elected member of the International Statistical Institute. Her primary research focus is on the design of screening experiments. She is coauthor (with D.T. Voss) of the textbook *Design and Analysis of Experiments* (Springer-Verlag, Berlin, Germany, 1999) and coeditor (with S.M. Lewis) of the book *Screening: Methods for Industrial Experimentation, Drug Discovery and Genetics* (Springer-Verlag, Berlin, Germany, 2006). Together with D. Majumdar, K. Challoner, and D.K. Lin, she co-organized the First Midwest Conference in Experimental Design, 2000, and the Conference for Future Directions in Experimental Design, 2003. She has held permanent or visiting positions at The Open University (United Kingdom), the University of Texas, the University of Wisconsin, and the University of Southampton (United Kingdom). She has served as chair of the Section on Physical and Engineering Sciences of the American Statistical Association.

Max Morris is professor and chair of the Department of Statistics at Iowa State University (ISU), Ames, Iowa, and holds a courtesy appointment in the Department of Industrial and Manufacturing Systems Engineering. He has received the Jack Youden Prize and the Frank Wilcoxon Prize from the American Society for Quality, was the 2002 recipient of the Jerome Sacks Award for Cross-Disciplinary Research from the National Institute of Statistical Sciences, and was selected as a fellow of the American Statistical Association in 1994. His research program focuses on the design and analysis of experiments, with special emphasis on those that involve computer models. He is the author of *Design of Experiments: An Introduction Based on Linear Models* (Chapman & Hall/CRC, Boca Raton, FL, 2010). He served as editor of *Technometrics* from 1996 to 1998. He held positions at the University of Texas Health Sciences Center at San Antonio, the Statistics Department of Mississippi State University, and Oak Ridge National Laboratory before joining the ISU faculty. He has also held faculty affiliate and consultant positions at Los Alamos National Laboratory and Ames Laboratory of the U.S. Department of Energy since 2000.

John Stufken is the Charles Wexler Professor in Statistics in the School of Mathematical and Statistical Sciences at Arizona State University, Phoenix, Arizona. He is a fellow of the American Statistical Association and the Institute of Mathematical Statistics and an elected member of the International Statistical Institute. His primary area of research interest is design and analysis of experiments, having published approximately 75 publications in this area. He is coauthor (with A. Hedayat and N.J.A. Sloane) of the Springer-Verlag book *Orthogonal Arrays: Theory and Applications*, published in 1999. He has served as executive editor for the *Journal of Statistical Planning and Inference* and as editor for *The American Statistician*. He was co-organizer (with A. Mandal) of the Design and Analysis of Experiments 2012 conference. He has previously held positions at the University of Georgia, Iowa State University, and the National Science Foundation and more recently was professor and head of the Department of Statistics at the University of Georgia. In 2011, he held a one-month appointment as the Rothschild Distinguished Visiting Fellow at the Isaac Newton Institute for Mathematical Sciences in Cambridge, United Kingdom.

Contributors

Anthony C. Atkinson

Department of Statistics
London School of Economics
London, United Kingdom

R.A. Bailey

School of Mathematics and Statistics
University of St. Andrews
St. Andrews, United Kingdom

William Becker

Econometrics and Applied Statistics Unit
Joint Research Centre
European Commission
Ispra, Italy

Stefanie Biedermann

School of Mathematics
University of Southampton
Southampton, United Kingdom

Derek Bingham

Department of Statistics and Actuarial
Science
Simon Fraser University
Burnaby, British Columbia, Canada

Mausumi Bose

Applied Statistics Unit
Indian Statistical Institute, Kolkata
Kolkata, India

Angela Dean

Department of Statistics
The Ohio State University
Columbus, Ohio

Holger Dette

Department of Mathematics
Ruhr University Bochum
Bochum, Germany

Aloke Dey

Theoretical Statistics and Mathematics Unit
Indian Statistical Institute, Delhi
New Delhi, India

Evangelos Evangelou

Department of Mathematical Sciences
University of Bath
Bath, United Kingdom

Nancy Flournoy

Department of Statistics
University of Missouri-Columbia
Columbia, Missouri

Heiko Grossmann

School of Mathematical Sciences
Queen Mary University of London
London, United Kingdom

Linda M. Haines

Department of Statistical Sciences
University of Cape Town
Rondebosch, South Africa

Klaus Hinkelmann

Department of Statistics
Virginia Tech
Blacksburg, Virginia

Jason Ming-Hung Kao

School of Mathematical and Statistical
Sciences
Arizona State University
Tempe, Arizona

André I. Khuri

Department of Statistics
University of Florida
Gainesville, Florida

C. Devon Lin

Department of Mathematics and Statistics
Queen's University
Kingston, Ontario, Canada

Abhyuday Mandal

Department of Statistics
University of Georgia
Athens, Georgia

Hugo Maruri-Aguilar

School of Mathematical Sciences
Queen Mary University of London
London, United Kingdom

Robert Mee

Department of Business Analytics and
Statistics
University of Tennessee
Knoxville, Tennessee

Leslie M. Moore

Statistical Sciences Group
Los Alamos National Laboratory
Los Alamos, New Mexico

John P. Morgan

Department of Statistics
Virginia Tech
Blacksburg, Virginia

Max D. Morris

Department of Statistics
Iowa State University
Ames, Iowa

Siuli Mukhopadhyay

Department of Mathematics
Indian Institute of Technology, Bombay
Mumbai, India

William I. Notz

Department of Statistics
The Ohio State University
Columbus, Ohio

Assaf P. Oron

Children's Core for Biomedical Statistics
Seattle Children's Research Institute
Seattle, Washington

Andrey Pepelyshev

School of Mathematics
Cardiff University
Cardiff, United Kingdom

Xianggui Qu (Harvey)

Department of Mathematics and Statistics
Oakland University
Rochester, Michigan

Andrea Saltelli

Econometrics and Applied Statistics Unit
Joint Research Centre
European Commission
Ispra, Italy

Rainer Schwabe

Institute for Mathematical Stochastics
Otto von Guericke University Magdeburg
Magdeburg, Germany

John Stufken

School of Mathematical and Statistical
Sciences
Arizona State University
Tempe, Arizona

Boxin Tang

Department of Statistics and Actuarial
Science
Simon Fraser University
Burnaby, British Columbia, Canada

Douglas P. Wiens

Department of Mathematical and Statistical
Sciences
University of Alberta
Edmonton, Alberta, Canada

Weng Kee Wong

Department of Biostatistics
School of Public Health
University of California, Los Angeles
Los Angeles, California

David C. Woods

Southampton Statistical Sciences Research
Institute
University of Southampton
Southampton, United Kingdom

Henry P. Wynn

Department of Statistics
London School of Economics
London, United Kingdom

Hongquan Xu

Department of Statistics
University of California, Los Angeles
Los Angeles, California

Min Yang

Department of Mathematics, Statistics and
Computer Science
University of Illinois at Chicago
Chicago, Illinois

Stanley Young

National Institute of Statistical Sciences
Research Triangle Park, North Carolina

Yaming Yu

Department of Statistics
University of California, Irvine
Irvine, California

Anatoly Zhigljavsky

School of Mathematics
Cardiff University
Cardiff, United Kingdom

Zhengyuan Zhu

Department of Statistics
Iowa State University
Ames, Iowa

Section I

General Principles

1

History and Overview of Design and Analysis of Experiments

Klaus Hinkelmann

CONTENTS

1.1	Introduction.....	4
1.2	Rothamsted Experiments.....	5
1.2.1	Broadbalk Experiment.....	5
1.2.2	Hoosfield Experiment.....	6
1.3	Statistical Foundations of Scientific Experimentation: Three Principles.....	7
1.3.1	Replication.....	7
1.3.2	Randomization.....	8
1.3.3	Blocking.....	9
1.3.4	Analysis of Covariance.....	10
1.3.5	Systematic and Trend-Free Designs.....	10
1.4	Latin Square Principle.....	11
1.4.1	Latin Square Design.....	11
1.4.2	Change-Over Designs.....	13
1.4.3	Orthogonal Latin Squares.....	15
1.5	Factorial Experiments.....	16
1.5.1	Rationale: Efficiency and Comprehensiveness.....	16
1.5.2	Confounding.....	18
1.5.3	Some General Results.....	21
1.5.4	Asymmetrical Factorials.....	22
1.5.5	Split-Plot Experiments.....	23
1.6	Incomplete Block Designs.....	23
1.6.1	Balanced Incomplete Block Designs.....	24
1.6.2	Pseudo-Factorial Designs.....	26
1.6.3	Partially Balanced Incomplete Block Designs.....	28
1.6.4	Youden Square Designs.....	29
1.6.5	Comparing Treatments with a Control.....	30
1.7	Fractional Factorial Designs.....	31
1.7.1	Fractional Replication.....	31
1.7.2	Characterization of Regular Fractional Factorials.....	34
1.7.3	Main Effect Plans.....	35
1.7.4	Supersaturated and Search Designs.....	36
1.8	Industrial and Technological Experimentation.....	37
1.8.1	Early Use of Experimental Design.....	37
1.8.2	Basic Ideas of Response Surface Methodology.....	37
1.8.3	Response Surface Designs.....	38

1.8.4	Mixture Experiments.....	40
1.8.5	Computer Experiments.....	41
1.9	Efficiency and Optimality of Designs.....	42
1.9.1	Relative Efficiency.....	42
1.9.2	Efficiency Factor.....	43
1.9.3	Optimality.....	43
1.10	Matters of Analysis and Computation.....	44
1.10.1	Nonorthogonal Designs.....	44
1.10.2	Complex Error Structures.....	45
1.10.3	Graphical Methods.....	46
1.10.4	Statistical Software Packages.....	47
1.11	Areas of Application	47
1.11.1	Genetics and Breeding.....	48
1.11.2	Agronomy.....	49
1.11.3	Medicine and Pharmacology.....	49
1.11.4	Marketing.....	51
1.11.5	Manufacturing and Product Development.....	51
1.12	Epilogue.....	52
	References.....	53

1.1 Introduction

“An experiment is simply a question put to nature in the hope of discovering some secret” (Russell 1926, p. 989). With this in mind, there is obviously no question that experimentation has played a major role in human life and human endeavors over centuries and that inquiring minds have thought about proper designing of experiments. We can, of course, only speculate as to what extent experimentation has taken place early on, since no records of the earliest such attempts are available. One of the earliest descriptions of a scientific, controlled experiment is of an experiment on the treatment of scurvy as carried out by James Lind on board the HM Bark *Salisbury* in 1747 (Lind 1753). This can be regarded as one of the first recorded experiments with a medical or biological application. The beginning of agricultural experimentation can be traced back to the long-term field experiments started by John Lawes and Henry Gilbert at Rothamsted, United Kingdom, 25 miles north of London, between 1843 and 1856, among them the famous Broadbalk winter wheat experiment investigating the effects on crop yields of inorganic compounds (Rothamsted Research 2006). Even though these experiments were not laid out according to modern design principles, they provided inspiration later to R.A. Fisher and Frank Yates to develop just these principles and many of today’s well-known experimental designs through their pioneering work at Rothamsted Experimental Station, beginning in the 1920s. We shall trace these historic developments and extensions thereof in this chapter, from both a theoretical and an applications-oriented perspective.

In Section 1.2, we describe the famous long-term field experiments at Rothamsted, which later became the Rothamsted Experimental Station. It is here that Fisher and Yates started their pioneering work. In Section 1.3, we discuss the cornerstones of modern experimental design, namely, the notions of replication, randomization, and blocking. We trace further

developments in the form of the Latin square principle (Section 1.4), factorial experiments (Section 1.5), incomplete block designs (IBDs) (Section 1.6), fractional factorial designs (Section 1.7), and response surface and mixture designs (Section 1.8). Considerations of efficiency and optimality are taken up in Section 1.9 and matters of analysis and computation in Section 1.10. We conclude this chapter with a brief discussion of areas of application (Section 1.11), with emphasis on extensions and modifications of familiar designs discussed earlier.

The field of experimental design is very large, and not all areas and developments can be covered or covered completely in one chapter. The reader, however, should get the general picture of how the area developed from its very beginnings to the most recent times. Details about more recent research on many topics in experimental design will be given in subsequent chapters of this handbook, and we shall make references to them throughout. In addition to articles and books cited in this chapter, a useful source of further references, particularly for pre-1950 publications, is the bibliography collected by Federer and Balaam (1973).

1.2 Rothamsted Experiments

Between 1843 and 1856, Lawes and Gilbert started nine long-term field experiments at Rothamsted. Their objective was to investigate the effects of various inorganic compounds, individually and in combination, on several agricultural crops. Of the nine experiments, they abandoned only one and changed the treatments in others during the first years based on the results they observed. The eight experiments were then continued more or less as originally planned. They are known as the *classical experiments* and represent the oldest, continuous agronomic experiments in the world. In the following, we shall consider the essential features of two of these experiments; for more details, see Rothamsted Research (2006).

1.2.1 Broadbalk Experiment

The Broadbalk winter wheat experiment was begun in 1843. The objective was to compare the effects of inorganic fertilizers supplying the elements of nitrogen (N), phosphate (P), potash (a nutrient containing potassium) (K), soda (a nutrient containing sodium) (Na), and magnesium (Mg) in various combinations with the effects of organic manures. The experiment was put on a permanent basis in 1852, consisting of a series of long, narrow plots across the entire field (for a picture of the layout, see Rothamsted Research 2006). The essential features of the scheme of the fertilizer (treatment) assignments are given in Table 1.1 as presented by Yates (1935). This clearly shows the beginning of an incomplete factorial arrangement (as discussed in Section 1.5).

Inspection of Table 1.1 indicates that the effects of the various compounds can be estimated—for the most part—from the difference of a single pair of plots. The same is true for comparing the effect of different combinations of compounds versus the effect of farm manure. In addition, the effect of increasing the level of nitrogen can be estimated from the plots using different levels of nitrogen, represented by Yates as 0, $\frac{1}{2}$, 1, $1\frac{1}{2}$, and 2.

TABLE 1.1
Fertilizer Assignment for Broadbalk Wheat Experiment

Plot No.	Treatment				
2.2	Farmyard manure				
3–4	No manure				
5	—	P	K	Na	Mg
6	1/2 N	P	K	Na	Mg
7	N	P	K	Na	Mg
8	1 1/2 N	P	K	Na	Mg
16	2 N	P	K	Na	Mg
10	N	—	—	—	—
11	N	P	—	—	—
13	N	P	K	—	—
12	N	P	—	3 ² /3 Na	—
14	N	P	—	—	2 ³ /4 Mg
2.1	—	—	K	Na	Mg

The experiment was modified in 1968, with one important change being the introduction of short-strawed cultivars of wheat and another being the subdivision of some sections in order to compare the yield of wheat grown continuously versus that of wheat grown in rotation after a 2-year break. Results from the experiment provide important information concerning the problem of food production for a growing world population, as illustrated in an article *The Economist* (Parker 2011). It shows that the Broadbalk experiment is a microcosm of the world’s current growing conditions, from the very poor in parts of Africa to the best in the developed parts of the world, represented by the plots that had no fertilizer or pesticides applied and those plots that received the best plants, fertilizers, fungicides, and best agronomic practices, respectively.

1.2.2 Hoosfield Experiment

The Hoosfield experiment on Spring Barley has been conducted (with some modifications) since 1852. This experiment, too, is, in the words of Yates (1935), a *complex experiment* as it uses a factorial structure for its fertilizer treatments but in a design that is quite different from the Broadbalk experiment: in four strips, it uses the nutrients/nutrient combinations 0 (nothing), P, KMgNa, and PKMgNa, and these are *crossed* with four forms of N (no nitrogen, ammonium sulfate [N₁], sodium nitrate [N₂], rape cake [N₃]) as shown in Figure 1.1 (adapted from Rothamsted Research 2006; see also Yates 1935).

In modern terminology, this is a 2 × 2 × 4 factorial arranged in a split-block design (see Hinkelmann and Kempthorne 2008) without replication, to which were added plots with farmyard manure and no manure (as indicated in Figure 1.1). Even though all 16 nutrient plots can be used for the estimation of the various effects of the nutrients, it is doubtful that they actually were. Also, a serious drawback to the arrangement is the possible fertility gradient along the length of the field. If the northernmost part of the field has the highest soil fertility, then the effect of rape cake (N₃) may turn out to be falsely declared significant. This problem does not arise in the Broadbalk experiment as each treatment there is applied over the entire length of the field.

					N ₃
					N ₂
FYM					N ₁
No manure					No N
	PKMgNa	KMgNa	P	No mineral	

FIGURE 1.1

Hoosfield experimental layout. (Adapted from Rothamsted Research, *Guide to the Classical and Other Long-Term Experiments, Datasets and Sample Archives*, Rothamsted Research, Harpenden, Herts, 2006.)

1.3 Statistical Foundations of Scientific Experimentation: Three Principles

The Rothamsted long-term experiments, of which we have described two earlier, have proved to be very important, not only for agricultural research but also for the development of the statistical principles of experimental design. This development, in the context of agronomic experimentation, is closely connected with the statisticians who worked at Rothamsted, foremost among are them R.A. Fisher and Frank Yates. Efforts to analyze the data from the experiments with a statistical point of view in mind illustrated several flaws in the original layouts and led Fisher to the introduction of important principles for designing experiments. In her biography of Fisher, Fisher Box (1978) described how her father started his work at Rothamsted in 1919 and studied and analyzed the data from the Broadbalk experiment. From the articles "Studies in Crop Variation," "...we may see how, concurrently with work on estimation, Fisher was led to the analysis of variance procedure and to the principles of experimental design. Within 7 years he had solved all major problems and placed in the hands of the experimenter both the techniques for conducting experiments and the mathematical and arithmetical procedures for making sense of the results" (Fisher Box 1978, p. 100).

1.3.1 Replication

It had long been recognized that the variation in soil fertility presented great difficulties for making valid comparisons between the various manuring treatments using designs as advocated in the early experiments. One way out of this dilemma was the introduction of *uniformity trials*, that is, trials without treatments, a year or even several years prior to the real experiment. This provided some measures of experimental error but was, in most cases, not a feasible method in practice. Then, the idea of *replication* was introduced,

typically two or three replications, in the form of either replicating the same experiment at a different location or replicating each treatment on the same field. In the latter case, the recommendation was to use a *balanced* or *systematic* plan (Russell 1926). For more than two replicates and the case of two treatments, say A and B, in strips, Russell wrote that “the arrangement

ABABABAB

is not sufficient because A is always left to B and will always come out better if the fertility of the land is falling off from the left to the right of the plots,” rather the arrangement should be

ABBAABBA

so that “each treatment is compared with itself on one side and the other treatment on the other side and the plots are balanced about the centre. Whichever way the fertility may be varying the comparisons can still be made” (Russell 1926, p. 996) (see also Section 1.3.5 for trend-free designs). Russell, who was the Director of Rothamsted Experimental Station, provided other examples of balanced designs for more complex experiments.

1.3.2 Randomization

The question about valid estimation of error was carried further by R.A. Fisher, who had joined Rothamsted Experimental Station in 1919. He was critical of the systematic arrangements and provided arguments that these arrangements might lead to wrongly estimated standard errors. He instead proposed what was then a revolutionary idea, namely, that of *randomization*: “One way of making sure that a valid estimate of error will be obtained is to arrange the plots deliberately at random, so that no distinction can creep in between pairs of plots treated alike and pairs treated differently; in such a case an estimate of error, derived in the usual way from the variations of plots treated alike, may be applied to test the significance of the observed difference between the averages of plots treated differently” (Fisher 1926, p. 506). Moreover he claimed that this estimate of error could be used in testing hypotheses about the treatment effects: “The estimate of error is valid, because, if we imagine a large number of different results obtained by different random arrangements, the ratio of the real to the estimated error, calculated afresh for each of these arrangements, will be actually distributed in the theoretical distribution by which the significance of the result is tested” (Fisher 1926, p. 507). The theoretical distribution refers, of course, to the normal distribution. Thus, Fisher suggested that randomized experiments could be analyzed as if the observations were independent and normally distributed (see Fisher Box, 1978), a result that was justified theoretically by Kempthorne (1952, 1955), who showed that the randomization test for testing hypotheses about treatment effects can be approximated by the *F*-test (see also Hinkelmann and Kempthorne 2008). This result is still not widely appreciated today as most textbooks postulate that the errors have to be normally distributed.

The role played by randomization in experimental design is still a topic of great interest and current research (see, e.g., Calinski and Kageyama 2000, 2003, and Chapter 3). Important advances, in particular in the context of multiple and structured randomizations, were made by, for example, Brien and Payne (1999) and Brien and Bailey (2006). Extensions of the Kempthorne approach to various forms of block designs (Section 1.3.3) were given

by Hinkelmann and Alcorn (1998). Applications of randomization in medical trials (see Section 1.11.3) are discussed by Rosenberger and Lachin (2002).

1.3.3 Blocking

In the context of various forms of replication of treatments (varieties), Fisher (1926) also considered different components of error when he alluded to errors that arise in the course of experimentation, which should be eliminated as much as possible through careful experimental practices, and errors that arise because of experimental conditions, which should be controlled through careful arrangement of the treatments. He illustrated this in terms of what he called the most flexible and useful type of arrangement, namely, the arrangement in *randomized blocks* (see Chapter 3). Even though the use of blocks, containing more or less uniform plots, was already part of agricultural experimentation, this nevertheless represented a formal statement of the *reduction of error through blocking*, or as Fisher (1925, 1970, p. 265) put it: "...it is still possible to eliminate much of the effect of soil heterogeneity, and so increase the accuracy of our observations, by laying restrictions on the order in which the strips are arranged." He then showed how—through the use of analysis of variance (ANOVA)—the total variation can be separated into three parts representing (1) local differences between blocks, (2) differences due to treatment, and (3) experimental errors (with variance σ_e^2), based upon independent randomizations in each block, contrary to the common practice at the time. This is illustrated in Table 1.2 for the *randomized complete block design* (RCBD) with t treatments in b blocks, such that each treatment occurs exactly once in each block, showing the three sources of variation and indicating how differences among the treatment effects ($\tau_j, j = 1, 2, \dots, t$, with $\Sigma\tau_j = 0$) can be tested. This ANOVA table is based on the linear model (see Chapter 2)

$$\text{observation} = \text{mean} + \text{block effect} + \text{treatment effect} + \text{error},$$

or, more formally,

$$y_{ij} = \mu + \beta_i + \tau_j + \epsilon_{ij}.$$

These three principles—*replication, randomization, and blocking*—represent the foundation of all of experimental design and upon which all subsequent developments are based, not just in agronomy but in all fields of experimentation including engineering, medicine, and marketing. These principles ensure the validity of estimating experimental error, provide a method of reducing experimental error, and establish a procedure for testing

TABLE 1.2
ANOVA for RCBD

Source	Degrees of Freedom (d.f.)	Sums of Squares (SS)	Mean Squares (MS)	Expected Mean Squares (E(MS))
Blocks (B)	$b - 1$	SS(B)	MS(B)	—
Treatments (T)	$t - 1$	SS(T)	MS(T)	$\sigma_e^2 + b\Sigma\tau_j^2/(t - 1)$
Error (E)	$(b - 1)(t - 1)$	SS(E)	MS(E)	σ_e^2
Total	$bt - 1$	SS(Total)		

hypotheses about treatment effects. And they lay the foundation for the development of other experimental designs.

1.3.4 Analysis of Covariance

In addition to blocking, “there is a second means by which precision may, in appropriate cases, be much increased by elimination of causes of variation which cannot be controlled” (Fisher 1935, 1971, p. 163). This method requires additional measurements, referred to as *concomitant variables* or *covariates*, on the experimental units, such as initial weight of the animals in a feeding trial, number of plants on an agronomic plot, and yield of orchard trees in the previous year. These measurements will then be used to *correct* or *adjust* the effects of the treatment. “The process of calculating the average apparent effect on the experimental value of the increase of one unit of the concomitant measurement is, in principle, extremely simple” (Fisher 1971, p. 171). Already in 1935, Fisher proposed to use regression methods to perform this adjustment, and he introduced the *analysis of covariance* as the appropriate procedure, showing that the extra variation accounts for one degree of freedom (d.f.) that is separated from the experimental error. If the linear relationship between covariates and final measurement is strong, this will lead then to a substantial reduction in the error variance.

1.3.5 Systematic and Trend-Free Designs

Although randomization is the fundamental principle in experimental design, there may be situations where other ways of arranging the treatments may be more appropriate (see Section 1.3.1). Cox (1951), for example, described a situation where several lots of wool are being treated by different treatments sequentially over several weeks. Because the wool deteriorates over time, there will be, superimposed on any treatment differences, a smooth trend due to aging. He then asked the question: “In what order should the treatments be applied?”

One way might be to divide the entire time period of the experiment into several blocks and use a RCBD. This would eliminate part of the trend. Another recommendation is to use the time points as covariates and analyze the data via analysis of covariance for a completely randomized design (CRD). Yet another method might be to repeat the treatments in the same order over the duration of the experiment and analyze the data as a CRD. The method proposed by Cox (1951) assumes a polynomial trend and then assigns the treatments so that their estimated effects are independent, or nearly independent of the trend. We call this an assignment of treatments orthogonal or nearly orthogonal to trend. For example, for $t = 2$ treatments, eight time periods, and a linear trend, the arrangement of the treatments might look as follows, where the first row represents the coefficients for the linear trend with equally spaced plots (time periods):

-7	-5	-3	-1	1	3	5	7
T_2	T_2	T_1	T_1	T_1	T_1	T_2	T_2

This shows that the coefficients for each treatment add up to zero, indicating the required orthogonality. The preceding linear trend coefficients represent the coefficients of the first-order orthogonal polynomial (see, e.g., Fisher and Yates 1938, 1957, or, e.g., Beyer 1991,

TABLE 1.3

Trend-Free Block Design

Block	Treatments						
1	1	2	3	4	5	6	7
2	6	4	2	7	5	3	1
3	7	5	3	1	6	4	2
P_1	-3	-2	-1	0	1	2	3

Draper and Smith 1998). In this connection, we should mention that Fisher (1921) introduced the use of orthogonal polynomials in statistics to investigate the trends of the effects from quantitative treatments.

The ideas exemplified above have been extended, leading to what are now called *trend-free designs*. Bradley and Yeh (1980) and Yeh and Bradley (1983) consider trend-free block designs and give necessary and sufficient conditions for their existence. An example is given in Table 1.3 for $t = 7$ in $b = 3$ blocks assuming a linear trend (P_1).

For further developments, see for example, Yeh et al. (1985); Jacroux et al. (1995); Jacroux (1998); and Lin and Stufken (1999).

1.4 Latin Square Principle

1.4.1 Latin Square Design

The fact that soil fertility was quite variable over an experimental field or, more generally, that experimental material often exhibited great variability became a major concern when designing experiments. Fisher (1926, p. 510) argued that, for variety trials and simple manurial trials, “the problem of designing economic and effective field experiments reduces to two main principles: (1) the division of the experimental area into plots as small as possible subject to the type of farm machinery used, and to adequate precautions against edge effect; (2) the use of arrangements [of the varieties or treatments] which eliminate a maximum fraction of the soil heterogeneity, and yet provide a valid estimate of the residual errors. Of these arrangements, by far the most efficient, as judged by experiments upon uniformity trial data, is that which the writer has named the Latin Square.” After the idea of blocking in one direction, leading to the randomized block design, Fisher thus introduced the idea of blocking in two directions, which led to various forms of *Latin square designs* (see Chapters 3 and 10).

Initially, systematic arrangements in a square were used, in particular in Denmark and Norway since 1872, in which the number of rows and columns were equal to the number of varieties. For example, for five varieties, one such arrangement is given in Table 1.4a, where the rows and columns represent the two directions of variability. This arrangement would, if possible, be repeated several times. This particular arrangement is, however, such that the distribution of the varieties (treatments) over the field is along diagonals. This may be a drawback in that such an arrangement may mimic a possible fertility trend. To counteract this, another systematic arrangement, known as the *Knut Vik square*, was introduced, which

TABLE 1.4
Systematic Latin Squares

(a)	A	B	C	D	E	(b)	A	B	C	D	E
	E	A	B	C	D		D	E	A	B	C
	D	E	A	B	C		B	C	D	E	A
	C	D	E	A	B		E	A	B	C	D
	B	C	D	E	A		C	D	E	A	B

for five varieties is given in Table 1.4b. It is obtained by moving each row forward two places instead of one. This provides an even distribution of the varieties over the field, but Fisher (1935) showed that this arrangement does not lead to valid estimation of error and called it a *biased design* (see also Kempthorne 1952, p. 330).

In order to eliminate these problems, Fisher insisted that the term *Latin square design* should be “applied only to a process of randomization by which one [square] is selected at random out of the total number of Latin squares possible” (Fisher 1926, p. 510) (see also Yates 1933b). For this purpose, he had enumerated all 4×4 , 5×5 , and 6×6 squares, as these were the most important ones from a practical point of view. He wrote that the Statistical Laboratory at Rothamsted would be able to “supply these, or other types of random arrangements, to intending experimenters.” In their *Statistical Tables for Biological, Agricultural and Medical Research*, Fisher and Yates (1938) then describe more practical randomizations for Latin square designs of various sizes.

Fisher (1935) showed how the analysis of variance, which he had introduced earlier (Fisher 1918, 1926), can be applied to estimate the experimental error for making comparisons between varieties or, more generally, between treatments (we shall use this terminology henceforth to reflect today’s much wider use of experimental design). The subdivision of the $s^2 - 1$ d.f. among the observations from an $s \times s$ Latin square design is given in Table 1.5.

It shows that the $2(s - 1)$ d.f. ascribable to differences between rows and columns and representing the components of heterogeneity of soil fertility, which have been eliminated in the field, are also eliminated from the estimate of error. Based on randomization theory, this error is then used to test hypotheses about treatments (see also Hinkelmann and Kempthorne, 2008). This method represents a *correction* of the faulty analysis of the systematic squares used earlier (see Fisher 1935).

TABLE 1.5
Subdivision of d.f. for $s \times s$ Latin Square

Rows	$s - 1$
Columns	$s - 1$
Treatments	$s - 1$
Error	$(s - 1)(s - 2)$
Total	$s^2 - 1$

1.4.2 Change-Over Designs

We have seen in previous sections how, in agronomic trials, the problem of the variability of soil fertility was dealt with through the introduction of block designs and Latin square designs. Using animals, for example in feeding trials, presented the researchers with similar problems as the variability among animals is usually quite large. Thus, rather than using each animal as the experimental unit, each animal was used as a block, and the individual treatments were randomly assigned over several time periods. These designs were called *switchover designs* (later also *change-over designs* or *crossover designs*; see Chapter 4). In the case of dairy cows, however, this practice would lead to other problems because of a drop in milk yield toward the end of the lactation period. These considerations have led to the use of Latin square designs with rows as periods and columns as cows. For three treatments, that is, different rations, Cochran et al. (1941) at the Iowa Agricultural Experiment Station used the Latin square given in Table 1.6a for several, say s , groups of three cows using the same or different time periods.

A possible partitioning of the d.f. in the analysis of variance (ANOVA) is given in Table 1.7. The effects included in the ANOVA are dictated by the underlying design, except for the groups \times rations effect, which is optional depending on the specific experimental situation encountered.

Such an arrangement would be sufficient if there are long enough rest periods between successive assignments of different rations so that there are no *carry-over effects* from a treatment in one period to that in the next. For shorter change-over periods, however, one may anticipate a carry-over effect of the ration given in the previous period for certain types of

TABLE 1.6
Latin Square SwitchOver Designs

(a)	Cow			(b)	Cow			
	Period	1	2		3	Period	1	2
	1	A	B	C	1	B	A	C
	2	B	C	A	2	A	C	B
	3	C	A	B	3	C	B	A

TABLE 1.7
ANOVA for SwitchOver Design

Source	d.f.	
	Same Periods	Different Periods
Groups	$s - 1$	$s - 1$
Cows within groups	$2s$	$2s$
Periods	2	$2s$
Rations	2	2
Groups \times rations	$2(s - 1)$	$2(s - 1)$
Error	$4s - 2$	$2s$
Total	$9s - 1$	$9s - 1$

The designs for t even are also known as *sequentially counterbalanced Latin squares* in the psychology literature and as *column complete squares* in the mathematics literature. For a discussion and overview of extensions of the class of Williams squares, see Isaac et al. (2001).

An important feature of the designs mentioned earlier is that the number of treatments, t , is the same as the number of periods (rows), p say. In some situations, it may, however, be desirable, for practical and/or ethical reasons, to use designs where $p < t$. A simple method to obtain such designs is to delete one or more periods from the full design described earlier (Patterson 1950). Other methods are discussed by Jones and Kenward (2003), using, for example, Youden square designs (see Section 1.6.4) or balanced incomplete block (BIB) designs (see Section 1.6.1).

In other situations, it may be possible and desirable to use designs where $p > t$ (for a discussion, see Jones and Kenward 2003). One important case is where $p = t + 1$. Such designs are used mainly for statistical reasons in order to remove the nonorthogonality between direct and residual effects from the balanced designs with $p = t$. This is particularly important if the residual effects themselves are of interest to the researcher. The nonorthogonality arises because each treatment is followed by every other treatment the same number of times, except by itself. More efficient designs can be constructed by simply adding as an extra period the last period of the original design as suggested by Lucas (1957) at North Carolina Agricultural Experiment Station. An alternative procedure is to have a preperiod the same as the first period of the original design. These designs were later called *strongly balanced designs for residual effects* (e.g., Hinkelmann and Kempthorne 2005). Omitting the observations from the preperiod in the analysis will lead to orthogonal designs.

For a discussion of other designs with $p > t$, see Jones and Kenward (2003).

Further variations of change-over designs with carry-over effects from the preceding period are designs that take into account second-order residual effects, that is, carry-over effects from two preceding periods. Such designs were considered by Williams (1950), who proposed a method of construction based on *Græco-Latin squares* (see Section 1.4.3).

1.4.3 Orthogonal Latin Squares

“Two Latin squares may be said to be orthogonal to each other if, when they are superimposed, every letter of one square occurs once with every letter of the other” (Bose 1938, p. 323). If the letters of one Latin square are Latin letters and the letters of the other Latin square are Greek letters, then superimposing these two Latin squares leads to what has been called a Græco-Latin square (see, e.g., Fisher and Yates 1938). *Græco-Latin square designs* can be used to eliminate variability in three directions, using rows, columns, and Greek letters as the levels of three blocking factors and Latin letters as the treatments.

For any Latin square of size s (except for $s = 2$ and 6), it is possible to construct a set of Latin squares, such that any two of them are mutually orthogonal (Bose et al. 1960). Fisher conjectured that if s is a prime or a prime power, then there exist $s - 1$ mutually orthogonal Latin squares. This conjecture was proved independently by Bose (1938) and Stevens (1939). They also provided methods of constructing such sets, which Bose referred to as *hyper-Græco-Latin squares* and Stevens as *completely orthogonalized Latin squares*. For s not a prime or prime power, the number of possible mutually orthogonal Latin squares depends on s and is at most $s - 1$.

TABLE 1.9
4×4 Hyper-Græco-Latin Square

	P_1	P_2	P_3	P_4
C_1	$A\alpha 1$	$B\beta 2$	$C\gamma 3$	$D\delta 4$
C_2	$C\beta 4$	$D\alpha 3$	$A\delta 2$	$B\gamma 1$
C_3	$D\gamma 2$	$C\delta 1$	$B\alpha 4$	$A\beta 3$
C_4	$B\delta 3$	$A\gamma 4$	$D\beta 1$	$C\alpha 2$

An example of a hyper-Græco-Latin square design for $s = 4$ was described by Box et al. (2005) for an experiment with a Martindale wear tester, which is used for testing the wearing quality of types of cloth. Four pieces of cloth can be tested simultaneously in one machine cycle. The response is the weight loss of the test piece after being rubbed against emery paper for 1000 revolutions of the machine. Specimens of four different types of cloth, representing the treatments (A, B, C, D), are mounted in four holders (1, 2, 3, 4). Each holder can be in any one of four different positions (P_1, P_2, P_3, P_4) on the machine. Each emery paper sheet ($\alpha, \beta, \gamma, \delta$) was cut into four quarters and each quarter was used to complete a single cycle (C_1, C_2, C_3, C_4) of 1000 revolutions. With holder, position, emery paper, and cycle as the blocking factors and type of cloth as the treatment factor, the layout of the experimental design (apart from randomization) is given in Table 1.9.

Since each factor accounts for three d.f., this design (with 16 observations) leaves no d.f. for error in the ANOVA table. Therefore, in order to make this experiment work, from a statistical point of view, one needs to replicate the design given in Table 1.9. This is discussed in detail by Box et al. (2005).

1.5 Factorial Experiments

“No aphorism is more frequently repeated in connection with field trials, than that we must ask Nature few questions, or, ideally, one question at a time. The writer is convinced that this view is wholly mistaken. Nature, he suggests, will best respond to a logical and carefully thought out questionnaire; indeed, if we ask her a single question, she will often refuse to answer until some other topic has been discussed” (Fisher 1926, p. 511).

1.5.1 Rationale: Efficiency and Comprehensiveness

With the preceding words, Fisher provides his rationale for carrying out *factorial experiments* rather than the usual *one-at-a-time experiments* (or single-question method) common at that time (see also Daniel 1973). He illustrates this with an agricultural field trial with winter oats (Eden and Fisher 1927) using three treatment factors: type of nitrogenous manure (M, S), amount of manure (0, 1, 2), and time of application (early, late). On the surface, this is a $2 \times 3 \times 2$ factorial, but all plots receiving the amount 0 are indistinguishable and act essentially as a control, C . This is therefore an experiment with 9 (instead of 12) different treatments, consisting of a $2 \times 2 \times 2$ factorial plus control (Kempthorne 1952 refers to such experiments as *partially factorial experiments*). The experiment was performed in 8 blocks of size 12 as shown in Table 1.10.

TABLE 1.10

Complex Experiment with Winter Oats

C	2 M Early	2 S Late	C	2 S Late	C	C	1 S Early
1 S Early	1 M Early	1 M Late	1 S Late	2 M Early	2 M Late	1 M Early	1 M Late
C	2 M Late	C	2 S Early	C	1 S Late	C	2 S Early
2 S Early	2 M Early	C	1 M Late	C	2 S Early	2 S Late	2 M Late
C	1 S Late	1 S Early	1 M Early	1 M Late	C	C	1 S Late
2 M Late	C	2 S Late	C	2 M Early	C	1 M Early	1 S Early
2 S Early	2 M Late	1 S Early	2 M Early	2 S Late	2 S Early	2 M Early	C
C	C	1 M Late	C	1 M Early	2 M Late	C	1 M Late
2 S Late	1 M Early	C	1 S Late	C	C	1 S Early	1 S Late
2 M Early	1 M Early	2 M Late	2 S Late	1 S Early	C	C	1 S Late
1 S Late	C	C	1 M Late	1 M Early	2 S Early	2 M Late	C
1 S Early	C	2 S Early	C	C	2 M Early	2 S Late	1 M Late

Source: Adapted from Fisher, R.A., *J. Ministry Agric.*, 33, 512, 1926.

Fisher used this experiment to point out that “any general difference between *M* and *S*, between early and late application, or ascribable to quantity of nitrogenous manure, can be based on 32 comparisons, each of which is affected only by such soil heterogeneity as exists between plots in the same block. To make these three sets of comparisons only, with the same accuracy, by single question methods, would require 224 plots, against our 96.”

The important point here is that each observation from the factorial experiment is used in several comparisons, whereas with the single-question method, one would need 32 plots for each level of each factor to achieve the same accuracy for each comparison (assuming the error variances are the same for both designs, say σ_e^2). For example, to assess the effect of *M* versus *S*, we consider in each block comparisons of the observations for the treatments (x, M, y) and (x, S, y) , where (x, y) represents each combination of the different forms (referred to also as *levels*) of the other two factors, amount of manure and time of application, respectively. Thus, we have four comparisons in each of the eight blocks. The average of these 32 comparisons represents an estimate of the differential effect of *M* versus *S*. The variance of this estimate is $2\sigma_e^2/32$ (for a more detailed analysis of the results of this experiment, see Eden and Fisher 1927). For the single-question method, we estimate the differential effect of *M* versus *S* by comparing the means of the observations from two treatments differing only with respect to the factor type of nitrogen, for example, (1, *M*, early) versus (1, *S*, early). Similarly, if we add the treatment (2, *S*, early), we can estimate the differential effect of the quantity 1 versus 2. And using the treatment (2, *S*, late), it is possible to estimate the effect of early versus late. In order for these estimates to have variance $\sigma_e^2/16$, we need 32 replications for each of the four treatments, whereas in the experimental setup of Table 1.10, each of the 8 factorial treatments is replicated only 8 times. Finally, inclusion of 32 replications (plots) for the control *C* leads to 160 plots (rather than the 224 plots mentioned but not explained by Fisher) as a minimal design using the single-question method, as compared to 96 plots for the factorial method.

Fisher (1935) refers to this as factorial experiments having greater *efficiency*: comparisons are made with the same precision with a fraction of the observations that would otherwise be necessary. Also, in addition to comparisons among individual factors, factorial experiments can investigate possible *interactions* among two or more factors. He refers to this as greater *comprehensiveness*, because this information cannot be obtained by the conventional

method of changing one factor at a time or only at considerable expense over time and under ideal conditions, which usually are not available for agricultural field experiments. Even for a variety trial, as Fisher points out, it will be advantageous and even necessary to include what he calls *subsidiary factors*, such as seed rate, drilling width, and amount of seed. The inclusion of additional factors provides for wider *inductive inference*.

Frank Yates had joined Fisher in 1931 at Rothamsted Experimental Station and became the head of statistics at Rothamsted Experimental Station in 1933 when Fisher was appointed to a chair at University College London. Jointly, they put forth powerful and what appeared to be convincing arguments for the use of factorial experiments (Fisher 1935; Yates 1935, 1937), not only in agriculture but in other areas of application as well (see Section 1.8.1), but they encountered strong resistance and criticism. Part of this criticism is best documented in the discussion of Yates' paper delivered before a session of the Royal Statistical Society (Yates 1935). Neyman (in Yates 1935, p. 235) wrote: "...I am inclined to think that before trusting it [complex experimentation] so entirely as Mr. Yates and many of the other speakers have done, it is desirable to produce some further evidence as to its validity." He criticized the use of *main effects* (represented by the comparison of the effects of two different levels of one factor averaged over the level combinations of all the other factors) as compared to *simple effects* (represented by the comparison of the effects of two levels of one factor for particular level combinations of the other factors, as used in the preceding single-question method) and the difficulty of detecting main effects and interactions because of insufficient number of replications. He thereby reiterated an earlier criticism of Wishart (1934), summarizing his remarks (p. 241) as follows: "...if the experimenter is inclined to believe in the absence of troublesome interactions and thus in the soundness of the method of complex experiments, I think it will be useful for him to realize that the method is based on belief. Further, he should recognize what curious answers may be given to the questionnaire if the number of replications is small and Nature chooses to be frivolous and not behave as the experimenter expects her to do." Another criticism concerned the use of large blocks because of the large number of treatments, which might lead to an increase in error. Before refuting the criticism by pointing to the higher efficiency of factorial experiments and reliance on a series of experiments rather than a single experiment as well as on systems of *confounding* (see Section 1.5.2), Yates made the following general remark (p. 243): "A general survey of the remarks immediately brings out one fact of considerable interest. Those speakers who are actually engaged in experimental work or are in close contact with experimental workers, appear completely satisfied with factorial design, while those who are not so engaged have raised several objections to the method. Now if the method of factorial design were as fundamentally unsound, misleading, and unreliable as the critics would have us believe, one would have expected some at least of those who had most used it to have discovered the fact, both from bitter experience and because they had at one time and another devoted a considerable amount of thought to it. But this is not the case. I am encouraged, therefore, to believe that the force of these criticisms is not perhaps as great as their volume would indicate." In the end, factorial designs and factorial experiments proved, of course, to be great achievements in the development of experimental design (see Chapters 7 through 9).

1.5.2 Confounding

Heterogeneity in soil fertility or in experimental material in general requires the experimenter to try to reduce experimental error as much as possible by using appropriate

methods of designing the experiment. Since it is an empirical fact that the plots in smaller blocks are more homogeneous than plots in larger blocks, it would be advantageous, in particular in factorial experiments with a large number of treatment combinations, to make use of smaller blocks wherever possible. The question how to do this was answered by Fisher (1935) with the introduction of the notion of *confounding*. The idea behind this notion is that very often higher-order interactions are nonexistent or if they exist, “there would be no immediate prospect of the fact being utilized” (Fisher 1971, p. 111). In such a case, blocks within a replicate would be created such that the contrasts between blocks shall be contrasts between such unimportant or nonexistent interactions. Fisher’s representation of the main effects and interactions in a simple algebraic form makes this process very simple.

To illustrate this point, we consider the 2^3 factorial, that is, three factors (A, B, C) at two levels (0, 1) each, with blocks of size 4. Following the notation introduced by Fisher and Yates (see, e.g., Yates 1937), the eight treatment combinations are denoted by (1), $a, b, ab, c, ac, bc,$ and abc , where the correspondence between treatment combinations and factor-level combinations is given here:

Treatment									
Combination		(1)	a	b	ab	c	ac	bc	abc
Factor level	A	0	1	0	1	0	1	0	1
	B	0	0	1	1	0	0	1	1
	C	0	0	0	0	1	1	1	1

Using the notation of Hinkelmann and Kempthorne (2005), we can write all the main effects and interactions symbolically (Fisher 1935) as follows:

$$A^\alpha B^\beta C^\gamma = \frac{[a + (-1)^\alpha][b + (-1)^\beta][c + (-1)^\gamma]}{4},$$

where α, β, γ equal 1 if the corresponding factor is included in the main effect/interaction or 0 otherwise. If we assume that the three-factor interaction ABC (i.e., $\alpha, \beta, \gamma = 1$), given by

$$4ABC = (a - 1)(b - 1)(c - 1) = abc - ab - ac + a - bc + b + c - 1,$$

is negligible, then the treatment combinations with a plus sign are assigned to units in one block and the treatment combinations with a minus sign are assigned to units in the other block; that is, we have

Block 1 : $a, b, c, abc,$

Block 2 : $ab, ac, bc, (1).$

Denoting the treatment combinations by $(x_1x_2x_3)$ with $x_i = 0, 1$ representing the *low* and *high* levels of factor $i(i = 1, 2, 3)$, the aforementioned blocks 1 and 2 can be expressed as

Block 1: 100, 010, 001, 111,

Block 2: 110, 101, 011, 000.

This arrangement, repeated in all replicates, allows the estimation of all main effects and interactions other than ABC with equal variance. We emphasize that the treatments need to be randomized independently in each block of each replicate.

This procedure is easily generalized to the general 2^f factorial with blocks of size 2^q or to the 3^f factorial and blocks of size 3^q for any $q < f$, and Fisher provided examples of that (Fisher 1935). Using group theoretical arguments, Fisher (1942, 1945) also provided a simple method of constructing systems of confounding. He introduced the notion of the *intra-block subgroup*, which represents the initial block from which all other blocks can be derived quite easily. He also introduced the notion of *partial confounding*, whereby, for example, one interaction is confounded in one replicate and another interaction is confounded in another replicate. To use the preceding example, suppose we were interested in at least some information about the ABC interaction, we could construct a design in which ABC is confounded in, say, replicate 1, and another interaction, say AB , is confounded in replicate 2. Then, in replicate 1, we would have the same blocks as given earlier, but in replicate 2, the block arrangement would be as follows:

Block 3: (1), ab , c , abc ,
 Block 4: a , b , ac , bc ,

or in alternative notation

Block 3: 000, 110, 001, 111,
 Block 4: 100, 010, 101, 011.

This design would allow estimation of all main effects and interactions, but the variance of the estimators for AB and ABC would be twice as large as the variances for the estimators for the remaining effects. This fact is also reflected in the ANOVA tables given in Table 1.11, where we consider the complete confounding system with $2r$ replications of blocks 1 and 2 and the partial confounding system with r replications of blocks 1, 2, 3, and 4, that is, $4r$ blocks altogether in each design. Comparison of the expected mean squares (E(MS)) shows again explicitly that the information on AB is less with partial confounding than with no

TABLE 1.11
 ANOVA Tables for Systems of Confounding

Source	Complete Confounding		Partial Confounding	
	d.f.	E(MS)	d.f.	E(MS)
Replication	$2r - 1$		$2r - 1$	
Blocks within reps	$2r$		$2r$	
A	1	$\sigma_e^2 + 4r[A]^2$	1	$\sigma_e^2 + 4r[A]^2$
B	1	$\sigma_e^2 + 4r[B]^2$	1	$\sigma_e^2 + 4r[B]^2$
AB	1	$\sigma_e^2 + 4r[AB]^2$	1	$\sigma_e^2 + 2r[AB]^2$
C	1	$\sigma_e^2 + 4r[C]^2$	1	$\sigma_e^2 + 4r[C]^2$
AC	1	$\sigma_e^2 + 4r[AC]^2$	1	$\sigma_e^2 + 4r[AC]^2$
BC	1	$\sigma_e^2 + 4r[BC]^2$	1	$\sigma_e^2 + 4r[BC]^2$
ABC	0	—	1	$\sigma_e^2 + 2r[ABC]^2$
Error	$6(2r - 1)$	σ_e^2	$5(2r - 1) + 2(r - 1)$	σ_e^2

confounding of AB , that is, $2r[AB]^2$ versus $4r[AB]^2$, where $[AB]^2$ (and similar expressions) refers to the squared value of the effect AB as defined earlier.

1.5.3 Some General Results

The preceding example illustrates the basic principle of partial confounding and the method of constructing such a *system of confounding*. This idea was then generalized to include situations where more than two interactions were to be confounded with blocks, either completely or partially. Fisher (1942) showed how combinatorics and group theory can be used to construct appropriate designs and how to use the intrablock subgroup to easily construct the other blocks for a given system of confounding. One important result is that if two interactions, say $ABCD$ and $CDEF$, are confounded with blocks, then also what we call now the *generalized interaction* $ABEF$ [= $(ABCD) \times (CDEF) = ABC^2D^2EF$ and deleting any letter raised to the power 2] (see also Section 1.7.1) is confounded with blocks. This imposes certain restrictions on the choice of appropriate interactions that can be confounded with blocks if one wants to avoid confounding main effects and two-factor interactions. Fisher showed that the block size determines how many factors can be accommodated without having to confound main effects and two-factor interactions. For a generalization of this result, see Hinkelmann and Kempthorne (2005, Chapter 11).

The mathematical theory of systems of confounding for symmetrical factorials, with each factor having p^s levels, with p being a prime number, was given by Bose (1947). His discussion is based essentially on geometrical arguments by identifying the treatment combinations as points in a Euclidean geometry and the main effects and interactions as pencils in lower-dimensional flats. A more algebraic presentation was provided by Kempthorne (1947, 1952), who introduced an overparameterized model for the true (or expected) yield of a treatment combination that is useful not only for constructing and evaluating systems of confounding but also for the analysis of data from the actual experiment. His method formalizes and extends the representation given by Yates (1937) (see also Cochran and Cox 1950, 1957), who referred, for example, in the 3^2 experiment to I - and J -components as part of the two-factor interaction, which Kempthorne calls AB and AB^2 , respectively.

An extension of the principle of confounding is to make use of the variance-reducing property of the Latin square and use systems of *double confounding*, whereby some effects are confounded with rows and some with columns. In fact, this method is not confined to arrays in a square but can be used also more generally in a rectangle (Fisher 1935; Yates 1937). As an example, consider the arrangement in Table 1.12 of a 2^5 factorial in a 4×8 rectangle, where ABC , CDE , and $ABDE$ are confounded with rows and AB , CD , $ABCD$, BDE , ADE , BCE , and ACE are confounded with columns.

The first row and the first column represent the intrablock subgroups, respectively, for the two aforementioned confounding systems.

TABLE 1.12

Double Confounding

	1	2	3	4	5	6	7	8
1 (1)	<i>ab</i>	<i>ace</i>	<i>bce</i>	<i>abde</i>	<i>de</i>	<i>bcd</i>	<i>acd</i>	
2	<i>abe</i>	<i>e</i>	<i>bc</i>	<i>ac</i>	<i>d</i>	<i>abd</i>	<i>acde</i>	<i>bcde</i>
3	<i>cde</i>	<i>abcde</i>	<i>ad</i>	<i>bd</i>	<i>abc</i>	<i>c</i>	<i>be</i>	<i>ae</i>
4	<i>abcd</i>	<i>cd</i>	<i>bde</i>	<i>ade</i>	<i>ce</i>	<i>abce</i>	<i>a</i>	<i>b</i>

1.5.4 Asymmetrical Factorials

Factorial designs considered in previous sections are referred to as *symmetrical factorials*, which means that each factor has the same number of levels. Among these, the factorials with 2 and 3 levels, that is, the 2^f and 3^f factorials, are of particular practical importance, but the theory was extended to p^f and s^f factorials, where p is a prime number and s is the power of a prime number (for details, see Hinkelmann and Kempthorne 2005). Practical situations often require the use of designs where different factors have different numbers of levels. These are referred to as *mixed* or *asymmetrical factorials*. Of particular interest are $2^f \times 3^g$ or $2^f \times 3^g \times 4^h$ factorials. To construct systems of confounding for such factorials is in general a much more complicated problem than that for symmetrical factorials (see Hinkelmann and Kempthorne 2005). Yates (1937) and Li (1944) provided extensive tables of what are now called *balanced factorial designs* (Shah 1958). Their feature is that for some main effects or interactions, part of the information is confounded with blocks over the entire design. This is different from the type of partial confounding discussed earlier, where in any given replicate, one (or more) interaction(s) is (are) completely confounded. As an example, consider the design in Table 1.13 for a $2 \times 2 \times 4$ factorial in blocks of size 4 (Li 1944). Here, the levels of the factors are denoted by 0, 1, 2, and 3.

The result of the arrangement of the treatment combinations for factors A , B , and C as given in Table 1.13 is that the interactions AC , BC , ABC are partially confounded with blocks. More specifically, these interactions are estimated with efficiency $2/3$, that is, $1/3$ of the information is lost due to confounding, whereas all other effects are estimated with full efficiency or efficiency 1, that is, there is no loss of information. Another interesting result is that this type of design represents a close connection between balanced factorial designs and *partially balanced incomplete block (PBIB) designs* (see Section 1.6). More specifically, using results established by Dean and John (1975) and Kshirsagar (1966), this is an EGD-PBIB(7) design given by Hinkelmann (1964) (for more details, see Hinkelmann and Kempthorne 2005). One method of constructing some of the balanced factorial designs is through the use of *pseudo-factors* (see Hinkelmann and Kempthorne 2005). For example, the four-level factor C in Table 1.13 can be represented by a 2^2 factorial with factors C_1 and C_2 say, and the design can be obtained by confounding AC_1 , BC_2 , and ABC_1C_2 in replication 1; AC_2 , BC_1C_2 , and ABC_1 in replication 2; and AC_1C_2 , BC_1 , and ABC_2 in replication 3.

These results give an indication of the versatility and usefulness of factorial designs not only in agricultural field trials but in other areas of scientific and industrial experimentation as well. Other aspects and more details of factorial designs are given in Chapters 7 through 11. Also, for further discussion of factorial designs and their applications in various scientific fields and industry, the reader is referred to a large number of textbooks, among

TABLE 1.13
Balanced Factorial Design

Replication 1				Replication 2				Replication 3			
Block				Block				Block			
1	2	3	4	1	2	3	4	1	2	3	4
000	010	100	110	000	010	100	110	000	010	100	110
011	001	111	101	111	101	011	001	101	111	001	011
102	112	002	012	012	002	112	102	112	102	012	002
113	103	013	003	103	113	003	013	013	003	113	103

them are Box et al. (2005), Cox (1958), Dean and Voss (1999), Giesbrecht and Gumpertz (2004), Hinkelmann and Kempthorne (2005, 2008), Montgomery (2012), Ryan (2007), and Wu and Hamada (2009).

1.5.5 Split-Plot Experiments

A particular type of factorial experiment, dictated by practical circumstances, was described by Yates (1935). It involved two factors, say A with a levels and B with b levels, where factor A could only be applied to large plots and factor B could be applied to smaller plots. For example, factor A may represent different types of soil preparation, and factor B may represent different varieties. The experiment was arranged such that the field was divided into $a \times r$ large plots to which the levels of factor A were randomly assigned such that each level of factor A occurred r times. Then each large plot was subdivided into b subplots to which the levels of factor B were assigned randomly. For this reason, such an experiment is referred to as a *split-plot experiment* consisting of *whole plots* for factor A and *split plots* or *subplots* for factor B . The special feature of this design is that there are two independent randomizations and factor A is applied in a CRD, whereas factor B is applied in a RCBD, with the levels of factor A representing the *blocks*. This leads to two different types of tests: one for the main effect A and one for the main effect B and the interaction AB , using the whole-plot error and the split-plot error, respectively (see also Chapter 8).

Different experimental situations have led to different arrangements of this type. For example, both factors may only be applied to large plots. Then the (rectangular) field could be divided into large blocks in both directions, with the levels of factor A being applied in one direction and the levels of factor B being applied in the other direction. Such a design was referred to as a *split-plot design in strips* or, more recently, as a *split-block design* (see Section 1.2.2) or a *strip-block design* in industrial experimentation (Vivacqua and Bisgaard 2004).

Another extension is to subdivide the split plots further into *split-split plots* that can accommodate a third factor C , say. Thus, there are now three independent randomizations and three types of tests: for (1) A , (2) B and AB , and (3) C , AC , BC , and ABC . The design is referred to as a *split-split-plot design*.

There are many more variations of these types of designs. For a listing and classification of such designs, see Hinkelmann and Kempthorne (2008). For more on split-plot-type designs, see also Chapter 8. For new aspects of constructing appropriate split-plot-type designs for industrial applications, see Vining (2012).

1.6 Incomplete Block Designs

In agricultural field experiments, it is generally possible to use blocks large enough to accommodate all treatments in an RCBD or Latin square design. This may come, however, at the expense of accuracy of estimation, because the variability in soil fertility may be quite large, leading to large experimental error. But “in some other experimental material the groupings which most effectively eliminate heterogeneity are definitely limited in numbers” (Yates 1936a, p. 122). Yates mentioned as examples of such limitation the number

of pigs in a litter, monozygotic twins, or the number of leaves on a plant. When the treatments have a factorial structure, this disadvantage may be overcome by employing systems of confounding (see Section 1.5.2). This leads to the use of *incomplete blocks* and to a reduction in the variance of the estimates of treatment comparisons. But what should one do in situations with a limited number of experimental units in natural blocks and a large number of treatments that do not have a factorial structure?

1.6.1 Balanced Incomplete Block Designs

One answer to the question of how to modify the randomized block type of arrangement in this situation was given by Yates (1936a, p. 122), who recommended a particular type of arrangement that he characterized as follows: "In this modified type of arrangement the number of experimental units per block is fixed, being less than the number of treatments, and the treatments are so allotted to the blocks that every two treatments occur together in a block equally frequently." He called this type of arrangement a *symmetrical incomplete randomized block arrangement*, which is now referred to as a BIB design (see Chapter 3). As an example, consider the (nonrandomized) design, given by Yates (1936a), in Table 1.14, with $t = 6$ treatments in $b = 10$ blocks of size $k = 3$ and with every two treatments occurring together $\lambda = 2$ times in a block.

Yates provided a number of possible designs but emphasized that in many cases the designs would have to be quite large to satisfy the basic equal frequency property. Many more BIB designs were found after they were introduced by Yates (for some more complete listings, see, e.g., Fisher and Yates 1938; Bose 1939; Cochran and Cox 1957; Raghavarao 1971; and Hinkelmann and Kempthorne 2005). Many of these designs can be obtained by statistical packages (see Section 1.10.4). Of particular interest are BIB designs with blocks of size 2 that are simply obtained by taking all possible combinations of size 2. Yates pointed out that these designs are more efficient than the designs that had been used theretofore by using one treatment as the control and comparing each other treatment against this control. The BIB designs with blocks of size 2 have found wide application in biology, genetics, psychology, and medicine (where they often are referred to as *left-right designs* because of the application of treatments to left and right body parts).

Yates also argued that experimental arrangements are not likely to be of practical value unless they are easy to analyze, and he showed that the analysis of data from a BIB design is indeed quite simple, even though the BIB design is a *nonorthogonal design* in the sense that the treatment effects are not independently estimated from the block effects. Using the *method of least squares* (see Chapter 2) for the customary linear model for observations from a block design, namely,

$$\text{observation} = \text{mean} + \text{block effect} + \text{treatment effect} + \text{error},$$

TABLE 1.14

BIB Design

	Block									
	1	2	3	4	5	6	7	8	9	10
Treatments	1	1	1	1	1	2	2	2	3	3
	2	2	3	4	5	3	4	5	4	4
	3	4	5	6	6	6	5	6	5	6

Yates derived explicit expressions for estimating linear functions of treatment effects and for estimating the error variance from the ANOVA table. One important feature of the BIB design is that all simple treatment comparisons are estimated with the same variance. He further showed that the *relative efficiency* (RE) of the BIB design relative to the RCBD with the same number of replications is given by

$$\frac{(1 - 1/k)\sigma_{\text{RCBD}}^2}{(1 - 1/t)\sigma_{\text{BIBD}}^2},$$

where k denotes the block size, t denotes the number of treatments, and σ_{RCBD}^2 and σ_{BIBD}^2 represent the error variances for the RCBD and the BIB design, respectively. The ratio $(1 - 1/k)/(1 - 1/t) = (k - 1)t/(t - 1)k$ he called the *efficiency factor* of the BIB design, which, since $k < t$, is less than one (see also Sections 1.9.1 and 1.9.2).

The analysis given by Yates (1936a) is based on *intrablock comparisons* between treatments, but he argued (Yates 1940) that information about treatment effects can also be obtained from *interblock comparisons* if block effects are considered to be random effects. He referred to this as *recovery of interblock information*. Calling σ_e^2 the intrablock variance and $\sigma_e^2 + k\sigma_b^2$ the interblock variance, where σ_b^2 is a measure of the variability among blocks, he defined the weights $w = 1/\sigma_e^2$ and $w' = 1/(\sigma_e^2 + k\sigma_b^2)$ and used them to obtain the *combined estimator* of treatment differences as the weighted combination of the corresponding intrablock and interblock estimators. The actual estimates of treatment comparisons are obtained by using estimates of w and w' , or of $\rho = w/w'$. Yates showed how this can be accomplished by using different forms of the ANOVA table as indicated in Table 1.15, where n denotes the total number of observations, b the number of blocks, and t the number of treatments (Yates and other writers later changed t to v to indicate varieties).

In Table 1.15, the sums of squares SS(B), SS(T/B) (i.e., the *sequential sums of squares* for the preceding observation model), and SS(Total) in the left part are obtained directly, and SS(E) is then obtained by subtraction. In the right part, SS(T) is obtained directly, SS(Total) and SS(E) are the same as on the left, and SS(B/T) is then obtained by subtraction, where SS(T) and SS(B/T) are the sequential sums of squares for the ordered model

$$\text{observation} = \text{mean} + \text{treatment effect} + \text{block effect} + \text{error}.$$

The important features of these ANOVA tables are that (1) the mean squares MS(E) and MS(B/T) can be used to estimate σ_e^2 and σ_b^2 , since

$$E[\text{MS}(E)] = \sigma_e^2 \text{ and } E[\text{MS}(B/T)] = \sigma_e^2 + [(n - t)/(b - 1)]\sigma_b^2,$$

TABLE 1.15
ANOVA Tables for Incomplete Block Designs

Source	d.f.	Sums of Squares		Source
Blocks ignoring treatments	$b - 1$	SS(B)	SS(B/T)	Blocks eliminating treatments
Treatments eliminating blocks	$t - 1$	SS(T/B)	SS(T)	Treatments ignoring blocks
Error	$n - b - t + 1$	SS(E)	SS(E)	Error
Total	$n - 1$	SS(Total)	SS(Total)	Total

and (2) MS(T/B)/MS(E) provides a test for the equality of treatment effects using intrablock information only. The estimation procedure for σ_b^2 is referred to as the *Yates procedure* and the estimator of σ_b^2 as the *Yates estimator* (see Section 1.10.1 for other approaches).

Yates (1940, p. 325) concluded with some advice concerning the usefulness of BIB designs: "In agricultural experiments the gain from the use of inter-block information will not in general be so great as in similar quasifactorial (lattice) designs [see Section 1.6.2], since complete replications cannot (except in special cases) be arranged in compact groups of blocks. For this reason also, cases will arise in which the use of ordinary randomized blocks will be more efficient than the use of incomplete blocks, whereas lattice designs can never be less efficient than ordinary randomized blocks. Their [BIB designs] greatest value is likely to be found in dealing with experimental material in which the block size is definitely determined by the nature of the material. A further use is in co-operative experiments in which each centre can only undertake a limited number of treatments."

The preceding discussion should give the reader some idea how the necessity for IBDs has evolved from practical considerations. Some further extensions are given in Sections 1.6.2 through 1.6.4. Beyond that, the notion of IBD has received considerable attention over the years, and more modern results are being presented in Chapter 3.

1.6.2 Pseudo-Factorial Designs

According to Yates (1936b, p. 424), "a simple method [of dealing with the problem of comparing a large number of varieties and] of keeping the block size small is to select one or more varieties as controls and to divide the rest into sets, each set being arranged with the controls in a number of randomized blocks. Unfortunately this method has the disadvantage that comparisons between varieties in different sets are of lower accuracy than comparisons between varieties in the same set." To alleviate this problem partly and to avoid the waste of experimental material by the controls, Yates (1936b) introduced what he called *pseudo-factorial* or *quasi-factorial designs*.

For the simplest quasi-factorial design with $t = p^2$ treatments (varieties), the treatments are labeled by a pair of symbols (xy) with $x, y = 0, 1, 2, \dots, p - 1 (=q)$ and arranged in a square or lattice (hence also the name *lattice design* even though this arrangement is not the design itself) as follows:

00	01	02	0q
10	11	12	1q
20	21	22	2q
			.	
			.	
			.	
q0	q1	q2	qq

Blocks of size p are then formed by first taking each row as a block (the *X* arrangement) and then taking the columns as blocks (the *Y* arrangement), thus creating two replications with p blocks each. In terms of the *pseudo-factors* X and Y , this means that the X main effects are confounded with blocks in one replicate and the Y main effects are confounded with blocks in the other replicate, but all the $(p - 1)^2$ d.f. for interactions are unconfounded.

This is, therefore, a system of partial confounding. We are, however, not interested in main effects and interactions but rather in differences between treatments (varieties). Now, all contrasts among treatments are estimable, but with different precisions. For this particular setup, Yates showed that there are two types of treatment comparisons, one type being estimated with full precision and the other type being estimated with half precision, that is, twice the variance of the former.

Still other arrangements are possible. Consider, for example, the case $p = 5$: we could superimpose on the 5×5 lattice a Latin square of size 5, and generate an additional replicate with blocks formed by having all the treatments with the same Latin letter in the same block. Since 5 is a prime number, we know that there exist 4 pairwise orthogonal Latin squares, and each can be used as described earlier. This, together with the X and Y arrangements, will lead to 6 replicates, each with 5 blocks of size 5. As a consequence of this arrangement, each simple treatment comparison is estimated with the same precision. For this reason, this design is called a *balanced lattice design*, which is a special case of a BIB design. If p is not a prime or prime power, then the number of possible replicates is restricted by the maximum number of existing mutually orthogonal Latin squares for p (see Section 1.4.3), say q , leading to at most $q + 1$ replicates.

Yates (1936b) proposed a wide variety of lattice designs (for a concise tabulation, see, e.g., Hinkelmann and Kempthorne 2005). In addition to two-dimensional square lattice designs discussed above, he discussed lattices for $t = p_1 p_2$ when the treatments are arranged in a rectangle with p_1 rows and p_2 columns, and the lattice design generated from this arrangement consists of p_1 blocks of size p_2 (using the rows as blocks) in one replicate and p_2 blocks of size p_1 (using the columns as blocks) in the second replicate. These designs, however, did not find wide application. Their drawback is that if p_1 and p_2 are quite different, then the error variances for the different block sizes may be quite different from each other, which leads to complicated analyses. Yates (1939) also proposed *three-dimensional* or *cubic lattices* with $t = p^3$ in blocks of size p . In addition to blocking in one direction only, he further considered blocking in two directions, that is, confounding in two directions, which led him to *lattice squares* and *lattice rectangles*. Extensions of many of the most commonly used lattice designs were given by Kempthorne and Federer (1948a,b) and Federer (1955).

Another extension of the lattice designs introduced by Yates was given by Harshbarger (1947, 1949, 1951). He considered the case of $t = p(p - 1)$, which is of considerable practical value since it fills the gap between $t = (p - 1)^2$ and $t = p^2$. Specifically, Harshbarger considered designs with blocks of size $p - 1$ and called these *rectangular lattice designs*. The treatments again are labeled (xy) with $x, y = 0, 1, \dots, p - 1$ and $x \neq y$ and arranged in a square array, illustrated here for $p = 5$.

—	01	02	03	04
10	—	12	13	14
20	21	—	23	24
30	31	32	—	34
40	41	42	43	—

The first replicate is then obtained by using the rows as blocks and the second replicate by using the columns as blocks. This design is called a *simple rectangular lattice* with $t = 20$ treatments in 10 blocks of size 4, each treatment being replicated twice. Harshbarger also

introduced a *triple rectangular lattice* by superimposing a Latin square of order p on this array and deleting the diagonal. For this to work, however, the Latin square must be one with a transversal along the main diagonal, that is, the cells along the main diagonal must contain all of the p Latin letters. The third replicate is then obtained by using the Latin letters to form the p blocks of size $p - 1$.

Lattice designs are, by their very nature, IBDs. Another property is that each replicate contains all the treatment combinations (varieties). They are referred to as *resolvable incomplete block designs* (see also Chapter 3). The property of *resolvability*, introduced by Bose (1942), is of particular importance in variety trials as argued by Patterson and Williams (1976, p. 84): "Practical field conditions dictate that all designs used for these trials are resolvable. Thus some important disease measurements are expensive and have to be restricted to one or two replications. Again, large trials cannot always be completely drilled or harvested in a single session. Use of resolvable designs allows these operations to be done in stages, with one or more complete replications dealt with at each stage."

In addition to the lattice designs described above, other resolvable designs were proposed by Bose and Nair (1962) and by David (1967) with his *cyclic designs*. To provide even more flexibility, Patterson and Williams (1976) introduced what they refer to as *alpha designs*. Their algorithm allows for the construction of resolvable designs for any number of treatments (varieties) t and block sizes k such that t is a multiple of k . Their designs include some but not all square and rectangular lattices.

1.6.3 Partially Balanced Incomplete Block Designs

Realizing that practically useful BIB designs exist only for a relatively small set of parameters (t, b, k, r), Bose and Nair (1939) introduced a rather rich class of IBDs that they called *partially balanced incomplete block designs*. The main difference compared to BIB designs is that now simple treatment differences are estimated with different precisions according to certain rules.

In its simplest and most intuitive and practicable form, there will be two different variances for the estimates of treatment differences. Using a well-defined *association scheme*, for each treatment the remaining treatments are divided into first associates and second associates. Each treatment has n_1 first associates and n_2 second associates, and each treatment occurs together with its first associates in the same block (of size k) λ_1 times and with its second associates λ_2 times, where λ_1 or λ_2 can be 0, but $\lambda_1 \neq \lambda_2$. Denoting the estimator for the treatment effect difference $\tau_u - \tau_v$ by $\hat{\tau}_u - \hat{\tau}_v$ any two treatments, T_u and T_v say, we then have $\text{Var}(\hat{\tau}_u - \hat{\tau}_v) = c_1\sigma_e^2$ if T_u and T_v are first associates or $c_2\sigma_e^2$ if they are second associates, with $c_1 \neq c_2$. The most important condition for this result to hold is the following: if any two treatments T_u and T_v are i th associates, then the number of treatments common to the j th associates of T_u and the k th associates of T_v is p_{jk}^i , independent of the i th associates T_u and T_v . To illustrate these concepts, we consider the following example (see Clatworthy 1973):

Suppose we have $t = 8$ treatments and $b = 6$ blocks of size $k = 4$. Using the *group-divisible association scheme*

1	5
2	6
3	7
4	8

where two treatments in the same row are first associates, and otherwise are second associates, we have $n_1=1$ and $n_2=6$. The matrices $P_i=(p_{jk}^i)$ are given by

$$P_1 = \begin{pmatrix} 0 & 0 \\ 0 & 6 \end{pmatrix}, \quad P_2 = \begin{pmatrix} 0 & 1 \\ 1 & 4 \end{pmatrix}.$$

A *group-divisible PBIB design* with $\lambda_1=3$ and $\lambda_2=1$ is given by (with the rows representing the blocks, before randomization of the treatments)

1	5	2	6
3	7	4	8
1	5	3	7
2	6	4	8
1	5	4	8
2	6	3	7

This design is also a *resolvable design*, because blocks 1 and 2, blocks 3 and 4, and blocks 5 and 6 each form a replicate of all 8 treatments.

Based on earlier results by Bose (1939) for BIB designs, Bose and Nair (1939) and subsequently many other contributors provided an array of mathematical tools for the construction of PBIB designs, such as geometrical configurations, finite Euclidean and projective geometries, the method of symmetrically repeated differences, triple systems, and orthogonal arrays (see Raghavarao, 1971; Street and Street, 1987; Hinkelmann and Kempthorne, 2005). Bose and Shimamoto (1952) studied in detail the *two-associate-class PBIB designs*, which, from a practical point of view, represent the most important PBIB designs. Extensive tables of these designs were prepared by Bose et al. (1954) and enlarged and revised by Clatworthy (1973).

Another large and important class of PBIB designs are the *cyclic PBIB designs* with between 2 and $t/2$ associate classes. They were introduced by Kempthorne (1953) and Zoellner and Kempthorne (1954) for designs with blocks of size $k=2$. They were further developed more generally by, for example, David (1963, 1965), David and Wolock (1965), John (1966, 1969), and Wolock (1964). An extensive table of these designs was prepared by John et al. (1972) (see also John and Williams 1995). For $k=2$ these designs also became known as *paired comparison designs*. Cyclic PBIB designs proved to be quite useful (see, e.g., Section 1.11.1) since (1) they are easy to construct, (2) they exist for a wide combination of design parameters, and (3) they are easy to analyze.

Further developments led to many other types of PBIB designs with more than two associate classes. Among them are the Kronecker product designs by Vartak (1955) (see also Surendran 1968), m -associate class group-divisible designs of Roy (1953–1954) and Raghavarao (1960), extended group-divisible designs by Vartak (1959) and Hinkelmann (1964), cubic designs by Raghavarao and Chandrasekhararao (1964), and hypercubic designs by Shah (1958) and Kusumoto (1965). For mathematical details of constructing IBDs, see Raghavarao (1971) and Street and Street (1987). Later developments are discussed in Chapter 3.

1.6.4 Youden Square Designs

The introduction of IBDs was an important step in conducting experiments with a large number of treatments. The use of small blocks, either intentionally or dictated by the type

TABLE 1.16

5 × 21 Youden Square

Pos.	Plant																				
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
1	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
2	B	F	K	M	S	A	O	C	D	E	L	H	U	P	I	R	J	Q	T	N	G
3	C	N	F	O	L	G	B	S	J	H	A	P	I	Q	R	M	T	K	U	E	D
4	D	J	U	F	Q	I	S	N	P	O	M	T	B	A	C	E	G	H	R	K	L
5	E	R	P	T	F	H	K	M	S	U	J	B	Q	O	L	G	C	D	A	I	N

of available experimental material, often leads to considerable reduction of experimental error. Experimenting with tobacco-mosaic virus at the Boyce Thompson Institute for Plant Research, Youden (1940, p. 228) wrote: “It is apparent that much depends upon an apt selection of an arrangement which will take the fullest possible advantage of either known or suspected similarities in the material or environment.” Using leaves of plants as experimental units, he specifically referred to similarities of the leaves from a single plant and similarities of the leaves in the same position on different plants. This led him (Youden 1937) to use special arrangements of BIB designs with plants as blocks, arranging the treatments such that each treatment occurred exactly once in each position. Such an arrangement is part of a $t \times t$ Latin square, specifically a $k \times t$ rectangle, where k is the block size of the BIB design. For this reason, these designs are referred to as *incomplete Latin square designs* or as *Youden squares*, the latter name being proposed by Fisher. They thus provide for elimination of heterogeneity in two directions, just as the Latin square design (see Chapters 3 and 10).

As an example, we consider the Youden square design of Table 1.16 that was actually used by Youden (1937) for 21 virus preparations on five leaves, each of 21 *Nicotiana glutinosa* L. plants.

Youden square designs exist for all *symmetrical BIB designs*, that is, all BIB designs with $t = b$ can be arranged as Youden square designs (Hartley and Smith 1948). A listing of Youden square designs is provided by Cochran and Cox (1957).

1.6.5 Comparing Treatments with a Control

In some situations, for example, variety trials, the emphasis of the experiment may be to compare *test treatments (varieties)* with a *control or standard*. Even though this can be accomplished with any of the available designs by having one of the treatments as the control, it seemed more appropriate to develop new designs that pay attention to the special role of the control and that are more efficient for this type of comparisons. To serve this purpose, several classes of such designs were developed.

An obvious choice was to reinforce or supplement an existing design, such as a complete block or IBD with one or more replications per block of the control. Such designs are referred to as *supplemented balance designs* (Hoblyn et al. 1954; Pearce 1960) or *reinforced BIB designs* (Cox 1958; Das 1958).

A more specific approach was taken by Bechhofer and Tamhane (1981) when they introduced the *balanced treatment incomplete block design (BTIB)*, which is defined as a design consisting of t test treatments and one control in b blocks of size $k < t + 1$ with the

following properties: (1) each test treatment occurs together with the control λ_0 times in a block and (2) any two test treatments occur together λ_1 times in a block. Tables of such designs were provided by Bechhofer and Tamhane (1985). A discussion of the properties of designs for comparing test treatments with controls is given by Hedayat et al. (1988) (see also Chapter 3).

Further extensions of these designs involve the use of *partially balanced treatment incomplete block designs* (see Rashed 1984; Jacroux 2003; Hinkelmann and Kempthorne 2005) and the inclusion of several controls (see Jacroux 2002, 2003).

1.7 Fractional Factorial Designs

“The use of a factorial set of treatment combinations has now become a widely accepted means of investigating, within a single experiment, the effects on the experimental material of several different treatment variants. The principles of factorial design have been elaborated primarily for agricultural field experiments, but have been found valuable in many other types of experimentation” (Finney 1945, p. 291). At the same time, there was some concern about the large number of treatment combinations necessary for such experiments and the practicality of carrying them out. Finney, in cooperation with Yates and Kempthorne, began to discuss ways of reducing the number of treatment combinations, and Finney (1945) introduced the notion of *fractional replication*, an important development that opened up the possibility of using different fractions of the complete factorial, based on different assumptions. This made the application of factorial designs much more practical and, in turn, led to a much wider use of factorial experiments, not only in the agricultural setting but particularly in industrial and manufacturing applications (see, e.g., Davies 1956; Box et al. 2005).

1.7.1 Fractional Replication

The basic idea for considering only a fraction of all possible treatment combinations was based on an argument often given by Fisher and Yates for justifying the use of nonreplicated complete factorials, namely, that higher-order interactions are (1) generally of no interest and (2) usually negligibly small, the latter supported by practical experience. The question then is how to select a subset of all treatment combinations without sacrificing too much information. Finney (1945) used the 2^f factorial to develop a suitable procedure (see also Chapters 7 through 9).

As an example, we consider the 2^3 factorial with factors A, B, C . The eight treatment combinations (observations) of the full factorial allow the estimation of all the main effects and interactions, representing seven estimable functions. With only a fraction of the treatment combinations, the number of estimable functions will necessarily be reduced. If the three-factor interaction ABC is assumed to be negligible, then it seems natural to consider a fraction consisting of the treatment combinations that have, say, negative signs in the symbolic representation of ABC (see Section 1.5.2) or, using the terminology of Kempthorne (1952), that enter negatively into ABC . These treatment combinations are 000, 110, 101, and 011, representing a $1/2$ fraction of the 2^3 factorial, denoted by 2^{3-1} . Since there are only four observations, only three effects/interactions can be estimated, and each effect/interaction

TABLE 1.17
 Estimation of Effects/Interactions in a 2^{3-1} Factorial

	Effects							
	Mean	A	B	AB	C	AC	BC	ABC
Treatments								
000	+	-	-	+	-	+	+	-
110	+	+	+	+	-	-	-	-
101	+	+	-	-	+	+	-	-
011	+	-	+	-	+	-	+	-

(except the mean and ABC) is a linear contrast of the observations associated with the four available treatment combinations. These contrasts can be obtained from Table 1.17, where we list the signs for each treatment combination in the symbolic representation of the various main effects and interactions.

It is obvious from Table 1.17 that in addition to the mean- ABC (i.e., the average of all the observations), we can estimate the following three combinations of effects: $A-BC$, $B-AC$, and $C-AB$. Thus, A is *confused* (using Finney’s terminology) or *aliased* with BC , B with AC , and C with AB . Designs with these properties are generally not of practical value, unless two-factor interactions are negligible. This example, however, illustrates the nature of fractional factorial designs, in that various effects are aliased with each other and assumptions have to be made in order to estimate the relevant effects.

Formally, the particular fraction as well as its *alias structure* is determined by the *identity relationship* or *defining relation*. For the 2^{3-1} factorial given earlier, this defining relation is given by

$$I = ABC \quad \text{or more precisely} \quad I = -ABC,$$

which says that the fraction consists of all the treatment combinations that enter negatively into ABC . Using this equation formally as a mathematical equation with I as the identity and replacing every letter raised to an even power by 1, we obtain the alias structure by multiplying each effect into both sides of the defining relation, such as

$$A = -A(ABC) = -A^2BC = -BC.$$

This indicates, as shown earlier, that $A-BC$ is estimable or, alternatively, that A is aliased with BC .

Finney (1945) (see also Kempthorne 1947) showed that this procedure can be used to obtain different fractions for the 2^f factorial. For example, a $1/4$ fraction of the 2^6 factorial, or a 2^{6-2} , can be obtained by specifying a suitable defining relation, such as

$$I = ABCD = CDEF,$$

which means that the fraction consists of all the treatment combinations that enter with the same sign into $ABCD$ and $CDEF$. This yields 16 treatment combinations, and it is easy to see that these 16 treatment combinations also enter with the same sign into the *generalized*

interaction of $ABCD$ and $CDEF$, which is formally obtained as $(ABCD) \times (CDEF) = ABEF$. With this the final defining relation is given by

$$I = ABCD = CDEF = ABEF.$$

Finney (1945) refers to the elements in the defining relation as the *alias subgroup*. The reason, of course, is that from this relationship, we obtain the alias structure, such as

$$\begin{aligned} A &= BCD = ACDEF = BEF, \\ AB &= CD = ABCDEF = EF. \end{aligned}$$

This can be continued until all 15 estimable functions of effects have been identified. The preceding equations indicate, for example, that the main effect A is aliased with two 3-factor interactions and one 5-factor interaction. If these can be assumed to be negligible, then A is estimable. This will be the case with all other main effects as well. On the other hand, the second equation shows that two-factor interactions are aliased with other two-factor interactions, and judicious assumptions will have to be made if one wants to estimate specific two-factor interactions (see Chapter 7).

These kinds of arguments were also extended by Finney (1945) and Kempthorne (1947) to 3^{f-8} fractions. As an example, consider the 3^{4-2} , that is, a 1/9th replicate of the 3^4 factorial consisting of nine treatment combinations. Using Kempthorne's (1947, 1952) notation, a possible defining relation might be

$$I = ABC^2 = ACD = (ABC^2)(ACD) = (ABC^2)(ACD)^2.$$

In this defining relation, ABC^2 and ACD are referred to as the *independent interactions*. There are now two generalized interactions, obtained by multiplying the first independent interaction into the second independent interaction and into the squared interaction. The multiplication is then carried out formally that leads to

$$I = ABC^2 = ACD = A^2BC^3D = A^3BC^4D^2.$$

In this equation, each power is reduced modulo 3, dropping each letter raised to the power 0 and using the convention that the first letter of each interaction should be raised to the first power, which is achieved by squaring the entire interaction term, as needed. Then the final defining relation is

$$I = ABC^2 = ACD = AB^2D^2 = BCD^2.$$

It can be shown (see Hinkelmann and Kempthorne 2005, Example 13.4) that with the resulting fraction, one can estimate the main effects for factors A , B , C , and D , assuming that all interactions are negligible. Since each main effect accounts for 2 d.f., the four main effects account for the 8 d.f. available from the nine observations.

The mathematics for the 3^{f-8} fractional factorial laid the foundation for dealing with *regular fractions* for the general s^f factorial, where s is the power of a prime number (see Kempthorne 1952; Hinkelmann and Kempthorne 2005). An important result was given by Rao (1947), who established the relationship between fractional factorials and orthogonal arrays (for further reading on orthogonal arrays, see Hedayat et al. 1999).

Employing orthogonal arrays relaxes the necessity to use a number of treatment combinations that is a power of s , which is a restriction when considering regular fractions. In any case, however, for the practical application, Finney (1945, p. 301) advises: "Fractional replication is a device which should always be used with caution, and skillful choice is needed to ensure that only one of the several aliases of each degree of freedom is at all likely to represent a real treatment effect." Examples of a wide range of applications using fractional factorial designs are given by Prvan and Street (2002).

1.7.2 Characterization of Regular Fractional Factorials

Since their introduction, fractional factorial designs have been used widely, in particular in industrial and manufacturing applications. Because there are many choices for designing fractional factorials, it seemed appropriate to provide some sort of classification for them that would reflect their properties with respect to the estimation of effects, in particular main effects and two-factor interactions. The first important step in this direction was made by Box and Hunter (1961a,b). They introduced the notion of the *resolution* of a design as the lowest order of any interaction included in the defining relationship. For example, if the defining relation contains a three-factor interaction and no main effects or two-factor interactions, then the design is said to be a resolution III design. The important consequence then is that some main effects are aliased with two-factor interactions. Hence, those main effects can only be estimated if those two-factor interactions (plus some other higher-order interactions) are negligible. In resolution IV designs, main effects are aliased with three-factor interactions, and two-factor interactions are aliased with other two-factor interactions. From an estimation point of view, resolution V designs are of major interest, since they allow the estimation of main effects and two-factor interactions, assuming that all other interactions are negligible. The drawback of these designs, of course, is that they cannot be constructed with very small fractions. For example, for the 2^8 factorial, only a 1/4th fraction with 64 treatment combinations can result in a resolution V design, whereas a 1/8th or 1/16th fraction with 32 or 16 treatment combinations, respectively, can be obtained as a resolution IV design. Chapter 7 provides a detailed account of the construction of resolution III, IV, and V two-level designs.

Another important concept to distinguish between competing fractional factorials was introduced by Fries and Hunter (1980). They wanted to distinguish between various designs of the same resolution for the same number of factors. As an example, consider the following 2^{8-3} fractions of resolution IV, d_1 and d_2 say, given by the defining relations

$$d_1 : I = ABCDEF = CDEG = ABFG = BDEH = ACFH = BCGH = ADEFGH,$$

$$d_2 : I = CDEF = ABDEG = ABCFG = ABCEH = ABDFH = CDGH = EFGH.$$

The defining relation contains 5 four-factor interactions for d_1 and 3 for d_2 . As a consequence, in d_1 24 of the 28 two-factor interactions are aliased with each other, and the remaining 4 are aliased with higher-order interactions, whereas in d_2 only 13 two-factor interactions are aliased with each other, and the remaining 15 are said to be *clear* (Wu and Chen 1992), i.e., aliased with three-factor or higher-order interactions. According to Fries and Hunter (1980), d_2 has less *aberration* than d_1 . Among all competing designs, the one

with the smallest number of four-factor interactions in the defining relation is said to be a *minimum aberration design* of resolution IV.

Still other criteria have been introduced to characterize and distinguish among competing designs, such as *estimation capacity*, *projections of a design*, *estimation index*, and *generalized minimum aberration*. These criteria are useful not only for constructing regular fractional factorials but also for constructing *nonregular fractional factorials*. Nonregular designs are more flexible with respect to number of treatment combinations than regular designs as they are not based on defining relations but mostly on orthogonal arrays (see Plackett–Burman designs in Section 1.7.3). They also accommodate situations where different factors have different numbers of levels (see, e.g., Addelman 1962). This leads often to rather complicated alias structures, however. An excellent discussion of the aforementioned criteria and the relationships among them pertaining to regular and nonregular fractional factorial designs, without and with blocking, has been given by Chen and Cheng (2012). For more details, see also Wu and Hamada (2009), Mukerjee and Wu (2006), and Chapter 9.

1.7.3 Main Effect Plans

Plackett and Burman (1946) posed the question how one should select n treatment combinations from a complete factorial such that the main effects will be estimated with the same precision as if attention had been concentrated on varying the levels of a single factor throughout the n treatment combinations. An example of their effort is the following design for 11 factors, each at 2 levels, in $n = 12$ assemblies (treatment combinations or runs), where the rows represent treatment combinations and the columns represent factors (the *low* and *high* levels are denoted by 0 and 1, respectively):

1	1	0	1	1	1	0	0	0	1	0
0	1	1	0	1	1	1	0	0	0	1
1	0	1	1	0	1	1	1	0	0	0
0	1	0	1	1	0	1	1	1	0	0
0	0	1	0	1	1	0	1	1	1	0
0	0	0	1	0	1	1	0	1	1	1
1	0	0	0	1	0	1	1	0	1	1
1	1	0	0	0	1	0	1	1	0	1
1	1	1	0	0	0	1	0	1	1	0
0	1	1	1	0	0	0	1	0	1	1
1	0	1	1	1	0	0	0	1	0	1
0	0	0	0	0	0	0	0	0	0	0

This design is referred to as a *main effect plan* or more precisely as a *saturated orthogonal main effect plan* (OMEPE). The construction is based on the existence of Hadamard matrices (see Hedayat et al. 1999). Plackett and Burman have shown that for the 2^f factorial such plans exist for $f = n - 1$ and $n = 4q$ ($q = 2, 3, \dots$), although it is not known whether they exist for all q (see Hedayat et al. 1999). Plackett and Burman have listed the actual designs for some values of q (see also Chapter 9).

Another method of constructing OMEPEs is based on Fisher's (1942, 1945) procedure for constructing the intrablock subgroup for systems of confounding. This subgroup is the

OMEF, called the Fisher plan (it is worth mentioning at this point that the mathematical ideas for regular fractional factorials and for confounding [Sections 1.5.2 and 1.5.3] are essentially the same). The OMEFs are resolution III designs. For asymmetrical factorials, Addelman and Kempthorne (1961) developed several methods of constructing OMEFs, essentially extending Fisher's method.

OMEFs have become a very important tool in exploratory studies, in particular when a large number of factors have to be screened. The use of OMEFs and other highly fractionated factorial designs (see, e.g., Tippett 1936) "has usually rested on an implicit hypothesis of what we call *factor sparsity*. This hypothesis is that in relation to the noise only a small proportion of the factors have effects that are large. The former will be called *active* factors" (Box and Meyer, 1986, p. 11). Factor sparsity typically also implies *effect sparsity*, in that only the main effects of the active factors and perhaps low-order interactions among them are important.

1.7.4 Supersaturated and Search Designs

For screening a very large number of factors, f say, even OMEFs may not be practical because they require too many treatment combinations. Moreover, assuming factor sparsity, as mentioned above, the main task then really becomes to select a small number, n say, with $n < f + 1$, from the totality of all possible treatment combinations. Such designs are called *supersaturated designs*. They are no longer orthogonal designs. Also, the effects can no longer be estimated simultaneously but, based on the assumption of effect sparsity and absence of factor interactions, some method of sequential regression analysis may be appropriate.

Different methods of constructing supersaturated designs have been proposed. Among them are designs found by computer search, first proposed by Booth and Cox (1962), for the following pairs of (f, n) : (16, 12), (20, 12), (24, 12), (24, 18), (30, 18), (36, 18), and (30, 24). For more recent results see Gupta et al. (2010). Other methods are based on Hadamard matrices (Lin 1993; Wu 1993; Butler et al. 2001; Gupta et al. 2008), on orthogonal arrays (Fang et al. 2000; Liu and Lin 2009; Gupta et al. 2008), on BIB designs (Nguyen 1996; Butler et al. 2001; Gupta et al. 2011), and on cyclic developments (Liu and Dean 2004; Georgiou and Koukouvinos 2006; Georgiou et al. 2009). For more details see Chapter 9, and for an overview of two-level supersaturated designs, see Kole et al. (2010).

The usefulness of OMEFs and supersaturated designs is based on the assumption that essentially all interactions are negligible. In some cases that may be a questionable assumption. This raises the question whether one can construct designs, which (1) consist of a small number of treatment combinations and (2) allow the exploration of a limited, but unspecified, number of low-level interactions. This led Srivastava (1975) to the development of what he called *search designs*.

The most important classes of search designs are the main-effect-plus-one or resolution III.1 plans, allowing the search of one active two-factor interaction, and the resolution V.1 plans, allowing the search for one active three-factor interaction (in addition to all two-factor interactions) (Srivastava and Gupta 1979; Srivastava and Ghosh 1976) from among the totality of all such interactions. Optimal two-level designs were identified by Ghosh and Tian (2006). A method for comparing search designs based on search probabilities was proposed by Ghosh and Teschmacher (2002) (see also Sarkar and Chatterjee 2010).

1.8 Industrial and Technological Experimentation

1.8.1 Early Use of Experimental Design

The experimental designs we have discussed so far were developed mainly in the context of agricultural experimentation. Many of these ideas were accepted, however, also in other fields of experimentation (see also Section 1.11), in particular in industrial production. Tippett (1935, 1936, 1938), after having visited Fisher at Rothamsted, advocated the use of experimental design when he joined the British Cotton Industry Research Association's Shirley Institute. This included the use of orthogonalized Latin squares (see Section 1.4.3) and factorial experiments. However, differences between agricultural and industrial experimentation and reluctance on the part of production managers to conduct complex experiments (see Hamaker 1955) led to a slow increase in the application of experimental design in industry. One of the differences between agriculture and industry is the speed with which results from experiments become available, generally much faster in industry than in biological and medical sciences. For this reason, smaller experiments were performed that required little time and needed less careful planning. In addition, since variables could be controlled to a large extent, the results were much less subject to experimental error, and hence, the analyses relied less on statistical arguments. An increase in reported experiments, however, occurred in the 1950s using available designs. For references to a large number of case studies, see Bisgaard (1992), and for examples of industrial applications, see Hamaker (1955), Dean and Voss (1999), and Box et al. (2005). The latter also emphasize graphical representation of experimental results and analyses, which is an important requirement for engineers and managers (Hamaker 1955).

Most of the designs we have described are intended for comparative experiments, that is, designs where the main interest is in comparing different qualitative or quantitative treatments with respect to their effects or performance. This is not always of interest in industrial experiments. There, one would like to investigate the relationship between treatments and responses in the form of a response surface, that is, between input or independent and output or dependent variables. Moreover, often one would like to use this relationship to find the factor-level combination that gives the optimal (highest or lowest) response. Methods that are directed toward this kind of investigation, using tools from experimental design and regression analysis, are commonly known as *response surface methodology* (RSM). Beginning with the seminal paper by Box and Wilson (1951), this started a new and fruitful direction in industrial experimentation.

1.8.2 Basic Ideas of Response Surface Methodology

The basic ideas of RSM were essentially formulated by Box and Wilson (1951) as they described the practical situation and presented methods on how to solve the problem of "experimental attainment of optimum conditions."

The response η of a process is assumed to depend on k standardized quantitative factors x_1, x_2, \dots, x_k , which can be measured or controlled exactly (we note here that we now use k for the number of factors rather than f in order to conform to the notation used in RSM). For the u th combination ($u = 1, 2, \dots, n$) of factor levels, we express this as

$$\eta_u = \phi(x_{1u}, x_{2u}, \dots, x_{ku}; \theta_1, \theta_2, \dots, \theta_q) = \phi(x_u; \theta),$$

where $\theta_1, \theta_2, \dots, \theta_q$ are unknown parameters. The true yield η at any given point in the *operational region* and the functional relationship ϕ are unknown. Instead, we have available only observed responses $y_u = y(x_{1u}, x_{2u}, \dots, x_{ku})$ and shall attempt to approximate $\phi(x; \theta)$ by a polynomial function $f(x_1, x_2, \dots, x_k; \beta_1, \beta_2, \dots, \beta_m) = f(x; \beta)$. This could be done in principle by having observations over a fine grid in the operational region, but that is not feasible from a practical point of view. Instead, we will be restricted to a relatively small number of points (which are often referred to as runs or experiments) that will typically be restricted to a smaller region, called the *region of interest*. Because of “unavoidable uncontrolled factors the observed response y_u varies in repeated observations, having mean η_u and variance σ^2 ” (Box and Wilson, 1951, p. 2). Box and Wilson argued that the experimental strategy to find the point $(x_1^0, x_2^0, \dots, x_k^0)$ in the experimental region, at which η is a maximum or minimum, is influenced by

1. The magnitude of the experimental error;
2. The complexity of the response surface;
3. Whether or not experiments may be conducted sequentially so that each experimental point may be designed using the knowledge gained from the previous runs (experimental points).

In their words (p. 2), “Since the experimental error is small, small changes can be determined accurately, and the experimenter may explore adequately a small *sub-region* in the whole experimental region with only a few experiments. Also since the experiments are sequential the possibility arises of using the results obtained in one sub-region to move to a second in which the response is higher. By successive application of such a procedure, a maximum or at least a near-stationary point of high response should be reached.” As such, a method they introduced is what they refer to as the *method of steepest ascent*. To apply the method of steepest ascent, they considered first-order polynomial models, and to investigate the response surface at a near-stationary region, they considered second-order polynomial models.

1.8.3 Response Surface Designs

Box and Wilson (1951) pointed out that their procedure leads to two types of errors: (1) experimental error in estimating the function $f(x; \beta)$ and (2) bias due to the inadequacy of $f(x; \beta)$ approximating $\phi(x; \theta)$. To minimize these errors, singly or jointly, is essentially the focus of response surface designs. Box and Hunter (1957) suggested the following basic requirements of such designs:

1. Assuming that a polynomial $f(x; \beta)$ of degree d approximates $\phi(x; \theta)$ sufficiently well, the design should allow $f(x; \beta)$ to be estimated with satisfactory precision.
2. The design should allow a check of whether the chosen estimate of $f(x; \beta)$ provides a satisfactory fit to the response surface or whether a different polynomial may be more appropriate.
3. The design should not contain an excessively large number of experimental points.
4. The design should lend itself to adequate blocking of the experimental points.
5. One should be able to amend the design in case the polynomial of degree d proves to be inadequate and a polynomial of degree $d + 1$ needs to be fitted.

To estimate the coefficients in first-order models, we use 2^k factorial designs, or if they are too large, we consider suitable fractional factorial designs. By suitable, we mean, of course, that the design allows the estimation of main effects. Alternative designs called *simplex designs* were introduced by Box (1952). In these designs, the design points are located at the vertices of a regular k -dimensional simplex.

For the estimation of the coefficients of second-order models, one could, of course, consider 3^k factorials or resolution V 3-level fractional factorials. But even resolution V designs contain an excessive number of points for the relatively small number of regression coefficients to be estimated. Typically, one would like to estimate the linear, quadratic, and linear \times linear effects, utilizing both d.f. for the main effects, but only one d.f. from each of the two-factor interactions. This may leave an excessive number of d.f. for testing lack of fit and for error. Designs more suitable for this purpose have been developed over the years.

As one such design, Box and Wilson (1951) introduced the *central composite design*, probably the most commonly used design in RSM. In this design, each factor is used at five different levels, but not all level combinations occur. Rather, the design consists of three parts, as illustrated in Figure 1.2 for $k = 2$ and 3:

1. A factorial or *cube* part consisting of a full 2^k factorial or a 2^{k-p} fraction of at least resolution V, each point being replicated, say, r_F times.
2. An axial or *star* part consisting of $2k$ points on the axis of each factor at a distance α from the center of the design, each point being replicated, say, r_A times.
3. One or more, say n_0 , replications of the center point $(0, 0, \dots, 0)$.

The total number of experimental runs then is $n = 2^{k-p}r_F + 2kr_A + n_0$. Box and Hunter (1957) added to this design the property of *rotatability*, which is automatically satisfied for $k = 1$, but in general imposes certain conditions on the value for α , with $\alpha = \sqrt{2}$ for $k = 2$, for example. Rotatability is a desirable property as it ensures that the variance of the predicted response is the same for all points that are equidistant from the design center.




FIGURE 1.2 Central composite designs. (a) $k = 2$ and (b) $k = 3$.

Another important class of designs in this context are the *Box–Behnken designs* introduced by Box and Behnken (1960). These designs use only three levels of each factor. They are constructed by combining ideas from 2^k factorial designs and IBDs (BIB or PBIB designs with k treatments and b blocks of size m), resulting in $b2^m$ experimental points plus n_0 center runs (for a general description, see Jo and Hinkelmann 1993). For further designs and modifications of the designs described above, see Box and Draper (2007), Khuri and Cornell (1996), Myers et al. (2009), and Chapter 5.

1.8.4 Mixture Experiments

A special and yet quite distinct application of RSM occurs in experiments with mixtures. In such cases, the output $\eta = \phi(x; \theta)$ does not depend on the actual input values, but rather on the proportions to each other, represented by $x = (x_1, x_2, \dots, x_k)$. This means that $x_1 + x_2 + \dots + x_k = 1$ or, alternatively, 100%. An example may be the blending of $k = 3$ gasoline and ethanol stocks to determine the fuel blend that gives the best gas mileage.

Following initial work by Claringbold (1955), a major impetus for research in this area was provided by Scheffé (1958). He introduced the (k, m) -*lattice design*, referring to a collection of uniformly spaced design points on a simplex (for this reason, the design is now called a *simplex-lattice design*, see Cornell 2002). The input values x_i are of the form $x_i = 0, 1/m, 2/m, \dots, 1$ subject to the condition given earlier. Thus, for example, the $(3, 2)$ -lattice consists of the design points $(x_1 x_2 x_3) = \{(100), (010), (001), (1/2 \ 1/2 \ 0), (1/2 \ 0 \ 1/2), (0 \ 1/2 \ 1/2)\}$. Other commonly used mixture designs are *simplex-centroid designs* and *axial designs* (for their definitions, see Cornell 1975, 2002).

Typically, the analysis of data from mixture experiments is based on polynomial models of first or second degree, that is,

$$y(x) = \beta_0 + \sum_{i=1}^k \beta_i x_i + \epsilon$$

or

$$y(x) = \beta_0 + \sum_{i=1}^k \beta_i x_i + \sum_{i=1}^k \beta_{ii} x_i^2 + \sum_{i < j} \beta_{ij} x_i x_j + \epsilon.$$

Because of the dependency among the x_i 's, the parameters associated with the various model terms are not unique. To remove the dependency among the x_i 's, we could write, for example, $x_k = 1 - \sum_{i=1}^k x_i$ and substitute this expression for x_k into the preceding models. For the first-degree model, this would lead to a model with parameters $\beta_i - \beta_k$, obscuring the effects of the individual components. This becomes even more complicated for the second-degree model. An alternative method of dealing with the dependence among the x_i 's is to use so-called canonical polynomials (Scheffé 1958). For the first-degree model, such a polynomial is obtained by writing the original model as

$$y(x) = \beta_0 \left(\sum_{i=1}^k x_i \right) + \sum_{i=1}^k \beta_i x_i + \epsilon$$

using the fact that $\sum x_i = 1$. Combining terms leads to the model

$$y(x) = \sum_{i=1}^k \beta_i^* x_i + \epsilon$$

with $\beta_i^* = \beta_0 + \beta_i$ ($i = 1, 2, \dots, k$). This model retains its symmetry and the β_i^* have a clear meaning. For more details, we refer to Cornell (2002), Piepel and Cooley (2009), Sinha et al. (2010), Zhang and Peng (2012), and Chapter 5.

In addition to the input variables x_1, x_2, \dots, x_k , the experiment may also include so-called process variables, say, z_1, z_2, \dots, z_q , which also affect the output y , but are not connected with the blending process itself. For example, in addition to the blending proportions of gasoline and ethanol stocks, the output may also be a function of the type of car and driving habits. This leads to obvious extensions of the aforementioned models, but at the same time is bringing blocking and factorial designs into the picture. As a result, designing and analyzing such experiments becomes more complicated. For a more general discussion, see Cornell (2002, 2011) as well as some specific results given, for example, by Prescott (2004a,b) and Goldfarb et al. (2004).

1.8.5 Computer Experiments

The experiments we have considered so far are *physical experiments*, which are carried out in the real world, that is, on a piece of land, in the laboratory, or in a factory, but even they can become quite complicated. We have seen earlier (see Section 1.7.4) how one can deal with, for example, large numbers of input variables. In addition to a very large number of factors, another complication may arise in the form of a complex relationship between input and output variables, something that in, say, physics or meteorology may be expressed mathematically in terms of a complicated set of differential equations. With the advent of fast and powerful computers, these types of situations have been dealt with through computer simulation in the form of *computer experiments*. See Chapter 16 for a history and overview of computer experiments.

A computer program, incorporating methods of obtaining a numerical solution to the problem at hand, is written to obtain from a set of input variables, say, $x = (x_1, x_2, \dots, x_k)$, the possibly vector-valued output variable, say, $y = f(x)$. According to McKay et al. (1979), the output y often is available in the form of a graph $y(t)$ as a function over time. From a practical point of view, the problem may arise of what input values to choose. This is discussed by McKay et al. (1979), who consider three methods that have intuitive appeal: (1) random sampling; (2) stratified sampling, making sure that all areas of the sample space of x are represented; and (3) *Latin hypercube sampling*, making sure that each input variable x_i ($i = 1, 2, \dots, k$) has all portions of its distribution represented by input variables. Each form of sampling represents a design for the computer experiment, and Latin hypercube sampling has become a preferred method (see Chapter 17). Still other designs are uniform designs or distance-based designs. For a general discussion, see Sacks et al. (1989); Santner et al. (2003); Lin et al. (2010); Morris (2012); and Chapters 16 through 19. For an interesting application to defense and homeland security, we refer to Sanchez et al. (2012).

1.9 Efficiency and Optimality of Designs

As different designs were introduced that became progressively more complicated in order to cope with different physical or industrial conditions, it was quite natural that methods would be developed to assess which design would be *better* or more appropriate for given experimental situations. To this end, different criteria would be developed, among them the notions of RE, efficiency factor, and optimality.

1.9.1 Relative Efficiency

In agricultural field trials, it is important to understand the nature of the existing heterogeneity so that, if necessary, effective blocking can be used. In order to evaluate the soil heterogeneity, the usual practice in the beginning of scientific experimentation was to perform what are referred to as *uniformity trials*, that is, trials where only one treatment (variety) was used so that the results were not affected by treatment effects. For obvious reasons, this is, of course, not always a practical or economical method to follow. Rather than performing a uniformity trial, Yates (1935) proposed the idea of actually performing the intended experiment and obtaining retrospectively some information on how *successful* the employed design had been. From a practical point of view, this information would be useful only if this type of experiment were to be repeated subsequently. In order to assess how successful a particular form of blocking had been, Yates (1935) introduced the notion of RE. Specifically, he considered the situation of an experiment using a RCBD and then, using the results from the experiment, assessed to what extent blocking had been successful as compared to using a CRD, that is, no blocking at all. This information could be used then for designing future experiments on the same field or, perhaps, similar fields.

Defining the RE of one design, d_1 , relative to another design, d_2 , as

$$\begin{aligned} \text{RE} (d_1 \text{ to } d_2) &= \frac{\text{Efficiency } d_1}{\text{Efficiency } d_2} \\ &= \frac{\text{Var}_{d_2}}{\text{Var}_{d_1}}, \end{aligned}$$

where Var_d refers to $\text{Var}(\sum c_i \hat{\tau}_i)$, the variance of an estimated standardized treatment contrast $\sum c_i \tau_i$ ($\sum c_i = 0, \sum c_i^2 = 1$) for design d , Yates (1935) obtained an estimator for RE(RCBD to CRD). Considering a conceptual uniformity trial and using the notation of Table 1.2 for the RCBD actually used, he obtained

$$\text{RE(RCBD to CRD)} = [(b - 1)\text{MS(B)} + b(t - 1)\text{MS(E)}]/[(bt - 1)\text{MS(E)}].$$

For a detailed derivation, see Hinkelmann and Kempthorne (2008), who called this expression the *estimated RE*.

The practical implication of RE is the relationship between the number of blocks in the RCBD and the number of replications, r , for the CRD when using the same experimental material. This relationship is given by $r = b \times \text{RE}$. It implies that if $\text{RE} > 1$, then the number of replications for a CRD to achieve the same efficiency as the RCBD would have to be RE times as large as the number of blocks, that is, the number of replications for the

RCBD as used in the trial. These ideas can be extended to designs with increasing blocking complexity.

1.9.2 Efficiency Factor

Another notion for comparing properties of different designs, particularly IBDs, is that of *efficiency factor*. It compares, apart from the residual variance σ_e^2 , the average variance of simple treatment comparisons for two competing designs (see Section 1.6.1). For example, comparing an RCBD with r blocks and an IBD, both with the same number of experimental units, the efficiency factor E_{IBD} , assuming that $\sigma_e^2(\text{RCBD}) = \sigma_e^2(\text{IBD})$, is given by

$$E_{\text{IBD}} = \text{av. variance } (\hat{\tau}_i - \hat{\tau}_j)_{\text{RCBD}} / \text{av. variance } (\hat{\tau}_i - \hat{\tau}_j)_{\text{IBD}},$$

where the average (av) is taken over all possible combinations of i and j ($i \neq j$).

If the IBD is a BIB design, then $E_{\text{BIB}} = (k-1)t/(t-1)k$ (Yates 1936), where t is the number of treatments and k is the block size. E_{BIB} is the upper bound for E , that is, the largest possible value for the efficiency factor of an IBD with t treatments and b blocks of size k (Kempthorne 1956). If the IBD is a two-associate class PBIB design, then there are two efficiency factors: E_1 for comparisons between treatments that are first associates and E_2 for comparisons between treatments that are second associates. The overall efficiency factor is then given by

$$E_{\text{PBIB}} = (n_1 E_1 + n_2 E_2) / (n_1 + n_2),$$

where n_1 and n_2 are the number of first and second associates, respectively (Bose and Nair 1939). The values for E_1 , E_2 , and E_{PBIB} are listed by Clatworthy (1973) for making comparisons among competing PBIB designs. For example, for two PBIB designs, say, PBIB(1) and PBIB(2), if $E_{\text{PBIB}(1)} > E_{\text{PBIB}(2)}$, we would most likely choose PBIB(1) over PBIB(2), unless practical considerations would suggest otherwise.

1.9.3 Optimality

The notion of an *optimal design* was first considered by Smith (1918) in the context of fitting polynomial models. She considered experiments with the aim of studying the relationship between treatments (inputs) and responses (outputs), as in response surface designs, for example. And it is in this context that the concept of *optimality* has had the greatest influence even though much of the development goes back to considering *comparative experiments*.

Based on the early results by Wald (1943) and Ehrenfeld (1953) for Latin square-type designs, Kiefer (1958, 1959, 1975a,b) set the stage for a much more rigorous and systematic investigation of optimal designs. He introduced different optimality criteria, which are based on different statistical properties concerning variances of the estimates of treatment contrasts. These criteria can be discussed in terms of either maximizing a suitable function of Fisher's information matrix or minimizing a corresponding function of the variance-covariance matrix, V_d , of a maximal set of orthonormal treatment contrast estimates, say, $P'\hat{\tau}$. The most important optimality criteria are D -optimality where the determinant of V_d is minimized, A -optimality where the trace of V_d is minimized, and E -optimality where the largest eigenvalue of V_d is minimized, with d being a design within a class of competing designs. The statistical meaning of these criteria is that D -optimality minimizes the generalized variance or volume of the confidence ellipsoid of $P'\hat{\tau}$, A -optimality minimizes the

average variance of the set $P\hat{\tau}$, and E-optimality minimizes the maximum variance of a single normalized contrast (for more details, see Chapters 2 and 3).

These three criteria do not always lead to the same optimal design, but Kiefer (1958) showed that *symmetric designs*, that is, designs with their information matrix given by $aI + bJ$ (with a, b being constants, I being the $t \times t$ identity matrix, J being a $t \times t$ matrix with unity elements), are A -, D -, and E -optimal when their information matrices have maximal trace. Examples of such designs include BIB designs, Latin square designs, and Youden designs. From this, additional conditions led to the definition of *universal optimality* (Kiefer 1975a) or *Kiefer optimality* (Pukelsheim 1993).

Since the introduction of A -, D -, and E -optimalities, other optimality criteria have been introduced to serve particular purposes, in particular in the context of response surface and regression-type designs. For a general discussion, see Pukelsheim (1993) and Atkinson et al. (2007).

1.10 Matters of Analysis and Computation

The properties of balancedness and orthogonality play an important part in the development of experimental designs. Not only are designs with these properties often optimal, but also the computations required for estimating contrasts of treatment effects and for obtaining the ANOVA table are quite simple, as they are based on the manipulation of treatment means, block means, etc., conforming to the underlying *linear model*. Concerning the notion of a linear model, it should be noted here that in the early work of Fisher and Yates, we do not find such models stated explicitly. Analyses became more computationally involved with the introduction of nonorthogonal designs, nonlinear models, and more complex error structures. We note, for example, how Cochran and Cox (1950, 1957) go into great detail to explain to the reader how the computations need to be done and how, in certain cases, one can simplify matters.

1.10.1 Nonorthogonal Designs

The first examples of what we now call *nonorthogonal designs* or *nonorthogonal linear models* occur with the method of increasing precision through *concomitant measurements* (or covariates) (Fisher 1935), that is, *analysis of covariance*, and with the occurrence of missing observations from an orthogonal design. In both situations, the computations become more complicated, and Fisher (1935) explained in some detail how the analysis needs to be performed. In the case of covariates, the estimates of the treatment effects (means) need to be adjusted as well as the sum of squares for error. In the case of missing observations, Fisher (1935) described an algebraic method of estimating the missing values for subsequent calculations and reducing the d.f. for the error sum of squares for each missing observation (see also Yates 1933a). The problem of missing observations was dealt with for the general situation by Coons (1957), who used (upon recommendation by M.S. Bartlett) analysis of covariance techniques for this purpose. For several missing values, this becomes computationally quite complicated using the calculators existing at that time.

A large class of nonorthogonal designs is that of IBDs, including BIB, PBIB, and BTIB designs as described in Section 1.6 and Chapter 3. As illustrated by Yates (1936a) for the

BIB design (see Section 1.6.1 and Table 1.15), there no longer exists a unique ANOVA table, but rather two ANOVA tables with what are now called *sequential sums of squares* (see, e.g., Hinkelmann and Kempthorne 2008, Section 4.7; see also Section 1.6.1). One of them is used to test hypotheses about treatment effects and to estimate σ_e^2 , and the other is used by Yates (1940) to estimate σ_b^2 or the ratio $\rho = \sigma_b^2 / \sigma_e^2$ for the recovery of interblock information. Although this estimation procedure is fairly simple to perform, the estimator itself is not the most efficient. Other methods have been proposed; see, for example, Shah (1964), Graybill and Weeks (1959), Graybill and Deal (1959), and Nelder (1968). The *best* estimation procedure, however, is a modified maximum likelihood procedure proposed by Patterson and Thompson (1971) for the general IBD, based on earlier work by Hartley and Rao (1967). The special feature of the Patterson–Thompson procedure is that it is applicable also to IBDs with block sizes not necessarily equal. Such cases may occur quite naturally in animal experiments when, for example, litters are considered to be blocks, which often are not of equal size, a situation considered also by Cunningham and Henderson (1968).

Patterson and Thompson (1971, p. 545, 546) describe the basic idea of their method as follows: “The contrasts among yields are divided into two sets: (i) contrasts between treatment totals; and (ii) contrasts with zero expectation, i.e. error contrasts. The method consists of maximizing the joint likelihood of all possible contrasts in set (ii). Contrasts in set (i) are excluded from the likelihood function on the grounds that, as long as treatment effects are regarded as unknown, fixed, as opposed to random, and without restraints, no contrast in set (i) can provide any information on error.” It is for this reason that the method has become known as *residual* or *restricted maximum likelihood* (REML). Computationally, it is more complex than the earlier methods, but as it has become part of several modern software packages (see Section 1.10.4), it has become the most used method (see Littell 2002). Its applicability goes much wider and beyond the analysis of experiments.

1.10.2 Complex Error Structures

The method of combining intra- and interblock information for the purpose of estimating treatment effects can be accomplished in one step by using for the observation in the ℓ th plot of the j th block a *mixed linear model* of the form

$$y_{j\ell} = \mu + \tau_{j\ell} + \beta_j + \epsilon_{j\ell},$$

where $j = 1, 2, \dots, b$; $\ell = 1, 2, \dots, k_j$; $\tau_{j\ell}$ denotes the treatment effect; β_j the block effect, which is assumed to be a random effect with mean 0 and variance σ_b^2 ; and $\epsilon_{j\ell}$ the error with mean 0 and variance σ_e^2 . This leads to a covariance structure for the observations within the j th block, with the variance–covariance matrix given by

$$V_j = \sigma_e^2 I_{k_j} + \sigma_b^2 J_{k_j}$$

(with I_{k_j} being the $k_j \times k_j$ identity matrix and J_{k_j} being the $k_j \times k_j$ matrix of unity elements) and zero correlation for observations in different blocks. This represents a covariance structure that, for purposes of the analysis, leads to the method of *generalized least squares*

requiring the estimation of two parameters, σ_e^2 and σ_b^2 . As an alternative method, this represents the prime example of REML estimation as developed by Patterson and Thompson (1971).

The combined analysis of an IBD as described above is an example of an analysis based on a *mixed linear model*. Another example is encountered sometimes in the context of agricultural field experiments. Even though blocking is being used to eliminate variability, additional variability, called *spatial variability*, may arise. Such variability can be eliminated through further conditions on the design (see Section 1.11.2) or can be dealt with through a special type of correlation structure for the error. One method was proposed by Papadakis (1937), referred to as the *nearest neighbor method*. Basically, the yields of neighboring plots in a single row, adjusted for treatment effects, are used as covariates. This method was improved and extended to two dimensions by Bartlett (1978). Further developments were proposed by Wilkinson et al. (1983) in the form of a *smooth trend plus independent error model*. In it, they assumed that trend effects were made small by taking second differences of the observations. Green et al. (1985) improved this method by using *least squares smoothing*. Gleeson and Cullis (1987) unified previous approaches to *spatial analysis* by advocating the use of the general class of *autoregressive integrated moving average* random processes using REML for variance and covariance parameter estimation. For a comprehensive review, see Stringer et al. (2012).

Another example of a complex error structure arises in the context of *repeated measures designs*. For example, in order to compare the effectiveness of different drugs to control blood pressure, different patients receive different drugs, and each patient's blood pressure is measured repeatedly over the course of a given time period. For a long time, these designs were routinely analyzed as split-plot designs (often referred to as *split-plot design in time*), even though it was realized that there may exist correlations between the repeated measures, in particular when they were close together in time or space (Finney 1990a). Geisser and Greenhouse (1958) showed that this analysis was correct if the correlation structure was that of *compound symmetry*, defined as the same correlation between any two measurements (as encountered in split-plot-type designs; see Section 1.5.5). Huynh and Feldt (1970) extended the Geisser–Greenhouse result to a more general correlation structure given by the Huynh–Feldt condition, whereby the (k, ℓ) element of the correlation matrix is given by

$$\sigma_{k\ell} = \lambda\delta_{k\ell} + \gamma_k + \gamma_\ell,$$

where $\delta_{k\ell} = 1$ if $k = \ell$ and $= 0$ otherwise, and γ_k and γ_ℓ are constants. These two conditions are quite special and not always realistic. The introduction of REML and its incorporation into software packages has made a routine application of different correlation structures, such as AR(1), or spatial power correlation or even unstructured correlation, possible (see Hinkelmann and Kempthorne 2008).

1.10.3 Graphical Methods

Graphical representations of data, either in their raw form or adjusted for certain extraneous effects, have always been important in the data analysis process. This has been made even easier with the advent of high-speed computers. In particular, *residual plots* are a useful tool for checking certain assumptions concerning the model that has been used for the analysis. In the case of spatial models, for example, such plots may reveal additional trends or

different trends in different blocks. They also may detect potential outliers that need to be checked.

Graphical methods may also be used when numerical methods are not available. Such a situation was considered by Daniel (1959) for the analysis of unreplicated saturated two-level experiments, when no d.f. are available for estimating the error variance. Consider a 2^{f-8} fractional factorial and suppose that the $m = 2^{f-8} - 1$ effects that can be estimated with equal precision are main effects and two-factor interactions. Even though we think that some of these effects are zero, we have no way of knowing which and we cannot test that hypothesis because we do not have an estimate of σ_e^2 . Daniel (1959) proposed a graphical procedure, the *half-normal plot*, to deal with this problem (see also Birnbaum 1959). The procedure consists of plotting the absolute values of the m estimates on the upper half of normal probability paper, which can now be done routinely with the help of computer packages. If, as assumed, most of the estimates are, in fact, estimating zero, then they should fall on a straight line through the origin. So-called guard rails (Daniel 1959; Zahn 1975) are used to judge the significance of those effects/interactions that deviate from this line. Then, from a new half-normal plot using only the nonsignificant effects, an estimate of σ_e^2 can be obtained (Zahn 1975). For a discussion of using half-normal plots in industrial applications, see Taylor (1994).

1.10.4 Statistical Software Packages

With the advent of modern computers, it did not take long before programs were written to perform standard statistical analyses. In the early 1970s, a group of statisticians at land-grant universities in the southern United States, called the University Statisticians at Southern Experiment Stations, began to undertake a project to write and compile such programs and make them generally available. Eventually, this project was transferred to a private company, and the resulting product became known as the Statistical Analysis System (SAS), with SAS Institute becoming the company's name. Many of the programs in SAS are either directly devoted to experimental design, in particular programs to construct particular designs, such as IBDs (PROC PLAN), systems of confounding, fractional factorials (PROC FACTEX), response surface designs (PROC RSREG), and optimal designs (PROC OPTEX), or indirectly as they can be used to analyze data from experimental and observational studies, such as the *General Linear Models Procedure* (PROC GLM) or the *Mixed Models Procedure* (PROC MIXED), which also incorporates REML estimation.

Concurrently with the development of SAS, a statistical software package was developed at Rothamsted Experimental Station under the direction of John Nelder, head of the statistics department. This program is called GENSTAT. Its primary objective was the analysis of data from designed experiments, but many other procedures were incorporated as well.

In addition to these two major software packages, other statistical packages with programs for designing and analyzing experiments are widely available. Among them are SPSS, BMDP, JMP, Minitab, Design-Expert, and R codes.

1.11 Areas of Application

As we have shown, many of the experimental designs have their origin in agricultural field experimentation. Over the course of time, however, these ideas have found application in

many other subject matter areas. Moreover, such applications as well as applications in some subdisciplines of agriculture have led to modifications of existing designs, new ideas, and new designs. In this section, we shall give some examples of such developments and references.

1.11.1 Genetics and Breeding

Diallel crosses, introduced by Schmidt (1919), represent an important tool in plant breeding and to a lesser extent in animal breeding. In such a program, a set of f females is crossed with a set of m males in all possible combinations. This mating design is also known as the North Carolina Design II (Comstock and Robinson 1952). For monoecious parents, where each parent has male *and* female reproductive units, this design has been discussed extensively by Griffing (1956). The goal is to assess the breeding potential of the individual lines through their *general combining abilities* (Sprague and Tatum 1942) and to investigate the type of gene actions present in the population through the estimation of genetic variance components. To this end, offspring from this mating design are grown in an environmental design, such as CRD, RCBD, and IBD, and the yield data of these offspring are then used to estimate relevant genetic parameters.

In some cases, it is either not practical or too costly to perform all matings. This has led to the development of incomplete mating designs, so-called partial diallel crosses (PDCs) (Gilbert 1958; Hinkelmann and Stern 1960; Kempthorne and Curnow 1961). The construction of PDCs has been based on an analogy with IBDs, in particular cyclic designs with blocks of size 2 (see Section 1.6.3), in the following way. For p parents, the complete diallel cross (CDC) (without self- and reciprocal crosses) leads to $p(p - 1)/2$ crosses. This corresponds to a BIB design with p treatments and $p(p - 1)/2$ blocks of size 2, the cross $i \times j$ corresponding to a block containing treatments i and j , and $\lambda = 1$. This suggests that a PDC corresponds to an m -associate PBIB design with p treatments, blocks of size 2, and parameters $\lambda_1, \lambda_2, \dots, \lambda_m$ either 1 (indicating that the corresponding associates are combined in a cross) or 0. The number of replicates per treatment is the number of times that an individual is being crossed, and the number of blocks is equal to the number of crosses (see Curnow 1963; Hinkelmann and Kempthorne 1963). The offspring from these crosses are grown in environmental designs, possibly IBDs (Agarwal and Das 1990; Gupta et al. 1995; Singh and Hinkelmann 1995). Several methods of constructing optimal mating-environmental designs have been proposed. For detailed discussion, see Singh et al. (2012).

The ideas of constructing PDCs have been extended to the construction of partial three-way and four-way crosses. For the case of three-way crosses, Hinkelmann (1965) established their correspondence with *generalized PBIB designs* as developed by Shah (1959) and provided appropriate construction methods.

Using microarray technology to study gene functions and biological processes, new designs had to be devised to perform and evaluate *microarray experiments*. Such experiments allow researchers to measure the expression levels of thousands of genes in biological samples. Microarray experiments are two-phase experiments (Kerr 2003). In the first phase, well-established designs are used, such as CRD and RCBD, depending on the question of interest. The second phase of a microarray experiment, which involves measuring the mRNA content of the tissue samples using microarrays (Nettleton 2012), has led to new design notions and to new uses of familiar designs, such as sequential designs (Durrieu and Briollais 2009) and factorial designs (Banerjee and Chatterjee 2008). An overview of appropriate designs for two-color microarray experiments can be found in Nettleton (2012).

1.11.2 Agronomy

Most agronomic field trials use one of the existing designs described in previous sections. There are, however, situations where special conditions require additional attention with regard to design and/or analysis. We refer here to trials where several crops are involved, either simultaneously or sequentially. Basically, they fall into one of three categories: *rotation experiments*, *competition experiments*, and *intercropping experiments*. Although different in several ways, these types of experiments bear a close relationship. For the most part, they are long-term experiments, and changes in external conditions may require changes in the experimental setup. Hence, “flexibility of design is therefore required” (Patterson 1965, p. 171).

Rotation experiments have existed for a long time, but their statistical aspects were discussed first by Cochran (1939) and Yates (1949, 1954), who gave definitions of various forms of such experiments. Design issues have to do with determining the number and type of replications, such as replications of crop phases within or over different years (Cochran 1939); the inclusion of auxiliary treatments, such as different nitrogen factors, for comparisons between rotations; and testing of hypotheses about the underlying causes of the effects of rotations (Patterson 1964). Special consideration must be given to plot sizes and shapes for the different crops (Dyke 1974). And there are a number of issues concerning the analysis (Yates 1954), such as whether or how to perform interim analyses, which type of variate to analyze, and how to deal with serial plot correlations (Patterson and Lowe 1970; Preece 1986). For details on designing and analyzing competition experiments, see also Street and Wilson (1985) and Goldberg and Scheiner (1993).

As Mead (1979) pointed out, competition experiments occur in different complexities, and in their most complex form, they are the same as intercropping experiments, that is, growing several crops (most often two) simultaneously on the same piece of land. And concerning intercropping, an age-old farming practice, Finney (1990b, p. 74) writes: “To applied statisticians, data from experiments on intercropping are intrinsically interesting; they introduce variants of problems of multivariate analysis, parameter estimation, and experimental design.” He continues (p. 81) to state three principles of design for research on intercropping with two crops, *A* and *B*, say: “(i) The treatments under test must include sole crops as well as intercrops unless *either* the merits of intercropping are already so firmly established that sole crops need not even be discussed *or* tradition so strongly favours intercropping that no research findings will change it; (ii) for there to be any basis for discussing the optimal ratio between the areas occupied by *A* and *B*, intercropping experiments must include at least two such ratios (say 2:1 and 1:2) in addition to the sole crops; (iii) if the argument for intercropping is to involve insurance against weather or other disasters to one component, experiments must be conducted in several years.” Suitable designs are discussed by Federer (1993, 1999) for two and three or more crops, respectively. Intercropping data are essentially multivariate, but the discussion continues whether the use of various forms of (univariate) *land equivalent ratios* (LERs) is appropriate, where an LER is a measure of the efficiency of an intercrop in terms of the land area required under sole cropping to give the yields from the individual crops (see Mead and Riley 1981).

1.11.3 Medicine and Pharmacology

“A *clinical trial* is an experiment testing medical treatments on human subjects” (Piantadosi 1997, p. 10), reiterating the statement by Hill (1963, p. 1044): “In the assessment of a treatment, medicine always has proceeded, and always must proceed, by way of experiment.

The experiment may merely consist in giving the treatment to a particular patient or series of patients, and of observing and recording what follows—with all the difficulty of interpretation, of distinguishing the *propter hoc* from the *post hoc*.” In principle, the designs for randomized controlled clinical trials are usually quite simple, but they are made more complicated by the fact that human beings are involved who have to give their consent to be included in the trial and that such trials are subject to regulatory conditions and monitored by data safety monitoring boards (see Wittes and Yu 2012). Also the use of randomization has not been without controversy (Rosenberger and Lachin 2002) as it involves ethical issues. Clinical trials become even more complex since they usually are of a sequential nature and often subject to monitoring. This leads to the use of *adaptive designs* and *group sequential methods* when repeated interim analyses lead to changes in the trial design (see Chow and Chang 2011; Leifer and Geller 2012; Yin 2012) and to the use of *adaptive randomization procedures*, such as Efron’s (1971) biased coin design (see Zhang and Rosenberger 2012 and Chapter 24).

A different kind of experiment, relative to those we have discussed so far, is that of *biological assay* or *bioassay*. This area is associated with the names of Bliss, Fieller, and Finney, with, among many other things, introducing the method of probit analysis (Bliss 1934), describing theory and practice of assays (Bliss and Cattell 1943), formulating what has become known as Fieller’s theorem for making inference in bioassay (Fieller 1940), and giving the first systematic account of the statistical principles for bioassay (Finney 1947). In the words of Finney (1952, 1971, p. 1) “an assay is [thus] a form of biological experiment, but the interest lies in comparing the potencies of treatments on an agreed scale instead of in comparing the magnitudes of effects of different treatments.” A stimulus at different doses is given to subjects, and the subsequent response of a measurable characteristic of the subject, as a function of the dose, is of interest for comparing the potencies. The potency is being obtained relative to a standard preparation of the stimulus in a direct or indirect assay. Typically the dose–response relationship is modeled as a dose–response regression, which may be linear or nonlinear. Designs are then required to estimate the dose–response relationship for two major types of assays: *parallel line assays* and *slope ratio assays*. For parallel line assays, one considers designs with k nonzero doses, each for the standard and test preparations, referred to as (k, k) designs. Analogous designs for slope ratio assays, referred to as $(1, k, k)$ designs, typically have additional observations at zero dose. Incomplete block designs and crossover designs play an important role in this context. An interesting example is a twin crossover design (Smith et al. 1944; Finney 1955), given in Table 1.18, where in a basic design four subjects are given two doses from two standard and test preparations, S_1 , S_2 , T_1 , and T_2 .

TABLE 1.18
Twin Crossover Design

Subject	Dose on Occasion	
	1	2
1	S_1	T_2
2	S_2	T_1
3	T_1	S_2
4	T_2	S_1

For a detailed discussion of available designs and types of analyses, see Finney (1978). For more recent developments concerning aspects of design, see for example, Smith and Ridout (2003), Hanusz and Rutkowska (2006), and Donev et al. (2008).

1.11.4 Marketing

Stated choice or *stated preference experiments* have become accepted tools in marketing research as well as in other areas where one wants to obtain information about products or services that do not yet exist and where people are presented with options to choose from. More specifically, “a stated choice experiment consists of a set of choice sets. Each choice set consists of two or more options (or alternatives). Each respondent is shown each choice set in turn and asked to choose the option they think is best (or worst) from among the options presented in the choice set” (Street and Burgess 2012, p. 332).

A particular, and perhaps most commonly used, choice experiment is the so-called generic stated preference choice experiment, in which it is assumed that all options in each choice set, other than the *none of these* option, if present, are described by the same set of attributes and each attribute can take one level from a finite set of possible levels. This last statement suggests that attributes correspond to the factors of a factorial design and the levels of the attributes correspond to the levels of those factors. An option in a choice set corresponds then to a particular factor-level combination. All possible options would therefore correspond to the complete set of factor-level combinations. It is, obviously, impracticable to present a responder with all those options. An obvious way to proceed then is to use fractional factorials, and this is what Street and Burgess (2012) use to construct optimal choice designs. This is made more complicated, however, since the attributes usually have different numbers of levels, which leads to asymmetric designs. For a detailed discussion of choice experiments, see Louviere et al. (2000), Street and Burgess (2007, 2012), and Chapter 22.

1.11.5 Manufacturing and Product Development

Statistical methods, in particular experimental designs, have long been used to determine *optimal* manufacturing conditions and properties; see, for example, RSM (Section 1.8.2). A new philosophy, however, was introduced by Taguchi (1986). He introduced what is now referred to as *robust parameter design*. The basic idea is that there are two types of factors that operate on a system: *control factors* and *noise factors*. Control factors are set at a certain level and remain unchanged during the process, whereas noise factors may change randomly during the process. Very often the noise factors are environmental factors and hence capable of causing unwanted variation in the outcome of the process. “The goal of robust parameter design is to design the process in such a way as to operate at the levels of the control variables, that is parameters, that make the process as insensitive as possible to the random fluctuations of the noise factors” (Robinson and Anderson-Cook 2012, p. 444).

There are different approaches to achieving that goal. Taguchi’s (1986) main influence has been to consider not only the mean of the process but also the variance. To this end, he combined the sample mean and sample variance into a single performance measure, which he called the *signal-to-noise ratio* (SNR). There are a number of possible SNR, and the appropriate choice depends on the goal of the experiment (see, e.g., Robinson and Anderson-Cook 2012). The SNR is then used to construct an optimality criterion.

Taguchi combined fractional factorials in what he called inner and outer arrays to produce optimal designs according to his criterion. This approach has been criticized for a number of reasons (see e.g., Box 1988; Nair 1992), and other approaches were developed with respect to the underlying design and the modeling of the data. Some of these approaches make use of response surface methods. In dual response modeling (Box 1988; Vining and Myers 1990), two separate models are developed: one for the process mean as determined by the control factors and the other for the process variance as determined by the noise factors. For this approach, modifications to the crossed array designs were suggested by Steiner and MacKay (2005) and Asilahijani et al. (2010). As an alternative to the dual-response modeling, Welch et al. (1990) and Myers et al. (1992) proposed to use combined arrays together with a single model response approach. Special forms of combined array designs are mixed resolution and composite designs (Borkowski and Lucas 1997). For a detailed discussion, see Ankenman and Dean (2003), Bursztyn and Steinberg (2006), and Robinson and Anderson-Cook (2012).

1.12 Epilogue

Experimental design is an important feature of empirical research and industrial investigations. We have shown how the development started with simple agronomic trials, which had practical intent and scientific foundation, but no statistical validity. Nevertheless, these experiments were the building blocks from which Fisher and Yates began to formulate and develop the basic principles of experimental design within a statistical framework. These principles are still valid today and form the foundation of experimental design. Throughout this chapter, we have commented and elaborated on many developments that have taken place over almost 100 years. Similar reviews have been presented by Cox and Reid (2000) along more technical lines and by Atkinson and Bailey (2001) distinguishing between agricultural and industrial experiments. Ideas of reducing experimental error through more sophisticated *error-control designs* (Hinkelmann and Kempthorne 2008); of increasing the inference space through more complex *treatment designs*, such as various types of factorial designs; and through the inclusion of subject-specific factors, referred to by Cox (1984) as *intrinsic factors*, have been implemented over the years. Important strides have also been made with respect to the analysis of data, in particular from nonorthogonal designs. We have touched on many of these aspects, but it is obviously impossible to capture all the developments in a single chapter. Many of the topics, however, will be discussed in subsequent chapters in much more detail, including some of the latest results. Throughout this introductory chapter, we have made references to these chapters as well as to pertinent contributions in the vast literature.

Other topics, which have not even been mentioned here, will be presented in some chapters of this handbook, for example, designs for nonlinear models (Chapters 12 through 14). In some cases they deal with fairly new topics. To tackle new problems, however, requires an understanding of past developments. At the same time we must keep in mind one reviewer's comment: "Challenges in design and analysis of experiments change with changing needs in the sciences, and we are in an era in which rapid changes in technology make some of today's hot challenges tomorrow's history." In any case, experimental design is a constantly evolving part of statistical science and practice.

References

- Addelman, S. (1962). Symmetrical and asymmetrical fractional factorial plans. *Technometrics*, **4**, 47–58.
- Addelman, S. and O. Kempthorne (1961). Some main effect plans and orthogonal arrays of strength two. *Ann. Math. Stat.*, **32**, 1167–1176.
- Agarwal, S. C. and M. N. Das (1990). Use of n -ary block designs in diallel cross evaluation. *J. Appl. Stat.*, **17**, 125–131.
- Ankenman, B. E. and A. M. Dean (2003). Quality improvement and robustness via design of experiments. In *Handbook of Statistics*, Vol. 22 (C. R. Rao and R. Khattree, eds.), pp. 263–317. Amsterdam, the Netherlands: Elsevier.
- Asilahijani, H., S. H. Steiner, and R. J. MacKay (2010). Reducing variation in an existing process with robust parameter design. *Qual. Eng.*, **22**, 30–45.
- Atkinson, A. C., and R. A. Bailey (2001). One hundred years of the design of experiments on and off the pages of *Biometrika*. *Biometrika*, **88**, 53–97.
- Atkinson, A. C., A. N. Donev, and R. D. Tobias (2007). *Optimum Experimental Designs, with SAS*. Oxford, U.K. Oxford University Press.
- Banerjee, T. and R. Mukerjee (2008). Optimal factorial designs for CDNA microarray experiments. *Ann. Appl. Stat.*, **2**, 366–385.
- Bartlett, M. S. (1978). Nearest neighbour models in the analysis of field experiments. *J. R. Stat. Soc., B*, **40**, 147–158.
- Bechhofer, R. E. and A. C. Tamhane (1981). Incomplete block designs for comparing treatments with a control: General theory. *Technometrics*, **23**, 45–57 (Corrigendum **24**, 171).
- Bechhofer, R. E. and A. C. Tamhane (1985). Tables of admissible and optimal balanced treatment incomplete block designs. *Sel. Tables Math. Stat.*, **8**, 41–139.
- Beyer, W. H. (ed.) (1991). *CRC Standard Probability and Statistics Tables and Formulae*. Boca Raton, FL: CRC Press, Inc.
- Birnbaum, A. (1959). On the analysis of factorial experiments without replication. *Technometrics*, **1**, 343–357.
- Bisgaard, S. (1992). Industrial use of statistically designed experiments: Case study references and some historical anecdotes. *Qual. Eng.*, **4**, 547–562.
- Bliss, C. I. (1934). The method of probits. *Science*, **79**, 38–39.
- Bliss, C. I. and McK. Cattell (1943). Biological assay. *Ann. Rev. Physiol.*, **5**, 479–539.
- Booth, K. H. V. and D. R. Cox (1962). Some systematic supersaturated designs. *Technometrics*, **4**, 489–495.
- Borkowski, J. J. and J. M. Lucas (1997). Designs of mixed resolution for process robustness studies. *Technometrics*, **39**, 63–70.
- Bose, R. C. (1938). On the application of the properties of Galois fields to the problem of construction of hyper-Græco-Latin squares. *Sankhyā*, **3**, 323–338.
- Bose, R. C. (1939). On the construction of balanced incomplete block designs. *Ann. Eugen.*, **9**, 353–399.
- Bose, R. C. (1942). A note on the resolvability of balanced incomplete block designs. *Sankhyā*, **6**, 105–110.
- Bose, R. C. (1947). Mathematical theory of the symmetrical factorial design. *Sankhyā*, **8**, 107–166.
- Bose, R. C., W. H. Clatworthy, and S. S. Shrikhande (1954). Tables of partially balanced incomplete block designs with two associate classes. *N. C. Agric. Exp. Station Tech. Bull.*, **107**, 1–255.
- Bose, R. C. and K. R. Nair (1939). Partially balanced incomplete block designs. *Sankhyā*, **4**, 337–372.
- Bose, R. C. and K. R. Nair (1962). Resolvable incomplete block designs with two replications. *Sankhyā A*, **24**, 9–24.
- Bose, R. C. and T. Shimamoto (1952). Classification and analysis of partially balanced incomplete block designs with two associate classes. *J. Am. Stat. Assoc.*, **47**, 151–184.

- Bose, R. C., S. S. Shrikhande, and E. T. Parker (1960). Further results on the construction of mutually orthogonal Latin squares and the falsity of Euler's conjecture. *Can. J. Math.*, **12**, 189–203.
- Box, G. E. P. (1952). Multi-factor designs of first order. *Biometrika*, **39**, 49–57.
- Box, G. E. P. (1988). Signal-to-noise ratios, performance criteria, and transformations. *Technometrics*, **30**, 1–17.
- Box, G. E. P. and D. W. Behnken (1960). Some new three level designs for the study of quantitative variables. *Technometrics*, **2**, 455–475 (Errata **3**, 131, 576).
- Box, G. E. P. and N. R. Draper (2007). *Response Surfaces, Mixtures, and Ridge Analysis* (2nd edition). New York: Wiley.
- Box, G. E. P. and J. S. Hunter (1957). Multi-factor experimental designs for exploring response surfaces. *Ann. Math. Stat.*, **28**, 195–241.
- Box, G. E. P. and J. S. Hunter (1961a). The 2^{k-p} fractional factorial designs. Part I. *Technometrics*, **3**, 311–351 (Errata **5**, 417).
- Box, G. E. P. and J. S. Hunter (1961b). The 2^{k-p} fractional factorial designs. Part II. *Technometrics*, **3**, 449–458.
- Box, G. E. P., J. S. Hunter, and W. G. Hunter (2005). *Statistics for Experimenters: Design, Innovation, and Discovery* (2nd edition). Hoboken, NJ: Wiley.
- Box, G. E. P. and R. D. Meyer (1986). An analysis for unreplicated fractional factorials. *Technometrics*, **28**, 11–18.
- Box, G. E. P. and K. B. Wilson (1951). On the experimental attainment of optimum conditions (with discussion). *J. R. Stat. Soc. B*, **13**, 1–45.
- Bradley, J. V. (1958). Complete counterbalancing of immediate sequential effects in a Latin square design. *J. Am. Stat. Assoc.*, **53**, 525–528.
- Bradley, R. A. and C.-M. Yeh (1980). Trend-free block designs: Theory. *Ann. Stat.*, **8**, 883–893.
- Brien, C. J. and R. A. Bailey (2006). Multiple randomizations (with discussion). *J. R. Stat. Soc. B*, **68**, 571–609.
- Brien, C. J. and R. W. Payne (1999). Tier, structure formulae and the analysis of complicated experiments. *Statistician*, **48**, 41–52.
- Bursztyn, D. and D. M. Steinberg (2006). Screening experiments for dispersion effects. In *Screening: Methods for Experimentation in Industry, Drug Discovery, and Genetics* (A. M. Dean and S. M. Lewis, eds.), pp. 21–47. New York: Springer.
- Butler, N. A., R. Mead, K. M. Eskridge, and S. G. Gilmour (2001). A general method of constructing $E(s^2)$ -optimal supersaturated designs. *J. R. Stat. Soc. B*, **63**, 621–632.
- Calinski, T. and S. Kageyama (2000). *Block Designs: A Randomization Approach*. Vol. 1: Analysis. New York: Springer.
- Calinski, T. and S. Kageyama (2003). *Block Designs: A Randomization Approach*. Vol. 2: Design. New York: Springer.
- Chen, H. H. and C.-S. Cheng (2012). Minimum aberration and related criteria for fractional factorial designs. In *Design and Analysis of Experiments*. Vol. 3: Special Designs and Applications (K. Hinkelmann, ed.), pp. 299–329. Hoboken, NJ: Wiley.
- Chow, S.-C. and M. Chang (2011). *Adaptive Design Methods in Clinical Trials* (2nd edition). London, U.K.: Chapman & Hall/CRC.
- Claringbold, P. J. (1955). Use of the simplex design in the study of joint action of related hormones. *Biometrics*, **11**, 174–185.
- Clatworthy, W. H. (1973). Tables of two-associate class partially balanced designs. *National Bureau of Standards Appl. Math. Ser.*, **63**.
- Cochran, W. G. (1939). Long-term agricultural experiments (with discussion). *J. R. Stat. Soc., Suppl.*, **6**, 104–148.
- Cochran, W. G., K. M. Autrey, and C. Y. Cannon (1941). A double change-over design for dairy cattle feeding experiments. *J. Dairy Sci.*, **24**, 937–951.
- Cochran, W. G. and G. M. Cox (1950). *Experimental Designs*. New York: Wiley.
- Cochran, W. G. and G. M. Cox (1957). *Experimental Designs* (2nd edition). New York: Wiley.

- Comstock, R. E. and H. F. Robinson (1952). Estimation of average dominance of genes. In *Heterosis* (J. W. Gowen, ed.), pp. 494–516. Ames, IA: Iowa State College Press.
- Coons, I. (1957). The analysis of covariance as a missing plot technique. *Biometrics*, **13**, 387–405.
- Cornell, J. A. (1975). Some comments on designs for Cox's mixture polynomial. *Technometrics*, **17**, 25–35.
- Cornell, J. (2002). *Experiments with Mixtures* (3rd edition). New York: Wiley.
- Cornell, J. (2011). *A Primer on Experiments with Mixtures*. Hoboken, NJ: Wiley.
- Cox, D. R. (1951). Some systematic experimental designs. *Biometrika*, **38**, 312–323.
- Cox, D. R. (1958). *Planning of Experiments*. New York: Wiley.
- Cox, D. R. (1984). Interaction (with discussion). *Int. Stat. Rev.*, **52**, 1–32.
- Cox, D. R. and N. Reid (2000). *The Theory of the Design of Experiments*. Boca Raton, FL: Chapman & Hall/CRC.
- Cunningham, E. P. and C. R. Henderson (1968). An iterative procedure for estimating fixed effects and variance components in mixed model situations. *Biometrics*, **24**, 13–25.
- Curnow, R. N. (1963). Sampling the diallel cross. *Biometrics*, **19**, 287–306.
- Daniel, C. (1959). Use of the half-normal plots in interpreting factorial two-level experiments. *Technometrics*, **1**, 311–341.
- Daniel, C. (1973). One-at-a-time plans. *J. Am. Stat. Assoc.*, **68**, 353–360.
- Das, M. N. (1958). On reinforced incomplete block designs. *J. Ind. Soc. Agric. Stat.*, **10**, 73–77.
- David, H. A. (1963). *The Method of Paired Comparisons*. London, U.K.: Griffin.
- David, H. A. (1965). Enumeration of cyclic paired comparison designs. *Am. Math. Monthly*, **72**, 241–248.
- David, H. A. (1967). Resolvable cyclic designs. *Sankhyā A*, **29**, 191–198.
- David, H. A. and F. W. Wolock (1965). Cyclic designs. *Ann. Math. Stat.*, **36**, 1526–1534.
- Davies, O. L. (ed.) (1956). *The Design and Analysis of Industrial Experiments*. London, U.K.: Oliver and Boyd.
- Dean, A. M. and J. A. John (1975). Single replication factorial experiments in cyclic designs: II. Asymmetrical arrangements. *J. R. Stat. Soc. B*, **37**, 72–76.
- Dean, A. and D. Voss (1999). *Design and Analysis of Experiments*. New York: Springer.
- Donev, A. N., R. Tobias, and F. Monadjemi (2008). Cost-cautious designs for confirmatory bioassay. *J. Stat. Plann. Infer.*, **138**, 3805–3812.
- Draper, N. R. and H. Smith (1998). *Applied Regression Analysis* (3rd edition). New York: Wiley.
- Durrieu, G. and L. Briollais (2009). Sequential design for microarray experiments. *J. Am. Stat. Assoc.*, **104**, 650–660.
- Dyke, G. V. (1974). *Comparative Experiments with Field Crops*. London, U.K.: Butterworths.
- Eden, T. and R. A. Fisher (1927). Studies in crop variation. IV. The experimental determination of the value of top dressing with cereals. *J. Agric. Sci.*, **17**, 548–562.
- Efron, B. (1971). Forcing a sequential experiment to be balanced. *Biometrika*, **58**, 403–417.
- Ehrenfeld, S. (1953). On the efficiency of experimental design. *Ann. Math. Stat.*, **26**, 247–255.
- Fang, K.-T., D. K. J. Lin, and C.-X Ma (2000). On the construction of multi-level supersaturated designs. *J. Stat. Plann. Infer.*, **86**, 239–252.
- Federer, W. T. (1955). *Experimental Design—Theory and Application*. New York: Macmillan.
- Federer, W. T. (1993). *Statistical Design and Analysis for Intercropping Experiments*. Vol. I: Two Crops. New York: Springer.
- Federer, W. T. (1999). *Statistical Design and Analysis for Intercropping Experiments*. Vol. II: Three or More Crops. New York: Springer.
- Federer, W. T. and L. N. Balaam (1973). *Bibliography on Experiment and Treatment Design Pre 1968*. New York: Hafner.
- Fieller, E. C. (1940). The biological standardization of insulin. *J. R. Stat. Soc. Suppl.*, **7**, 1–64.
- Finney, D. J. (1945). The fractional replication of factorial experiments. *Ann. Eugen.*, **12**, 291–301 (Correction **15**, 276).
- Finney, D. J. (1947). The principles of biological assay. *J. R. Stat. Soc. Suppl.*, **9**, 46–91.
- Finney, D. J. (1952). *Statistical Method in Biological Assay*. London, U.K.: Griffin.

- Finney, D. J. (1955). *Experimental Design and Its Statistical Basis*. Chicago, IL: University of Chicago Press.
- Finney, D. J. (1971). *Statistical Method in Biological Assay* (3rd edition). New York: Hafner.
- Finney, D. J. (1990a). Repeated measurements: What is measured and what repeats? *Stat. Med.*, **9**, 639–644.
- Finney, D. J. (1990b). Intercropping experiments, statistical analysis, and agricultural practice. *Exp. Agric.*, **26**, 73–81.
- Fisher, R. A. (1918). The correlation between relatives on the supposition of Mendelian inheritance. *Trans. R. Soc. Edinburgh*, **52**, 399–433.
- Fisher, R. A. (1921). Studies in crop variation. *J. Agric. Sci.*, **11**, 107–135.
- Fisher, R. A. (1925). *Statistical Methods for Research Workers*. Edinburgh, U.K.: Oliver and Boyd.
- Fisher, R. A. (1926). The arrangement of field experiments. *J. Ministry Agric.*, **33**, 503–513.
- Fisher, R. A. (1935). *Design of Experiments*. Edinburgh, U.K.: Oliver and Boyd.
- Fisher, R. A. (1942). The theory of confounding in factorial experiments in relation to the theory of groups. *Ann. Eugen.*, **11**, 341–353.
- Fisher, R. A. (1945). A system of confounding for factors with more than two alternatives, giving completely orthogonal cubes and higher powers. *Ann. Eugen.*, **12**, 283–290.
- Fisher, R. A. (1970). *Statistical Methods for Research Workers* (14th edition). Darien, CT: Hafner.
- Fisher, R. A. (1971). *Design of Experiments* (9th edition). New York: Hafner.
- Fisher, R. A. and F. Yates (1938). *Statistical Tables for Biological, Agricultural and Medical Research*. Edinburgh, U.K.: Oliver and Boyd.
- Fisher, R. A. and F. Yates (1957). *Statistical Tables for Biological, Agricultural and Medical Research* (5th edition). New York: Hafner Publishing & Co.
- Fisher Box, J. (1978). *R. A. Fisher: The Life of a Scientist*. New York: Wiley.
- Fries, A. and W. G. Hunter (1980). Minimum aberration 2^{k-p} designs. *Technometrics*, **22**, 601–608.
- Geisser, S. and S. Greenhouse (1958). An extension of Box's results on the use of the F -distribution in multivariate analysis. *Ann. Math. Stat.*, **29**, 885–891.
- Georgiou, S. D., D. Draguljić, and A. M. Dean (2009). An overview of two-level supersaturated designs with cyclic structure. *J. Stat. Theory Pract.*, **3**, 489–504.
- Georgiou, S. and C. Koukouvinos (2006). Multi-level k -circulant supersaturated designs. *Metrika*, **64**, 209–220.
- Giesbrecht, F. G. and M. L. Gumpertz (2004). *Planning, Construction, and Statistical Analysis of Comparative Experiments*. Hoboken, NJ: Wiley.
- Ghosh, S. and L. Teschmacher (2002). Comparison of search designs using search probabilities. *J. Stat. Plann. Infer.*, **104**, 439–458.
- Ghosh, S. and Y. Tian (2006). Optimum two-level fractional factorial plans for model identification and discrimination. *J. Multivariate Anal.*, **97**, 1437–1450.
- Gilbert, N. (1958). Diallel cross in plant breeding. *Heredity*, **12**, 477–492.
- Gleason, A. C. and B. R. Cullis (1987). Residual maximum likelihood (REML) estimation of a neighbour model for field experiments. *Biometrics*, **43**, 277–288.
- Goldberg, D. E. and S. M. Scheiner (1993). ANOVA and ANCOVA: Field competition experiments. In *Design and Analysis of Ecological Experiments* (S. M. Scheiner and J. Gurevitch, eds.), pp. 69–93. London, U.K.: Chapman & Hall.
- Goldfarb, H. B., C. M. Anderson-Cook, C. M. Borror, and D. C. Montgomery (2004). Fraction of design space plots for assessing mixture and mixture-process design. *J. Qual. Technol.*, **36**, 169–179.
- Graybill, F. A. and R. B. Deal (1959). Combining unbiased estimators. *Biometrics*, **15**, 543–550.
- Graybill, F. A. and D. L. Weeks (1959). Combining inter-block and intra-block information in balanced incomplete block designs. *Ann. Math. Stat.*, **30**, 799–805.
- Green, P., C. Jennison, and A. Seheult (1985). Analysis of field experiments by least squares smoothing. *J. R. Stat. Soc. B*, **47**, 299–315.
- Griffing, B. (1956). Concepts of general and specific combining ability in relation to diallel crossing systems. *Aust. J. Biol. Sci.*, **9**, 463–493.

- Gupta, S., A. Das, and S. Kageyama (1995). Single replicate orthogonal block designs for circulant partial diallel crosses. *Commun. Stat. Theory Methods* **24**, 2601–2607.
- Gupta, S., K. Hisano, and L. B. Morales (2011). Optimal k -circulant supersaturated designs. *J. Stat. Plann. Infer.*, **141**, 782–786.
- Gupta, V. K., R. Parsad, L. M. Bhar, and B. Kole (2008). Supersaturated designs for asymmetrical factorial experiments. *J. Stat. Theory Pract.*, **2**, 95–108.
- Gupta, V. K., P. Singh, B. Kole, and R. Parsad (2010). Computer aided construction of efficient multi-level supersaturated designs. *J. Stat. Theory Pract.*, **4**, 221–231.
- Hamaker, H. C. (1955). Experimental design in industry. *Biometrics*, **11**, 257–286.
- Hanusz, Z. and A. Rutkowska (2006). Comparison of several test preparations with one standard in multivariate bioassay. *Biom. Lett.*, **43**, 99–107.
- Harshbarger, B. (1947). Rectangular lattices. *Virginia Agric. Exp. Station Memoir*, **1**, 1–26.
- Harshbarger, B. (1949). Triple rectangular lattices. *Biometrics*, **5**, 1–13.
- Harshbarger, B. (1951). Near balance rectangular lattices. *Virginia J. Sci.*, **2**, 13–27.
- Hartley, H. O. and C. A. B. Smith (1948). The construction of Youden squares. *J. R. Stat. Soc. B*, **10**, 262–263.
- Hartley, H. O. and J. N. K. Rao (1967). Maximum likelihood estimation for the mixed analysis of variance model. *Biometrika*, **54**, 93–108.
- Hedayat, A. S., M. Jacroux, and D. Majumdar (1988). Optimal designs for comparing test treatments with controls (with discussion). *Stat. Sci.*, **3**, 462–491.
- Hedayat, A. S., N. J. A. Sloane, and J. Stufken (1999). *Orthogonal Arrays: Theory and Applications*. New York: Springer.
- Hill, A. B. (April 1963). Medical ethics and controlled trials. *Br. Med. J.*, **20**, 1043–1049.
- Hinkelmann, K. (1964). Extended group divisible partially balanced incomplete block designs. *Ann. Math. Stat.*, **35**, 681–695.
- Hinkelmann, K. (1965). Partial diallel crosses. *Sankhyā A*, **27**, 173–196.
- Hinkelmann, K. and J. S. Alcorn (1998). Randomization analysis of replicated complete block designs. *J. Combin. Inform. Syst. Sci.*, **23**, 317–332.
- Hinkelmann, K. and O. Kempthorne (1963). Two classes of group divisible partial diallel crosses. *Biometrika*, **50**, 281–291.
- Hinkelmann, K. and O. Kempthorne (2005). *Design and Analysis of Experiments, Vol. 2: Advanced Experimental Design*. Hoboken, NJ: Wiley.
- Hinkelmann, K. and O. Kempthorne (2008). *Design and Analysis of Experiments, Vol. 1: Introduction to Experimental Design (2nd edition)*. Hoboken, NJ: Wiley.
- Hinkelmann, K. and K. Stern (1960). Kreuzungspläne zur Selektionszüchtung von Waldbäumen. *Silvae Genet.*, **9**, 121–133.
- Hoblyn, T. N., S. C. Pearce, and G. H. Freeman (1954). Some considerations in the design of successive experiments in fruit plantations. *Biometrics*, **10**, 503–515.
- Huynh, H. and L. S. Feldt (1970). Conditions under which mean square ratios in repeated measurements design have exact F -distributions. *J. Am. Stat. Assoc.*, **65**, 1582–1589.
- Isaac, P. D., A. M. Dean, and T. Ostrom (2001). Sequentially counterbalanced Latin squares. *Stat. Appl.*, **3**, 25–46.
- Jacroux, M. (1998). On the determination and construction of E -optimal block designs in the presence of trends. *J. Stat. Plann. Infer.*, **67**, 331–342.
- Jacroux, M. (2002). On the determination and construction of A - and MV -optimal block designs for comparing a set of test treatments to a set of standard treatments. *J. Stat. Plann. Infer.*, **106**, 191–204.
- Jacroux, M. (2003). Some MV -optimal group divisible type block designs for comparing a set of test treatments with a set of standard treatments. *J. Stat. Plann. Infer.*, **113**, 597–615.
- Jacroux, M., D. Majumdar, and R. K. Shah (1995). Efficient block designs in the presence of trends. *Stat. Sin.*, **5**, 605–615.
- Jo, J. and K. Hinkelmann (1993). Some properties of Box-Behnken designs. *J. Combin. Inform. Syst. Sci.*, **18**, 273–287.

- John, J. A. (1966). Cyclic incomplete block designs. *J. R. Stat. Soc. B*, **28**, 345–360.
- John, J. A. (1969). A relationship between cyclic and PBIB designs. *Sankhyā B*, **31**, 535–540.
- John, J. A. and E. R. Williams (1995). *Cyclic and Computer Generated Designs* (2nd edition). London, U.K.: Chapman & Hall.
- John, J. A., F. W. Wolock, and H. A. David (1972). Cyclic designs. *Appl. Math. Ser.*, **62**, National Bureau of Standards, Washington, DC.
- Jones, B. and M. G. Kenward (2003). *Design and Analysis of Cross-Over Trials* (2nd edition). London, U.K.: Chapman & Hall/CRC.
- Kempthorne, O. (1947). A simple approach to confounding and fractional replication in factorial experiments. *Biometrika*, **34**, 255–272.
- Kempthorne, O. (1952). *Design and Analysis of Experiments*. New York: Wiley.
- Kempthorne, O. (1953). A class of experimental designs using blocks of two plots. *Ann. Math. Stat.*, **24**, 76–84.
- Kempthorne, O. (1955). The randomization theory of experimental inference. *J. Am. Stat. Assoc.*, **50**, 946–967.
- Kempthorne, O. (1956). The efficiency factor of an incomplete block design. *Ann. Math. Stat.*, **27**, 846–849.
- Kempthorne, O. and R. N. Curnow (1961). The partial diallel cross. *Biometrics*, **17**, 229–250 (Correction: **18**, 128).
- Kempthorne, O. and W. T. Federer (1948a). The general theory of prime-power lattice designs. I. Introduction and designs for p^n varieties in blocks of p plots. *Biometrics*, **4**, 54–79.
- Kempthorne, O. and W. T. Federer (1948b). The general theory of prime-power lattice designs. II. Designs for p^n varieties in blocks of p^s plots and in squares. *Biometrics*, **4**, 109–121.
- Kerr, M. K. (2003). Design considerations for efficient and effective microarray studies. *Biometrics*, **59**, 822–828.
- Khuri, A. I. and J. A. Cornell (1996). *Response Surfaces: Designs and Analyses* (2nd edition). New York: Marcel Dekker.
- Kiefer, J. (1958). On the nonrandomized optimality and randomized nonoptimality of symmetrical designs. *Ann. Math. Stat.*, **29**, 675–699.
- Kiefer, J. (1959). Optimum experimental designs (with discussion). *J. R. Stat. Soc. B*, **21**, 272–319.
- Kiefer, J. (1975a). Construction and optimality of generalized Youden designs. In *A Survey of Statistical Design and Linear Models* (J. N. Srivastava, ed.), pp. 333–353. Amsterdam, the Netherlands: North-Holland.
- Kiefer, J. (1975b). Balanced block designs and generalized Youden squares. *Ann. Stat.*, **3**, 109–118.
- Kole, B., J. Gangwani, V. K. Gupta, and R. Parsad (2010). Two-level supersaturated designs: A review. *J. Stat. Theory Pract.*, **4**, 589–608.
- Kshirsagar, A. M. (1966). Balanced factorial designs. *J. R. Stat. Soc. B*, **28**, 559–567.
- Kusumoto, K. (1965). Hyper-cubic designs. *Wakayama Med. Rep.*, **9**, 123–132.
- Leifer, E. S. and N. L. Geller (2012). Monitoring randomized clinical trials. In *Design and Analysis of Experiments*. Vol. 3: Special Designs and Applications (K. Hinkelmann, ed.), pp. 213–249. Hoboken, NJ: Wiley.
- Li, J. C. R. (1944). Design and statistical analysis of some confounded factorial experiments. *Iowa Agric. Exp. Station Res. Bull.*, **333**, 452–492.
- Lin, D. K. J. (1993). A new class of supersaturated designs. *Technometrics*, **35**, 28–31.
- Lin, C. D., D. Bingham, R. R. Sitter, and B. Tang (2010). A new and flexible method for constructing designs for computer experiments. *Ann. Stat.*, **38**, 1460–1477.
- Lin, W.-C. and J. Stufken (1999). On finding trend-free block designs. *J. Stat. Plann. Infer.*, **78**, 57–70.
- Lind, J. (1753). *A Treatise of the Scurvy*. In *Three Parts*. Edinburgh, U.K.: Sands, Murray and Cochran.
- Littell, R. C. (2002). Analysis of unbalanced mixed model data: A case study comparison of ANOVA versus REML/GLS. *J. Agric., Biol. Environ. Stat.*, **7**, 472–490.
- Liu, M.-Q. and D. K. J. Lin (2009). Construction of optimal mixed-level super-saturated designs. *Stat. Sinica*, **19**, 197–211.
- Liu, Y. and A. Dean (2004). K -circulant supersaturated designs. *Technometrics*, **46**, 32–43.

- Louviere, J., D. Hensher, and J. Swait (2000). *Stated Choice Methods: Analysis and Applications*. Cambridge, U.K.: Cambridge University Press.
- Lucas, H. L. (1957). Extra-period Latin square change-over design. *J. Dairy Sci.*, **40**, 225–239.
- McKay, M. D., R. J. Beckman, and W. J. Conover (1979). A comparison of three methods of selecting values of input variables in the analysis of output from a computer code. *Technometrics*, **21**, 239–245.
- Mead, R. (1979). Competition experiments. *Biometrics*, **35**, 41–54.
- Mead, R. and J. Riley (1981). A review of statistical ideas relevant to intercropping research (with discussion). *J. R. Stat. Soc. A*, **144**, 462–509.
- Montgomery, D. C. (2012). *Design and Analysis of Experiments* (8th edition). Hoboken, NJ: Wiley.
- Morris, M. D. (2012). Computer experiments. In *Design and Analysis of Experiments*. Vol. 3: Special Designs and Applications (K. Hinkelmann, ed.), pp. 379–411. Hoboken, NJ: Wiley.
- Mukerjee, R. and C.-F. Wu (2006). *A Modern Theory of Factorial Designs*. New York: Springer.
- Myers, R. H., A. I. Khuri, and G. G. Vining (1992). Response surface alternatives to the Taguchi robust parameter design approach. *Am. Stat.*, **46**, 131–139.
- Myers, R. H., D. C. Montgomery, and C. M. Anderson-Cook (2009). *Response Surface Methodology: Process and Product Optimization Using Designed Experiments* (3rd edition). New York: Wiley.
- Nair, V. N. (1992). Taguchi's parameter design: A panel discussion. *Technometrics*, **34**, 127–161.
- Nelder, J. A. (1968). The combination of information in generally balanced designs. *J. R. Stat. Soc. B*, **30**, 303–311.
- Nettleton, D. (2012). Design of gene expression microarray experiments. In *Design and Analysis of Experiments*. Vol. 3: Special Designs and Applications (K. Hinkelmann, ed.), pp. 73–108. Hoboken, NJ: Wiley.
- Nguyen, N.-K. (1996). An algorithmic approach to constructing supersaturated designs. *Technometrics*, **38**, 69–73.
- Papadakis, J. S. (1937). Méthode statistique pour des expériences sur champ. *Bull. Inst. Amél. Plante à Salonique*, **23**, 1–30.
- Parker, J. (February 24, 2011). The 9-billion people question. *Economist*, **398** (8722) 3–5.
- Patterson, H. D. (1950). The analysis of change-over trials. *J. Agric. Sci.*, **40**, 375–380.
- Patterson, H. D. (1964). Theory of cyclic rotation experiments. *J. R. Stat. Soc. B*, **26**, 1–45.
- Patterson, H. D. (1965). The factorial combination of treatments in rotation experiments. *J. Agric. Sci.*, **65**, 171–182.
- Patterson, H. D. and B. I. Lowe, (1970). The errors of long-term experiments. *J. Agric. Sci.*, **74**, 53–60.
- Patterson, H. D. and R. Thompson (1971). Recovery of inter-block information when block sizes are unequal. *Biometrika*, **58**, 545–554.
- Patterson, H. D. and E. R. Williams (1976). A new class of resolvable incomplete block designs. *Biometrika*, **63**, 83–92.
- Pearce, S. C. (1960). Supplemental balance. *Biometrika*, **47**, 263–271 (Corrigenda **48**, 475).
- Piantadori, S. (1997). *Clinical Trials—A Methodologic Perspective*. New York: Wiley.
- Piepel, G. F. and S. K. Cooley (2009). Automated method for reducing Scheffé linear mixture experiment models. *Qual. Technol. Quant. Manage.*, **6**, 255–270.
- Plackett, R. L. and J. P. Burman (1946). The design of optimum multifactorial experiments. *Biometrika*, **33**, 305–325.
- Preece, D. A. (1986). Some general principles of crop rotation experiments. *Exp. Agric.*, **22**, 187–198.
- Prescott, P. (2004a). A class of designs for mixture experiments based on augmented pairs. *Commun. Stat. Theory Methods*, **33**, 2759–2785.
- Prescott, P. (2004b). Modelling in mixture experiments including interactions with process variables. *Qual. Technol. Quant. Manage.*, **1**, 87–103 (Correction: **2**, 123–123).
- Prvan, T. and D. J. Street (2002). An annotated bibliography of application papers using certain classes of fractional factorial and related designs. *J. Stat. Plann. Infer.*, **106**, 245–269.
- Pukelsheim, F. (1993). *Optimal Design of Experiments*. New York: Wiley.
- Raghavarao, D. (1960). A generalization of group divisible designs. *Ann. Math. Stat.*, **31**, 756–771.

- Raghavarao, D. (1971). *Construction and Combinatorial Problems in Design of Experiments*. New York: Wiley.
- Raghavarao, D. and K. Chandrasekhararao (1964). Cubic designs. *Ann. Math. Stat.*, **35**, 389–397.
- Rao, C. R. (1947). Factorial experiments derivable from combinatorial arrangements of arrays. *J. R. Stat. Soc. B*, **9**, 128–139.
- Rashed, D. (1984). Designs for multiple comparisons of control versus test treatments. Unpublished PhD dissertation. Virginia Polytechnic Institute and State University, Blacksburg, VA.
- Robinson, T. and C. M. Anderson-Cook (2012). Robust parameter design. In *Design and Analysis of Experiments*. Vol. 3: Special Designs and Applications (K. Hinkelmann, ed.), pp. 443–469. Hoboken, NJ: Wiley.
- Rosenberger W. F. and J. M. Lachin (2002). *Randomization in Clinical Trials: Theory and Practice*. New York: Wiley.
- Rothamsted Research (2006). *Guide to the Classical and Other Long-Term Experiments, Datasets and Sample Archives*. Harpenden, U.K.: Rothamsted Research.
- Roy, P. M. (1953–54). Hierarchical group divisible incomplete block designs with m -associate classes. *Sci. Cult.*, **19**, 210–211.
- Russell, E. J. (1926). Field experiments: How they are made and what they are. *J. Ministry Agric.*, **32**, 989–1001.
- Ryan, T. P. (2007). *Modern Experimental Design*. Hoboken, NJ.: Wiley.
- Sacks, J., W. J. Welch, T. J. Mitchell, and H. P. Wynn (1989). Design and analysis of computer experiments. *Stat. Sci.*, **4**, 409–435.
- Sanchez, S. M., T. W. Lucas, P. J. Sanchez, C. J. Nannini, and H. Wan (2012). Designs for large-scale simulation experiments, with applications to defense and homeland security. In *Design and Analysis of Experiments*. Vol. 3: Special Designs and Applications (K. Hinkelmann, ed.), pp. 413–441. Hoboken, NJ: Wiley.
- Santner, T. J., B. J. Williams, and W. I. Notz (2003). *The Design and Analysis of Computer Experiments*. New York: Springer.
- Sarkar, A. and K. Chatterjee (2010). An MEP.2 plan in 3^n factorial experiment and its probability of correct identification. *J. Stat. Plan. Infer.*, **140**, 3531–3539.
- Scheffé, H. (1958). Experiments with mixtures. *J. R. Stat. Soc. B*, **20**, 344–360.
- Schmidt, J. (1919). La valeur de l'individu à titre de générateur appréciée suivant méthode du croisement diallél. *Compt. Rend. Lab. Carlsberg*, **14(6)**, 1–33.
- Shah, B. V. (1958). On balancing in factorial experiments. *Ann. Math. Stat.*, **29**, 766–779.
- Shah, B. V. (1959). A generalization of partially balanced incomplete block designs. *Ann. Math. Stat.*, **30**, 1041–1050.
- Shah, K. R. (1964). The use of inter-block information to obtain uniformly better estimators. *Ann. Math. Stat.*, **35**, 1064–1078.
- Singh, M., S. Gupta, and R. Parsad (2012). Genetic crosses experiments. In *Design and Analysis of Experiments*. Vol. 3: Special Designs and Applications (K. Hinkelmann, ed.), pp. 1–71. Hoboken, NJ: Wiley.
- Singh, M. and K. Hinkelmann (1995). Partial diallel crosses in incomplete blocks. *Biometrics*, **51**, 1302–1314.
- Sinha, B. K., P. Das, N. K. Mandal, and M. Pal (2010). Parameter estimation in linear and quadratic mixture models: A review. *Pak. J. Stat.*, **26**, 77–96.
- Smith, D. M. and M. S. Ridout (2003). Optimal designs for criteria involving log(potency) in comparative binary bioassays. *J. Stat. Plann. Infer.*, **113**, 617–632.
- Smith, K. (1918). On the standard deviation of adjusted and interpolated values of an observed polynomial function and its constants, and the guidance they give towards a proper choice of the distribution of observations. *Biometrika*, **12**, 1–85.
- Smith, K. W., H. P. Marker, E. C. Fieller, and W. A. Broom (1944). An extended cross-over design and its use in insulin assay. *Q. J. Pharm. Pharmacol.*, **17**, 108–117.
- Sprague, G. F. and L. A. Tatum (1942). General vs. specific combining ability in single crosses of corn. *Am. Soc. Agron.*, **34**, 923–932.

- Srivastava, J. N. (1975). Designs for searching non-negligible effects. In *A Survey of Statistical Design and Linear Models* (J. N. Srivastava, ed.), pp. 507–519. Amsterdam, the Netherlands: North-Holland.
- Srivastava, J. N. and S. Ghosh (1976). A series of balanced 2^m factorial designs of resolution V which allow search and estimation of one extra unknown effect. *Sankhyā B*, **38**, 280–289.
- Srivastava, J. N. and B. C. Gupta (1979). Main effect plans for 2^m factorials which allow search and estimation of one unknown effect. *J. Stat. Plann. Infer.* **3**, 259–265.
- Steiner, S. H. and R. J. MacKay (2005). *Statistical Engineering*. Milwaukee, WI: ASQ Press.
- Stevens, W. L. (1939). The completely orthogonalized Latin square. *Ann. Eugen.*, **9**, 82–93.
- Street, A. P. and D. J. Street (1987). *Combinatorics of Experimental Design*. Oxford, U.K.: Clarendon.
- Street, D. J. and L. Burgess (2007). *The Construction of Optimal Stated Choice Experiments: Theory and Methods*. Hoboken, NJ: Wiley.
- Street, D. J. and L. Burgess (2012). Designs for choice experiments for the multinomial logit model. In *Design and Analysis of Experiments*. Vol. 3: Special Designs and Applications (K. Hinkelmann, ed.), pp. 331–378. Hoboken, NJ: Wiley.
- Street, D. J. and W. H. Wilson (1985). Balanced designs for two-variety competition experiments. *Utilitas Math.*, **28**, 113–120.
- Stringer, J. K., A. B. Smith, and B. R. Cullis (2012). Spatial analysis of agricultural field experiments. In *Design and Analysis of Experiments*. Vol. 3: Special Designs and Applications (K. Hinkelmann, ed.), pp. 109–136. Hoboken, NJ: Wiley.
- Surendran, P. U. (1968). Association matrices and the Kronecker product of designs. *Ann. Math. Stat.*, **39**, 676–680.
- Taguchi, G. (1986). *Introduction to Quality Engineering*. White Plains, NY: UNIPUB/Kraus International.
- Taylor, G. A. R. (1994). Analysis of experiments by using half-normal plots. *The Statistician*, **43**, 529–536.
- Tippett, L. H. C. (1935). Some applications of statistical methods to the study of variation of quality of cotton yarn (with discussion). *J. R. Stat. Soc. B*, **2**, 27–62.
- Tippett, L. H. C. (1936). Applications of statistical methods to the control of quality in industrial production. *Trans. Manchr. Stat. Soc.*, 1–32.
- Tippett, L. H. C. (1938). The statistical principles of experimental design with particular reference to experimentation in textile factories. In *Proceedings of the Industrial Statistics Conference*. New York: Pitman Publishing Co.
- Vartak, M. N. (1955). On the application of the Kronecker product of matrices to statistical designs. *Ann. Math. Stat.*, **26**, 420–438.
- Vartak, M. N. (1959). The non-existence of certain PBIB designs. *Ann. Math. Stat.*, **30**, 1051–1062.
- Vining, G. G. (2012). Split-plot response surface designs. In *Design and Analysis of Experiments*. Vol. 3: Special Designs and Applications (K. Hinkelmann, ed.), pp. 471–500. Hoboken, NJ: Wiley.
- Vining, G. G. and Myers, R. H. (1990). Combining Taguchi and response surface philosophies: A dual response approach. *J. Qual. Technol.*, **22**, 38–45.
- Vivacqua, C. A. and S. Bisgaard (2004). Strip-block experiments for process improvement and robustness. *Qual. Eng.*, **16**, 671–675.
- Wald, A. (1943). On the efficient design of statistical investigations. *Ann. Math. Stat.*, **14**, 134–140.
- Welch, W. J., T. K. Yu, S. M. Kang, and J. Sacks. (1990). Computer experiments for quality control by parameter design. *J. Qual. Technol.*, **22**, 38–45.
- Wilkinson, G. N., S. R. Eckert, T. W. Hancock, and O. Mayo (1983). Nearest neighbour (NN) analysis of field experiments. *J. R. Stat. Soc. B*, **45**, 151–178.
- Williams, E. J. (1949). Experimental designs balanced for the estimation of residual effects of treatments. *Aust. J. Sci. Res. A*, **2**, 149–168.
- Williams, E. J. (1950). Experimental designs balanced for pairs of residual effects. *Aust. J. Sci. Res. A*, **3**, 351–363.
- Wishart, J. (1934). Statistics in agricultural research. *J. R. Stat. Soc. B*, **1**, 26–61.

- Wittes, J. and Z.-F. Yu (2012). Design and analysis of randomized clinical trials. In *Design and Analysis of Experiments*. Vol. 3: Special Designs and Applications (K. Hinkelmann, ed.), pp. 165–211. Hoboken, NJ: Wiley.
- Wolock, F. W. (1964). Cyclic designs. Unpublished PhD dissertation, Virginia Polytechnic Institute, Blacksburg, VA.
- Wu, C. F. J. (1993). Construction of supersaturated designs through partially aliased interactions. *Biometrika*, **80**, 661–669.
- Wu, C. F. J. and Y. Chen (1992). A graph-aided method for planning two-level experiments when certain interactions are important. *Technometrics*, **34**, 162–175.
- Wu, C. F. J. and M. S. Hamada (2009). *Experiments: Planning, Analysis, and Optimization* (2nd edition). Hoboken, NJ: Wiley.
- Yates, F. (1933a). The analysis of replicated experiments when the field results are incomplete. *Empire J. Exp. Agric.*, **1**, 129–142.
- Yates, F. (1933b). The formation of Latin squares for use in field experiments. *Empire J. Exp. Agric.*, **1**, 235–244.
- Yates, F. (1935). Complex experiments (with discussion). *J. R. Stat. Soc. Suppl.*, **2**, 181–247.
- Yates, F. (1936a). Incomplete randomized blocks. *Ann. Eugen.*, **7**, 121–140.
- Yates, F. (1936b). A new method of arranging variety trials: Quasi-Latin squares. *J. Agric. Sci.*, **26**, 424–455.
- Yates, F. (1937). The design and analysis of factorial experiments. *Imp. Bur. Soil Sci. Technol. Commun.*, **35**, 1–95.
- Yates, F. (1939). The recovery of inter-block information in variety trials arranged in three-dimensional lattices. *Ann. Eugen.*, **9**, 136–156.
- Yates, F. (1940). The recovery of inter-block information in balanced incomplete block designs. *Ann. Eugen.*, **10**, 317–325.
- Yates, F. (1949). The design of rotation experiments. *Commonw. Bur. Soil Sci. Tech. Commun.*, **46**.
- Yates, F. (1954). The analysis of experiments containing different crop rotations. *Biometrics*, **10**, 324–346.
- Yeh, C.-M. and R. A. Bradley (1983). Trend-free block designs: Existence and construction results. *Commun. Stat. Theory Methods*, **12**, 1–24.
- Yeh, C.-M., R. A. Bradley, and W. I. Notz (1985). Nearly trend-free block designs. *J. Am. Stat. Assoc.*, **80**, 985–992.
- Yin, G. (2012). *Clinical Trial Design: Bayesian and Frequentist Adaptive Methods*. New York: Wiley.
- Youden, W. J. (1937). The use of incomplete block replications in estimating tobacco-mosaic virus. *Contr. Boyce Thompson Inst.*, **9**, 41–48.
- Youden, W. J. (1940). Experimental designs to increase accuracy of greenhouse studies. *Contr. Boyce Thompson Inst.*, **11**, 219–228.
- Zahn, D. A. (1975). Modifications of revised critical values for the half-normal plot. *Technometrics*, **17**, 189–200.
- Zhang, C. and H. Peng (2012). D-optimal designs for quadratic mixture canonical polynomials with spline. *Stat. Probab. Lett.*, **82**, 1095–1101.
- Zhang, L. and W. Rosenberger (2012). Adaptive randomization in clinical trials. In *Design and Analysis of Experiments*. Vol. 3: Special Designs and Applications (K. Hinkelmann, ed.), pp. 251–281. Hoboken, NJ: Wiley.
- Zoellner, J. A. and O. Kempthorne (1954). Incomplete block designs with blocks of two plots. *Iowa State Coll. Agric. Exp. Station Res. Bull.*, 171–180.

2

Introduction to Linear Models

Linda M. Haines

CONTENTS

2.1	Introduction.....	63
2.2	Formal Definitions.....	64
2.3	Linear Model of Full Rank.....	65
2.3.1	Independent Errors with Constant Variance.....	65
2.3.1.1	Theory.....	65
2.3.1.2	Design.....	68
2.3.2	Heteroscedastic and Correlated Error Terms.....	74
2.4	Linear Model Not of Full Rank.....	75
2.4.1	Theory.....	75
2.4.2	Examples and Design.....	78
2.4.2.1	Block Designs.....	79
2.4.2.2	Factorial Experiments.....	82
2.5	Linear Mixed Model.....	83
2.5.1	Theory.....	84
2.5.1.1	Estimation.....	84
2.5.1.2	Prediction.....	86
2.5.2	Examples and Design.....	88
2.5.2.1	Observations on Individuals.....	88
2.5.2.2	Kriging and Other Models.....	91
2.6	Conclusions.....	92
	Acknowledgments.....	93
	References.....	93

2.1 Introduction

The *linear model* is exceptionally well documented throughout the statistical literature. The theory is elegant, and applications in science, engineering, medicine, education, social science, commerce and industry abound. Accounts are included in introductory texts such as that of Kutner et al. (2004), the book devoted to regression by Draper and Smith (1998) and matrix algebra-based texts such as the classic books by Scheffé (1959), Searle (1971), Graybill (1976) and Rao (2001) and the more recent books by Ravishanker and Dey (2002) and Khuri (2010). The theory of the linear model falls naturally into two parts, that for the linear model of full rank and that for the linear model not of full rank. While certain concepts relating to these two forms of the linear model intersect, the models nevertheless merit

separate attention. The aim in this chapter, therefore, is to give concise accounts of each separately, emphasizing those aspects which have an impact on design and introducing some key concepts relating to the construction of optimal designs.

The *linear mixed model* is, technically, an extension of the linear model which includes *random effects*. In practice, however, the linear mixed model offers very much more in terms of flexibility and scope than this extension might at first suggest. The basic ideas underpinning the model are rooted in animal genetics but today have far-reaching application in particular in medicine, biology and education. There are a number of excellent texts devoted to the linear mixed model, many of which also include accounts of the linear model itself. These include the books by Duchateau et al. (1998), Pinheiro and Bates (2000) and Verbeke and Molenberghs (2000) and the strongly theoretical classic texts by Searle et al. (1992) and McCulloch and Searle (2001). In this chapter, the basic theory underpinning the linear mixed model is introduced, and elements of design, specifically relating to two main areas of application, are then explored.

2.2 Formal Definitions

The linear model can be specified quite straightforwardly by the equations

$$y_i = f'(x_i)\beta + \epsilon_i, \quad i = 1, \dots, n, \quad (2.1)$$

where y_i is the i th observation taken at the k explanatory variables specified by the vector $x_i = (x_{i1}, x_{i2}, \dots, x_{ik})$, $f(x_i)$ is a $(p + 1) \times 1$ vector of functions of those variables with $\beta = (\beta_0, \beta_1, \dots, \beta_p)$ a conformable vector of unknown parameters, and ϵ_i is an error term which has mean zero and is not necessarily independent of other error terms. The functions $f(x_i)$ have a leading element of 1 associated with the intercept or mean parameter β_0 and in addition comprise elements corresponding to the explanatory variables $x_{i1}, x_{i2}, \dots, x_{ik}$ and, quite commonly, to scalar functions of those variables such as $x_{i1}x_{i2}$, x_{i1}^2 and x_{i2}^3 . The explanatory variables themselves take values from a design space \mathcal{X} which is determined by the nature of the experiment. The specification of this design space and the attendant selection of explanatory variables together provide the essential framework for a designed experiment.

Equation (2.1) is generally referred to as the observational equation and can be readily assembled in matrix form as

$$y = X\beta + \epsilon,$$

where $y = (y_1, \dots, y_n)'$, the matrix X is $n \times (p + 1)$ with i th row $f'(x_i)$ and $\epsilon = (\epsilon_1, \dots, \epsilon_n)'$ is distributed as a random vector with mean $\mathbf{0}$ and variance matrix Σ . For the linear model of full rank, the matrix X is of full column rank, that is, $\text{rank}(X) = p + 1$ if $n \geq p + 1$, and for the linear model not of full rank, X is not of full column rank, that is, $\text{rank}(X) < p + 1$. Linear models not of full rank are defined intrinsically, at least in general. However, such models can also arise in practice, most commonly when data collection does not follow a well-designed experiment.

The linear mixed model extends the linear model to incorporate random effects and can be specified quite simply as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u} + \boldsymbol{\epsilon},$$

where \mathbf{u} is a $q \times 1$ vector of random effects with mean $\mathbf{0}$ and variance matrix \mathbf{G} , independent of the error term $\boldsymbol{\epsilon}$, and \mathbf{Z} is an $n \times q$ matrix of functions of explanatory variables whose values are taken from a design space \mathcal{Z} . Designs for this model are thus specified by the matrices \mathbf{X} and \mathbf{Z} . However, it is not uncommon for the explanatory variables defining the matrix \mathbf{Z} to be fixed, that is, determined by the structure of the problem, and thus for design issues to be dictated solely by the specification of the matrix \mathbf{X} .

2.3 Linear Model of Full Rank

The theory of the linear model of full rank with error variables that are independently and identically distributed provides the basis for theory relating to more general settings. This theory, together with some attendant design issues, is therefore discussed in detail as follows, and extensions to the model with errors that are heteroscedastic or correlated or both are then considered.

2.3.1 Independent Errors with Constant Variance

2.3.1.1 Theory

Consider the linear model of full rank with error variables independently and normally distributed with common variance σ^2 . Then the observations \mathbf{y} are distributed as $\mathbf{y} \sim N(\mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{I})$, where \mathbf{I} is the identity matrix, and the log-likelihood is given by

$$l = l(\boldsymbol{\beta}, \sigma^2; \mathbf{y}) = -\frac{n}{2} \ln(2\pi) - \frac{n}{2} \ln(\sigma^2) - \frac{1}{2\sigma^2} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta})' (\mathbf{y} - \mathbf{X}\boldsymbol{\beta}).$$

The maximum likelihood estimators (MLEs) of the unknown parameters $\boldsymbol{\beta}$ and σ^2 can be obtained by solving the likelihood equations $\frac{\partial l}{\partial \boldsymbol{\beta}} = \mathbf{0}$ and $\frac{\partial l}{\partial \sigma^2} = 0$ (Khuri 2010). Thus, the MLE for $\boldsymbol{\beta}$, written $\hat{\boldsymbol{\beta}}$, is given by the solution to the normal equations $\mathbf{X}'\mathbf{X}\hat{\boldsymbol{\beta}} = \mathbf{X}'\mathbf{y}$. Since \mathbf{X} has full column rank, the matrix $\mathbf{X}'\mathbf{X}$, of order $p + 1$ and $\text{rank}(\mathbf{X}'\mathbf{X}) = \text{rank}(\mathbf{X}) = p + 1$, is nonsingular and thus

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}.$$

The MLE of σ^2 , denoted $\hat{\sigma}^2$, follows immediately as

$$\hat{\sigma}^2 = \frac{(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})' (\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})}{n} = \frac{\mathbf{y}' (\mathbf{I} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}') \mathbf{y}}{n}.$$

Note that it follows from the form of the likelihood that $\widehat{\boldsymbol{\beta}}$ is also the least squares estimator of $\boldsymbol{\beta}$. In addition, it can be shown, by an appropriate Gauss–Markov theorem, that $\widehat{\boldsymbol{\beta}}$ is the best linear unbiased estimator (BLUE) of $\boldsymbol{\beta}$ (Khuri 2010).

Under the assumption of normality, the distributions of $\widehat{\boldsymbol{\beta}}$ and $\widehat{\sigma}^2$ are given by

$$\widehat{\boldsymbol{\beta}} \sim N(\boldsymbol{\beta}, \sigma^2(\mathbf{X}'\mathbf{X})^{-1}) \quad \text{and} \quad \frac{n\widehat{\sigma}^2}{\sigma^2} \sim \chi_{n-(p+1)}^2$$

respectively. Thus, $\widehat{\boldsymbol{\beta}}$ is unbiased, whereas $\widehat{\sigma}^2$ is biased since $E(\widehat{\sigma}^2) = \frac{(n-(p+1))}{n}\sigma^2$. For this reason, the least squares estimator

$$s^2 = \frac{(\mathbf{y} - \mathbf{X}\widehat{\boldsymbol{\beta}})'(\mathbf{y} - \mathbf{X}\widehat{\boldsymbol{\beta}})}{n - (p + 1)}$$

is commonly used in place of $\widehat{\sigma}^2$ to estimate σ^2 . Further, key properties of $\widehat{\boldsymbol{\beta}}$ and $\widehat{\sigma}^2$ follow immediately. In particular, $\widehat{\boldsymbol{\beta}}$ and $\widehat{\sigma}^2$ are sufficient statistics for $\boldsymbol{\beta}$ and σ^2 and in addition are independently distributed.

Interest, particularly when constructing good designs, often focuses on linear functions of the regression parameters $\boldsymbol{\beta}$. It follows from the principle of invariance that a function $\mathbf{c}'\boldsymbol{\beta}$, where \mathbf{c} is a constant vector conformable with $\boldsymbol{\beta}$, has MLE $\mathbf{c}'\widehat{\boldsymbol{\beta}}$ and, from the distribution of $\widehat{\boldsymbol{\beta}}$, that

$$\mathbf{c}'\widehat{\boldsymbol{\beta}} \sim N(\mathbf{c}'\boldsymbol{\beta}, \sigma^2\mathbf{c}'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{c}).$$

More specifically, the fitted values are given by $\widehat{\mathbf{y}} = \mathbf{X}\widehat{\boldsymbol{\beta}}$ and are distributed as

$$\widehat{\mathbf{y}} \sim N(\mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'),$$

where the matrix $\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$ is referred to as the *hat matrix* or the orthogonal projection matrix, succinctly denoted \mathbf{P}_X . Similarly, the *predicted value* of a response \widehat{y}_0 at a specified setting of the explanatory variables \mathbf{x}_0 is given by $\widehat{y}_0 = \mathbf{f}'(\mathbf{x}_0)\widehat{\boldsymbol{\beta}}$ with distribution

$$\widehat{y}_0 \sim N(\mathbf{f}'(\mathbf{x}_0)\boldsymbol{\beta}, \sigma^2\mathbf{f}'(\mathbf{x}_0)(\mathbf{X}'\mathbf{X})^{-1}\mathbf{f}(\mathbf{x}_0)).$$

The information matrix for $\boldsymbol{\beta}$ and σ^2 at the design specified by the explanatory variable settings $\mathbf{x}_1, \dots, \mathbf{x}_n$, and hence by the matrix \mathbf{X} , plays a crucial role in the design of experiments for the linear model of full rank. In the present context, the information matrix is given by $-E\left[\frac{\partial^2 l}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'}\right]$ where $\boldsymbol{\theta} = (\boldsymbol{\beta}, \sigma^2)$ and thus by

$$\begin{bmatrix} \frac{1}{\sigma^2}(\mathbf{X}'\mathbf{X}) & 0 \\ 0 & \frac{n}{2\sigma^4} \end{bmatrix}$$

(McCulloch and Searle 2001). Clearly, $\widehat{\boldsymbol{\beta}}$ is an efficient estimator since $\text{Var}(\widehat{\boldsymbol{\beta}}) = \sigma^2(\mathbf{X}'\mathbf{X})^{-1}$ attains the Rao–Cramér lower bound. However, it can be shown that $\widehat{\sigma}^2$ and indeed s^2 are not efficient for σ^2 .

Certain aspects of inference for the parameters $\boldsymbol{\beta}$ and σ^2 have an impact on design, albeit somewhat indirectly, and these are outlined here. Thus consider testing the null hypothesis $H_0 : \mathbf{H}\boldsymbol{\beta} = \mathbf{h}$ against the alternative $H_A : \mathbf{H}\boldsymbol{\beta} \neq \mathbf{h}$ where \mathbf{H} is a $q \times (p+1)$ matrix of full row rank and \mathbf{h} is a vector of constants such that $\mathbf{h} \in \mathcal{C}(\mathbf{H})$ where \mathcal{C} is the column space. Then the likelihood ratio test is based on the test statistic

$$\frac{(\mathbf{H}\widehat{\boldsymbol{\beta}} - \mathbf{h})' (\mathbf{H}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{H}')^{-1} (\mathbf{H}\widehat{\boldsymbol{\beta}} - \mathbf{h})}{qs^2}$$

which follows an F distribution with degrees of freedom q and $n - (p+1)$ under H_0 and a noncentral F distribution with noncentrality parameter

$$\lambda = (\mathbf{H}\boldsymbol{\beta} - \mathbf{h})' (\mathbf{H}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{H}')^{-1} (\mathbf{H}\boldsymbol{\beta} - \mathbf{h}) / \sigma^2$$

under H_A . For a linear function $\mathbf{c}'\boldsymbol{\beta}$, the test statistic for $H_0 : \mathbf{c}'\boldsymbol{\beta} = c_0$ against $H_A : \mathbf{c}'\boldsymbol{\beta} \neq c_0$ simplifies to

$$\frac{(\mathbf{c}'\widehat{\boldsymbol{\beta}} - c_0)^2}{s^2(\mathbf{c}'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{c})} \sim t_{n-(p+1)}^2$$

under H_0 , and the associated $100(1 - \alpha)\%$ confidence interval for $\mathbf{c}'\boldsymbol{\beta}$ can be summarized as

$$\mathbf{c}'\widehat{\boldsymbol{\beta}} \pm t_{n-(p+1); \frac{\alpha}{2}}^* s \sqrt{\mathbf{c}'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{c}}$$

where $t_{n-(p+1); \frac{\alpha}{2}}^*$ is the $100(1 - \frac{\alpha}{2})$ th percentile of the $t_{n-(p+1)}$ distribution. Special cases include $\mathbf{c} = \mathbf{e}_j$ corresponding to the j th individual parameter, with \mathbf{e}_j the $(p+1) \times 1$ vector with j th element equal to 1 and all other elements 0, and $\mathbf{c} = \mathbf{f}(x_i)$ corresponding to the fitted value of the i th observation. If several such linear functions are compared simultaneously, then intervals based on a multiple comparison correction, such as that of Bonferroni, should be invoked. For the parameters themselves, the test statistic for $H_0 : \boldsymbol{\beta} = \mathbf{0}$ reduces to the analysis of variance–based form

$$\frac{\mathbf{y}'\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}}{(p+1)s^2}$$

where $\mathbf{y}'\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}$ is the regression sum of squares and the attendant $100(1 - \alpha)\%$ confidence region for $\boldsymbol{\beta}$ is the ellipsoid specified by

$$(\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}})' \mathbf{X}'\mathbf{X} (\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}) \leq (p+1)s^2 F_{p+1, n-(p+1); \alpha}^*$$

where $F_{p+1, n-(p+1); \alpha}^*$ is the $100(1 - \alpha)$ th percentile of the $F_{p+1, n-(p+1)}$ distribution.

Explanatory variables quite often fall naturally into, or are considered to be in, two or more groups. This setting is explored here for two groups with extensions to more than two groups being immediate. Thus consider the linear model of full rank partitioned as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} = \mathbf{X}_1 \boldsymbol{\beta}_1 + \mathbf{X}_2 \boldsymbol{\beta}_2 + \boldsymbol{\epsilon} \quad (2.2)$$

where $\mathbf{X} = [\mathbf{X}_1 \ \mathbf{X}_2]$ and $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$ are $p_1 \times 1$ and $p_2 \times 1$ vectors of unknown parameters with $p_1 + p_2 = p$ and with associated, conformable model matrices \mathbf{X}_1 and \mathbf{X}_2 , respectively. Note that the mean or intercept β_0 is assumed to be incorporated into the parameter vector $\boldsymbol{\beta}_1$. The results for the full model hold with obvious modification, but it is worth considering the information matrix and confidence regions for $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$ a little more carefully. Thus the information matrix for the parameters $\boldsymbol{\beta}' = (\boldsymbol{\beta}'_1, \boldsymbol{\beta}'_2)$ has the form

$$\frac{1}{\sigma^2} \begin{bmatrix} \mathbf{X}'_1 \mathbf{X}_1 & \mathbf{X}'_1 \mathbf{X}_2 \\ \mathbf{X}'_2 \mathbf{X}_1 & \mathbf{X}'_2 \mathbf{X}_2 \end{bmatrix},$$

and the information for $\boldsymbol{\beta}_1$ alone, that is, ignoring $\boldsymbol{\beta}_2$, is clearly proportional to $\mathbf{X}'_1 \mathbf{X}_1$. It now follows quite straightforwardly that the information for $\boldsymbol{\beta}_2$ eliminating $\boldsymbol{\beta}_1$, that is, with the parameter $\boldsymbol{\beta}_1$ unknown (Pawitan 2001), is summarized in the Schur complement $\mathbf{M}^*_{22} = \mathbf{X}'_2 \mathbf{X}_2 - \mathbf{X}'_2 \mathbf{X}_1 (\mathbf{X}'_1 \mathbf{X}_1)^{-1} \mathbf{X}'_1 \mathbf{X}_2$. Furthermore, the test of the hypothesis $H_0 : \boldsymbol{\beta}_2 = \boldsymbol{\beta}_{20}$ against $H_A : \boldsymbol{\beta}_2 \neq \boldsymbol{\beta}_{20}$ is based on the statistic

$$\frac{(\widehat{\boldsymbol{\beta}}_2 - \boldsymbol{\beta}_{20})' \mathbf{M}^*_{22} (\widehat{\boldsymbol{\beta}}_2 - \boldsymbol{\beta}_{20})}{p_2 s^2},$$

and the associated $100(1 - \alpha)\%$ confidence region for $\boldsymbol{\beta}_2$ is given by

$$(\boldsymbol{\beta}_2 - \widehat{\boldsymbol{\beta}}_2)' \mathbf{M}^*_{22} (\boldsymbol{\beta}_2 - \widehat{\boldsymbol{\beta}}_2) \leq p_2 s^2 F^*_{p_2, n-(p+1); \alpha},$$

where $F^*_{p_2, n-(p+1); \alpha}$ is the $100(1 - \alpha)$ th percentile of the $F_{p_2, n-(p+1)}$ distribution.

2.3.1.2 Design

In the context of the linear model of full rank, a design represents an allocation of n_i observations to the i th distinct explanatory variable setting x_i for $i = 1, \dots, t$ where $\sum_{i=1}^t n_i = n$, and is denoted

$$\xi_n = \begin{Bmatrix} x_1, & x_2, & \dots, & x_t \\ n_1, & n_2, & \dots, & n_t \end{Bmatrix}. \quad (2.3)$$

Note that the explanatory variables $x_i, i = 1, \dots, t$, are generally referred to as the support points of the design or, collectively, as the support of the design. The design ξ_n defines the matrix \mathbf{X} , with n_i rows of \mathbf{X} corresponding to the setting or functions of the setting $x_i, i = 1, \dots, t$. Optimal designs, that is, designs which are in some sense optimal according

to a specified criterion, are now sought. The immediate challenge therefore is to formulate criteria which are meaningful and applicable.

In linear regression, optimality criteria relating to the precise estimation of the parameters β and to linear functions of those parameters are of key importance and are commonly based on the information matrix $X'X/\sigma^2$ for β , with σ^2 an unknown nuisance parameter. It is sometimes possible to order the information matrices for β , or more precisely the matrices $X'X$, for different designs according to their nonnegative definite differences, that is, according to the Loewner ordering, but such an ordering is not in general readily identified (Pukelsheim 1993; Yang et al. 2011). Designs based on criteria which are appropriate scalar functions of the information matrix are therefore sought.

In terms of the parameters β themselves, the criteria of D - and A -optimality have dominated the literature. Specifically, D -optimal designs are those designs which minimize the *generalized variance* of the parameter estimates $\hat{\beta}$ or, equivalently, the volume of the confidence ellipsoid for β and thus which maximize the determinant $\det(X'X)$, and A -optimal designs are those designs which minimize the average variance of the estimates $\hat{\beta}$ and thus the criterion $\text{trace}((X'X)^{-1})$. D_S -optimal designs provide a useful counterpoint to D -optimal designs and are those designs which minimize the generalized variance of a subset of the parameters β . Thus for the parameters β_2 introduced earlier in the partitioned model (2.2), a D_S -optimal design is a design which maximizes the determinant $\det(M_{22}^*)$. In particular, the D_S -optimality criterion plays a key role in design for model checking (Atkinson et al. 2007).

The precise estimation of linear functions of the regression parameters, that is, functions of the form $c'\beta$ where c is a vector of constants, is of central interest. The variance of $c'\hat{\beta}$, the estimate of $c'\beta$, is given by $\sigma^2 c'(X'X)^{-1}c$ (Section 2.3.1.1), and the definition of a c -optimal design as a design which minimizes this variance follows immediately. The criteria of D_{A-} , C - and E -optimality are built on c -optimality in an obvious way. Specifically, a set of linear functions of β can be assembled row-wise as $T\beta$ where T is an appropriate matrix of constants, and the estimate $T\hat{\beta}$ has variance matrix $\sigma^2 T(X'X)^{-1}T'$. Thus D_{A-} and C -optimal designs can be introduced as those designs which minimize the determinant and the trace of the matrix $T(X'X)^{-1}T'$, respectively. The variances of the linear functions embedded in $T\beta$ depend on the scaling of the rows of T , and this can be countered by introducing E -optimal designs as those designs which minimize the maximum variance $\sigma^2 c'(X'X)^{-1}c$ subject to the normalizing constraint $c'c = 1$.

The predicted values \hat{y}_0 (Section 2.3.1.1) are, of course, estimates of the linear functions of the parameters of the specific form $f'(x_0)\beta$, and they play a key role in any regression analysis. As a consequence, optimality criteria specific to the precise estimation of the predicted responses have been formulated. In particular, designs minimizing the maximum variance of $f'(x_0)\hat{\beta}$, that is, designs minimizing the maximum of the function $f'(x_0)(X'X)^{-1}f(x_0)$, over the design space \mathcal{X} are termed G -optimal designs. In addition, designs minimizing the average variance of the predicted responses over the settings in a specified subset of the design space, say \mathcal{X}_0 , that is, designs minimizing the criterion

$$\int_{x_0 \in \mathcal{X}_0} f'(x_0)(X'X)^{-1}f(x_0) d\mu(x_0) = \text{trace} \left[(X'X)^{-1} \int_{x_0 \in \mathcal{X}_0} f(x_0)f'(x_0) d\mu(x_0) \right]$$

where μ is a probability measure on \mathcal{X}_0 , are termed V -optimal if the subspace \mathcal{X}_0 comprises a finite number of points and I -optimal otherwise.

The criteria just discussed are aptly referred to as the *alphabetic optimality criteria*. There are in fact many such criteria other than those already noted, but it is possible to cut across this large family and to identify classes of criteria based on form and properties. Thus the D -, D_S - and D_A -optimality criteria are all determinant based, the A -, c -, C -, V - and I -optimality criteria all have the form $\text{trace}((X'X)^{-1}L)$ for some matrix L and are termed L -optimality criteria (Fedorov 1972), and the E - and G -optimality criteria are both minimax. In addition, certain of these criteria can be expressed in terms of the eigenvalues of the matrix $X'X$, denoted $\lambda_i, i = 1, \dots, p + 1$. Specifically, D -optimal designs maximize the product $\prod_{i=1}^{p+1} \lambda_i$, A -optimal designs minimize the sum $\sum_{i=1}^{p+1} \frac{1}{\lambda_i}$ and E -optimal designs maximize the minimum eigenvalue. Finally, it is readily seen that D - and D_S -optimal designs are invariant to linear transformations of the rows of X , that is, transformations of the form XA where A is an appropriate nonsingular matrix of order $p + 1$. Specifically, in the case of D -optimality, the determinant of the information matrix for the regression parameters associated with such a transformation is given by $\det(A'X'XA)$ and, since A is constant and nonsingular, is proportional to $\det(X'X)$. This invariance does not, at least in general, extend to other criteria.

The discussion thus far has focused on designs with integer numbers of points allocated to the explanatory variable settings, that is, designs of the form given in (2.3). Such designs are used in practice and are referred to as exact designs. There is however no unifying mathematical theory associated with optimality for exact designs. To compound matters, the construction of optimal exact designs is invariably combinatorial in nature and necessarily computationally challenging. In addition, the optimality or otherwise of a candidate design cannot be readily demonstrated. In the late 1950s, Kiefer broke through the exact design mind-set by introducing the notion of an approximate design and by developing the elegant theory that is associated with such designs. His ideas are formalized in two papers, a sole-author discussion paper in 1959 (Kiefer 1959) and a paper with Wolfowitz in 1960 (Kiefer and Wolfowitz 1960), and are now presented here, albeit somewhat briefly.

An approximate design is defined, quite straightforwardly, as a probability measure on the design space \mathcal{X} . If the measure comprises a finite set of points, then this definition can be relaxed to give the design

$$\xi = \left\{ \begin{array}{cccc} x_1, & x_2, & \dots, & x_t \\ w_1, & w_2, & \dots, & w_t \end{array} \right\}, \quad (2.4)$$

with weights allocated to the distinct settings, that is, with $0 < w_i < 1$ for $i = 1, \dots, t$ and with $\sum_{i=1}^t w_i = 1$. The formulation relates directly back to that of the exact design (2.3) but with the proportions of observations allocated to the settings not necessarily being ratios of integers. It now follows that the information matrix for β at an approximate design ξ based on a finite set of points can be introduced as

$$M(\xi) = X'WX = \sum_{i=1}^t w_i f(x_i) f'(x_i)$$

where W is a diagonal matrix of the weights, that is, $W = \text{diag}(w_1, \dots, w_t)$, and the nuisance parameter σ^2 is taken without loss of generality to be 1. Furthermore, optimality criteria based on this information matrix, and expressed as $\Psi(M(\xi))$, or more succinctly and depending on the context as $\Psi(\xi)$, can then be easily formulated.

The theory associated with optimal approximate designs is now introduced. Convex criteria play a key role and are defined as those criteria which satisfy the inequality

$$\Psi[\alpha\xi_1 + (1 - \alpha)\xi_2] \geq \alpha\Psi(\xi_1) + (1 - \alpha)\Psi(\xi_2)$$

where ξ_1 and ξ_2 are approximate designs of the form (2.4) and $0 \leq \alpha \leq 1$. Then it can be shown, by invoking Carathéodory's theorem (Silvey 1980; Rockafellar 1996), that an approximate design which minimizes a criterion that is convex over the space of all such designs, Ξ , can always be formulated as a design based on a finite number of support points and furthermore that an upper bound of $p(p + 1)/2 + 1$ can be placed on that number. This powerful result then leads directly into the formulation and proof of a suite of general equivalence theorems (Fedorov 1972; Silvey 1980; Atkinson et al. 2007). The theorems are presented here in broad form and without proof. Specifically, consider a convex criterion $\Psi(\xi)$ with a directional derivative at ξ in the direction of the point x defined by

$$\phi(x, \xi) = \lim_{\alpha \rightarrow 0} \frac{\Psi(M[(1 - \alpha)\xi + \alpha x]) - \Psi(M(\xi))}{\alpha}.$$

Then the theorem stated as follows holds, provided the design space \mathcal{X} is compact.

Theorem 2.1 (A General Equivalence Theorem) *The following three conditions characterize a Ψ -optimal approximate design ξ^* :*

1. *The design ξ^* minimizes $\Psi(\xi)$ over all approximate designs $\xi \in \Xi$.*
2. *The design ξ^* maximizes the minimum of the directional derivative $\phi(x, \xi)$ over all $x \in \mathcal{X}$.*
3. *The minimum of $\phi(x, \xi)$ is given by $\phi(x, \xi^*) = 0$ and is attained at the support points of the design.*

The impact of the theorem is immediate and far reaching. One of the key features, and arguably the most important, is that the conditions of the theorem provide a mechanism for confirming the global optimality or otherwise of a candidate optimal design. In addition, and to fix ideas, it is worth considering the special cases of D -, L - and E -optimality in a little more detail. Thus the D -optimality criterion $\det(M(\xi))$ can be reexpressed as the convex criterion $\Psi_D(\xi) = -\ln(\det(M(\xi)))$, and the associated directional derivative of $\Psi_D(\xi)$ at a design ξ in the direction of a point x is given by

$$\phi_D(x, \xi) = (p + 1) - f'(x)M(\xi)^{-1}f(x)$$

where $f'(x)M(\xi)^{-1}f(x)$ is proportional to the variance of the predicted response $f'(x)\widehat{\beta}$ and is termed the standardized variance function, written $d(x, \xi)$. It thus follows that the condition $\phi_D(x, \xi) \geq 0$ for all $x \in \mathcal{X}$ for the design ξ to be D -optimal, which arises from the equivalence theorem, is equivalent to the variance-based condition $d(x, \xi) \leq p + 1$ for all $x \in \mathcal{X}$ and thus that approximate D -optimal designs are also G -optimal. In other words, within the context of approximate designs, the D -optimality criterion which is convex and tractable and the G -optimality criterion which is minimax and difficult to manipulate are in fact equivalent.

The L -optimality criterion is based on minimizing $\Psi_L(\xi) = \text{trace}(\mathbf{M}(\xi)^{-1}\mathbf{L})$ and, as noted earlier, includes A -, c -, C -, I - and V -optimality. The criterion can be readily shown to be convex, and the attendant directional derivative of $\Psi_L(\xi)$ at a design ξ in the direction of a point x is given by

$$\phi_L(x, \xi) = \text{trace}(\mathbf{M}(\xi)^{-1}\mathbf{L}) - f'(x)\mathbf{M}(\xi)^{-1}\mathbf{L}\mathbf{M}(\xi)^{-1}f(x).$$

The general equivalence theorem stated earlier can thus be invoked, thereby extending the notion of L -optimality to approximate designs. However, directional derivatives of the form $\phi_L(x, \xi)$ do not enjoy such a ready interpretation as that of $\phi_D(x, \xi)$.

The E -optimality criterion, which is based on minimizing the maximum of $c'\mathbf{M}(\xi)^{-1}c$ over all c such that $c'c = 1$, is clearly minimax, and as a consequence, the attendant equivalence theorem must be delicately formulated. Specifically, the theorem states that an approximate design ξ is E -optimal if and only if there exists a matrix $E = \sum_{i=1}^{p_0} \alpha_i z_i z_i'$ such that

$$\phi_E(x, \xi) = \lambda_{\min} - x'Ex = \lambda_{\min} - \sum_{i=1}^{p_0} \alpha_i (x'z_i)^2 \geq 0 \quad \text{for all } x \in \mathcal{X}$$

where the terms α_i are positive numbers which sum to 1 and the vectors z_i comprise the set of p_0 normalized eigenvectors of $\mathbf{M}(\xi)$ corresponding to the smallest eigenvalue λ_{\min} for $i = 1, \dots, p_0$ and $1 \leq p_0 \leq p + 1$. The condition on $\phi_E(x, \xi)$ is subtle, somewhat circular and not in general algebraically tractable, reflecting the fact that the E -criterion is minimax and not amenable to differentiation, at least in a more conventional sense.

Finally, the notion of design *efficiency* merits special attention. Specifically, the D - and L -efficiency of a design ξ_A relative to an optimally better design ξ_B are defined as the ratios

$$\left[\frac{\det(\mathbf{M}(\xi_A))}{\det(\mathbf{M}(\xi_B))} \right]^{\frac{1}{p+1}} \quad \text{and} \quad \frac{\text{trace}(\mathbf{M}(\xi_B)^{-1}\mathbf{L})}{\text{trace}(\mathbf{M}(\xi_A)^{-1}\mathbf{L})}$$

respectively, where ξ_A and ξ_B can be either approximate or exact designs. These efficiencies are, in essence, ratios of variances, and their inverses can thus be interpreted as the numbers of replications of design ξ_A required for that design to have the same variance as that of design ξ_B . In addition, if an approximate optimal design is available for a particular setting, then the design can be used as a benchmark, that is, as the base design ξ_B in the calculation of efficiencies.

The criteria introduced here form the building blocks for many other regression-based criteria, such as compound and constrained criteria, and criteria for augmented and sequential designs. These somewhat more complicated criteria are covered in detail both in Atkinson et al. (2007) and in later chapters of the handbook and thus, for the sake of brevity, are not discussed here.

Some ideas relating to the construction of optimal designs are now outlined. More complete details are given in Atkinson et al. (2007). Thus, for certain model settings, it is possible to obtain explicit expressions for approximate optimal designs by invoking algebraic arguments based on the relevant equivalence theorem. Some fascinating examples are given in Fedorov (1972), Silvey (1980) and Atkinson et al. (2007) and indeed throughout the design literature. However, more generally, optimal or near-optimal approximate designs must

be found numerically using an appropriate constrained optimization procedure and the global optimality or otherwise of the resultant candidate design then confirmed, again numerically.

The construction of exact optimal designs is more challenging, and arguably more intriguing, than that of approximate optimal designs. The essential problem is combinatorial, and indeed Welch (1982) has shown that the construction of exact D -optimal designs is NP-hard. In certain cases, it may be feasible to obtain the requisite globally optimal exact design by, for example, complete enumeration, branch-and-bound search or subtle algebraic arguments. More generally, however, recourse must be made to heuristic algorithms. The latter fall broadly into two classes, greedy algorithms such as those based on Fedorov-type exchanges, candidate-free coordinate exchange and multiplicative updating of the weights and nature-inspired algorithms including simulated annealing, genetic algorithms and swarm intelligence. Further issues relating to algorithmic searches for optimal designs are covered in Chapter 21.

A small example serves to fix ideas relating to design construction. Suppose that a cubic polynomial is to be fitted to observations on the interval $[-1, 1]$ in such a way that the regression parameters are estimated as precisely as possible. Suppose further that only five measurements can be taken at locations which can be specified to one decimal place. What can the designer recommend? First, the approximate D -optimal design follows immediately from a suite of results based on elegant arguments involving differential equations (Fedorov 1972) and puts equal weights on the four support points $-1, -\frac{1}{\sqrt{5}}, \frac{1}{\sqrt{5}}$ and 1 . This design provides a benchmark for efficiencies but is not practically meaningful in that the five observations cannot be distributed equally across the four support points of the approximate design. Second, exact D -optimal designs comprising five points taken on the interval $[-1, 1]$ can be constructed numerically using a suitable constrained optimization routine and, somewhat surprisingly, comprise the four support points of the approximate design with any one of those points repeated twice. The loss in efficiency relative to the benchmark approximate D -optimal design is small, just less than 5%. Finally, five-point D -optimal designs based on measurements taken on a grid with spacings of 1 decimal place can be readily found by complete enumeration and comprise the points $-1, -.4, .4, .5$ and 1 or, by symmetry, the points $-1, -.5, -.4, .4$ and 1 . The loss of efficiency for these designs relative to the five point D -optimal design based on continuous measurements is less than 0.5% and is surprisingly small.

Finally, an account of the linear model of full rank with independent error terms of constant variance would not be complete without mention of response surface methodology (RSM). The main thrust of RSM is to explore the surface generated by responses to several variables, usually coded, by fitting first- and second-order polynomial models which provide an approximate rather than a mechanistic model for the data. More specific aims are to examine the nature of the surface and to locate an appropriate optimum. The theory which underpins RSM is that of the linear model of full rank, but complex and often conflicting desiderata influence design. Thus the alphabetic criteria described previously which are critically important in design construction within the regression context are subsumed in a wealth of other criteria (Myers 1999). Specifically, RSM is used extensively in industry, and as a consequence, criteria for good design are strongly driven by practical concerns and include, *inter alia*, issues such as the nature of the variances of predicted responses across the design region, cost-effectiveness, robustness to missing values and estimation of pure error. Many of these requirements are not always easy to quantify and can be example specific. There are a number of excellent texts on the topic of RSM, including those by

Box et al. (2005), Khuri (2006) and Myers et al. (2009), and Chapter 5 provides a review of recent design-related developments.

2.3.2 Heteroscedastic and Correlated Error Terms

Settings in which the error terms of the linear model of full rank have a variance matrix of general form, that is, $\text{Var}(\epsilon) = \Sigma$, are easy to formulate but fraught with problems relating to the analytic and algorithmic construction of the attendant optimal designs.

Suppose first that the errors are independent and that the variance of an observation at the generic setting x is of known form $\sigma^2/\lambda(x)$, where $\lambda(x)$ is termed the efficiency function (Fedorov 1972). Then it is easy to see that the model can be reformulated as a linear model with explanatory variables specified by $\sqrt{\lambda(x)}f(x)$ and with error terms of constant variance σ^2 . Thus all the results pertaining to the linear model of full rank described in the previous subsection accrue. In particular, in the case of polynomial models in a single explanatory variable, the theory of differential equations and related orthogonal polynomials can often be invoked in order to derive analytic expressions for D -optimal designs, and a synthesis of the results is given in the paper by Dette and Trampisch (2010).

Suppose now that the observations are correlated, that is, $\Sigma = \sigma^2\mathbf{R}$, with the correlation matrix \mathbf{R} a function of specified explanatory variables and of known form. This setting is readily formulated and is nicely exemplified by a single time series with autocorrelated errors. The analytic derivation of the associated optimal designs is however challenging, and, to compound matters, numerical construction requires cunningly devised and specialized algorithms. An introduction to this area of design construction is given in Atkinson et al. (2007), and a comprehensive account, together with some cutting-edge developments, is presented in Chapter 6.

Suppose, more generally, that the error variables are normally distributed with a variance matrix Σ which is of known form but is now a function of unknown parameters as well as explanatory variables. Then the log-likelihood can be expressed in the usual way as

$$-\frac{n}{2} \ln(2\pi) - \frac{1}{2} \ln \det(\Sigma) - \frac{1}{2} (\mathbf{y} - \mathbf{X}\beta)^T \Sigma^{-1} (\mathbf{y} - \mathbf{X}\beta)$$

and the information matrix for the parameters derived taking into account the parameter dependencies of Σ . In so doing, however, three scenarios which depend on the interplay between the unknown parameters in Σ and the regression parameters β should be recognized. Specifically, if Σ depends on unknown parameters θ which do not include β , that is, $\Sigma = \Sigma(\theta)$, it can be shown, by invoking some subtle matrix algebra, that the information matrix for β and θ is block diagonal and is given by

$$\begin{bmatrix} \mathbf{X}^T \Sigma^{-1} \mathbf{X} & \mathbf{0} \\ \mathbf{0}' & \frac{1}{2} \left\{ \text{trace} \left(\Sigma^{-1} \frac{\partial \Sigma}{\partial \theta_r} \Sigma^{-1} \frac{\partial \Sigma}{\partial \theta_s} \right) \right\} \end{bmatrix}, \quad (2.5)$$

where the notation (B_{rs}) is used to denote a generic matrix \mathbf{B} with rs th element B_{rs} . In addition, if Σ depends on β alone, that is, $\Sigma = \Sigma(\beta)$, the information matrix for β is then

$$\mathbf{X}^T \Sigma^{-1} \mathbf{X} + \frac{1}{2} \left\{ \text{trace} \left(\Sigma^{-1} \frac{\partial \Sigma}{\partial \beta_r} \Sigma^{-1} \frac{\partial \Sigma}{\partial \beta_s} \right) \right\},$$

and if Σ depends on unknown parameters θ and also, and separately, on the regression parameters β , that is, $\Sigma = \Sigma(\beta, \theta)$, then the off-diagonal block matrices in (2.5) are no longer zero and are given by

$$\frac{1}{2} \left\{ \text{trace} \left(\Sigma^{-1} \frac{\partial \Sigma}{\partial \beta_r} \Sigma^{-1} \frac{\partial \Sigma}{\partial \theta_s} \right) \right\}.$$

These formulae simplify considerably in the case of independent observations, and full details are given in the paper by Atkinson and Cook (1995). If the observations are correlated, however, the major application of the setting is to be found within the context of the linear mixed model and is discussed in Section 2.5.

2.4 Linear Model Not of Full Rank

The theory of the linear model not of full rank is somewhat more challenging and more intricate than that of the linear model of full rank and relies heavily on the notion and properties of the generalized or g -inverses of a matrix. Furthermore, this reliance permeates into the formulation of applications based on the model. Note therefore that a g -inverse of an $n \times k$ matrix A , not necessarily square, is a $k \times n$ matrix, denoted A^- , which satisfies the condition $AA^-A = A$. It can be shown that such a matrix always exists but that it is not necessarily unique. Note also that a unique g -inverse of A , termed the Moore–Penrose inverse and denoted A^+ , does in fact exist and satisfies the three conditions $A^+AA^+ = A^+$, $(AA^+)' = AA^+$ and $(A^+A)' = A^+A$, in addition to the base condition $AA^+A = A$. The Moore–Penrose inverse finds powerful application in the theory of the linear model not of full rank and can be readily computed by invoking the singular value decomposition. Further details relating to g -inverses are included in, for example, the books by Pringle and Rayner (1971) and Ben-Israel and Greville (2003).

2.4.1 Theory

Consider again the linear model $\mathbf{y} = \mathbf{X}\beta + \epsilon$ with $\epsilon \sim N(\mathbf{0}, \sigma^2\mathbf{I})$ but now with the matrix \mathbf{X} not of full column rank, that is, $\text{rank}(\mathbf{X}) = r < p + 1 \leq n$. The normal equations

$$\mathbf{X}'\mathbf{X}\beta^0 = \mathbf{X}'\mathbf{y},$$

with the vector β^0 representing a solution to these equations, follows immediately by minimizing the sum of squares

$$S(\beta) = (\mathbf{y} - \mathbf{X}\beta)'(\mathbf{y} - \mathbf{X}\beta)$$

or, equivalently, maximizing the likelihood with respect to the parameter β . However, $\text{rank}(\mathbf{X}'\mathbf{X}) = \text{rank}(\mathbf{X}) = r < p + 1$ and thus $\mathbf{X}'\mathbf{X}$ is singular, and the solution β^0 is not unique (Khuri 2010). Indeed, there are infinitely many such solutions. To be specific, the normal

equations are consistent since $\mathcal{C}(X'X) = \mathcal{C}(X')$, and $X'y \in \mathcal{C}(X'X)$, and a general solution to them can be expressed as

$$\beta^0 = (X'X)^- X'y + [I - (X'X)^- X'X]z \quad (2.6)$$

where z is an arbitrary vector (Pringle and Rayner 1971).

Since there are infinitely many solutions to the normal equations, it is clear that it is not sensible to consider least squares or, equivalently, maximum likelihood estimation of the parameters β . The question therefore immediately arises as to whether it is meaningful to estimate any function of the parameters β . The issue can be approached by defining a linear function of β , written $c'\beta$, to be estimable if and only if $c'\beta^0$ is unique, that is, invariant to the choice of β^0 . A further question then arises as to whether any such linear functions exist. In fact, estimable functions are readily characterized as is shown in the following theorem.

Theorem 2.2 *The linear function $c'\beta$ is estimable if and only if $c' \in \mathcal{R}(X)$ or, equivalently, if and only if $c' \in \mathcal{R}(X'X)$, where \mathcal{R} denotes the row space.*

The proof is short and elegant and is therefore included here.

Proof: Suppose that β_1^0 and β_2^0 are two different solutions to the normal equations. Then $X'X\beta_1^0 = X'y = X'X\beta_2^0$ and thus $X'X(\beta_1^0 - \beta_2^0) = 0$, that is, $(\beta_1^0 - \beta_2^0) \in [\mathcal{R}(X'X)]^\perp$, where \mathcal{W}^\perp denotes the orthogonal complement of an inner product space \mathcal{W} . Now $c'\beta$ is estimable if and only if $c'\beta_1^0 = c'\beta_2^0$, that is, $c'(\beta_1^0 - \beta_2^0) = 0$, and thus if and only if $c' \in \mathcal{R}(X'X) = \mathcal{R}(X)$.

Thus, $c'\beta^0$ is taken to be the least squares or, equivalently, the MLE of the estimable function $c'\beta$. Furthermore, it can be shown that the function $c'\beta$ is estimable if and only if there exists an unbiased estimator of $c'\beta$ which is linear in y . In fact, the existence of a linear unbiased estimator for $c'\beta$ can be used as the basis for an alternative, and often more tractable, definition of estimability. It can also be proved that $c'\beta^0$ is the BLUE of $c'\beta$, thereby establishing a Gauss–Markov result (Khuri 2010). Finally, note that a set of functions $T\beta$ is termed estimable if and only if each individual function so defined is estimable and thus if and only if $\mathcal{R}(T) \subseteq \mathcal{R}(X)$. The maximum number of linearly independent estimable functions is clearly $\text{rank}(X) = r$.

Many estimable functions can be identified by inspection of the rows of X or $X'X$. Thus, the fitted values $X\beta$ are clearly estimable, the least squares estimate of σ^2 ,

$$s^2 = \frac{(y - X\beta^0)'(y - X\beta^0)}{n - r}$$

is thus unique and a full analysis of variance can be formulated. However, identifying estimable functions by eye is not always straightforward, and the subtle problem of establishing whether or not a given function $c'\beta$ is estimable then arises. One approach is to use the result

$$c'\beta \text{ is estimable if and only if } c' = c'(X'X)^-(X'X)$$

to establish estimability (Pringle and Rayner 1971). More generally, symbolic algebra packages, as, for example, Mathematica (Wolfram Research 2014), provide matrix operations such as Gaussian elimination that are performed exactly and can be used to confirm or refute the estimability of a set of functions $T\beta$.

A solution to the normal equations is required in order to calculate the estimate $c'\beta^0$ of an estimable function $c'\beta$. The most obvious approach is to set $z=0$ in the general solution (2.6) and to invoke $\beta^0 = (X'X)^-X'y$ where $(X'X)^-$ is any generalized inverse of $X'X$. There are a number of ways in which such a generalized inverse can be found. For example, if $X'X$ is partitioned so that the leading diagonal matrix X'_1X_1 is nonsingular with rank $r = \text{rank}(X'X)$, then a g -inverse of $X'X$ is given by

$$\begin{bmatrix} (X'_1X_1)^{-1} & \mathbf{0} \\ \mathbf{0}' & \mathbf{0} \end{bmatrix}.$$

More generally, symbolic algebra packages such as Mathematica (Wolfram Research 2014) and GAP (GAP 2014) provide explicit expressions, and programming languages such as R (R Development Core Team 2014) and MATLAB[®] (MATLAB 2014) provide numeric results for the Moore–Penrose inverse of a matrix.

A second approach to finding a solution to the normal equations is to impose a constraint $H\beta^0 = h$ directly on the equations themselves. The matrix H must necessarily be a row complement of $X'X$ or equivalently of X , that is, H must be a $q \times (p+1)$ matrix of rank $p+1-r$ which satisfies $\mathcal{R}(X) \cap \mathcal{R}(H) = \emptyset$. In certain cases, such as the completely randomized design, it is possible to impose the constraints by inspection, that is, by hand, but more generally, a solution to the normal equations based on H is given by $\beta^0 = (X'X + H'H)^{-1}X'y$. A particularly useful choice of H is an orthogonal complement of X of full row rank, that is, a row complement of X satisfying $H'X = 0$ with $q = p+1-r$.

The matrix $(X'X + H'H)^{-1}$ is a g -inverse of $X'X$. In fact, more generally, the two approaches to finding a solution to the normal equations, that of finding a g -inverse of $X'X$ and that of imposing a constraint on the normal equations, can be shown to be equivalent. However, the two approaches are subtly different in implementation and are valuable in somewhat different contexts.

In many settings for which the linear model not of full rank is appropriate, the explanatory variables fall naturally into two or more groups. The partitioned form of the model and the attendant model matrix therefore play a key role in the design and analysis of an experiment. Some notions relating to the normal equations and to estimability for the model comprising two groups of explanatory variables are therefore introduced here. Specifically, consider again the partitioned model $y = X\beta = X_1\beta_1 + X_2\beta_2 + \epsilon$ but now with the matrix $X = [X_1 \ X_2]$ not of full rank. Then the normal equations can be expressed succinctly as

$$\begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix} \begin{bmatrix} \beta_1^0 \\ \beta_2^0 \end{bmatrix} = \begin{bmatrix} X'_1y \\ X'_2y \end{bmatrix}$$

where $S_{11} = X'_1X_1$, $S_{22} = X'_2X_2$, and $S_{12} = X'_1X_2$ with $S_{21} = S'_{12}$. Solutions to the normal equations can be obtained by a *sweep-out* procedure, that is, by premultiplying the equations by the *sweep* matrix

$$F = \begin{bmatrix} I & \mathbf{0} \\ -S_{21}S_{11}^{-1} & I \end{bmatrix}$$

to give

$$\begin{bmatrix} S_{11} & S_{12} \\ \mathbf{0} & S_{22} - S_{21}S_{11}^{-1}S_{12} \end{bmatrix} \begin{bmatrix} \beta_1^0 \\ \beta_2^0 \end{bmatrix} = \begin{bmatrix} X_1' \mathbf{y} \\ X_2' \mathbf{y} - S_{21}S_{11}^{-1}X_1' \mathbf{y} \end{bmatrix}$$

(Pringle and Rayner 1971). Thus, information on the parameters β_2 for a particular design is contained in the Schur complement $S_{22}^* = S_{22} - S_{21}S_{11}^{-1}S_{12}$, and it follows that the linear function $c_2'\beta_2$ is estimable if and only if $c_2' \in \mathcal{R}(S_{22}^*)$. Similarly, if the coefficients β_1 are of interest, then clearly information on β_1 is contained in the matrix $S_{11}^* = S_{11} - S_{12}S_{22}^{-1}S_{21}$, and $c_1'\beta_1$ is estimable if and only if $c_1' \in \mathcal{R}(S_{11}^*)$. Note, however, that, unless the off-diagonal matrix S_{12} is zero so that the effects relating to β_1 and β_2 are in that sense orthogonal, information and estimability for the parameter vectors β_1 and β_2 jointly must be approached with care. For example, the linear function $c_1'\beta_1 + c_2'\beta_2$ is estimable if and only if $(c_1', c_2') \in \mathcal{R}([X_1 \ X_2])$.

Details relating to the linear model not of full rank are intricate, and it is therefore tempting to reformulate the model as one of full rank. Such a reparametrization can be introduced formally as follows. Consider finding a basis set of estimable functions assembled as $\beta^* = T\beta$ with T an $r \times (p + 1)$ matrix of full row rank, and consider setting $X\beta = U\beta^*$ where U is a matrix of full column rank. Then it can be shown that the model $y = U\beta^* + \epsilon$ is "equivalent in every way to the original model" (Pringle and Rayner 1971). The construction of U follows by observing that $X = UT$ so that $XT' = UTT'$ and hence that $U = XT'(TT')^{-1}$. Also, if $U'U$ is diagonal, then the transformation is said to be an orthogonal full rank reparametrization. Such reparametrizations can be useful, for example, in the case of 2^k factorial experiments discussed in Chapter 7, but more generally, the essence and ready interpretation of the model is lost.

The problems encountered in developing the theory of the linear model not of full rank are further exacerbated if the error terms are assumed to be correlated. The research questions which arise are challenging and impinge on notions of blocking for spatial data. Some issues, more particularly in relation to design, are discussed in the recent paper by Kunert et al. (2010) and references therein.

The theory of the linear model not of full rank can also be developed using projection matrices. Specifically, the hat matrix $X(X'X)^{-1}X'$ defines the estimation space and the matrix $I - X(X'X)^{-1}X'$ the error space of the model, and arguments based on projections onto these matrices can then be invoked. This approach provides a fascinating and powerful counterpoint to the one presented here, and full details are available in the books by Bailey (2008) and by Morris (2011).

2.4.2 Examples and Design

In order to fix ideas, examples from two major areas of application of the linear model not of full rank, namely, block designs for varietal trials and factorial experiments, are now introduced. In addition, since it is not easy to formulate a general framework for design, many issues relating specifically to design are explored within the context of these examples. The concepts presented here provide a basis for an understanding of the later chapters on block designs, that is, Chapters 3, 4, and 10, and on factorial experiments, that is, Chapters 7 through 9. Books on block designs include the classic texts of Cochran and Cox (1957)

and Shah and Sinha (1989) and the more recent texts by John and Williams (1995) and Dey (2010), and those on factorial experiments include the classic text by Raktoc et al. (1981) and the recent book by Wu and Hamada (2009) and provide excellent further detail.

2.4.2.1 Block Designs

Consider first a block design with t treatments replicated r_i times, $i = 1, \dots, t$, and the treatments allocated to b blocks of size k_j , $j = 1, \dots, b$. The total number of observations is therefore $n = \sum_{i=1}^t r_i = \sum_{j=1}^b k_j$. Assume that the structure and allocation are such that a treatment occurs at most once in each block and thus that the design is a binary block design. Then the allocation of treatments to blocks can be summarized succinctly by the $t \times b$ incidence matrix $N = \{n_{ij}\}$ where $n_{ij} = 1$ if treatment i occurs in block j and 0 otherwise. Observe that the sum of the rows of the matrix N is the vector of block sizes, that is, $\mathbf{k} = N'\mathbf{1} = (k_1, \dots, k_b)'$, and that the sum of the columns of N is the vector of treatment replications, that is, $\mathbf{r} = N\mathbf{1} = (r_1, \dots, r_t)'$. For ease of exposition, the theory of block designs developed here is restricted to binary block designs but can be readily extended to block designs for which treatments are replicated within blocks.

If the i th treatment occurs in the j th block, the attendant observation y_{ij} can be modelled as

$$y_{ij} = \mu + \tau_i + \beta_j + \epsilon_{ij}$$

where μ is the mean effect, τ_i the i th treatment effect, β_j the j th block effect, and ϵ_{ij} is an error term distributed as $N(0, \sigma^2)$ independently of all other such terms. The model can also be expressed in matrix form as

$$\mathbf{y} = \mathbf{1}\mu + \mathbf{X}_1\boldsymbol{\tau} + \mathbf{X}_2\boldsymbol{\beta} + \boldsymbol{\epsilon}$$

where the matrices \mathbf{X}_1 and \mathbf{X}_2 are the unit/treatment and unit/block incidence matrices, respectively, $\boldsymbol{\tau} = (\tau_1, \dots, \tau_t)'$, $\boldsymbol{\beta} = (\beta_1, \dots, \beta_b)'$ and $\boldsymbol{\epsilon} \sim N(\mathbf{0}, \sigma^2\mathbf{I})$. Note that the matrix \mathbf{X}_1 is $n \times t$, the matrix \mathbf{X}_2 is $n \times b$ and both matrices have rows with one element equal to 1 and all other elements 0. The incidence matrix itself is then given by $N = \mathbf{X}_1'\mathbf{X}_2$. Two linear relationships between the columns of the $n \times (t + b + 1)$ matrix $\mathbf{X} = [\mathbf{1} \ \mathbf{X}_1 \ \mathbf{X}_2]$ can immediately be identified since $\mathbf{X}_1\mathbf{1} - \mathbf{1} = \mathbf{0}$ and $\mathbf{X}_2\mathbf{1} - \mathbf{1} = \mathbf{0}$. Thus $\text{rank}(\mathbf{X}) \leq t + b - 1$ and the linear model for the block design is not of full rank.

The normal equations for this setting are readily derived and can be expressed succinctly as

$$\begin{bmatrix} n & \mathbf{r}' & \mathbf{k}' \\ \mathbf{r} & \mathbf{R} & \mathbf{N} \\ \mathbf{k} & \mathbf{N}' & \mathbf{K} \end{bmatrix} \begin{bmatrix} \mu^0 \\ \boldsymbol{\tau}^0 \\ \boldsymbol{\beta}^0 \end{bmatrix} = \begin{bmatrix} G \\ \mathbf{t} \\ \mathbf{b} \end{bmatrix}$$

where $\mathbf{R} = \text{diag}(r_1, \dots, r_t)$, $\mathbf{K} = \text{diag}(k_1, \dots, k_b)$, G is the sum of the observed values, that is, $G = \mathbf{1}'\mathbf{y}$, and $\mathbf{t} = \mathbf{X}_1'\mathbf{y}$ and $\mathbf{b} = \mathbf{X}_2'\mathbf{y}$ are the vectors of treatment and block totals respectively (John and Williams 1995). The singularity of the normal equations follows from the fact that $\mathbf{X}'\mathbf{X}$ is of order $t + b + 1$ but with $\text{rank}(\mathbf{X}'\mathbf{X}) = \text{rank}(\mathbf{X}) \leq t + b - 1$. It is convenient to identify the singularities more explicitly. Thus, consider reducing the normal equations

to a more tractable form by invoking the *sweep-out* procedure introduced earlier, that is, by multiplying both sides of the equations by the nonsingular *sweep* matrix

$$F = \begin{bmatrix} n^{-1} & \mathbf{0}' & \mathbf{0}' \\ \mathbf{0} & I & -NK^{-1} \\ \mathbf{0} & -N'R^{-1} & I \end{bmatrix}$$

to give the set of equations

$$\begin{bmatrix} 1 & n^{-1}r' & n^{-1}k' \\ \mathbf{0} & R - NK^{-1}N' & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & K - N'R^{-1}N \end{bmatrix} \begin{bmatrix} \mu^0 \\ \boldsymbol{\tau}^0 \\ \boldsymbol{\beta}^0 \end{bmatrix} = \begin{bmatrix} \bar{y} \\ \mathbf{t} - NK^{-1}\mathbf{b} \\ \mathbf{b} - N'R^{-1}\mathbf{t} \end{bmatrix}$$

where $\bar{y} = G/n$. Thus, $(R - NK^{-1}N')\mathbf{1} = \mathbf{0}$ and $(K - N'R^{-1}N)\mathbf{1} = \mathbf{0}$ thereby concentrating the singularities in $X'X$ in two matrices. Furthermore, solutions for the treatment effects can be obtained from the equation $C\boldsymbol{\tau}^0 = \mathbf{q}$ where $C = R - NK^{-1}N'$ and $\mathbf{q} = \mathbf{t} - NK^{-1}\mathbf{b}$, and all information regarding these effects is summarized in the matrix C , termed the treatment information matrix.

The question now arises as to which linear functions of the parameters μ , $\boldsymbol{\tau}$, and $\boldsymbol{\beta}$ are estimable given that the normal equations possess at least two singularities. The answer to this question is inextricably tied up with the notion of connectedness. Specifically, a block and a treatment are said to be associated if the treatment occurs in the specified block, and two treatments are said to be connected if they are linked by a chain of associated treatments and blocks (Dey 2010). A design is said to be connected if every pair of treatments is connected. It can then be shown that, for a connected design, $\text{rank}(C) = \text{rank}(R - NK^{-1}N') = t - 1$ and $\text{rank}(K - N'R^{-1}N) = b - 1$ and hence that the matrix $X'X$ has $\text{rank } t + b - 1$. Here, attention is restricted to connected block designs.

Since the treatment information matrix C is of order t with $C\mathbf{1} = \mathbf{0}$ and $\text{rank}(C) = t - 1$, it is now immediately clear that all treatment contrasts of the form $c'\boldsymbol{\tau}$ with $c'\mathbf{1} = 0$, that is, with $c' \in \mathcal{R}(C)$, are estimable. Similarly, all block contrasts are estimable. Furthermore, a solution to the equations $C\boldsymbol{\tau}^0 = \mathbf{q}$ is given by $\boldsymbol{\tau}^0 = C^- \mathbf{q}$ where C^- is a g -inverse of C . Specifically, note that the vector $\mathbf{1}$ is an orthogonal row complement of C and thus that a particularly tractable g -inverse of C is given by $(C + aJ)^{-1}$ where a is any nonzero constant and $J = \mathbf{1}\mathbf{1}'$. Note also that the unique Moore–Penrose inverse of C is given by $C^+ = (C + t^{-1}J)^{-1} - t^{-1}J$ (Caliński and Kageyama 2000). The least squares estimator of the contrast $c'\boldsymbol{\tau}$ can now be obtained as $c'C^- \mathbf{q}$. In addition, the vector \mathbf{q} can be expressed as $\mathbf{q} = (X_1' - NK^{-1}X_2')\mathbf{y}$ with $\mathbf{y} \sim N(\mathbf{0}, \sigma^2 I)$, and it then follows by some straightforward algebra that $\text{Var}(\mathbf{q}) = \sigma^2 C$ and hence, since $c' \in \mathcal{R}(C)$, that $\text{Var}(c'\boldsymbol{\tau}^0) = \sigma^2 c' C^- c$.

The setting just described provides a template for block designs more generally. For example, the randomized complete block design can be recovered by taking blocks of size t , while row-and-column designs, that is, designs with two block structures that are arranged, at least figuratively, as rows and columns, can be readily formulated through direct extensions of the observational and normal equations. In addition, the design setting provides sufficient flexibility to demonstrate how designs, which are in some sense optimal according to specified criteria, can be accommodated in terms of the allocation of treatments to blocks. This idea is now explored a little further.

The primary aim of any block design for varietal and related trials is to compare treatments once the effects of the blocks have been discounted. It is natural therefore to focus on design criteria which are based on the variances of the estimated pairwise treatment differences. An immediately appealing optimality criterion in this regard is that of equality of the variances, that is, variance balance. As an example, consider the balanced incomplete block design (BIBD), which is equi-replicate with treatments replicated r times, which has blocks of the same size k and which is such that the number of times each pair of treatments occurs in the same block is the same and is equal to λ . Then, intuitively, it is clear that the variances of estimated pairwise treatment differences are equal, but it is informative to demonstrate this explicitly. Thus, it is readily shown from earlier derivations and from the fact that any treatment shares a block with $r(k-1) = \lambda(t-1)$ other treatments that

$$C = rI - \frac{1}{k} [(r - \lambda)I + \lambda J] = \frac{r(k-1) + \lambda}{k} I - \frac{\lambda}{k} J = \frac{\lambda t}{k} I - \frac{\lambda}{k} J.$$

It now follows that $\frac{k}{\lambda t} I$ is a g -inverse of C and hence that the variances of all estimated pairwise treatment differences are equal and are given by $2k\sigma^2/(\lambda t)$. The BIBD is not a panacea for variance balance and good design however. Specifically the three conditions $rt = bk$, $\lambda(t-1) = r(k-1)$ and $b \geq t$ (Fisher's inequality) are necessary but not sufficient for a BIBD to exist. For example, it can be shown that a BIBD with $t = b = 22$, $r = k = 7$ and $\lambda = 2$ does not exist. As a consequence, there is a large body of mathematically intricate research devoted to broad classes of designs which, *inter alia*, go part way to meeting the requirement of variance balance (John 1971; John and Williams 1995; Bailey 2004). Block designs with independent responses, including the BIBD and other related designs, are discussed in comprehensive detail in Chapter 3.

The structural restrictions associated with BIBDs and related designs lead naturally into a consideration of alphabetic optimality criteria rather than the more stringent criterion of variance balance. Thus, A - and MV -optimal designs are those designs which minimize the average of the variances of the estimated pairwise treatment differences and the maximum of those variances respectively and are well documented and widely used (John and Williams 1995). In the case of A -optimality, the criterion reduces to minimizing trace($BC^{-1}B'$), where the $\binom{t}{2} \times t$ matrix B comprises rows with elements 0, 1 and -1 which define appropriately all pairwise treatment differences, and thus, since $B'B = tI - J$, to trace($C^{-1}(tI - J)$). Furthermore, if the Moore–Penrose inverse is chosen as the g -inverse of C , then $C^+ \mathbf{1} = \mathbf{0}$ and the A -optimality criterion reduces quite simply to minimizing trace(C^+) (Bailey 2009). The A -optimality criterion can be expressed in terms of the $t-1$ nonzero eigenvalues of the matrix C , denoted $\lambda_1, \dots, \lambda_{t-1}$, as $\sum_{i=1}^{t-1} \frac{1}{\lambda_i}$. Other criteria based solely on these eigenvalues can also be formulated but are not as widely used or as practically meaningful as the A -optimal and MV -optimal criteria. For example, D -optimal designs maximize the product $\prod_{i=1}^{t-1} \lambda_i$, E -optimal designs maximize the minimum nonzero eigenvalue of C and (M, S) -optimal designs minimize the sum of squares $\sum_{i=1}^{t-1} \lambda_i^2$ over the set of all designs for which the sum $\sum_{i=1}^{t-1} \lambda_i$ is a maximum. Further discussion relating to optimality criteria which are both meaningful and widely used in block designs is given in Chapter 3.

Some attention should also be given to the notion of orthogonality of block designs. Specifically a block design is said to be orthogonal if estimates of treatment contrasts are the same whether or not the block effects are included in the model. For example, the randomized complete block design is orthogonal. More generally, many row-and-column designs

such as the Latin square and the Graeco-Latin square designs are also orthogonal and enjoy all the powerful properties that so accrue.

Finally, it should be emphasized that there are a number of *standard* or *textbook* block designs which are precisely specified and which enjoy certain desirable properties but which are not sufficiently flexible, in terms of allocation of treatments to blocks, to accommodate a range of optimality criteria. Such designs include Latin square designs, Graeco-Latin squares and Youden squares and the randomized complete block designs, amongst others. The analysis of these designs is rooted in the theory of the linear model not of full rank described earlier, and the implementation and attendant analyses are comprehensively delineated in a number of intermediate texts such as those of Dean and Voss (1999) and Montgomery (2012).

2.4.2.2 Factorial Experiments

Consider first the case of two factors A and B with numbers of levels a and b , respectively. Suppose that observations are taken at combinations of the factor levels specified by an $a \times b$ incidence matrix N , with the number of observations in the ij th cell, that is, taken at level i of A and level j of B , given by n_{ij} , with $n_{ij} \geq 0$, $i = 1, \dots, a$ and $j = 1, \dots, b$. Then, if $n_{ij} > 0$, the observations in the ij th cell can be modelled as

$$y_{ijk} = \mu + A_i + B_j + \epsilon_{ijk} \quad \text{for } k = 1, \dots, n_{ij},$$

where μ is the mean effect, A_i and B_j are the effects of the factors A and B at levels i and j , respectively, and the error term ϵ_{ijk} is distributed as $N(0, \sigma^2)$ independently of all other such error terms. Technically this model is the same as that introduced for the block design earlier in this section, and similar considerations in terms of connectedness and estimability apply. For example, in a connected two-factor model, all contrasts of the effects A_i , $i = 1, \dots, a$, and, separately, of the effects B_j , $j = 1, \dots, b$, are estimable. However, it should be emphasized that underlying issues relating to the experiment itself, in particular the choice of factors and levels and the randomization and replication of treatment combinations, are crucial and are comprehensively delineated in Chapter 1. As an aside, it is interesting to note that the effects of A and B are not necessarily orthogonal but that there is a class of designs for which orthogonality holds. Specifically, it can be shown that the effects of the factors A and B are orthogonal if and only if the proportionality condition

$$\frac{n_{i0}n_{0j}}{n} = n_{ij} \quad \text{with} \quad n_{i0} = \sum_{j=1}^b n_{ij} \quad \text{and} \quad n_{0j} = \sum_{i=1}^a n_{ij},$$

holds for all $i = 1, \dots, a$ and $j = 1, \dots, b$. The two-factor model can be extended to accommodate interaction between the factors and to models with several factors with and without interaction, but, for the sake of brevity, full details are not presented here. The primary challenge in the design of factorial experiments is encountered when blocking is required or when the numbers of treatment combinations is large and only a fraction can be implemented.

The theory which underpins s^n factorial experiments, that is, designs with n factors at s levels where s is a prime number, is rooted in Galois field theory and is very different to that for the general multifactor model. Many of the key features relating to these experiments,

in particular blocking and fractional factorials, are introduced in Chapter 1 and discussed later in comprehensive detail in Chapter 7. More general results, including those for mixed level factorial designs such as $2^m 4^n$ and $2^m 3^n$ factorials, are given in a number of texts, for example, Wu and Hamada (2009). The aim here therefore is to introduce a small example in order to demonstrate, albeit briefly, how the basic concepts of blocking in s^n factorial experiments emanate from the formulation of an appropriate linear model not of full rank.

Consider a 3^2 design, that is, a design comprising two factors A and B each at 3 levels, 0, 1 and 2. Then the responses to the nine treatment combinations, coded as ij for $i, j = 0, 1, 2$, can be modelled as

$$y_{ij} = \mu + A_i + B_j + AB_{i+j} + AB_{i+2j}^2 + \epsilon_{ij}$$

where A_i and B_j are main effects and AB_{i+j} and AB_{i+2j}^2 are interaction terms with arithmetic for the subscripts performed modulo 3. Thus, for example, $y_{12} = \mu + A_1 + B_2 + AB_0 + AB_2^2 + \epsilon_{12}$. It then follows immediately from the form of the observational equations for a single replicate that the model is not of full rank but that the treatment means and the contrasts for individual treatment effects, for example, $A_0 - A_2$ and $AB_0 - 2AB_1 + AB_2$, are estimable. Suppose now that replicates of a 3^2 experiment are to be arranged in 3 blocks of size 3. The set of nine treatment combinations can be partitioned into the three subsets $\{00, 11, 22\}$, $\{10, 21, 02\}$ and $\{20, 01, 12\}$ which are associated with the interaction terms AB_0^2, AB_1^2 and AB_2^2 , respectively. These subsets can then be taken as specifying the blocks for each replicate, with the interaction AB^2 thereby confounded with blocks. Contrasts involving A , B , and AB , separately, remain estimable. However, while contrasts of treatment means associated with the same level of AB^2 are estimable, the treatment means themselves are not.

2.5 Linear Mixed Model

The linear mixed model was introduced earlier but, for clarity, is defined again more compactly here. The model can be formulated as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}\mathbf{u} + \boldsymbol{\epsilon}$$

where \mathbf{y} is the $n \times 1$ vector of responses, $\boldsymbol{\beta}$ is a $(p + 1) \times 1$ vector of unknown fixed effects with attendant model matrix \mathbf{X} , \mathbf{u} is a $q \times 1$ vector of random effects with attendant model matrix \mathbf{Z} , and $\boldsymbol{\epsilon}$ is an $n \times 1$ vector of error terms. The random effects are taken to be distributed as $\mathbf{u} \sim r.v.(\mathbf{0}, \mathbf{G})$, where *r.v.* indicates a random variable, the error terms as $\boldsymbol{\epsilon} \sim r.v.(\mathbf{0}, \mathbf{R})$, and \mathbf{u} and $\boldsymbol{\epsilon}$ are taken to be independent. The variance-covariance matrix of \mathbf{y} is assumed to depend on a vector of variance components $\boldsymbol{\theta}$ through the matrices \mathbf{G} and \mathbf{R} and is expressed succinctly as $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}(\boldsymbol{\theta}) = \mathbf{Z}\mathbf{G}\mathbf{Z}' + \mathbf{R}$. Note that for ease of exposition, the matrix \mathbf{X} is taken here to be of full column rank. Matrices \mathbf{X} not of full column rank can be accommodated by introducing *g*-inverses and incorporating notions relating to estimability.

2.5.1 Theory

2.5.1.1 Estimation

Suppose that the random effects \mathbf{u} and the errors $\boldsymbol{\epsilon}$ are normally distributed. Then the distribution of \mathbf{y} is normal and is given by

$$\mathbf{y} \sim N(\mathbf{X}\boldsymbol{\beta}, \mathbf{ZGZ}' + \mathbf{R}),$$

and the log-likelihood thus follows in the usual way as

$$l = l(\boldsymbol{\beta}, \boldsymbol{\theta}; \mathbf{y}) = -\frac{n}{2} \ln(2\pi) - \frac{1}{2} \ln(\det(\boldsymbol{\Sigma})) - \frac{1}{2} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta})' \boldsymbol{\Sigma}^{-1} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta}).$$

If the variance components $\boldsymbol{\theta}$ are known, then the MLEs of the fixed effects $\boldsymbol{\beta}$ follow immediately from the likelihood equations $\frac{\partial l}{\partial \boldsymbol{\beta}} = \mathbf{0}$ as

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}' \boldsymbol{\Sigma}^{-1} \mathbf{X})^{-1} \mathbf{X}' \boldsymbol{\Sigma}^{-1} \mathbf{y} \sim N(\boldsymbol{\beta}, (\mathbf{X}' \boldsymbol{\Sigma}^{-1} \mathbf{X})^{-1}), \quad (2.7)$$

and all the results relating to the generalized least squares setting accrue. However, if the parameters $\boldsymbol{\theta}$ are not known, then the situation is more complicated in that $\boldsymbol{\theta}$ must be estimated and the distribution of the estimate of $\boldsymbol{\beta}$ is, at least in general, intractable. Specifically, consider solving the ML equations $\frac{\partial l}{\partial \boldsymbol{\beta}} = \mathbf{0}$ and $\frac{\partial l}{\partial \boldsymbol{\theta}} = \mathbf{0}$ simultaneously. Observe that this approach yields

$$\begin{aligned} \frac{\partial l}{\partial \boldsymbol{\theta}} &= -\frac{1}{2} \frac{\partial \ln(\det(\boldsymbol{\Sigma}))}{\partial \boldsymbol{\theta}} - \frac{1}{2} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta})' \frac{\partial \boldsymbol{\Sigma}^{-1}}{\partial \boldsymbol{\theta}} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta}) \\ &= -\frac{1}{2} \text{trace} \left(\frac{\partial \boldsymbol{\Sigma}}{\partial \boldsymbol{\theta}} \boldsymbol{\Sigma}^{-1} \right) + \frac{1}{2} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta})' \boldsymbol{\Sigma}^{-1} \frac{\partial \boldsymbol{\Sigma}}{\partial \boldsymbol{\theta}} \boldsymbol{\Sigma}^{-1} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta}) \\ &= \mathbf{0} \end{aligned}$$

with $\boldsymbol{\beta}$ a function of $\boldsymbol{\theta}$, that is, $\boldsymbol{\beta}(\boldsymbol{\theta}) = (\mathbf{X}' \boldsymbol{\Sigma}^{-1} \mathbf{X})^{-1} \mathbf{X}' \boldsymbol{\Sigma}^{-1} \mathbf{y}$. Thus, in order to obtain solutions to the ML equations for $\boldsymbol{\theta}$, the expression for $\boldsymbol{\beta} = \boldsymbol{\beta}(\boldsymbol{\theta})$ can be *plugged* into $\frac{\partial l}{\partial \boldsymbol{\theta}} = \mathbf{0}$ to yield

$$\text{trace} \left(\frac{\partial \boldsymbol{\Sigma}}{\partial \boldsymbol{\theta}} \boldsymbol{\Sigma}^{-1} \right) = \mathbf{y}' \mathbf{P} \frac{\partial \boldsymbol{\Sigma}}{\partial \boldsymbol{\theta}} \mathbf{P} \mathbf{y} \quad (2.8)$$

where \mathbf{P} is the symmetric matrix expressed succinctly as

$$\mathbf{P} = \mathbf{P}(\boldsymbol{\theta}) = \boldsymbol{\Sigma}^{-1} - \boldsymbol{\Sigma}^{-1} \mathbf{X} (\mathbf{X}' \boldsymbol{\Sigma}^{-1} \mathbf{X})^{-1} \mathbf{X}' \boldsymbol{\Sigma}^{-1}$$

and depends on $\boldsymbol{\theta}$. These equations are, in general, highly nonlinear in $\boldsymbol{\theta}$ and must be solved numerically. However, the solutions are problematic in that they do not always comply with the inference region, that is, with constraints on $\boldsymbol{\theta}$ to ensure that variances are nonnegative, and more generally that $\boldsymbol{\Sigma}$ itself is nonnegative definite. The MLEs for $\boldsymbol{\theta}$ are thus necessarily awkward to handle.

Many problems associated with the MLEs for the variance components θ are rooted in the fact that these estimates in turn incorporate estimates of the fixed effects β . To counter this, Patterson and Thompson (1971) introduced the notion of restricted or residual maximum likelihood estimation, that is, REML. Specifically a vector $w = K'y$ with $K'X = 0$, which comprises the maximum number of linearly independent error contrasts, is introduced and thus $w \sim N(0, K'\Sigma K)$. The vector w contains no information about β and has degrees of freedom $n - (p + 1)$, that is, $\text{rank}(K'\Sigma K) = n - (p + 1)$. The restricted log-likelihood for θ given the error contrasts $K'y$ can now be expressed by

$$l_R(\theta; K'y) = -\frac{(n - (p + 1))}{2} - \frac{1}{2} \ln(\det(K'\Sigma K)) - \frac{1}{2} y'K(K'\Sigma K)^{-1}K'y.$$

Furthermore, and crucially, it can then be shown by some subtle but rather tedious algebra that $l_R(\theta; K'y)$ is invariant to the specific choice of the matrix K and is given, up to an additive constant, by

$$-\frac{1}{2} \ln(\det(\Sigma)) - \frac{1}{2} \ln(\det(X'\Sigma^{-1}X)) - \frac{1}{2} y'Py,$$

(Harville 1977; Searle et al. 1992; Khuri 2010). The REML estimators for θ , written $\hat{\theta}_{REML}$, can then be obtained by solving the likelihood equation $\frac{\partial l_R(\theta; K'y)}{\partial \theta} = 0$ subject to the constraints of the inference region. In general, REML estimates of θ are to be preferred to the MLE since they eliminate the fixed effects and possess appropriate degrees of freedom. Strictly, there are no estimates for β associated with the REML approach, but clearly estimates of β can be obtained by plugging the estimator $\hat{\theta}_{REML}$ into the expression for $\hat{\beta}$.

Overall therefore, the estimate of the fixed effects parameters β is of the form

$$\hat{\beta} = (X'\Sigma^{-1}(\hat{\theta})X)^{-1} X'\Sigma(\hat{\theta})^{-1} y,$$

where the estimate of θ , denoted generically $\hat{\theta}$, is obtained either by ML or by REML. This estimate is readily computed, but its distribution is extremely complicated. In particular, the variance of $\hat{\beta}$ is of immediate interest, but small sample approximations are somewhat intractable (Kackar and Harville 1984). The estimates of the fixed effects β and the variance components θ , both ML and REML, are however based on likelihood, and the associated variance matrix can therefore be approximated asymptotically by the inverse of the Fisher information matrix. Furthermore, since the parameters β are not included in the variance components, it follows immediately that the information matrix for β and θ is given by the block diagonal matrix (2.5), that is, by

$$\begin{bmatrix} X'\Sigma^{-1}X & 0 \\ 0 & \frac{1}{2} \left\{ \text{trace} \left(\Sigma^{-1} \frac{\partial \Sigma}{\partial \theta_r} \Sigma^{-1} \frac{\partial \Sigma}{\partial \theta_s} \right) \right\}_{r,s} \end{bmatrix}.$$

This matrix depends on the matrix X , on the matrix Z through the relation $\Sigma = ZGZ' + R$ and, of course, on the variance components θ . This dependence in turn impacts on the construction of optimal designs based on information and is explored a little further in Section 2.5.2.

2.5.1.2 Prediction

In many applications based on the linear mixed model, particularly those in animal genetics, interest focuses on prediction of quantities involving the random effects. There are in fact a number of approaches to the problem of prediction, and the more informative ones are summarized briefly here. Full details are available in the seminal article by Robinson (1991) and in the books by McCulloch and Searle (2001) and Ruppert et al. (2003).

A particularly straightforward approach to prediction concentrates on the random effects u alone assuming, at least initially, that β and θ are known. The aim is to derive the best predictor of u , say \tilde{u} , where \tilde{u} is best in the sense that it minimizes the mean squared prediction error (MSPE), that is, $E[(\tilde{u} - u)'(\tilde{u} - u)]$. Thus the required predictor is given by the conditional expectation $E(u|y)$. Furthermore, if the random effects u and the error terms ϵ are normally distributed, then

$$\begin{bmatrix} y \\ u \end{bmatrix} \sim N \left(\begin{bmatrix} X\beta \\ 0 \end{bmatrix}, \begin{bmatrix} \Sigma & ZG \\ GZ' & G \end{bmatrix} \right),$$

and it follows immediately that $\tilde{u} = E(u|y) = GZ'\Sigma^{-1}(y - X\beta)$ with $\Sigma = ZGZ' + R$. As a counterpoint to this, the best linear predictor of u can be derived without making any distributional assumptions about u and ϵ and coincides with \tilde{u} . In practice, the fixed effects β are invariably unknown. Thus, assuming that the variance components θ are known, a naive *practical predictor* for u can be introduced as $\tilde{u} = GZ'\Sigma^{-1}(y - X\hat{\beta}) = GZ'Py$ where $\hat{\beta}$, the ML estimator of β given in (2.7), is plugged in for β .

A more stringent and cohesive approach to prediction which incorporates both the fixed effects β and the random effects u is that of best linear unbiased prediction (BLUP), where *best* is again interpreted in the sense of minimizing the MSPE. The approach has no distributional requirements and can be regarded as the prediction equivalent of BLUE. Specifically, consider the best linear predictor of a linear function of the form $a'X\beta + c'Zu$, where a and c are specified vectors, subject to the constraint that the linear function is unbiased, that is, $E(a'X\beta + c'Zu) = a'X\beta$. Then it follows by some intricate algebra that the BLUP of β is given by $\tilde{\beta} = (X'\Sigma^{-1}X)^{-1}X'\Sigma^{-1}y$, coinciding with the MLE of β under normality and with the BLUE of β , and the BLUP of u is given by $\tilde{u} = GZ'\Sigma^{-1}(y - X\tilde{\beta})$, coinciding with the naive predictor (McCulloch and Searle 2001).

An alternative approach to deriving the BLUPs, following Henderson et al. (1959), is to take the random effects u and the error terms ϵ to be normally distributed and to introduce the pseudo-p.d.f.

$$\begin{aligned} f(y, u; \beta, \theta) &= f(y|u)f(u), \\ &= \frac{1}{(2\pi)^{\frac{n}{2}} \det(\Sigma)^{\frac{1}{2}}} \exp \left\{ -\frac{1}{2}(y - X\beta - Zu)' \Sigma^{-1}(y - X\beta - Zu) \right\} \\ &\quad \times \frac{1}{(2\pi)^{\frac{q}{2}} \det(G)^{\frac{1}{2}}} \exp \left\{ -\frac{1}{2}u'G^{-1}u \right\} \end{aligned}$$

which treats the random effects u as observed. Then a hierarchical log-likelihood can be formulated as $l_H(\beta, \theta; y, u) = \ln(f(y, u; \beta, \theta))$. Estimates of β and predictors of u can be obtained by solving

$$\frac{\partial l_H(\boldsymbol{\beta}; \mathbf{y}, \mathbf{u})}{\partial \boldsymbol{\beta}} = \mathbf{0} \quad \text{and} \quad \frac{\partial l_H(\boldsymbol{\beta}; \mathbf{y}, \mathbf{u})}{\partial \mathbf{u}} = \mathbf{0},$$

and thus by solving the equations

$$\begin{bmatrix} \mathbf{X}'\mathbf{R}^{-1}\mathbf{X} & \mathbf{X}'\mathbf{R}^{-1}\mathbf{Z} \\ \mathbf{Z}'\mathbf{R}^{-1}\mathbf{X} & \mathbf{Z}'\mathbf{R}^{-1}\mathbf{Z} + \mathbf{G}^{-1} \end{bmatrix} \begin{bmatrix} \tilde{\boldsymbol{\beta}} \\ \tilde{\mathbf{u}} \end{bmatrix} = \begin{bmatrix} \mathbf{X}'\mathbf{R}^{-1}\mathbf{y} \\ \mathbf{Z}'\mathbf{R}^{-1}\mathbf{y} \end{bmatrix},$$

which are termed the mixed model equations. The immediate solutions to these equations for the BLUPs $\tilde{\boldsymbol{\beta}}$ and $\tilde{\mathbf{u}}$ are of a different algebraic form to those given previously but can be shown, by some careful manipulation, to be the same (McCulloch and Searle 2001). The mixed model equations are inherently attractive, and solving them directly is computationally efficient.

In order to draw inferences about the fixed effects and the random effects, it is necessary to derive appropriate variances for the BLUPs, $\tilde{\boldsymbol{\beta}}$ and $\tilde{\mathbf{u}}$. The variance components $\boldsymbol{\theta}$ are, at least initially, assumed to be known. Since $E(\tilde{\mathbf{u}}) = \mathbf{0}$ and $\mathbf{X}'\mathbf{P} = \mathbf{0}$ where, as before, $\mathbf{P} = \boldsymbol{\Sigma}^{-1} - \boldsymbol{\Sigma}^{-1}\mathbf{X}(\mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{X})^{-1}\mathbf{X}'\boldsymbol{\Sigma}^{-1}$, it follows that

$$\text{Var} \begin{bmatrix} \tilde{\boldsymbol{\beta}} \\ \tilde{\mathbf{u}} \end{bmatrix} = \begin{bmatrix} (\mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{X})^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{GZ}'\mathbf{PZG} \end{bmatrix}$$

(Searle et al. 1992). However, while the BLUP $\tilde{\boldsymbol{\beta}}$ coincides with the MLE for $\boldsymbol{\beta}$ and predicts or estimates the fixed effects, in contrast, the BLUP $\tilde{\mathbf{u}}$ is a predictor of the random effects \mathbf{u} . Thus the prediction error variance $\text{Var}(\tilde{\mathbf{u}} - \mathbf{u})$ is more meaningful than the variance $\text{Var}(\tilde{\mathbf{u}})$ and is to be preferred. The prediction error variance matrix for $\tilde{\boldsymbol{\beta}}$ and $\tilde{\mathbf{u}}$ is therefore given by

$$\text{Var} \begin{bmatrix} \tilde{\boldsymbol{\beta}} \\ \tilde{\mathbf{u}} - \mathbf{u} \end{bmatrix} = \begin{bmatrix} (\mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{X})^{-1} & -(\mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{X})^{-1}\mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{ZG} \\ -\mathbf{GZ}'\boldsymbol{\Sigma}^{-1}\mathbf{X}(\mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{X})^{-1} & \mathbf{G} - \mathbf{GZ}'\mathbf{PZG} \end{bmatrix}.$$

Furthermore, if a linear function of the form $g = g(\boldsymbol{\beta}, \mathbf{u}) = \mathbf{a}'\mathbf{X}\boldsymbol{\beta} + \mathbf{c}'\mathbf{Z}\mathbf{u}$, with BLUP given by $\tilde{g} = g(\tilde{\boldsymbol{\beta}}, \tilde{\mathbf{u}}) = \mathbf{a}'\mathbf{X}\tilde{\boldsymbol{\beta}} + \mathbf{c}'\mathbf{Z}\tilde{\mathbf{u}}$, is of interest, then the prediction error variance follows immediately as

$$\begin{aligned} \text{Var}(\tilde{g} - g) &= \text{Var}(\mathbf{a}'\mathbf{X}\tilde{\boldsymbol{\beta}} + \mathbf{c}'\mathbf{Z}(\tilde{\mathbf{u}} - \mathbf{u})) \\ &= \mathbf{c}'\mathbf{ZGZ}'\mathbf{c} - \mathbf{d}'\boldsymbol{\Sigma}^{-1}\mathbf{d} + (\mathbf{a}' - \mathbf{d}'\boldsymbol{\Sigma}^{-1})\mathbf{X}(\mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{X})^{-1}\mathbf{X}'(\mathbf{a} - \boldsymbol{\Sigma}^{-1}\mathbf{d}) \end{aligned}$$

where $\mathbf{d}' = \mathbf{c}'\mathbf{ZGZ}'$ and this expression can be used in inference. Similarly if the BLUP of $E(\mathbf{y}|\mathbf{u})$, that is, the vector of fitted values $\tilde{\mathbf{y}} = \mathbf{X}\tilde{\boldsymbol{\beta}} + \mathbf{Z}\tilde{\mathbf{u}}$, is of interest, then the associated prediction error variance matrix is given by

$$\begin{aligned} \text{Var}(\tilde{\mathbf{y}} - E[\mathbf{y}|\mathbf{u}]) &= \mathbf{X}(\mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{X})^{-1}\mathbf{X}' + \mathbf{ZGZ}' - \mathbf{ZGZ}'\mathbf{PZGZ}' \\ &\quad - \mathbf{X}(\mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{X})^{-1}\mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{ZGZ}' - \mathbf{ZGZ}'\boldsymbol{\Sigma}^{-1}\mathbf{X}(\mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{X})^{-1}\mathbf{X}'. \end{aligned}$$

The discussion on prediction thus far is predicated on the assumption that the variance components θ are known. However, this is rarely the case, and ML or REML estimators of θ must be incorporated into the prediction process. The obvious expedient of plugging an estimator of θ , generically $\hat{\theta}$, into expressions for the BLUPs is therefore adopted. Specifically the estimated BLUP of the random effects u , termed the EBLUP, is introduced as

$$\tilde{u} = Z'G(\hat{\theta})\Sigma^{-1}(\hat{\theta})(y - X\hat{\beta}),$$

but its distribution is not tractable. To compound matters, the plug-in estimate of the prediction error variance of the BLUP of the linear function $g = g(\beta, u) = a'X\beta + c'Zu$, termed the empirical mean squared prediction error (EMSPE) and obtained quite simply by plugging $\hat{\theta}$ into the expression for the prediction error variance $\text{Var}(\tilde{g} - g)$ with θ known, tends to underestimate the prediction error variance of the EBLUP. Approximations to this plug-in prediction error variance which improve on its properties and in particular the bias have been developed by Kackar and Harville (1984), Prasad and Rao (1990) and Harville and Jeske (1992) but are somewhat complicated to implement.

2.5.2 Examples and Design

In order to delve a little more deeply into the ideas which underpin the linear mixed model, particularly issues relating to design, two broad areas of application are introduced. Specifically, attention is focused first on settings for which multiple observations are taken on individual units and second on models, such as those for kriging and splines, which can be cast as linear mixed models.

2.5.2.1 Observations on Individuals

Consider a group of K individuals, for example, patients in a clinical trial. Suppose that d_i observations are taken on the i th individual, $i = 1, \dots, K$, giving a total of $n = \sum_{i=1}^K d_i$ observations. Then the linear mixed model for the i th individual can be specified as

$$y_i = X_i\beta + Z_iu_i + \epsilon_i,$$

where y_i is the $d_i \times 1$ vector of responses, β is a $(p + 1) \times 1$ vector of unknown fixed effects common to all K individuals, u_i is a $q \times 1$ vector of random effects specific to the i th individual, X_i and Z_i are conformable matrices of explanatory variables or functions of those variables associated with the fixed and the random effects, respectively, and ϵ_i is a $d_i \times 1$ vector of error terms. The random effects are taken to be distributed as $u_i \sim N(\mathbf{0}, G)$ and represent between-individual error, the error terms are assumed to be distributed as $\epsilon_i \sim N(\mathbf{0}, R_i)$ and represent within-individual error and u_i and ϵ_i are taken to be independent both within and between individuals. In addition, the matrices G and $R_i, i = 1, \dots, K$, are assumed to depend on a vector of variance components θ common to all K individuals. Clearly therefore, the response y_i is distributed as $N(X_i\beta, \Sigma_i)$ with the variance matrix $\Sigma_i = Z_iGZ_i' + R_i$ a function of the variance components $\theta, i = 1, \dots, K$.

The models for the individuals can be assembled as the linear mixed model

$$y = X\beta + Zu + \epsilon,$$

where, since observations between individuals are independent, $\mathbf{X}' = (\mathbf{X}'_1, \dots, \mathbf{X}'_K)$, $\mathbf{Z} = \text{diag}(\mathbf{Z}_1, \dots, \mathbf{Z}_K)$, $\mathbf{u}' = (\mathbf{u}'_1, \dots, \mathbf{u}'_K)$ and $\boldsymbol{\epsilon}' = (\boldsymbol{\epsilon}'_1, \dots, \boldsymbol{\epsilon}'_K)$ and thus

$$\mathbf{y} \sim N(\mathbf{X}\boldsymbol{\beta}, \boldsymbol{\Sigma}) \quad \text{where } \boldsymbol{\Sigma} = \text{diag}(\boldsymbol{\Sigma}_1, \dots, \boldsymbol{\Sigma}_K).$$

Furthermore, the information matrix for the fixed effects $\boldsymbol{\beta}$ is readily derived as $\sum_{i=1}^K \mathbf{X}'_i \boldsymbol{\Sigma}_i^{-1} \mathbf{X}_i$ and that for the variance components as

$$\frac{1}{2} \sum_{i=1}^K \left\{ \text{trace} \left(\boldsymbol{\Sigma}_i^{-1} \frac{\partial \boldsymbol{\Sigma}_i}{\partial \theta_r} \boldsymbol{\Sigma}_i^{-1} \frac{\partial \boldsymbol{\Sigma}_i}{\partial \theta_s} \right) \right\}_{r,s}.$$

Note that both these information matrices depend on the variance components $\boldsymbol{\theta}$. Thus, in the design phase, if $\boldsymbol{\theta}$ is unknown, then a best guess may be adopted and designs, termed locally optimal designs, constructed (Chernoff 1953).

The general setting for the linear mixed model just described is widely used in the modelling of longitudinal data, that is, data comprising observations made on a number of individuals over time. To fix ideas, suppose that a simple linear regression model with a random intercept and independent error terms is appropriate for modelling observations on each individual. Then the model for the i th individual is given by

$$\mathbf{y}_i = \beta_0 \mathbf{1} + \beta_1 \mathbf{t}_i + u_i \mathbf{1} + \boldsymbol{\epsilon}_i, \quad (2.9)$$

where \mathbf{y}_i is a $d_i \times 1$ vector of observations taken at time points specified by the $d_i \times 1$ vector \mathbf{t}_i , β_0 and β_1 are unknown fixed effects parameters, u_i is a random term distributed as $N(0, \sigma_u^2)$, thereby introducing the random intercept term $\beta_0 + u_i$, and the error terms $\boldsymbol{\epsilon}_i$ are distributed as $N(\mathbf{0}, \sigma_\epsilon^2 \mathbf{I})$. It now follows that $\boldsymbol{\Sigma}_i = \sigma_\epsilon^2 \mathbf{I}_{d_i} + \sigma_u^2 \mathbf{J}_{d_i}$, that

$$\boldsymbol{\Sigma}_i^{-1} = \frac{1}{\sigma_\epsilon^2} \left(\mathbf{I} - \frac{\gamma}{1 + d_i \gamma} \mathbf{J} \right)$$

where $\gamma = \sigma_u^2 / \sigma_\epsilon^2$ and hence the information matrix for the parameters β_0 and β_1 is given by

$$\mathbf{M}(\mathbf{t}_i) = \frac{1}{\sigma_\epsilon^2 (1 + d_i \gamma)} \begin{bmatrix} d_i & \mathbf{1}' \mathbf{t}_i \\ \mathbf{t}'_i \mathbf{1} & (1 + d_i \gamma) \mathbf{t}'_i \mathbf{t}_i - \gamma (\mathbf{1}' \mathbf{t}_i)^2 \end{bmatrix}$$

(Debushe and Haines 2008). The total information matrix is therefore $\sum_{i=1}^K \mathbf{M}(\mathbf{t}_i)$. The question now arises as to how to design such a longitudinal study. For example, suppose that 100 individuals are available, exactly 3 observations from the time points 0, 1, ..., 9 and 10 can be taken without repetition on each individual, and that the precise estimation of the fixed effects β_0 and β_1 is of interest. Should observations be taken on 100 individuals at the 3 time points (0, 5, 10) or on 30 individuals at the time points (0, 2, 10), 40 at (0, 5, 10) and 30 at (0, 8, 10), or should some other allocation be adopted? In fact, it can be shown, using the approach outlined as follows, that for all values of the variance ratio γ , the requisite design puts 50 individuals at the time points (0, 1, 10) and 50 at (0, 9, 10).

Questions such as the one just posed can be formalized and answered within the general longitudinal setting by introducing the notion of a population design. Broadly, a population design is a design in which groups of individuals are allocated to distinct design settings. More formally, n_i individuals are allocated to the distinct design setting (X_i, Z_i) which comprises d_i observations, $i = 1, \dots, t$. Thus the number of distinct design settings is t , and the total number of individuals in the design is now $K = \sum_{i=1}^t n_i d_i$. The design can be represented as

$$\xi_n = \left\{ \begin{matrix} (X_1, Z_1), & (X_2, Z_2) & \dots, & (X_t, Z_t) \\ n_1, & n_2, & \dots, & n_t \end{matrix} \right\},$$

and the attendant information matrix is given by $\sum_{i=1}^t n_i X_i' \Sigma_i^{-1} X_i$. The ideas which underpin optimal design for the linear model of full rank delineated in Section 2.3.1 can now be invoked within the context of population designs, but with points x_i replaced by distinct design settings (X_i, Z_i) for $i = 1, \dots, t$. In particular, alphabetic optimality criteria can be introduced, the notion of an approximate design formulated and exact designs constructed numerically using exchange-type algorithms. The only point of departure is that the information matrix in the linear mixed model setting depends sensitively on the variance components. A best guess for these components suffices, but care must be exercised for random coefficient models other than the random intercept model (Longford 1994). For example, if the slope and intercept in the simple linear regression model have random effects with variance matrix G , then the specification of G depends crucially on the location of the time points.

Two further points of interest in relation to the modelling of longitudinal data using linear mixed models should be noted. First, there have only been a few reports in the literature on designs, which in some sense maximize the information on the variance components. These include a review article by Khuri (2000) and insightful and informative papers by Giovagnoli and Sebastiani (1989), Liu et al. (2007), and Loeza-Serrano and Donev (2014). Second, nonlinear models are often used to model individual responses in longitudinal studies. In the context of design, these models are commonly linearized, and thus, many of the considerations for the linear mixed model outlined previously apply, except that the variance components now include the regression parameters (Mentre et al. 1997).

Block designs for which the block effects are assumed to be random can be effectively cast in the linear mixed model framework. For ease of exposition, the notation introduced here complies with that used for the block designs discussed in Section 2.4.2 and deviates somewhat from the notation for the linear mixed model used earlier. Thus consider a block design comprising b blocks of equal size k and t treatments replicated r_i times, $i = 1, \dots, t$. Then the $k \times 1$ vector of observations in the j th block, denoted y_j , can be modelled as

$$y_j = \mu \mathbf{1} + A_j \tau + u_j \mathbf{1} + \epsilon_j,$$

where μ is the mean, τ is the $t \times 1$ vector of fixed treatment effects, A_j is the $k \times t$ unit/treatment incidence matrix for the j th block, and the error term ϵ_j is distributed as $N(\mathbf{0}, \sigma_e^2 \mathbf{I})$. In addition, the term u_j represents the random effect for the j th block and is assumed to be distributed as $N(0, \sigma_u^2)$ independently of all other random effects and error terms. Thus it follows, using arguments similar to those invoked for the random intercept model (2.9) introduced earlier, that the information matrix for the treatment effects τ is given by

$$\frac{1}{\sigma_e^2(1+k\gamma)} \begin{bmatrix} bk & \sum_{j=1}^b \mathbf{1}'A_j \\ \sum_{j=1}^b A_j'\mathbf{1} & (1+k\gamma) \sum_{j=1}^b A_j'A_j - \gamma \sum_{j=1}^b A_j'\mathbf{1}\mathbf{1}'A_j \end{bmatrix}$$

where $\gamma = \sigma_u^2/\sigma_e^2$. A little reflection then shows that this matrix can be expressed as

$$\frac{1}{\sigma_e^2(1+k\gamma)} \begin{bmatrix} bk & \mathbf{r}' \\ \mathbf{r} & (1+k\gamma)\mathbf{R} - \gamma\mathbf{N}\mathbf{N}' \end{bmatrix}$$

in the usual block design notation. Furthermore, by invoking the sweep procedure and performing some tedious algebra, the information matrix for the treatments eliminating the fixed mean effect can be shown to be given, up to a multiplying constant, by

$$C_R = (\mathbf{R} - \frac{1}{k}\mathbf{N}'\mathbf{N}) + \frac{1}{1+k\gamma} \left(\frac{1}{k}\mathbf{N}'\mathbf{N} - \frac{1}{bk}\mathbf{r}\mathbf{r}' \right)$$

(Bailey 2009).

From a design perspective, similar considerations apply to the treatment information matrix C_R with random block effects to those which apply to the corresponding matrix C for block designs with fixed block effects introduced in Section 2.4.2. Thus, if the design is connected and the ratio γ is known or a best guess is available, then A -, MV -, D - and E -optimal designs can be constructed. From a broader perspective, the interpretation of the random block design itself is of interest. Thus the model with $\gamma = 0$ corresponds to no block effects and in the limit as γ approaches infinity to fixed block effects. More generally, the model combines intra- and interblock information in a seamless manner and can be extended for example to split-plot designs for which traditional methods of analysis have proved awkward (Goos 2002). Finally, note that this example differs from that for longitudinal studies, first, in that the model matrices associated with the fixed effects are not of full rank and, second, in that population designs are not relevant since blocks tend not to be repeated.

2.5.2.2 Kriging and Other Models

The universal kriging model for spatial data can be formulated as a linear mixed model, and certain advantages accrue from this insight. Specifically, consider the model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{u} + \boldsymbol{\epsilon}$$

where \mathbf{y} is an $n \times 1$ vector of observations taken at n locations in a spatial field, $\boldsymbol{\beta}$ is a $(p + 1) \times 1$ vector of unknown parameters, and \mathbf{X} is an $n \times (p + 1)$ matrix comprising explanatory variables and functions of those variables at the specified locations. The error structure is captured in the terms \mathbf{u} and $\boldsymbol{\epsilon}$, with \mathbf{u} taken from a stationary random field, usually Gaussian, with correlations dependent on distances between the locations at which observations are taken. More specifically, $\mathbf{u} \sim N(\mathbf{0}, \sigma^2\mathbf{R}_c)$ where \mathbf{R}_c is an appropriate correlation matrix and $\boldsymbol{\epsilon}$ is assumed to be distributed as $N(\mathbf{0}, \tau^2\mathbf{I})$ independently of \mathbf{u} . Note that, within the spatial context, the variances τ^2 and σ^2 are termed the nugget and the partial

sill, respectively. The terms \mathbf{u} are interpreted here as random effects. The variance of \mathbf{y} is thus given by $\boldsymbol{\Sigma} = \tau^2\mathbf{I} + \sigma^2\mathbf{R}_c$ and the estimation of the parameters σ^2 , τ^2 and any parameters embedded in the matrix \mathbf{R}_c follows in the usual way, with REML estimates generally favoured over MLEs.

Interest then centres on the prediction of an unknown observation y_0 at a given location with explanatory variables \mathbf{x}_0 already recorded there. The derivation of the BLUP for y_0 is a little different to that for the random effects \mathbf{u} themselves and in fact represents the more general case of prediction (Harville 1991). The resultant predictor is given by

$$\tilde{y}_0 = \mathbf{x}'_0\hat{\boldsymbol{\beta}} + \mathbf{v}'_0\boldsymbol{\Sigma}^{-1}(\mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}})$$

and the prediction or kriging variance, $\text{Var}(\tilde{y}_0 - y_0)$, by

$$(\sigma^2 + \tau^2) - \mathbf{v}'_0\boldsymbol{\Sigma}^{-1}\mathbf{v}_0 + (\mathbf{x}'_0 - \mathbf{v}'_0\boldsymbol{\Sigma}^{-1}\mathbf{X})(\mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{X})^{-1}(\mathbf{x}_0 - \mathbf{X}'\boldsymbol{\Sigma}^{-1}\mathbf{v}_0)$$

where \mathbf{v}_0 is the vector of covariances of the unknown observation y_0 with the known observations \mathbf{y} . Two key problems in design for spatial data can be identified, that of deciding at which location to next take an observation in the random field and that of which node to remove and which to add in a monitoring network. Design criteria which accommodate these goals can be formulated in terms of the kriging variance. For example, if the problem relates to deciding on the location of the next observation, then the location for which the kriging variance is a maximum can be sought.

Models for spatial data are of course not limited simply to the kriging model. There is a vast literature on spatial statistics, including the classic text of Cressie (1993) and the more recent books by Diggle and Ribeiro Jr. (2007) and Cressie and Wikle (2011). Design issues are specifically discussed in the text by Müller (2007) and the recent volume edited by Mateu and Müller (2013), and cutting edge issues relating to optimal design for spatial data are presented in Chapter 15. It should also be noted that kriging models can be extended to random fields in more than two dimensions and that such models without a nugget effect, that is, without the error terms, are used extensively in the design and analysis of computer experiments. More details, particularly relating to design, are available in the books by Santner et al. (2003) and by Fang et al. (2006) and in Chapters 16 through 19.

Finally, it should be noted that there are many settings ranging from ill-posed inverse problems, through to the Kalman filter and smoothing splines, which can also be mapped onto the linear mixed model. Some are mentioned briefly in the discussion of the paper by Robinson (1991), while the notion that a smoothing spline can be formulated as a linear mixed model is carefully explored in the seminal text by Ruppert et al. (2003). It is however arguable as to whether full advantage of these mappings has been taken within the context of the design of experiments.

2.6 Conclusions

The theory underpinning the linear model of full rank, the linear model not of full rank and the linear mixed model has been introduced, and notions relating to the design of experiments for these models have been developed and interwoven into the text. The discussion

is compact, and the interested reader can delve more deeply into the theorems, proofs and design concepts through the many excellent texts that are readily available and are included in the references which follow. The aim of this chapter is to provide a sound basis and some necessary building blocks for an understanding of the specialized chapters which follow.

Acknowledgments

Much of this chapter emanates from lectures given by Professor Peter Clarke and by the late Professor Arthur Rayner at the former University of Natal in Pietermaritzburg, South Africa. I am indebted to them both for this superb grounding. I would also like to thank the University of Cape Town and the National Research Foundation (NRF) of South Africa, grant (UID) 85456, for financial support. Any opinion, finding and conclusion or recommendation expressed in this material is that of the author and the NRF does not accept liability in this regard.

References

- Atkinson, A. C. and Cook, R. D. (1995), D-optimum designs for heteroscedastic linear models, *Journal of the American Statistical Association*, 90, 204–212.
- Atkinson, A. C., Donev, A. N., and Tobias, R. (2007), *Optimum Experimental Designs, with SAS*, Oxford University Press, Oxford, U.K.
- Bailey, R. A. (2004), *Association Schemes: Designed Experiments, Algebra and Combinatorics*, Cambridge University Press, Cambridge, U.K.
- Bailey, R. A. (2008), *Design of Comparative Experiments*, Cambridge University Press, Cambridge, U.K.
- Bailey, R. A. (2009), Designs for dose–escalation trials with quantitative responses, *Statistics in Medicine*, 28, 3721–3738.
- Ben-Israel, A. and Greville, T. N. E. (2003), *Generalized Inverses: Theory and Applications*, Springer, New York.
- Box, G. E. P., Hunter, J. S., and Hunter, W. G. (2005), *Statistics for Experimenters: Design, Innovation, and Discovery*, 2nd edn., Wiley, Hoboken, NJ.
- Caliński, T. and Kageyama, S. (2000), *Block Designs: A Randomization Approach*, Volume 1: Analysis, Springer, New York.
- Chernoff, H. (1953), Locally optimal designs for estimating parameters, *The Annals of Mathematical Statistics*, 24, 586–602.
- Cochran, W. G. and Cox, G. M. (1957), *Experimental Designs*, Wiley, New York.
- Cressie, N. and Wikle, C. K. (2011), *Statistics for Spatio-Temporal Data*, Wiley, Hoboken, NJ.
- Cressie, N. A. C. (1993), *Statistics for Spatial Data*, Wiley, New York.
- Dean, A. M. and Voss, D. (1999), *Design and Analysis of Experiments*, Springer, New York.
- Debushe, L. K. and Haines, L. M. (2008), V- and D-optimal population designs for the simple linear regression model with a random intercept term, *Journal of Statistical Planning and Inference*, 138, 1116–1130.
- Dette, H. and Trampisch, M. (2010), A general approach to D-optimal designs for weighted univariate polynomial regression models (with discussion), *Journal of the Korean Statistical Society*, 39, 1–42.
- Dey, A. (2010), *Incomplete Block Designs*, World Scientific, Hackensack, NJ.
- Diggle, P. J. and Ribeiro Jr, P. J. (2007), *Model-Based Geostatistics*, Springer, New York.

- Draper, N. R. and Smith, H. (1998), *Applied Regression Analysis*, Wiley, New York.
- Duchateau, L., Janssen, P., and Rowlands, G. (1998), *Linear Mixed Models. An Introduction with Applications in Veterinary Research*, International Livestock Research Institute, Nairobi, Kenya.
- Fang, K.-T., Li, R., and Sudjianto, A. (2006), *Design and Modeling for Computer Experiments*, Chapman & Hall/CRC, Boca Raton, FL.
- Fedorov, V. V. (1972), *Theory of Optimal Experiments*, Academic Press, New York.
- GAP (2014), *GAP—Groups, Algorithms, and Programming, Version 4.7.6*, The GAP Group.
- Giovagnoli, A. and Sebastiani, P. (1989), Experimental designs for mean and variance estimation in variance components models, *Computational Statistics and Data Analysis*, 8, 21–28.
- Goos, P. (2002), *The Optimal Design of Blocked and Split-Plot Experiments*, Springer, New York.
- Graybill, F. A. (1976), *Theory and Application of the Linear Model*, Duxbury Press, North Scituate, MA.
- Harville, D. A. (1977), Maximum likelihood approaches to variance component estimation and to related problems, *Journal of the American Statistical Association*, 72, 320–338.
- Harville, D. A. (1991), Discussion of the paper “That BLUP is a good thing: The estimation of random effects” by G. K. Robinson, *Statistical Science*, 6, 35–39.
- Harville, D. A. and Jeske, D. (1992), Mean squared error of estimation or prediction under a general linear model, *Journal of the American Statistical Association*, 87, 724–731.
- Henderson, C. R., Kempthorne, O., Searle, S. R., and Von Krosigk, C. N. (1959), The estimation of environmental and genetic trends from records subject to culling, *Biometrics*, 15, 192–218.
- John, J. A. and Williams, E. R. (1995), *Cyclic and Computer Generated Designs*, Chapman & Hall/CRC, London, U.K.
- John, P. W. M. (1971), *Statistical Design and Analysis of Experiments*, MacMillan, New York.
- Kackar, R. N. and Harville, D. A. (1984), Approximations for standard errors of estimators of fixed and random effects in mixed linear models, *Journal of the American Statistical Association*, 79, 853–862.
- Khuri, A. I. (2000), Designs for variance components estimation: Past and present, *International Statistical Review*, 68, 311–322.
- Khuri, A. I. (ed.) (2006), *Response Surface Methodology and Related Topics*, World Scientific, Hackensack, NJ.
- Khuri, A. I. (2010), *Linear Model Methodology*, Chapman & Hall/CRC, Boca Raton, FL.
- Kiefer, J. (1959), Optimum experimental designs (with discussion), *Journal of the Royal Statistical Society: Series B*, 21, 272–319.
- Kiefer, J. and Wolfowitz, J. (1960), The equivalence of two extremum problems, *Canadian Journal of Mathematics*, 12, 363–366.
- Kunert, J., Martin, R. J., and Eccleston, J. (2010), Optimal block designs comparing treatments with a control when the errors are correlated, *Journal of Statistical Planning and Inference*, 140, 2719–2738.
- Kutner, M. H., Nachtsheim, C. J., Neter, J., and Li, W. (2004), *Applied Linear Statistical Models*, 5th edn., McGraw-Hill, New York.
- Liu, Q., Dean, A. M., and Allenby, G. M. (2007), Design for hyperparameter estimation in linear models, *Journal of Statistical Theory and Practice*, 1, 311–328.
- Loeza-Serrano, S. and Donev, A. N. (2014), Construction of experimental designs for estimating variance components, *Computational Statistics and Data Analysis*, 71, 1168–1177.
- Longford, N. T. (1994), *Random Coefficient Models*, Clarendon Press, Oxford, U.K.
- Mateu, J. and Müller, W. G. (eds.) (2013), *Spatio-Temporal Design: Advances in Efficient Data Acquisition*, Wiley, Chichester, U.K.
- MATLAB (2014), *Version 8.4 (R2014b)*, The MathWorks Inc, Natick, MA.
- McCulloch, C. E. and Searle, S. R. (2001), *Generalized, Linear, and Mixed Models*, Wiley.
- Mentre, F., Mallet, A., and Baccar, D. (1997), Optimal design in random-effects regression models, *Biometrika*, 84, 429–442.
- Montgomery, D. (2012), *Design and Analysis of Experiments*, 8th edn., Wiley, New York.
- Morris, M. D. (2011), *Design of Experiments: An Introduction Based on Linear Models*, Chapman & Hall/CRC, Boca Raton, FL.

- Müller, W. G. (2007), *Collecting Spatial Data: Optimum Design of Experiments for Random Fields*, Springer, Berlin, Germany.
- Myers, R. H. (1999), Response surface methodology—current status and future directions (with discussion), *Journal of Quality Technology*, 31, 30–74.
- Myers, R. H., Montgomery, D. C., and Anderson-Cook, C. M. (2009), *Response Surface Methodology: Process and Product Optimization Using Designed Experiments*, Wiley, Hoboken, NJ.
- Patterson, H. D. and Thompson, R. (1971), Recovery of inter-block information when block sizes are unequal, *Biometrika*, 58, 545–554.
- Pawitan, Y. (2001), *In All Likelihood: Statistical Modelling and Inference Using Likelihood*, Oxford University Press, Oxford, U.K.
- Pinheiro, J. C. and Bates, D. M. (2000), *Mixed-Effects Models in S and S-PLUS*, Springer, New York.
- Prasad, N. G. N. and Rao, J. N. K. (1990), The estimation of the mean squared error of small-area estimators, *Journal of the American Statistical Association*, 85, 163–171.
- Pringle, R. M. and Rayner, A. A. (1971), *Generalized Inverse Matrices with Applications to Statistics*, Griffin, London, U.K.
- Pukelsheim, F. (1993), *Optimal Design of Experiments*, Wiley, New York.
- R Development Core Team (2014), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-900051-07-0.
- Raktoe, B. L., Hedayat, A., and Federer, W. T. (1981), *Factorial Designs*, Wiley, New York.
- Rao, C. R. (2001), *Linear Statistical Inference and Its Applications*, 2nd edn., Wiley, New York.
- Ravishanker, N. and Dey, D. K. (2002), *A First Course in Linear Model Theory*, Chapman & Hall/CRC, Boca Raton, FL.
- Robinson, G. K. (1991), That BLUP is a good thing: The estimation of random effects (with discussion), *Statistical Science*, 6, 15–51.
- Rockafellar, R. T. (1996), *Convex Analysis*, Princeton University Press, Princeton, NJ.
- Ruppert, D., Wand, M. P., and Carroll, R. J. (2003), *Semiparametric Regression*, Cambridge University Press, Cambridge, U.K.
- Santner, T. J., Williams, B. J., and Notz, W. I. (2003), *The Design and Analysis of Computer Experiments*, Springer, New York.
- Scheffé, H. (1959), *The Analysis of Variance*, Wiley, New York.
- Searle, S. R. (1971), *Linear Models*, Wiley, New York.
- Searle, S. R., Casella, G., and McCulloch, C. E. (1992), *Variance Components*, Wiley, New York.
- Shah, K. R. and Sinha, B. K. (1989), *Theory of Optimal Designs*, Springer, New York.
- Silvey, S. D. (1980), *Optimal Design: An Introduction to the Theory for Parameter Estimation*, Chapman and Hall, London, U.K.
- Verbeke, G. and Molenberghs, G. (2000), *Linear Mixed Models for Longitudinal Data*, Springer, New York.
- Welch, W. J. (1982), Computer aided design of experiments, PhD thesis, Imperial College London, London, U.K.
- Wolfram Research (2014), *Mathematica Edition: Version 10.0*, Wolfram Research, Inc., Champaign, IL.
- Wu, C. F. J. and Hamada, M. (2009), *Experiments: Planning, Analysis, and Optimization*, 2nd edn., Wiley, Hoboken, NJ.
- Yang, M., Zhang, B., and Huang, S. (2011), Optimal designs for generalized linear models with multiple design variables, *Statistica Sinica*, 21, 1415–1430.

Section II

Designs for Linear Models

3

Blocking with Independent Responses

John P. Morgan

CONTENTS

3.1	Blocking: The Basics.....	99
3.2	Cost of Blocking.....	105
3.3	Universal Optimality and BIBDs.....	110
3.4	Block Designs That Approximate BIBD Structure.....	113
3.4.1	Regular Graph Designs.....	114
3.4.2	Group Divisible Designs.....	116
3.4.3	Nearly Balanced Incomplete Block Designs.....	122
3.4.4	E-Optimal Block Designs.....	125
3.5	Small Block Designs.....	129
3.5.1	Block Size Two.....	129
3.5.2	Two Blocks.....	132
3.5.3	Minimally Connected Designs.....	133
3.6	Multiple Blocking Factors.....	134
3.6.1	Resolvable Block Designs.....	136
3.6.1.1	Affine Resolvable Designs.....	137
3.6.1.2	Other Resolvable Designs.....	138
3.6.2	Other Nested Block Designs.....	139
3.6.3	Row–Column Designs.....	139
3.6.4	Designs with Nesting and Crossing.....	142
3.6.4.1	Semi-Latin Squares.....	143
3.6.4.2	Nested Row–Column Designs.....	144
3.7	Control Treatments and Weighted Optimality.....	145
3.7.1	Design for TvC Experiments.....	148
3.7.2	Lesser Weight on the Control.....	151
3.8	Discussion and Directions.....	151
	References.....	153

3.1 Blocking: The Basics

Dean and Voss (1999) describe an experiment to assess the impact of visual context on human subjects' ability to reproduce a straight line. The experimental procedure has a subject look at a picture of a 5 cm straight line, then draw freehand a line of the same length on a sheet of paper. The recorded response is the length of the segment drawn. This basic procedure is repeated six times for each subject, with each repetition drawn on a

separate sheet. Visual context is varied with the six pictures, each having one of six, distinct, preprinted borders. The goal of the experiment is to understand differences, if any, in responses, depending on the six bordering treatments. That understanding is hampered in so far as responses vary for reasons having nothing to do with the treatments.

Variability is apparent in the line-segment experiment in the fact that very few, if any, individuals can consistently draw straight lines of precisely the same length even under identical conditions. Observe, however, that the lengths of lines drawn by the same person are likely to vary less than the lengths of lines drawn by different people. That is, for all of the reasons that humans vary in their drawing skill, subject-to-subject variability will be greater than within-subject variability, and possibly much greater. This suggests that it would be advantageous, if possible, to isolate the larger subject-to-subject variability from other, intrinsic sources of variability in response, including that due to visual contexts, so that context effects become more readily discernible. This is the purpose of *blocking*: an often powerful technique for isolating a source of variability in experimental material so that it does not interfere with inferences on comparisons among treatments. As discussed in Chapter 1, blocking was originally introduced and developed as an experimental principle by R. A. Fisher and F. Yates; see particularly Sections 1.3 and 1.6.

Formally, a *block* is a set of experimental units sharing a common value of some characteristic thought to represent a potentially major source of variation in the response. A set of experimental units is *blocked* if they have been arranged into disjoint blocks. A *block design* is a choice of treatments to be used in each block, followed by a random assignment of those treatments to the units in each block. It is quite common, though not required, to have the same number of experimental units and thus to use the same number of treatments, in each block. Only equisized blocks are considered in this chapter.

In the line-segment experiment described earlier, fourteen subjects were recruited. Each was presented with six line pictures, one for each of the six border treatments in a random order, and asked to draw the line on a blank paper. The blocks are the fourteen subjects, and the units within the blocks are the six presentations of pictures. Using b for the number of blocks, v for the number of treatments, and k for the number of units in a block, this experiment has $v = k = 6$ and $b = 14$. The name of this particular design is *randomized complete block design* (RCBD). The blocks are said to be “complete” because each comprises a single replicate of the v treatments (no treatment left out of the block and none repeated).

Blocking is a design choice. Alternatively, the line-segment experiment could have been run with a *completely randomized design* (CRD). With fourteen subjects and six presentations to each, there are a total of $14 \times 6 = 84$ experimental units. A CRD would start by preparing fourteen pictures with each border, then assigning them at random to the 84 units, where “at random” means that each of the $84!/(14!)^6$ distinct assignments is equally likely. The CRD is attempting through the randomization to distribute, without distinction, *all* nuisance sources of variation across the treatments, paying no special attention to subject to subject or any other particular source of variability. Blocking, on the other hand, serves to isolate one of those sources that, pre-experimentation, we have identified as potentially major, so that it does not affect our estimation of treatment differences. For the line-segment experiment, the separate random assignments of treatments to units within each subject (block) keep the major subject-to-subject variability separated from treatments while seeking to spread within-subject variability equally across the treatments. This is reflected in the analysis of variance (ANOVA) skeletons presented in Table 3.1. The sum of squares for treatments is the same in both ANOVAs. What changes is the removal of variability due to blocks from the error sum of squares. The F -test for comparing treatments, as well as all estimation for comparisons of treatments, is subject to the variability estimated

TABLE 3.1
ANOVAs for the Line-Segment Experiment

CRD		RCBD	
Source	d.f.	Source	d.f.
Treatments	5	Blocks = Subjects	13
Error	78	Treatments	5
		Error	65
Total	83	Total	83

by MSE. Reducing that variability increases the power of the test and the precision of those estimates. This is the primary purpose of blocking. For further exposition on the important role played by randomization, see Chapter 1. The randomization employed is a primary characteristic distinguishing block designs from CRDs.

The probability distribution employed to randomize treatment assignment within blocks also generates a statistical model for the responses; see Hinkelmann and Kempthorne (2008) for an introduction to and development of this theory. For our purposes, it is sufficient to simply state the model

$$y_{ju} = \mu + \tau_{d[j,u]} + \beta_j + \epsilon_{ju}. \tag{3.1}$$

where, y_{ju} is the response from unit u in block j , and the design function $d[j, u]$ identifies which of the v treatments have been assigned to that unit. The response is written in (3.1) as the sum of a general mean μ , an effect $\tau_{d[j,u]}$ of the treatment applied to the unit in question, an effect β_j of the block in which the unit lies, and a random error term ϵ_{ju} associated with that unit and its measurement. Though there are many other options, unless otherwise specified, the ϵ_{ju} in this section are taken to be uncorrelated, mean zero random variables, with common variance σ_E^2 .

The block effects β_j serve to reduce the variability to which treatment comparisons are exposed, but are not the target of inference. That target is the set of treatment effects τ_i , with the design goal of maximizing precision of treatment comparisons. To see how this is done, rewrite (3.1) in matrix form as

$$\mathbf{y} = \mu \mathbf{1} + \mathbf{X}_d \boldsymbol{\tau} + \mathbf{Z} \boldsymbol{\beta} + \boldsymbol{\epsilon}, \tag{3.2}$$

where for the total n of experimental units employed, \mathbf{X}_d is the $n \times v$ unit/treatment incidence matrix corresponding to the design function $d[j, u]$, that is, \mathbf{X}_d has a one in position (u, j) if treatment j is assigned to unit u ; $\boldsymbol{\tau} = (\tau_1, \dots, \tau_v)'$ is the $v \times 1$ vector of treatment effects; \mathbf{Z} is the $n \times b$ unit/block incidence matrix; and $\boldsymbol{\beta} = (\beta_1, \dots, \beta_b)'$ is the $b \times 1$ vector of treatment effects. Write $\mathbf{P}_Z = \mathbf{Z}(\mathbf{Z}'\mathbf{Z})^{-1}\mathbf{Z}'$ for the projector onto the column space of \mathbf{Z} . Then standard linear model theory (see Chapter 2) says that the least squares equations for estimation of $\boldsymbol{\tau}$ are $\mathbf{X}'_d(\mathbf{I} - \mathbf{P}_Z)\mathbf{X}_d \hat{\boldsymbol{\tau}} = \mathbf{X}'_d(\mathbf{I} - \mathbf{P}_Z)\mathbf{y}$. The coefficient matrix $\mathbf{C}_d = \mathbf{X}'_d(\mathbf{I} - \mathbf{P}_Z)\mathbf{X}_d$ is the *information matrix* for estimation of $\boldsymbol{\tau}$. It serves this role: if $\mathbf{c}'\boldsymbol{\tau}$ is any estimable function of the treatment effects, then $\text{Var}_d(\widehat{\mathbf{c}'\boldsymbol{\tau}}) = (\mathbf{c}'\mathbf{C}_d^{-1}\mathbf{c})\sigma_E^2$ for any generalized inverse (see Chapter 2) \mathbf{C}_d^{-1} of \mathbf{C}_d . A good design d will be one that makes the variances $(\mathbf{c}'\mathbf{C}_d^{-1}\mathbf{c})\sigma_E^2$ small in some overall sense (over all comparisons \mathbf{c} of interest). The purpose of blocking the units was to minimize, in so far as possible, the error variance σ_E^2 . Given the blocks and

whatever they accomplish in variance reduction, it is the choice of which treatments are assigned within each block (i.e., the selection of d) that determines \mathbf{X}_d and, hence, \mathbf{C}_d and the quadratic forms $\mathbf{c}'\mathbf{C}_d^-\mathbf{c}$. Good choice of d will minimize, again in an overall sense, these quadratic forms.

Before proceeding, here are a few pertinent facts about \mathbf{C}_d . Obviously, \mathbf{C}_d is symmetric, and since $(\mathbf{I} - \mathbf{P}_Z)$ is idempotent, $\mathbf{C}_d = \mathbf{X}'_d(\mathbf{I} - \mathbf{P}_Z)\mathbf{X}_d$ is nonnegative definite. As it is the coefficient matrix in the least squares equations, its row space consists of all vectors \mathbf{c} such that $\mathbf{c}'\boldsymbol{\tau}$ is estimable. Easily checked is $\mathbf{C}_d\mathbf{1} = \mathbf{0}$, so this $v \times v$ matrix has rank at most $v - 1$, and the estimability of $\mathbf{c}'\boldsymbol{\tau}$ requires $\mathbf{c}'\mathbf{1} = 0$. Hence, only contrasts of treatments are estimable, and all contrasts are estimable if, and only if, $\text{rank}(\mathbf{C}_d) = v - 1$. Accordingly, the terms $(\mathbf{c}'\mathbf{C}_d^-\mathbf{c})\sigma_E^2$ targeted by the design problem are termed *contrast variances*.

A design d is said to be *connected* if its information matrix \mathbf{C}_d has rank $v - 1$ and otherwise is *disconnected*. Because disconnected designs are unable to provide estimates of every treatment contrast, they are inappropriate unless a lower dimensional model for the treatment effects is in play, such as is employed with fractional factorial designs (Chapter 7). This chapter focuses on estimating all treatment contrasts, so considers only connected designs. Letting $\mathcal{D} = \mathcal{D}(v, b, k)$ denote the class of all connected designs for v treatments in b blocks of k units per block, the goal is to determine a "best" design in \mathcal{D} .

Needed now is a way to operationalize the idea of "best." As will be seen, there are many ways to do this, some more statistically meaningful than others and with varying relevance depending on experimenter goals. The basic idea is to select a function, let's call it Φ , that maps the information matrix \mathbf{C}_d to the real numbers \mathfrak{R} in a manner that summarizes variances of contrast estimates. Smaller values of Φ will correspond to designs that make variances smaller in the sense that Φ measures. We begin with a general framework for functions of this type, then explore particular functions that are useful in a variety of situations.

Formally, Φ is defined on the class \mathcal{C} of all symmetric, nonnegative definite matrices of rank $v - 1$ with zero row sums. Let $\Phi : \mathcal{C} \rightarrow \mathfrak{R}$. Φ corresponds to a standard optimality criterion if

- (i) Φ respects the nonnegative definite ordering: for any $\mathbf{C}_1, \mathbf{C}_2 \in \mathcal{C}$ such that $\mathbf{C}_1 - \mathbf{C}_2$ is nonnegative definite, $\Phi(\mathbf{C}_1) \leq \Phi(\mathbf{C}_2)$,
 - (ii) Φ is invariant to treatment permutation: $\Phi(\mathbf{PCP}') = \Phi(\mathbf{C})$ for each $\mathbf{C} \in \mathcal{C}$ and every permutation matrix $\mathbf{P}_{v \times v}$,
 - (iii) Φ is convex: $\Phi(\alpha\mathbf{C}_1 + (1 - \alpha)\mathbf{C}_2) \leq \alpha\Phi(\mathbf{C}_1) + (1 - \alpha)\Phi(\mathbf{C}_2)$ for all $\mathbf{C}_1, \mathbf{C}_2 \in \mathcal{C}$.
- (3.3)

A design minimizing Φ is said to be Φ -optimal.

It can be shown (e.g., Morgan 2007b) that $\mathbf{C}_1 - \mathbf{C}_2$ is nonnegative definite, implying that $\mathbf{c}'\mathbf{C}_1^-\mathbf{c} \leq \mathbf{c}'\mathbf{C}_2^-\mathbf{c}$ for all contrast vectors \mathbf{c} . Thus, property (3.3)(i) says that Φ summarizes variances of estimators in a way that uniformly smaller variances produce smaller values of Φ , this being a minimal requirement if Φ is to be statistically meaningful. Property (3.3)(ii) is equality of interest in all treatments, for permutation invariance says Φ produces the same summary measure regardless of how the treatments are interchanged. The convexity property (3.3)(iii) is a technical requirement that is advantageous for optimality arguments and that, happily, turns out to be an attribute of all the commonly employed, statistically reasonable criteria.

The notion of equal treatment interest, incorporated through (3.3)(ii), is appropriate for many, but not all, experiments. One notable class of exceptions is experiments having a

control treatment, with differing emphasis on contrasts involving, and not involving, the control. Criteria that incorporate unequal treatment interest are taken up in Section 3.7.

Let N_d be the $v \times b$ matrix whose (i, j) entry is n_{dij} = the number of units in block j assigned treatment i by design d . Further write $r_d = (r_{d1}, \dots, r_{dv})'$ where $r_{di} = \sum_j n_{dij}$ is the total number of units assigned treatment i , and let $R_d = \text{Diag}(r_d)$ be the diagonal matrix of the replication numbers r_{di} . Then the information matrix C_d can be expressed as

$$C_d = R_d - \frac{1}{k} N_d N_d' \tag{3.4}$$

Having selected the treatments to be used in any given block, they will be assigned to that block's units with a random device. The specific realization of the randomization does not change (3.4) and so does not affect the quality of the information that the experiment will produce.

The (i, i') entry of $N_d N_d'$ in (3.4), labeled $\lambda_{dii'}$ = $\sum_j n_{dij} n_{di'j}$, is called a *treatment concurrence number*. For $i \neq i'$ this is the number of pairs of units within blocks receiving the pair of treatments i, i' . As C_d is determined by the replication numbers r_{di} and the concurrence numbers $\lambda_{dii'}$, the search for an optimal design can be understood through these counts. As we shall see, depending on the criterion selected, different values for these counts, that is, different designs, can be judged best. But first we need to define some specific criteria.

The inverses of the $v - 1$ positive eigenvalues $e_{d1} \leq e_{d2} \leq \dots \leq e_{d,v-1}$ of C_d are termed *canonical variances*; aside from σ_E^2 , they are variances for the estimators of contrasts specified by their corresponding normalized eigenvectors. Many, though not all, of the commonly employed optimality criteria are functions of the canonical variances. Among these (see Table 3.2 and also Chapter 2) are the A , E , and D criteria. Each has a particular statistical meaning, as follows. Up to a design-independent constant, A measures (i) the average of the variances for all $v(v - 1)/2$ pairwise comparisons $\widehat{\tau}_i - \widehat{\tau}_{i'}$ and also (ii) the average of the variances when estimating any $v - 1$ orthonormal contrasts. The E -value for a design is proportional to the maximal estimator variance over all normalized contrasts, $\max_c \text{Var}_d(\widehat{c}'\tau)/c'c$. The D -value for a design is proportional, under a normality assumption, to the volume of the confidence ellipsoid for estimating any $v - 1$ orthonormal contrasts. Also displayed in Table 3.2 is MV , a standard criterion in the sense of (3.3) that is not solely a function of the eigenvalues. Like E , the MV criterion seeks to minimize the impact of the worst case, but with respect to pairwise contrasts rather than all contrasts. Of the measures in Table 3.2, the D criterion is the least popular for assessing block designs, though it is widely used in other design settings. Other interpretations of A are revealed in Morgan and Stallings (2014).

TABLE 3.2
Optimality Criteria

Criterion	Function
A	$\Phi_A = \sum_{i=1}^{v-1} \frac{1}{e_{di}}$
E	$\Phi_E = \frac{1}{e_{d1}}$
D	$\Phi_D = \prod_{i=1}^{v-1} \frac{1}{e_{di}}$
MV	$\Phi_{MV} = \max_{h \in H} h' C_d^- h$

Note: H is the collection of normalized contrast vectors with two nonzero coordinates.

A , D , and E are special instances of the Φ_t family of criteria specified by minimizing $\Phi_t(C_d) = \left[\sum_i e_{di}^{-t} / (v - 1) \right]^{1/t}$: take $t = 1$, $t \rightarrow 0$, and $t \rightarrow \infty$, respectively. All Φ_t criteria are standard criteria in the sense of (3.3). While other values of t do not admit simple statistical interpretations like those given earlier, it is interesting to see that the eigenvalue-based criteria in Table 3.2 are all part of the same continuum.

The criteria in Table 3.2 often, but not always, agree on what design is best. The three examples that follow illustrate several aspects of this situation. Each displays an *incomplete block design* (IBD), in which the block size k is smaller than the number of treatments, and so a block cannot receive a complete replicate of the treatments. The optimality results stated in the examples will be justified later in this chapter. Here and throughout this chapter, block designs will be displayed with blocks as columns.

Example 3.1

This design, for $(v, b, k) = (6, 10, 3)$, is optimal with respect to all of the criteria in Table 3.2.

1	2	3	4	5	1	2	3	4	5
2	3	4	5	1	2	3	4	5	1
6	6	6	6	6	4	5	1	2	3

In this design, all the replication numbers are equal ($r_{di} \equiv 5$) and all the concurrence numbers are equal ($\lambda_{dii'} \equiv 2, i \neq i'$).

Example 3.2

This design, for $(v, b, k) = (7, 28, 5)$, is E -optimal but not A -optimal.

1	1	1	1	1	1	1	1	1	1	1	1	1	1
2	2	2	2	2	2	2	2	2	2	2	2	3	3
3	3	3	3	4	4	4	4	4	4	5	5	4	4
4	4	5	5	5	5	5	5	6	6	6	6	5	5
6	7	6	7	6	6	7	7	7	7	7	7	6	6

1	1	1	1	1	1	2	2	2	2	2	2	2	2
3	3	3	3	3	3	3	3	3	3	3	3	3	3
4	4	4	4	5	5	4	4	4	4	4	4	5	5
5	5	6	6	6	6	5	5	5	5	6	6	6	6
7	7	7	7	7	7	6	6	7	7	7	7	7	7

In this design, all the replication numbers are equal ($r_{di} \equiv 20$). However, the concurrence numbers are not all the same, ranging from 12 (e.g., λ_{d12}) to 14 (e.g., λ_{d17}).

Example 3.3

This design, for $(v, b, k) = (5, 7, 3)$, is E -optimal and MV -optimal but neither A -optimal nor D -optimal.

1	1	1	1	2	2	2
1	3	3	4	3	3	4
2	4	5	5	4	5	5

This design has two different replication numbers ($r_{d1} = 5$, other $r_{di} = 4$), while all the concurrence numbers are equal ($\lambda_{dii'} \equiv 2, i \neq i'$).

Example 3.3 is distinguished from both Examples 3.1 and 3.2 by having a treatment (in this case, treatment 1) assigned to more than one experimental unit within a block. Examples 3.1 and 3.2 have all $n_{dij} \in \{0, 1\}$ and consequently are said to be *binary* block designs. Example 3.3 provides an example of a *nonbinary* block design. This and the other combinatorial features demonstrated by the preceding example designs will arise in the search for general properties that make for good block designs. Because the combinatorial structure (through r_{di} 's and λ_{dii} 's) determines the information structure (C_d), much of the statistical theory for block designs is aimed at identifying design structures that can be proven optimal or, failing that, near optimal.

3.2 Cost of Blocking

Variance reduction through blocking does not come for free. One aspect of this is seen in Table 3.1, where the error degrees of freedom for the RCBD is 13 fewer than for the CRD. If a block design uses b blocks, then $b - 1$ degrees of freedom are "stolen" from error relative to the unblocked CRD, resulting in a less precise estimate of the remaining error variance σ_E^2 . This impacts, for instance, widths of confidence intervals for treatment contrasts. Fortunately, this tends to be a small cost in all but the smallest designs (those providing few error degrees of freedom) and will be more than outweighed by the reduction in σ_E^2 achieved by a well-chosen blocking variable. There is a more substantial cost that is the focus of this section: treatment information in the data that are lost to blocks.

To frame the idea, look again at the design in Example 3.1. Consider the data contrast $y_6. - y_{1.}$, which is the difference in the totals for the sixth and first blocks. Were the units not blocked, $y_6. - y_{1.}$ would contribute to the estimate for $\tau_4 - \tau_6$. Because of the blocking, its expected value is $\tau_4 - \tau_6 + 3(\beta_6 - \beta_1)$, which is uninformative for $\tau_4 - \tau_6$. The treatment information in $y_6. - y_{1.}$, lost to the blocking, is said to be *confounded with blocks*. More generally, the expected value of an arbitrary data contrast $\sum_j c_j y_j.$ of block totals contains the term $\sum_j c_j \beta_j.$, so it cannot contribute to estimating treatment contrasts.

Starting with \mathbf{y} in model (3.2) and the blocks projection matrix \mathbf{P}_Z , consider the transformed data $\mathbf{y}^* = (\mathbf{I} - \mathbf{P}_Z)\mathbf{y}$ having model $\mathbf{y}^* = \mathbf{X}_d^* \boldsymbol{\tau} + \boldsymbol{\epsilon}^*$. Least squares estimation for $\boldsymbol{\tau}$ using the data \mathbf{y}^* leads to exactly the same information matrix (3.4) for $\boldsymbol{\tau}$ as that found when starting with the original data \mathbf{y} . Since \mathbf{P}_Z transforms \mathbf{y} to block averages, this shows that estimation of $\boldsymbol{\tau}$ uses only data information that is orthogonal to the block totals. Any treatment information found in the b block totals, and as explained in the preceding paragraph, is confounded with blocks, also said to *lie in the blocks stratum*. Confounding represents the primary cost of blocking: treatment information lost to blocks.

Confounding occurs whenever not every treatment is equally replicated in every block, so whenever incomplete blocks are used, but never with complete blocks. In a RCBD, every block contains each treatment once, so no contrast of block totals contains treatment information. Aside from the small cost of lost error degrees of freedom, a RCBD estimates treatment differences just as efficiently as a CRD even if achieving no reduction in σ_E^2 . Efficiency gains with a RCBD are typically substantial if the blocking actually delivers on its variance-reduction goal.

With an IBD, confounding means that less of the data are used in estimating treatment contrasts. Consequently, should blocking be ineffective and not reduce σ_E^2 , an IBD can

produce results that are markedly inferior to what would have been achieved without blocking at all. Prudent use of incomplete blocks requires understanding how much information is confounded and so the amount of variance reduction required to justify that use. This understanding can be gained through a modification of the treatment information matrix (3.4).

The *efficiency matrix* $C_{d,\text{eff}}$, here also shown in spectral form, is a transformed version of the information matrix

$$C_{d,\text{eff}} = \mathbf{R}_d^{-1/2} C_d \mathbf{R}_d^{-1/2} = \sum_{i=1}^{v-1} \varepsilon_{di} \mathbf{f}_{di} \mathbf{f}'_{di}. \tag{3.5}$$

The eigenvalues ε_{di} are termed the (*canonical*) *efficiency factors* for design d . Obviously, efficiency factors are positive (for connected designs), and it can be shown that they obey an upper bound of 1. The eigenvectors \mathbf{f}_{di} define the *basic contrast* vectors $\mathbf{c}_{di} = \mathbf{R}_d^{1/2} \mathbf{f}_{di}$. Aside from the constant σ_E^2 , the variance of the estimator for basic contrast i is $\text{Var}_d(\widehat{\mathbf{c}}'_{di} \boldsymbol{\tau}) = 1/\varepsilon_{di}$.

The efficiency matrix is also defined for unblocked designs. For a CRD with the same replication numbers r_d employed by the preceding IBD, the treatment information matrix is $\mathbf{R}_d - \frac{1}{n} \mathbf{r}_d \mathbf{r}'_d$ and so

$$\begin{aligned} C_{\text{CRD},\text{eff}} &= \mathbf{R}_d^{-1/2} \left[\mathbf{R}_d - \frac{1}{n} \mathbf{r}_d \mathbf{r}'_d \right] \mathbf{R}_d^{-1/2} \\ &= \mathbf{I} - \frac{1}{n} \mathbf{R}_d^{1/2} \mathbf{J} \mathbf{R}_d^{1/2}, \end{aligned} \tag{3.6}$$

where \mathbf{J} is the all-one matrix. The \mathbf{f}_{di} in (3.5) are also eigenvectors of (3.6), but with corresponding eigenvalues identically 1, for $C_{\text{CRD},\text{eff}} \mathbf{f}_{di} = \mathbf{f}_{di} - \frac{1}{n} \mathbf{R}_d^{1/2} \mathbf{J} \mathbf{c}_{di} = \mathbf{f}_{di} - \frac{1}{n} \mathbf{R}_d^{1/2} \mathbf{0} = \mathbf{f}_{di}$. Thus, and again aside from σ_E^2 , the basic contrast $\mathbf{c}'_{di} \boldsymbol{\tau}$ for the IBD d is estimated by a CRD having the same replication numbers as d with variance $\text{Var}_{\text{CRD}}(\widehat{\mathbf{c}}'_{di} \boldsymbol{\tau}) = 1$. The ratio of this variance to that produced by d , assuming both designs produce the same value for σ_E^2 (blocking achieves no reduction), is

$$\text{Eff}_d(\mathbf{c}'_{di} \boldsymbol{\tau}) = \text{Var}_{\text{CRD}}(\widehat{\mathbf{c}}'_{di} \boldsymbol{\tau}) / \text{Var}_d(\widehat{\mathbf{c}}'_{di} \boldsymbol{\tau}) = \varepsilon_{di}. \tag{3.7}$$

It follows from (3.7) that efficiency factor ε_{di} measures the proportion of unconfounded information available with d on the basic contrast $\mathbf{c}'_{di} \boldsymbol{\tau}$. Said another way, $1 - \varepsilon_{di}$ is the proportional reduction in σ_E^2 needed for the IBD d to estimate this contrast as well as an unblocked design. If d reduces σ_E^2 by more than $100(1 - \varepsilon_{di})\%$, then it will estimate this contrast better than a CRD.

It was mentioned earlier that there is no confounding with a RCBD. This is equivalent to saying that for a RCBD, all efficiency factors are 1. Readers may wish to check that the efficiency matrix for a RCBD is identical to that in (3.6) for an equally replicated CRD.

The basic contrasts are linearly independent and so form a basis for all treatment contrasts. Ordering the efficiency factors $\varepsilon_{d1} \leq \varepsilon_{d2} \leq \dots \leq \varepsilon_{d,v-1}$, it follows that the proportion of unconfounded information available on every contrast is at least ε_{d1} and no contrast uses more than $100\varepsilon_{d,v-1}\%$ of the information in the unblocked design. A single, summary

measure of the fraction of available information is provided by the harmonic mean of the efficiency factors

$$\frac{v - 1}{\sum_{i=1}^{v-1} \frac{1}{\epsilon_{di}}}, \tag{3.8}$$

called the *average efficiency factor*. It is the inverse of the average variance of the basic contrast estimators. The average efficiency factor is a convenient measure for comparing designs with an eye to minimizing confounding. Keep in mind, however, that it is expressed relative to an unblocked design with the same replication vector. Both individual and average efficiency factors are relative measures, valid for design comparison *only* among designs having the same r_d . Even then, full understanding of the average efficiency factor requires a broader context, to be explored in Section 3.7.

Examples of efficiency factors, calculated for the example designs in Section 3.1, are displayed in Table 3.3. The blocking in Example 3.2 must reduce error variance by about 7% if it is to perform as well, on average, as an unblocked design with 20 replicates. The design in Example 3.1 has all efficiency factors equal to 80%, so should error variance be reduced by more than 20%, every contrast will be estimated more efficiently than with an unblocked, 5-replicate competitor. Exact expressions for (3.8) as functions of v and other design parameters are known in a few cases; see Sections 1.6.1 and 1.9.2.

Many block designs used in practice are equally replicated, having all $r_{di} = bk/v \equiv r$. In this case, efficiency factors are proportional to eigenvalues of the information matrix C_d , $\epsilon_{di} = e_{di}/r$, the average efficiency factor (3.8) is inversely proportional to the A -optimality value, and the smallest efficiency factor is inversely proportional to the E -optimality value (see Table 3.2). We will return to the relationship between efficiency factors and optimality criteria in Section 3.7, when the interpretation of (3.8) for the Example 3.3 design can also be given.

Table 3.4 lists efficiencies for the most efficient IBDs, for up to ten treatments and up to ten replicates, for those (v, b, k) for which bk is a multiple of v . Included is the average efficiency factor, termed the A -efficiency, and the lowest efficiency factor ϵ_{d1} , termed the E -efficiency. Note that if blocking reduces variation by more than $(1 - \epsilon_{d1})100\%$, the design will be better than a CRD for estimating *every* treatment contrast. How these best IBDs are determined is the topic of Sections 3.3 and 3.4.

Before ending this section, it should be mentioned that if block effects are random, $\beta \sim (0, \sigma_B^2 I)$, then the blocks stratum information $P_Z y$ can contribute to treatment estimation. The procedure, known as analysis with *recovery of interblock information*, will be mentioned briefly in later sections. Owing to the additional variance term σ_B^2 , blocks stratum information is typically poor. That is why, as pursued in this chapter, blocking schemes are chosen to maximize the within-block information based on y^* and as captured in (3.4).

TABLE 3.3
Eigenvalues and Efficiency Factors for Example Designs

Example	Eigenvalues e_{di}	Efficiency Factors ϵ_{di}	Average Efficiency
3.1	4,4,4,4,4	$\frac{4}{5}, \frac{4}{5}, \frac{4}{5}, \frac{4}{5}, \frac{4}{5}$	$\frac{4}{5}$
3.2	18.4,18.4,18.4,18.4,18.8,19.6	0.92,0.92,0.92,0.92,0.94,0.98	0.933
3.3	$3\frac{1}{3}, 3\frac{1}{3}, 3\frac{1}{3}, 3\frac{1}{3}$	$\frac{5}{6}, \frac{5}{6}, \frac{5}{6}, \frac{7}{10}$	0.796

TABLE 3.4
Efficiencies for Best IBDs

v	b	k	r	A_{eff}	E_{eff}	v	b	k	r	A_{eff}	E_{eff}	v	b	k	r	A_{eff}	E_{eff}
3	3	2	2	0.750	0.750	6	10	3	5	0.800	0.800	8	24	3	9	0.759	0.741
3	6	2	4	0.750	0.750	6	12	3	6	0.794	0.722	8	4	4	2	0.808	0.500
3	9	2	6	0.750	0.750	6	14	3	7	0.797	0.762	8	6	4	3	0.838	0.667
3	12	2	8	0.750	0.750	6	16	3	8	0.798	0.750	8	8	4	4	0.850	0.750
3	15	2	10	0.750	0.750	6	18	3	9	0.798	0.778	8	10	4	5	0.854	0.800
4	4	2	2	0.600	0.500	6	20	3	10	0.800	0.800	8	12	4	6	0.854	0.833
4	6	2	3	0.667	0.667	6	3	4	2	0.882	0.750	8	14	4	7	0.857	0.857
4	8	2	4	0.643	0.500	6	6	4	4	0.894	0.813	8	16	4	8	0.855	0.813
4	10	2	5	0.662	0.600	6	9	4	6	0.897	0.875	8	18	4	9	0.856	0.833
4	12	2	6	0.667	0.667	6	12	4	8	0.899	0.875	8	20	4	10	0.856	0.815
4	14	2	7	0.659	0.571	6	15	4	10	0.900	0.900	8	8	5	5	0.913	0.880
4	16	2	8	0.662	0.625	6	6	5	5	0.960	0.960	8	16	5	10	0.914	0.888
4	18	2	9	0.667	0.667	6	12	5	10	0.960	0.960	8	4	6	3	0.949	0.889
4	20	2	10	0.663	0.600	7	7	2	2	0.375	0.188	8	8	6	6	0.951	0.905
4	4	3	3	0.889	0.889	7	14	2	4	0.545	0.375	8	12	6	9	0.952	0.926
4	8	3	6	0.889	0.889	7	14	2	4	0.542	0.400	8	8	7	7	0.980	0.980
4	12	3	9	0.889	0.889	7	21	2	6	0.583	0.583	9	9	2	2	0.300	0.117
5	5	2	2	0.500	0.345	7	28	2	8	0.573	0.500	9	18	2	4	0.511	0.375
5	10	2	4	0.625	0.625	7	35	2	10	0.577	0.500	9	27	2	6	0.545	0.500

(Continued)

TABLE 3.4 (CONTINUED)
Efficiencies for Best IBDs

v	b	k	r	A_{eff}	E_{eff}	v	b	k	r	A_{eff}	E_{eff}	v	b	k	r	A_{eff}	E_{eff}						
5	15	2	6	0.611	0.532	7	35	2	10	0.577	0.510	9	36	2	8	0.563	0.563	10	20	4	8	0.831	0.781
5	20	2	8	0.625	0.625	7	7	3	3	0.778	0.778	9	45	2	10	0.555	0.473	10	25	4	10	0.832	0.800
5	25	2	10	0.620	0.569	7	14	3	6	0.778	0.778	9	6	3	2	0.667	0.500	10	4	5	2	0.850	0.600
5	5	3	3	0.815	0.709	7	21	3	9	0.778	0.778	9	9	3	3	0.727	0.667	10	6	5	3	0.878	0.800
5	10	3	6	0.833	0.833	7	7	4	4	0.875	0.875	9	12	3	4	0.750	0.750	10	8	5	4	0.883	0.800
5	15	3	9	0.831	0.792	7	14	4	8	0.875	0.875	9	15	3	5	0.740	0.631	10	10	5	5	0.886	0.815
5	5	4	4	0.938	0.938	7	7	5	5	0.931	0.870	9	18	3	6	0.742	0.667	10	12	5	6	0.886	0.833
5	10	4	8	0.938	0.938	7	14	5	10	0.933	0.904	9	21	3	7	0.745	0.714	10	14	5	7	0.888	0.857
6	6	2	2	0.429	0.250	7	7	6	6	0.972	0.972	9	24	3	8	0.750	0.750	10	16	5	8	0.887	0.836
6	9	2	3	0.556	0.500	8	8	2	2	0.333	0.146	9	27	3	9	0.747	0.684	10	16	5	8	0.887	0.850
6	12	2	4	0.577	0.500	8	12	2	3	0.488	0.333	9	30	3	10	0.747	0.700	10	18	5	9	0.889	0.889
6	15	2	5	0.600	0.600	8	16	2	4	0.538	0.500	9	9	4	4	0.834	0.750	10	20	5	10	0.888	0.855
6	18	2	6	0.588	0.500	8	20	2	5	0.545	0.438	9	18	4	8	0.844	0.844	10	20	5	10	0.888	0.860
6	21	2	7	0.587	0.500	8	24	2	6	0.560	0.500	9	9	5	5	0.896	0.840	10	5	6	3	0.918	0.833
6	24	2	8	0.592	0.563	8	28	2	7	0.571	0.571	9	18	5	10	0.900	0.900	10	10	6	6	0.924	0.889
6	27	2	9	0.595	0.556	8	32	2	8	0.565	0.500	9	3	6	2	0.923	0.750	10	15	6	9	0.926	0.926
6	30	2	10	0.600	0.600	8	36	2	9	0.562	0.477	9	6	6	4	0.933	0.875	10	10	7	7	0.951	0.916
6	4	3	2	0.769	0.667	8	40	2	10	0.563	0.500	9	9	6	6	0.936	0.917	10	5	8	4	0.971	0.938
6	6	3	3	0.784	0.667	8	8	3	3	0.747	0.667	9	12	6	8	0.938	0.938	10	10	8	8	0.972	0.943
6	8	3	4	0.788	0.667	8	16	3	6	0.756	0.688	9	15	6	10	0.937	0.908	10	10	9	9	0.988	0.988

Notes: Where two lines are listed for the same (v, b, k, r) , the first has higher A -efficiency (sometimes this is not detectable to three decimal places) and the second has higher E -efficiency. When A -efficiency = E -efficiency, the best design is a BIBD.

3.3 Universal Optimality and BIBDs

In a landmark paper, Kiefer (1975) tackled the question of how a design might be optimal with respect to *all* criteria defined by (3.3). He found this result:

Theorem 3.1 *For the class of designs \mathcal{D} with corresponding information matrices $\mathcal{C} = \{C_d : d \in \mathcal{D}\}$ and having $C_d \mathbf{1} = \mathbf{0}$ for all d , suppose there is a $d^* \in \mathcal{D}$ for which*

$$\text{trace}(C_{d^*}) = \max_{d \in \mathcal{D}} \text{trace}(C_d); \tag{3.9}$$

$$C_{d^*} = \alpha \left(I - \frac{1}{v} J \right). \tag{3.10}$$

Then d^ is universally optimal, that is, d^* is optimal with respect to every criterion satisfying the properties (3.3).*

Proof: Let Φ be any of the criteria satisfying (3.3), let $d \in \mathcal{D}$ be any competing design, and let \mathcal{P} be the class of all $v \times v$ permutation matrices P . Then invoking (3.3)(ii) and (3.3)(iii),

$$\Phi(C_d) = \Phi(PC_dP') = \frac{1}{v!} \sum_{P \in \mathcal{P}} \Phi(PC_dP') \leq \Phi \left(\frac{1}{v!} \sum_{P \in \mathcal{P}} PC_dP' \right).$$

Writing \bar{C}_d for $\frac{1}{v!} \sum_{P \in \mathcal{P}} PC_dP'$, we have $\bar{C}_d = \alpha_d(I - \frac{1}{v}J)$, where $\alpha_d = \text{trace}(C_d)/(v - 1) \leq \alpha$. Consequently,

$$\Phi(C_d) \leq \Phi(\bar{C}_d) = \Phi \left(\alpha_d \left(I - \frac{1}{v} J \right) \right) \leq \Phi \left(\alpha \left(I - \frac{1}{v} J \right) \right) = \Phi(C_{d^*}),$$

the last inequality by virtue of (3.3)(i), since $C_{d^*} - \bar{C}_d = (\alpha - \alpha_d)(I - \frac{1}{v}J)$ is nonnegative definite. □

Theorem 3.1 applies to block designs and to any class of designs where each member of \mathcal{C} is symmetric, nonnegative definite, and of rank deficiency one with zero row sums. The proof demonstrates why all three of the properties in (3.3) are important.

A design d having information matrix of form (3.10), $C_d = \alpha(I - \frac{1}{v}J)$, enjoys an additional, useful quality that does not depend on variance magnitude, and so is not directly assessed by criteria (3.3). Let c be any normalized contrast vector ($c'c = 1$). A generalized inverse of C_d is $\frac{1}{\alpha}I$ so that $\text{Var}_d(c'\tau) = \sigma_E^2/\alpha$ does not depend on the contrast selected. This property, known as *variance balance*, says that all treatment comparisons are estimated with the same precision. Variance balance simplifies interpretation of experimental results. It is the exact expression of equal treatment interest.

Example 3.4

The design in Example 3.1 has constant replication 5, constant concurrence 2, and, hence from (3.4), information matrix

$$C_d = 5I - \frac{1}{3}[2(J - I) + 5I] = 4\left(I - \frac{1}{6}J\right).$$

This matrix has the complete symmetry required by (3.10) and, from Lemma 3.1, has the maximal trace property (3.9) relative to all designs for $v = 6$ treatments in $b = 10$ blocks of size $k = 3$. Thus, by Theorem 3.1, the design is universally optimal over $\mathcal{D}(6, 10, 3)$.

Block designs meeting the Theorem 3.1 conditions are easily characterized.

Lemma 3.1 *A block design d with block size $k \leq v$ satisfies (3.9) and (3.10) if, and only if,*

- (i) d is binary;
- (ii) d is equireplicate, that is, r_{di} is constant in i ;
- (iii) d is equiconcurrent, that is, $\lambda_{dii'}$ is constant in $i \neq i'$.

Setting $k = v$, Lemma 3.1 says the RCBDs are universally optimal. This should be no surprise; not only is the use of each treatment once in each block intuitively appealing, it has already been seen in Section 3.2 that RCBDs carry the same treatment information as an equally replicated CRD. It is easily shown that, among CRDs with a fixed number of experimental units, those with equal replication are universally optimal.

An IBD satisfying the conditions of Lemma 3.1 is called a *balanced incomplete block design* (BIBD). One example is the design in Example 3.1.

Given the strong result of Theorem 3.1, BIBDs are, in a sense, “ideal” IBDs; they are recommended for use whenever available. Their shortfall is that in the universe of all possible (v, b, k) , BIBDs are rare. Let the integer r denote the common replication for any equally replicated IBD, and let the integer λ denote the common treatment concurrence for a BIBD. Here are three necessary conditions for a BIBD to exist:

Lemma 3.2 *If there is a BIBD in $\mathcal{D}(v, b, k)$, then (i) $r = bk/v$, (ii) $\lambda = bk(k - 1)/v(v - 1)$, and (iii) $b \geq v$.*

The first two conditions of Lemma 3.2 follow immediately from Lemma 3.1. The third, known as *Fisher’s inequality*, was proven in Fisher (1940).

Lemma 3.2 places strong restrictions on the triples (v, b, k) . Thinking again of Example 3.1, fix $v = 6$ and $k = 3$. Then (i) and (ii) of Lemma 3.2 jointly tell us that b must be a multiple of 10. Only in 1/10 of the design classes $\mathcal{D}(6, b, 3)$ is a BIBD even possible.

In addition, the conditions of Lemma 3.2, while necessary, are not always sufficient to guarantee existence of a BIBD. One of the outstanding open problems in combinatorial mathematics is to determine when those conditions are in fact sufficient. The impressive work of H. Hanani, summarized in Hanani (1975), establishes, with one exception, sufficiency for $k \leq 5$. For that one exception, $(v, b, k) = (15, 21, 5)$, Chowla and Ryser (1950)

have shown that 15 treatments cannot be placed in 21 blocks of 5 treatments each so that (i) through (iii) of Lemma 3.1 simultaneously hold; this BIBD cannot be constructed.

For $k = 6$, the Lemma 3.2 conditions are sufficient whenever $\lambda > 2$ and when $\lambda = 2$ except for $(v, b, k) = (21, 28, 6)$. There are still several open cases for $\lambda = 1$. See Chapter II.3 of Colbourn and Dinitz (2007) for further details on this and larger k . This is a problem that grows more difficult as k increases.

BIBDs were first proposed as experimental designs by Yates (1936a). Since that time, a mountainous, and still expanding, body of work on BIBD construction has evolved, with many contributions from statisticians and mathematicians alike. A good, recent survey of known BIBD results can be found in Chapter II.1 of Colbourn and Dinitz (2007). BIBDs are cataloged online at, and can be downloaded from, *www.designtheory.org*.

The conditions of Theorem 3.1 can also be met when $k > v$. Define a *multiply complete* (for short, *m-complete*) block to be a block of $k = mv$ experimental units of which m are randomly assigned to each treatment. Universally optimal designs in $\mathcal{D}(v, b, mv)$ are the *generalized complete block designs* comprised solely of *m-complete* blocks. For $k > v$ but not an integer multiple of v , let $k_0 = k \pmod{v}$ and set $m = (k - k_0)/v$. If $k_0 > 1$ and there is a BIBD $d_0 \in \mathcal{D}(v, b, k_0)$, then there is a universally optimal design $d \in \mathcal{D}(v, b, k)$ found by appending an *m-complete* block to each block of d_0 . If $k_0 = 1$ and $b = sv$ for integer $s \geq 1$, then append each treatment to s blocks of a generalized complete block design in $\mathcal{D}(v, b, k - 1)$. These optimal designs with $k_0 > 0$ are collectively referred to as *Kiefer's balanced block designs* (KBBDs).*

Universal optimality is defined in Theorem 3.1 as optimality with respect to all criteria (3.3). In light of Theorem 3.1, it is natural to ask if this optimality can be had for any block design if the conditions of the theorem are not met. The answer, at least in a restricted sense, is yes.

Theorem 3.2 (Yeh 1988) *Fix $k = v - 1$. For every $m \geq 1$, there is a BIBD in $\mathcal{D}(v, b = mv, k)$, which is universally optimal. For other values of b , the following designs are universally optimal over the subclass of binary designs in $\mathcal{D}(v, b, k)$:*

- (i) $2 \leq b \leq v - 1$: any binary design d with $\max_i r_{di} = b$ and $\min_i r_{di} = b - 1$.
- (ii) $b = mv + 1$ and $m \geq 1$: a BIBD (v, mv, k) to which one binary block is added.
- (iii) $b = mv + b_0$ for $2 \leq b_0 \leq v - 1$ and $m \geq 1$: a BIBD (v, mv, k) to which the blocks of a design from (i) with b_0 blocks is added.

It is an open question as to whether, for any of the designs in Theorem 3.2(i)–(iii), universal optimality can be extended beyond the binary subclass. It certainly cannot always be done. Morgan and Jin (2007) show that for $b = 2$ the designs displayed in Table 3.5 are uniquely *E*-optimal and so *E*-better than any binary competitor.

Theorem 3.2 designs with $b \equiv \pm 1 \pmod{v}$ fall within two design classes that will be mined for optimality in the following sections. They can be classified both as GGDD(2)s (Section 3.4.2) and NBBDD(1)s (Section 3.4.3).

* These designs are more commonly referred to as *balanced block designs* per Kiefer's original terminology. The new name proposed here distinguishes them from the many other block designs that also have the balance property, cf. Preece (1982).

TABLE 3.5

E-Optimal Design
in $\mathcal{D}(v, 2, v - 1)$,
 $v \geq 7$

1	1
2	2
3	3
\vdots	\vdots
$v - 3$	$v - 3$
$v - 2$	$v - 1$
$v - 2$	v

3.4 Block Designs That Approximate BIBD Structure

The vast majority of block designs used in practice have block size no larger than the number of treatments. Setting $k = v$, Theorem 3.1 tells us that the RCBDs are the designs of choice. With IBDs, once one leaves the BIBDs, there is no simple answer. This section explores the considerable theory that has grown, and continues to grow, around the IBD problem.

Section 3.3 presents two key facts concerning BIBDs: in terms of criteria (3.3) they are optimal in the strongest sense possible, and the triples (v, b, k) for which they exist are rare in the universe of all such triples. To add perspective to the latter, for $3 \leq v \leq 15$ and $2 \leq r \leq 30$, there are 1021 integer triples (v, b, k) where $r = bk/v$ is an integer; of these, only 253 have $b \geq v$ and integer λ (see Lemma 3.2). Even restricting to equal replication, the BIBDs leave much more unsolved than solved in the search for optimal block designs.

Yet the BIBDs are suggestive of how one might proceed. BIBDs are defined by conditions (i) through (iii) of Lemma 3.1. If these conditions produce universal optimality (as they do), then designs that are in some sense “close” to achieving those conditions should be good candidates for, if not universal optimality, individual optimalities of interest, such as those in Table 3.2. That is the tack taken in this section.

The symbols r and λ were assigned specific meanings in Section 3.3; now they will be defined in such a way as to preserve the earlier meanings as special cases. The starting point is the block design setting (v, b, k) with $k < v$. Associated with this setting,

$$\begin{aligned}
 r &= \lfloor \frac{bk}{v} \rfloor &&= \text{maximal achievable minimum replication,} \\
 p &= bk - vr &&= \text{plot excess over those needed for equal replication } r, \\
 \lambda &= \lfloor \frac{r(k-1)}{v-1} \rfloor &&= \text{maximal achievable minimum concurrence for a} \\
 &&&\text{treatment having replication } r, \\
 q &= r(k - 1) - \lambda(v - 1) &&= \text{total concurrence excess, for a treatment having} \\
 &&&\text{replication } r, \text{ over constant concurrence } \lambda.
 \end{aligned}
 \tag{3.11}$$

These *auxiliary parameters* will be used throughout this section. The necessary conditions Lemma 3.2(i),(ii) for BIBD existence may be restated as $p = q = 0$, termed the *BIBD setting*.

The concern in this section is for those settings where at least one of the plot excess p and the concurrence excess q is not zero.

Because equal replication is so common in practice, greatest attention has been paid to settings having $p = 0$, and this is where we begin. Now equal replication is possible, and “closeness” to a BIBD can be formulated in terms of deviation of the $\lambda_{dii'}$ ’s from λ . If not a BIBD setting, then necessarily $1 \leq q \leq v - 2$.

3.4.1 Regular Graph Designs

With $p = 0$ and whatever the value of q , let d be a binary, equireplicate design. Using $(x_{ii'})_{i \neq i'}$ to denote a matrix that is zero on the diagonal, the information matrix C_d for design d can be written as

$$\begin{aligned} C_d &= r\mathbf{I} - \frac{1}{k}N_dN_d' = r\mathbf{I} - \frac{r}{k}\mathbf{I} - \frac{1}{k}(\lambda_{dii'})_{i \neq i'} \\ &= \frac{r(k-1)}{k}\mathbf{I} - \frac{1}{k}(\lambda)_{i \neq i'} - \frac{1}{k}[(\lambda_{dii'})_{i \neq i'} - (\lambda)_{i \neq i'}] \\ &= \left(\frac{v\lambda + q}{k}\right)\mathbf{I} - \frac{\lambda}{k}\mathbf{J} - \frac{1}{k}\mathbf{\Delta}_d, \end{aligned} \tag{3.12}$$

where the *discrepancy matrix* $\mathbf{\Delta}_d$ has elements $\delta_{dii'}$, called *discrepancies*, defined by

$$\delta_{dii'} = \begin{cases} 0 & \text{if } i = i', \\ \lambda_{dii'} - \lambda & \text{if } i \neq i'. \end{cases} \tag{3.13}$$

Any discrepancy matrix is symmetric with constant row sums of q .

A BIBD has all $\lambda_{dii'} = \lambda$ and so $\mathbf{\Delta}_d = 0_{v \times v}$. One idea for what it means for a design d to be “close” to a BIBD is for it to be binary and equireplicate (meeting the BIBD conditions Lemma 3.2(i),(ii)), and for it to have $\mathbf{\Delta}_d$ close to the zero matrix in some sense. This motivates the *regular graph designs*.

Definition 3.1 Let the class of designs $\mathcal{D}(v, b, k)$ have $p = 0$ and $q > 0$. A design $d \in \mathcal{D}$ is a regular graph design if

- (i) d is binary,
- (ii) d is equireplicate, and
- (iii) the $v(v - 1)$ discrepancy values $\delta_{dii'}$ are all in $\{0, 1\}$.

The name “regular graph design” comes from the fact that $\mathbf{\Delta}_d$ for such a design is the adjacency matrix of a regular graph with v vertices. There is an edge connecting vertices i and i' in this graph if $\delta_{dii'} = 1$, but no edge if $\delta_{dii'} = 0$. The “regular” means that each vertex is connected to the same number of other vertices, which is q . In graph terminology, q is called the *degree* or the *valency* of the regular graph.

Within the binary, equireplicate class, imposing Definition 3.1(iii) is equivalent to minimizing $\text{trace}(\mathbf{\Delta}_d' \mathbf{\Delta}_d) = \text{trace}(\mathbf{\Delta}_d^2)$, a natural way to minimize deviation from the complete symmetry condition (3.10) in Theorem 3.1. Alternatively, RGDs can be viewed as firstly

maximizing $\text{trace}(C_d)$ (i.e., (i) of Definition 3.1) and, given that, then minimizing $\text{trace}(C_d^2)$ (i.e., (ii) and (iii)). Designs with this (M, S) -property* have minimal dispersion of their eigenvalues among designs with maximal sum of eigenvalues. Take note that these ideas of closeness to a BIBD give priority to (3.9) over (3.10).

John and Mitchell (1977) conjectured that if $\mathcal{D}(v, b, k)$ contains at least one regular graph design, then the A -, D -, and E -optimal designs must each be a regular graph design. If their conjecture is true, then so long as at least one regular graph design exists, all other designs can be ignored when looking for an optimal design. For $k > 2$, *no counterexample is known for their conjecture for either A -optimality or D -optimality*. On the other hand, it has not been proven, either. Counterexamples are known for the E -criterion (see Section 3.4.4).

Example 3.5

This regular graph design, for $(v, b, k, q) = (7, 7, 5, 2)$, is A -, D -, and E -optimal among regular graph designs:

1	1	1	1	1	2	3
2	2	2	2	4	3	4
3	3	3	5	5	4	5
4	4	6	6	6	5	6
5	7	7	7	7	6	7

Its discrepancy matrix is

$$\Delta = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}.$$

If one could simply enumerate all designs in $\mathcal{D}(v, b, k)$, then optimality with respect to any criterion could be trivially determined. The John–Mitchell conjecture for regular graph designs brings the enumerative problem within reach, up to a point. In fact, it changes the design enumeration problem into a graph enumeration problem, as follows. Each graph is a discrepancy matrix, so enumerating all regular graphs of degree q is equivalent to enumerating information matrices of the form (3.12), regardless of whether or not the corresponding designs exist. The matrices can then be ranked according to the desired optimality criterion, the best determined, and only then is it necessary to find a corresponding design. Should such a design not exist, then one moves on to the next best graph, and so on.

John and Mitchell (1977) followed this strategy for up to $v = 12$ treatments and $r = 10$ replicates. Formal optimality tools can then be employed to show that all of their designs, including the design in Example 3.5, are A - and D -optimal over the entire class $\mathcal{D}(v, b, k)$; see Theorem 3.6. Their stopping point was driven by the size of the graph enumeration problem, which grows rapidly with v (and q). Over 35 years later, tremendous strides have been made in computational power, yet little progress has been gained on the graph enumeration

* (M, S) -property has been often called (M, S) -optimality, but as it does not directly measure magnitude of contrast variances, the former term is more appropriate.

problem for $v > 16$. Recent progress on optimal RGDs for selected cases with v up to 18 may be found in Cakiroglu (2013).

3.4.2 Group Divisible Designs

A different approach to approximating the BIBD information structure can be based on giving first priority to approaching the complete symmetry condition (3.10) corresponding to Lemma 3.1(iii). How this “closeness” idea should be implemented can be seen in how the proof of Theorem 3.1 employs the convexity (3.3)(iii) of optimality functions. That usage is a particular case of the technique of *matrix averaging*.

Definition 3.2 Let P_1, \dots, P_t be t permutation matrices of order v . An averaging of the information matrix C_d is $\bar{C}_d = \frac{1}{t} \sum_{i=1}^t P_i C_d P_i'$. \bar{C}_d is called an average matrix of C_d .

Lemma 3.3 (Constantine 1981) Suppose \bar{C}_d is an average matrix of C_d . Then \bar{C}_d is nonnegative definite with zero row sums and the same trace as C_d , and $\Phi(\bar{C}_d) \leq \Phi(C_d)$ for any Φ satisfying (3.3).

For given trace of C_d , Theorem 3.1 maximizes information by averaging over *all* treatment permutations. When complete symmetry is not attainable in an actual design, an information matrix might be more advantageously improved by averaging over subsets of treatment permutations. If a corresponding design for the “partially averaged” matrix can be found, then it could be a promising candidate for optimality. The structure \bar{C}_d produced by averaging within all subsets of a treatments partition is that of a *generalized group divisible design* (GGDD).

Definition 3.3 The design $d \in D(v, b, k)$ is a GGDD(s) if the treatments in d can be divided into s mutually disjoint sets V_1, \dots, V_s of size v_1, \dots, v_s such that the elements $c_{dii'}$ of C_d satisfy

- (i) For $g = 1, \dots, s$ and all $i \in V_g$, $c_{dii} = r_{di} - \lambda_{dii}/k = c_g$, where c_g depends on the set V_g but not otherwise on the treatment i ;
- (ii) For $g, h = 1, \dots, s$ and all $i \in V_g$ and $i' \in V_h$, with $i \neq i'$ if $g = h$, $-kc_{dii'} = \lambda_{dii'} = \gamma_{gh}$, where γ_{gh} depends on the sets V_g and V_h but not otherwise on the treatments i and j .

It is convenient to assume that the treatment subsets V_1, \dots, V_s in the preceding definition are arranged so that the c_g 's are in nonincreasing order: $c_1 \geq c_2 \geq \dots \geq c_s$. If $|V_g| = 1$, then γ_{gg} is defined to be zero. This version of the GGDD definition is due to Srivastav and Morgan (1998), with the idea in less general form dating back to Adhikary (1965).

Example 3.6

The following design $d \in \mathcal{D}(6, 11, 3)$ is a GGDD(2).

1	1	1	1	1	1	2	2	3	3	4
2	2	2	2	3	4	3	5	4	5	5
3	4	5	6	6	5	4	6	6	5	6

The information matrix for d is

$$C_d = \frac{1}{3} \begin{pmatrix} 12 & -4 & -2 & -2 & -2 & -2 \\ -4 & 12 & -2 & -2 & -2 & -2 \\ -2 & -2 & 10 & -2 & -2 & -2 \\ -2 & -2 & -2 & 10 & -2 & -2 \\ -2 & -2 & -2 & -2 & 10 & -2 \\ -2 & -2 & -2 & -2 & -2 & 10 \end{pmatrix}.$$

The groups are $V_1 = \{1, 2\}$ and $V_2 = \{3, 4, 5, 6\}$. The concurrence parameters are $\gamma_{11} = 4$ and $\gamma_{12} = \gamma_{22} = 2$.

Definition 3.3 requires neither binarity nor equal replication, though both will frequently be employed. For instance, a binary, equireplicate GGDD(1) is just a BIBD, giving us our first example of optimality for GGDDs. The optimal design in Example 3.3, which is also a GGDD(1), is a case where neither additional property is employed. Nor are those properties possessed by the designs in Table 3.5, which are E -optimal GGDD(3)s (see Theorem 3.16).

A binary, equireplicate GGDD, with equal group sizes and with the further property that γ_{gh} is constant in $g \neq h$, goes by the simpler name *group divisible design* (GDD). It is customary with GDDs to drop the γ notation entirely and refer instead to the *in-group concurrence* λ_1 and the *out-group concurrence* λ_2 . If $|\lambda_1 - \lambda_2| = 1$, then a GDD is a RGD. These designs were formally defined by Bose and Shimamoto (1952).

Some GDDs enjoy a very broad optimality that, though not as strong as universal optimality, does encompass all three of the eigenvalue-based criteria listed in Table 3.2. In terms of eigenvalues, complete symmetry (3.10) is equivalent to $e_{d1} = e_{d2} = \dots = e_{d,v-1}$. Approximating complete symmetry (as through the GDD notion) will be valuable from an optimality perspective if it corresponds to an optimal relaxation of eigenvalue equality. The *type 1* optimality criteria provide a framework for investigating the form of this relaxation.

Definition 3.4 Let f be a real-valued function on \mathfrak{R} and define Φ_f on the class of information matrices \mathcal{C} by $\Phi_f(C_d) = \sum_{i=1}^{v-1} f(e_{di})$. Minimizing Φ_f is a type 1 optimality criterion if

- (i) f is continuously differentiable on $(0, \mathcal{T})$,
- (ii) $f' < 0, f'' > 0, f''' < 0$ on $(0, \mathcal{T})$,
- (iii) f is continuous at 0 and $\lim_{x \rightarrow 0^+} f(x) = +\infty$,

where $\mathcal{T} = \max_{\mathcal{C}} \text{trace}(C_d)$.

Taking $f(x) = 1/x$ gives the A criterion and $f(x) = -\log(x)$ the D criterion. All of the Φ_t criteria are included in the type 1 class, including (in the limit) E . The type 1 class does not include the MV criterion, for its value is not determined solely by the e_{di} .

Theorems 3.3 through 3.5 to follow each provide eigenvalue conditions for attainment of optimality with respect to type 1 criteria. They do this by employing various combinations of moment and extremal bounds for the set of $v - 1$ eigenvalues. Though not all are needed in this subsection, these theorems are best understood as a group, so are given here. Applications of these results extend to many situations other than GDDs, some of which will be explored in later subsections.

The task of minimizing Φ_f can be cast as an optimization problem independent of information matrices and their eigenvalues. The goal is to

$$\text{minimize } \sum_{i=1}^{v-1} f(x_i)$$

for f of Definition 3.4 and positive variables $x_1 \leq x_2 \leq \dots \leq x_{v-1}$. Where the minimum is attained will depend on constraints ruling the x_i . We will consider (in various combinations) the following:

$$\begin{aligned} & (i) \quad \sum x_i = T_1 \text{ for some positive } T_1. \\ & (ii) \quad \sum x_i^2 \geq T_2 \text{ for some } T_2 \text{ satisfying } \frac{T_1^2}{v-1} \leq T_2 \leq T_1^2. \\ & (iii) \quad x_1 \leq T_0 \text{ for some } T_0 \text{ satisfying} \end{aligned} \tag{3.14}$$

$$\begin{aligned} & (a) \quad T_0 \leq (T_1 - \sqrt{\frac{v-1}{v-2}P})/(v-1) \text{ where } P = \sqrt{T_2 - \frac{T_1^2}{v-1}}, \\ & (b) \quad (T_1 - T_0)^2 \geq T_2 - T_0^2 \geq (T_1 - T_0)^2/(v-2). \end{aligned}$$

In the design optimality application, constraints (3.14)(i) through (iii) correspond, respectively, to fixing the trace value $\text{trace}(\mathbf{C}_d) = \sum_i e_{di}$, placing a lower bound on the trace square $\text{trace}(\mathbf{C}_d^2) = \sum_i e_{di}^2$, and placing an upper bound on the smallest eigenvalue e_{d1} .

Theorem 3.3 (Cheng 1978) *The minimum of $\sum f(x_i)$ subject to conditions (3.14)(i) and (ii) is at*

$$x_1 = x_2 = \dots = x_{v-2} = \frac{T_1 - \sqrt{\frac{v-1}{v-2}P}}{v-1} \quad \text{and} \quad x_{v-1} = \frac{T_1 + \sqrt{(v-1)(v-2)P}}{v-1},$$

where P is as defined in (3.14)(iii).

For Theorem 3.3 applied to the design optimality problem, the first moment of the eigenvalues is fixed, such as when restricting to binary designs, and their second moment is bounded from below. The minimum of Φ_f is achieved at the second moment bound if the smallest $v - 2$ eigenvalues can be made equal. The next result incorporates information on the smallest eigenvalue.

Theorem 3.4 (Jacroux 1985) *The minimum of $\sum f(x_i)$ subject to all of the conditions (3.14)(i) through (iii) is at $x_1 = T_0$,*

$$x_2 = \dots = x_{v-2} = \frac{(T_1 - T_0) - \sqrt{\frac{v-2}{v-3}P_0}}{v-2}, \quad \text{and} \quad x_{v-1} = \frac{(T_1 - T_0) + \sqrt{(v-2)(v-3)P_0}}{v-2},$$

where $P_0 = \left[(T_2 - T_0^2) - \frac{(T_1 - T_0)^2}{v-2} \right]^{\frac{1}{2}}$.

In Theorem 3.4 the first moment of the eigenvalues is again fixed and again their second moment is bounded from below, but now there is also an upper bound for the minimum eigenvalue. The optimum is achieved at the second moment bound if we can get the smallest eigenvalue to reach its upper bound, and the remaining eigenvalues have a structure as in Theorem 3.3. Theorem 3.3 can also be refined by incorporating information on the largest eigenvalue, as shown next.

Theorem 3.5 (Cheng and Bailey 1991) *The minimum of $\sum f(x_i)$ subject to conditions (3.14)(i) and (ii), and the additional constraint $x_{v-1} \leq T_v$, is at any set of x_i 's for which there are exactly two distinct values and for which*

$$\sum_i x_i^2 = T_2 \quad \text{and} \quad x_{v-1} = T_v.$$

For any fixed trace of the information matrix, Theorems 3.3 through 3.5 show how the “equality of eigenvalues” implied by (3.10) can be advantageously relaxed in situations where it is not achievable. We now investigate the application of these results.

In the following, a design is termed *generalized optimal* if it is optimal with respect to all type 1 criteria. Working with a variant of Theorem 3.3 (see the original paper for details), Cheng (1978) proved generalized optimality for GDD(2)s having $\lambda_2 = \lambda_1 + 1$. This remains one of the most far-reaching optimality results, outside the realm of universal optimality explored in Section 3.3, which has been established for block designs. Otherwise, with few exceptions (see the next paragraph and following Theorem 3.9), it has been necessary to either restrict the class $\mathcal{D}(v, b, k)$ of competing designs, to restrict to a particular criterion, or both. The GDDs have been fertile ground in both regards.

In a very early paper, Takeuchi (1961) established *E*- and *MV*-optimality of the GDD(*s*)s having $\lambda_2 = \lambda_1 + 1$ for any $s \geq 2$. A GDD is said to be *singular* if $r = \lambda_1$ and *semiregular* if $r > \lambda_1$ and $rk = v\lambda_2$. Employing Theorem 3.5, Cheng and Bailey (1991) established generalized optimality within the binary, equireplicate subclass of $\mathcal{D}(v, b, k)$, for singular GDDs with $\lambda_2 = \lambda_1 - 1$ and semiregular GDDs with $\lambda_2 = \lambda_1 + 1$. Using majorization (see Bhatia 1997; and Section 3.5.3) techniques to establish inferiority of nonbinary and unequally replicated competitors, Bagchi and Bagchi (2001) were in many cases able to extend these generalized optimality results to the full class $\mathcal{D}(v, b, k)$.

The *dual* of an IBD d with treatment/block incidence matrix N_d is the design \tilde{d} (say) for which $N_{\tilde{d}} = N'_d$. If d is binary with equal replication r , then \tilde{d} is binary with replication $\tilde{r} = k$ and $C_{\tilde{d}} = kI - \frac{1}{\tilde{r}}N'_dN_d$ (compare Equation 3.4). Owing to the fact, true of any matrix times its transpose, that N'_dN_d has the same nonzero eigenvalues as $N_dN'_d$, the duals of the singular and semiregular GDDs mentioned earlier are also generalized optimal among binary, equireplicate competitors. The Cheng and Bailey (1991) paper contains several other results of this flavor, found by identifying classes of designs whose eigenvalues satisfy Theorem 3.5.

As noted earlier, GDDs are RGDs whenever $|\lambda_1 - \lambda_2| = 1$. GDDs are attractive targets for optimality arguments along the lines given earlier because their information matrices produce only two distinct eigenvalues. While other RGDs need not have this distinction, they do all possess (*M, S*)-property, so are approximating the requirements of Theorem 3.1 but without necessarily having a “nice pattern” for their eigenvalues like found in Theorems 3.3 through 3.5. Theorem 3.4 can nevertheless be effectively exploited for the RGD-optimality

problem. A two-armed approach handles the binary and nonbinary subclasses of $\mathcal{D}(v, b, k)$ separately. Henceforth denote the binary subclass by $\mathcal{B}(v, b, k)$.

In the theorem to follow, T_0 and T_0^* are upper bounds for e_{d1} , found from Lemma 3.5 and Corollary 3.1 (in Section 3.4.4) applied to non-RGDs. The smallest trace($C_{\bar{d}}^2$) that can be attained by a binary, non-RGD is the value used for T_2 . These values are plugged into Theorem 3.4 and a result of Kunert (1985) to produce non-RGD lower bounds for $\Phi_f(C_{\bar{d}})$, displayed as the right-hand side of the inequalities (3.15) for binary designs and (3.16) for nonbinary designs.

Theorem 3.6 (Jacroux 1985) *Let $\bar{d} \in \mathcal{D}(v, b, k)$ for $v > k \geq 3$ be a RGD. Write*

$$T_0 = \max \left\{ \frac{[r(k-1) - (\lambda + 2)]v}{(v-2)k}, \frac{[r(k-1) + \lambda - 1]}{k}, \frac{(r-1)(k-1)v}{(v-1)k} \right\},$$

$$T_0^* = \max \left\{ \frac{[r(k-1) - 2]v}{(v-1)k}, \frac{(r-1)(k-1)v}{(v-1)k} \right\},$$

$$T_1 = b(k-1), \quad T_2 = \text{tr}(C_{\bar{d}}^2) + \frac{4}{k^2}, \quad \text{and} \quad P_0 = \left[(T_2 - T_0^2) - \frac{(T_1 - T_0)^2}{v-2} \right]^{\frac{1}{2}}.$$

If $T_0 \leq [(T_1 - T_0) - \sqrt{\frac{v-2}{v-3}}P_0]/(v-2)$ and

$$\sum_{i=1}^{v-1} f(e_{\bar{d}i}) < f(T_0) + (v-3)f\left(\frac{(T_1 - T_0) - \sqrt{\frac{v-2}{v-3}}P_0}{v-2}\right) + f\left(\frac{(T_1 - T_0) + \sqrt{(v-2)(v-3)}P_0}{v-2}\right), \tag{3.15}$$

then a Φ_f -optimal design in $\mathcal{B}(v, b, k)$ must be a RGD. If moreover $T_0^* \leq [T_1 - \frac{2}{k} - T_0^*]/(v-2)$ and

$$\sum_{i=1}^{v-1} f(e_{\bar{d}i}) < f(T_0^*) + (v-2)f\left(\frac{T_1 - \frac{2}{k} - T_0^*}{v-2}\right), \tag{3.16}$$

then a Φ_f -optimal design in all of $\mathcal{D}(v, b, k)$ must be a RGD.

If these conditions hold, then all non-RGDs are ruled out with respect to Φ_f , but Theorem 3.6 does not say that the design \bar{d} must be Φ_f -optimal. Rather, it says that *some* RGD will be the optimal design; one must still complete the enumeration described following Example 3.5. Jacroux (1985) employed Theorem 3.6 to establish *A*- and *D*-optimality, over all of $\mathcal{D}(v, b, k)$, of the best RGDs determined by John and Mitchell (1977), lending strong credence to the John–Mitchell conjecture. Theorem 3.6 and its application to RGDs are an excellent example of how good theory can bring a complicated design problem within computational reach.

To this point, all of the designs shown to be optimal in this section have been binary. Even more, they have all had (M, S) -property, though this is not a requirement of the GGDDs in Definition 3.3. Indeed, there is no known example of an A -optimal IBD that is not binary. The situation is different with respect to the minimax criteria E and MV . E -optimality will be explored in more detail in Section 3.4.4; the remainder of this subsection concentrates on the MV criterion.

Many of the MV -optimality arguments found in the literature build on the simple inequality (3.17) first proven by Takeuchi (1961). It is applicable to any class of information matrices \mathcal{C} as identified prior to (3.3). For $C_d \in \mathcal{C}$, $C_d = (c_{di'v})$,

$$\frac{\text{Var}_d(\widehat{\tau}_i - \widehat{\tau}_{i'})}{\sigma_E^2} \geq \frac{4}{c_{dii} + c_{di'v} - 2c_{di'v}}. \tag{3.17}$$

The proofs of Theorems 3.7 and 3.8, not shown here, depend critically on (3.17). Again notice the emphasis on symmetry. The Theorem 3.7 result is for GGDD(1)s.

Theorem 3.7 (Morgan and Uddin 1995) *For the class of designs $\mathcal{D} = \mathcal{D}(v, b, k)$, suppose there is a $d^* \in \mathcal{D}$ for which*

$$\min_i \sum_{i' \neq i} c_{d^*i'i'} = \max_{d \in \mathcal{D}} \min_i \sum_{i' \neq i} c_{di'i'}, \text{ and} \tag{3.18}$$

$$C_{d^*} = \alpha \left(I - \frac{1}{v} J \right). \tag{3.19}$$

Then d^ is MV -optimal in \mathcal{D} . Moreover, if $d \in \mathcal{D}$ and $C_d \neq C_{d^*}$, then d^* is MV -superior to d .*

The second condition (3.19) of Theorem 3.7 is identical to the second condition (3.10) of Theorem 3.1. The two theorems are distinguished by their first conditions (3.18) and (3.9), with (3.18) allowing the possibility that trace may not be maximized. When the two conditions of Theorem 3.7 can be met without simultaneously meeting those of Theorem 3.1, *maximum trace designs cannot be MV -optimum*, for the asymmetry they necessarily entail results in higher variances for some elementary contrast estimators.

Theorem 3.7 gives us our first insight into how nonbinaryity can advantageously enter into the IBD-optimality problem. When equal replication is not possible (plot excess $p > 0$), judicious use of nonbinaryity can improve the approximation to complete symmetry and with this improve an optimality measure, particularly those (MV and E) that are most sensitive to an extreme contrast variance. Morgan and Srivastav (2002) have proven that the MV conclusions of Theorem 3.7 also hold for the E -criterion.

Like Theorem 3.1, which requires $(p, q) = (0, 0)$, the scope of Theorem 3.7 is fairly small in the (v, b, k) universe, needing (as can be shown) $(p, q) = (1, 0)$. The nonbinary design in Example 3.3 is MV -optimal over $\mathcal{D}(5, 7, 3)$ by Theorem 3.7. The A -optimal design in $\mathcal{D}(5, 7, 3)$ is binary.

Nonbinaryity can also enter into MV -optimal GGDDs with $s > 1$, as demonstrated next. Though neither binarity nor equal replication is required, Theorem 3.8(i) does imply binarity when $p = 0$. For $r = \lfloor bk/v \rfloor$, the c in (i) is $c = r(k - 1)/(v - 1)$.

Theorem 3.8 (Srivastav and Morgan 1998) *Let $d^* \in \mathcal{D}(v, b, k)$ be a GGDD(s) for which*

- (i) $c_g \geq c$ for $1 \leq g \leq s - 1$ and $c_s = c$,
- (ii) $\gamma_{ss} \geq \lfloor \frac{pk+2r(k-1)}{2(v-1)} \rfloor$,
- (iii) $\gamma_{gg} \geq \gamma_{ss}$ for $1 \leq g \leq s - 1$,
- (iv) γ_{gh} is constant in $g \neq h$,
- (v) $\frac{(\gamma_{gg}+(v-1)\gamma_{1s})}{v\gamma_{1s}(kc_g+\gamma_{gg})} + \frac{(\gamma_{hh}+(v-1)\gamma_{1s})}{v\gamma_{1s}(kc_h+\gamma_{hh})} \leq \frac{2}{kc+\gamma_{ss}}$ for $1 \leq g, h \leq s$ with $g \neq h$.

If $p \leq v - 2$, then d^* is MV -optimal in $\mathcal{D}(v, b, k)$.

The nonbinary design in Example 3.6 is MV -optimal by Theorem 3.8. Again, nonbinaryity is coming into play here when equal replication is not possible, provided it allows the group divisible structure to be achieved, so approximating (3.10) at the expense of (3.9). Condition (i), however, limits the extent of nonbinaryity. Many results of a similar flavor to Theorem 3.8 are known; see Srivastav and Morgan (1998) and the references therein.

The last two theorems beg the question of whether an optimal design in any of the senses in Table 3.2 can be nonbinary when equal replication is possible (i.e., when $p = 0$). The conjectured answer is *no*. Morgan (2009) examined this question for GGDD(1)s, also known as variance-balanced designs. Consider two designs d_1 and d_2 . If $C_{d_1} - C_{d_2} \neq 0$ is nonnegative definite, then d_1 dominates d_2 in the sense explained following (3), in which case d_2 is inadmissible. Let $\bar{d} \in \mathcal{D}(v, b, k)$ with $p = 0$ be a nonbinary GGDD(1), implying that $e_{\bar{d}i} \equiv \bar{e}$ is constant in i . If there is an E -optimal $d \in \mathcal{D}(v, b, k)$ that is binary and equireplicate, then $C_d - C_{\bar{d}} = C_d - \bar{e}(\mathbf{I} - \frac{1}{v}\mathbf{J})$ is nonnegative definite and \bar{d} is inadmissible. Morgan (2009) showed that of the 230 GGDD(1)s with $v \leq 15$ and $p = 0$ listed by Billington and Robinson (1983), 228 are inadmissible by virtue of comparison with E -optimal designs; the two remaining cases were left unresolved due to the E -optimal designs not being known.

3.4.3 Nearly Balanced Incomplete Block Designs

Theorems 3.7 and 3.8 demonstrate that when equal replication is not possible, nonbinary designs with sufficient balance are sometimes MV - or E -optimal (Sections 3.4.2 and 3.4.4). This has not been found to be the case for A -optimality, leading to a need to identify more flexible, binary design classes that leave room for measured steps away from symmetry.

Definition 3.5 *A nearly balanced incomplete block design with concurrence range ψ , or NBBD(ψ), is a design $d \in \mathcal{D}(v, b, k)$ satisfying*

- (i) Each $n_{dij} = 0$ or 1 (the design is binary),
- (ii) Each $r_{di} = r$ or $r + 1$ (replication numbers as equal as possible),

- (iii) $\max_{i \neq i', j \neq j'} |\lambda_{dii'} - \lambda_{djj'}| = \psi,$
- (iv) d that minimizes $\text{trace}(C_d^2)$ over all designs satisfying (i) through (iii).

Definition 3.5 is from Morgan and Srivastav (2000). The NBBD(1)s were earlier studied by Jacroux (1985), Jacroux (1991), and Cheng and Wu (1981), the last also considering a subclass of the NBBD(2) designs. The definition is tailored toward (M, S) -property, while allowing through (iii) that the concurrence numbers $\lambda_{dii'}$ need not be as close as possible.

Most of the designs found optimal in Sections 3.4.1 and 3.4.2 are NBBD(1)s for plot excess $p = 0$. We begin here with $p > 0$ and $\psi = 1$ but will also see that $\psi = 1$ is not always achievable and that contrary to the intuition from BIBDs, minimizing ψ does not necessarily lead to the best design (see Example 3.8). Two simple methods for constructing NBBD(1)s when $p > 0$ begin with a BIBD $d \in \mathcal{D}(v, b, k)$. For any b_0 such that $b_0 k < v$, one may (i) delete b_0 disjoint blocks from d or (ii) append b_0 disjoint blocks on the v treatments to d . The resulting NBBDs are binary GGDD($b_0 + 1$)s.

If $p > 0$ so that equal replication is not possible, the NBBD(1)s, being structurally close to RGDs, are promising candidates for A - and D -optimality. Relative to a NBBD(1), other binary designs must have larger $\text{trace}(C_d^2)$, opening the door for an adaptation of Theorem 3.4 for NBBD(1)s similar to that done with Theorem 3.6 for RGDs. That is, one can bound the E -value and the trace-square value for binary, non-NBBD(1)s, use these bounds to calculate a bound for the A -value (say) of non-NBBD(1)s via Theorem 3.4, and then compare this bound to the A -value of an actual NBBD(1) (compare (3.15) in Theorem 3.6). If the A -value of the NBBD(1) is smaller, then an A -optimal design in $\mathcal{B}(v, b, k)$ must be a NBBD(1). Analogous steps lead to an inequality similar to (3.16) for ruling out nonbinary competitors. See Jacroux (1991) for details, where many NBBD(1)s with $k > 2$ are shown to be A -optimal and refinements to Theorem 3.4 are also incorporated.

Some settings (v, b, k) do not allow $\psi = 1$ and larger concurrence ranges need to be considered. This comes about in either of two basic ways. First, if binarity forces some λ_{dij} to be less than the concurrence parameter λ defined in (3.11), then the concurrence range is necessarily at least two. An example is $\mathcal{D}(7, 28, 5)$ of Example 3.2. Second, requiring all $\lambda_{dij} \geq \lambda$ will, in some settings, force some $\lambda_{dij} \geq \lambda + 2$ as shown in Lemma 3.4.

Lemma 3.4 (Cheng and Wu 1981) *Let $d \in \mathcal{B}(v, b, k)$ have replication numbers $r_{di} \geq r$ for all i and concurrence numbers $\lambda_{dij} \geq \lambda$ for all $i \neq j$. If the concurrence excess q satisfies (i) $q > v - k$ or (ii) $q \leq v - k$ and $p(k - p) > (v - 2p)q$, then $\lambda_{dij} \geq \lambda + 2$ for some $i \neq j$.*

In similar fashion as explained earlier for NBBD(1)s, for settings where $\psi < 2$ is not possible, Theorem 3.4 can be adapted for establishing optimality of NBBD(2)s. This was done by Morgan and Srivastav (2000) who, among other findings, discovered instances of A -optimal NBBD(3)s. Examples 3.7 and 3.8 are based on their results.

Example 3.7

The following member of $\mathcal{D}(5, 7, 3)$ is a NBBD(2) with minimal concurrence $\lambda_{d12} = 1 = \lambda - 1$.

1	1	1	1	2	2	2
2	3	3	4	3	3	4
3	4	5	5	4	5	5

This design is A - and D -optimal, but neither E - nor MV -optimal. The best design for the latter two criteria is nonbinary—see Example 3.3.

Example 3.8

Up to isomorphism, $\mathcal{D}(9, 11, 5)$ contains two NBBD(2)s, here labeled d_1, d_2 , and one NBBD(3), d_3 .

$d_1 :$	1	1	1	1	1	1	1	2	2	3	4
	2	2	2	2	3	3	5	3	4	5	5
	3	3	4	5	4	4	6	6	5	6	7
	4	7	6	8	6	7	7	7	6	8	8
	5	9	8	9	9	8	9	8	7	9	9

$d_2 :$	1	1	1	1	1	1	1	2	2	3	3
	2	2	2	2	3	3	4	4	5	4	5
	3	3	4	6	4	6	5	6	7	6	6
	4	5	5	7	5	7	8	7	8	8	7
	9	8	6	9	7	8	9	8	9	9	9

$d_3 :$	1	1	1	1	1	1	1	2	3	3	4
	2	2	2	2	2	3	3	5	4	5	5
	3	3	4	4	7	4	4	6	7	6	6
	5	6	5	6	8	5	6	7	8	7	8
	9	8	7	9	9	8	7	8	9	9	9

All three designs have the same $\text{trace}(C_d^2)$ value. Of the NBBD(2)s, d_1 is A - and D -better, while d_2 is E - and MV -better. The NBBD(3), however, is better than the NBBD(2)s on all four of these criteria and is A - and D -optimal over all of $\mathcal{D}(9, 11, 5)$.

The approach taken in this subsection is also useful for the *irregular BIBD settings*: those triples (v, b, k) where the necessary conditions for existence of a BIBD (Lemma 3.2) are satisfied, but no BIBD exists. In an irregular BIBD setting, a binary design must have some λ_{dij} less than λ , suggesting NBBD(2)s as likely candidates for optimality. Since $p = 0$, equal replication is possible, and any binary, equireplicate design d will have discrepancy matrix Δ_d with row and column sums of zero. Define the total discrepancy δ_d of such a matrix to be $\delta_d = \sum_{i \neq j} |\delta_{dij}|/2$. Reck and Morgan (2005) enumerated the 51 feasible nonzero submatrices for all Δ_d for NBBD(2)s and NBBD(3)s having $\delta_d \leq 5$ and established that should a corresponding design d exist for any one of these, then the A -optimal and D -optimal designs in an irregular BIBD setting with $r \leq 41$ must be a NBBD with concurrence range and total discrepancy in these ranges. For $r \leq 41$, there are 497 BIBD settings that are either known to be irregular or for which the existence question is open. The smallest of these is the basis for Example 3.9, taken from Reck and Morgan (2005).

Example 3.9

An A - and D -optimal design in $\mathcal{D}(15, 21, 5)$. The total discrepancy of this design is $\delta_d = 4$.

1	1	1	1	1	1	1	2	2	2	2	2
2	2	3	4	5	7	8	3	4	5	7	8
3	6	6	10	9	13	10	11	6	10	9	9
4	7	9	13	11	14	11	13	12	12	11	10
5	8	12	14	15	15	12	15	15	14	14	13
3	3	3	3	4	4	4	5	5			
4	5	6	8	5	6	7	6	6			
7	7	9	12	8	8	9	11	7			
10	8	10	14	9	11	12	12	10			
11	13	14	15	15	14	13	13	15			

Including $\Delta_d = 0$, there are seven discrepancy matrices that produce better A -values than that of this design, but in $\mathcal{D}(15, 21, 5)$, there is no corresponding design for any of them.

In a later paper, the same authors (Morgan and Reck 2007a) determined an E -optimal design in $\mathcal{D}(15, 21, 5)$. Interestingly, that NBB(2) has total discrepancy 7, so its discrepancy matrix is not among the 51 candidate matrices for the aforementioned A -optimality. Moreover, they found little relationship between the A - and E -rankings of discrepancy matrices. These results point to the difficult interplay between combinatorial and optimality aspects of designs that can arise when “nice” designs do not exist. To date $\mathcal{D}(15, 21, 5)$ is the only irregular BIBD setting where optimal designs have been determined.

3.4.4 E-Optimal Block Designs

There has been a wealth of research on E -optimality of block designs, particularly for settings with plot excess $p = 0$ and more generally for those that fall within the NBB(1) framework. This is partly because, in working with only the smallest eigenvalue e_{d1} of C_d , the technical problems tend to be more readily tractable than with other eigenvalue-based criteria. But it is also due to its pragmatic appeal of keeping any contrast from being estimated poorly. When E -optimal designs differ from those that are A - or D -optimal, it is not uncommonly due to the latter estimating a few contrasts poorly and many slightly better.

The E -criterion asks that e_{d1} be maximized, so that E -optimality arguments accordingly require upper bounds for this smallest, positive eigenvalue of C_d . The basic lemma by which these are obtained (cf. result 6.59(b) in Seber 2008) is stated next.

Lemma 3.5 For any C_d and any nonzero vector h ,

$$e_{d1} \leq \frac{h' C_d h}{h' (I - \frac{1}{v} J) h}.$$

Different choices for h lead to different bounds, with utility depending on the particular design problem faced. Corollary 3.1 displays the most frequently employed.

Corollary 3.1 *The smallest positive eigenvalue e_{d1} of C_d satisfies*

- (i) $e_{d1} \leq (c_{dii} + c_{d'i'i} - 2c_{dii'})/2, i \neq i' = 1, \dots, v;$
- (ii) $e_{d1} \leq \sum_{i \in M} \sum_{i' \in M} v c_{dii'} / m(v - m)$ for any $m = |M|, M \subset \{1, 2, \dots, v\}.$

The bound (i) is based on pairwise contrasts, compare (3.17). Application of Corollary 3.1 leads easily to bounds for investigating E -optimality of regular graph designs.

Theorem 3.9 (Cheng 1980; Jacroux 1980) *Let $d \in \mathcal{D}(v, b, k)$ with $k < v$ and $p = 0$:*

- (i) $e_{d1} \leq \max\{[r(k - 1) - \lambda - 2]v / [(v - 2)k], [r(k - 1) + \lambda - 1] / k\}$ if d is equireplicate but not a RGD.
- (ii) $e_{d1} \leq (r - 1)(k - 1)v / [(v - 1)k]$ if d is not equireplicate.

Thus, if the E -best RGD $d^* \in \mathcal{D}(v, b, k)$ satisfies

$$e_{d^*1} \geq \max\{[r(k - 1) - \lambda - 2]v / [(v - 2)k], [r(k - 1) + \lambda - 1] / k, (r - 1)(k - 1)v / [(v - 1)k]\},$$

then d^* is E -optimal over all of $\mathcal{D}(v, b, k)$.

Cheng (1980) and Jacroux (1980), whose broadly overlapping articles were published back to back, applied Theorem 3.9 and related arguments to between them establish E -optimality for a substantial number of smaller designs, including 159 of the 209 parameter sets considered by John and Mitchell (1977). Cheng (1980) further proved E -optimality for $GDD(\frac{v}{2})$ s having $\lambda_1 = \lambda_2 + 1, \lambda_2 \geq 1$ as well as for the duals of these designs and for the duals of BIBDs and of GDDs with $\lambda_2 = \lambda_1 + 1$. The duals of BIBDs, known as *linked block designs*, in some cases enjoy generalized optimality (Bagchi and Bagchi 2001).

GDDs are members of a class of binary, equireplicate IBDs known as *partially balanced incomplete block designs* (PBIBDs) with two associate classes. Any 2-class PBIBD d has exactly two concurrence numbers $\lambda_{dii'} \in \{\lambda_1, \lambda_2\}$, and the pairwise contrast variances $\text{Var}_d(\widehat{\tau_i} - \widehat{\tau_{i'}})$ take just two distinct values; see Dey (2010) or Clatworthy (1973) (which includes an extensive tabling of smaller designs) for a good introductory treatment, including definition of *triangular parameters* for these designs. PBIBDs were extensively studied in the statistics community from their introduction by Bose and Nair (1939) until interest was overtaken by increasing emphasis on optimality from the mid-1970s onward. In effect, the PBIBDs were too broad a class from an optimality perspective, containing both optimal and markedly inefficient members. PBIBDs are most competitive when, as in the optimality results for the aforementioned GDDs, they have concurrence range of one, in which case they have (M, S) -property. Constantine (1982), relying heavily on Corollary 3.1(ii), was able to establish E -optimality of 2-class PBIBDs with $(\lambda_1, \lambda_2) = (1, 0)$ or $(0, 1)$ whenever an additional inequality on the triangular parameters is met. This included all 2-class PBIBDs with $(\lambda_1, \lambda_2) = (1, 0)$ and $b < v$; see the paper for details and full results. Cheng (2012) has recently combined more detailed knowledge of regular graphs with the bounds of Theorem 3.9 to prove further E -optimality results for RGDs, including E -optimality over equireplicate designs for all those PBIBDs with $\lambda_2 = \lambda_1 + 1$ known as either triangular type or L_2 type.

Expanding on the John and Mitchell (1977) cataloging of RGDs and later proofs of their optimality over all competitors, there has been progress on determining *all* E -optimal block designs for settings with $p = bk - vr = 0$ in a practical range. The basic idea is first to determine the E -best design in the equireplicate subclass \mathcal{B}^* of $\mathcal{B}(v, b, k)$ and then to prove its optimality over all of $\mathcal{D}(v, b, k)$. The first step requires C_d to be of the form (3.12), in which case e_{d1} can be expressed in terms of a_{dv} , the largest eigenvalue (other than q) of the discrepancy matrix Δ_d :

$$e_{d1} = \frac{v\lambda + q}{k} - \frac{a_{dv}}{k}. \tag{3.20}$$

With (3.20), E -optimization over \mathcal{B}^* is translated to an eigenvalue problem for discrepancy matrices as defined by (3.12) and (3.13). That is, it is translated to the class of symmetric, zero-diagonal, integer matrices with row sums q . John and Mitchell (1977) restricted this class to contain only 0/1 matrices, making a full search feasible for up to $v = 12$. Morgan (2007b) adapted Lemma 3.5 to remove that restriction, allowing the full class to be studied, but without a full enumeration, for $v \leq 15$. Relative to regular graph discrepancy matrices, bounding via Lemma 3.5 can initially eliminate any discrepancy matrix Δ with values corresponding to $\lambda_{dii'}$ that are too far from λ to allow optimality. As v grows, this still leaves far too many matrices to evaluate, a problem that can be overcome with another application of Lemma 3.5. Let $\Delta_d^{(t)}$ be any $t \times t$ principal minor of Δ_d , and let H_t be the set of normalized t -vectors. Then

$$e_{d1} \leq \frac{v\lambda + q}{k} - \frac{1}{k} \max_{h \in H_t} \left(\frac{v}{v - (h'\mathbf{1})^2} \right) h'(\Delta_d^{(t)} - \frac{q}{v}\mathbf{J})h. \tag{3.21}$$

Inequality (3.21) allows competitive Δ matrices to be built up one row/column at a time (take $t = 3, \dots, v$), at each step eliminating those substructures that are bounded below a predetermined value (such as e_{d1} for a known RGD). In this way, Morgan (2007b) was able to determine E -best discrepancy matrices, including many with values outside $\{0, 1\}$. These matrices may be classified into type 1, those with values in $\{0, 1\}$ only; type 2, values ranging from 0 to 2; type 3, values ranging from -1 to 1; and type 4, values ranging from -1 to 2; see Table 3.6. There are many cases where RGDs are not uniquely E -optimal (compare Theorem 3.9). When no Δ matrix of type 1 is E -best, no RGD can be E -optimal, provided a design exists for the E -best matrix.

Example 3.10

The E -best discrepancy matrix for $v = 7, q = 5$ has minimal a_v -value in (3.20) of $a_v = 1$. This matrix, call it Δ_{opt} , is

$$\Delta_{opt} = \begin{pmatrix} 0 & -1 & -1 & 1 & 1 & 1 & 1 \\ -1 & 0 & -1 & 1 & 1 & 1 & 1 \\ -1 & -1 & 0 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & -1 & 0 & 0 \\ 1 & 1 & 1 & -1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 & -1 \\ 1 & 1 & 1 & 0 & 0 & -1 & 0 \end{pmatrix}.$$

Δ_{opt} is potentially the basis for E -optimal design in $\mathcal{D}(7, 7 + 21m, 5)$ for each $m \geq 0$. The discrepancy matrix for the E -best RGD, displayed in Example 3.5, has $a_v = 1.247$, a clearly

TABLE 3.6

Number of E-Best Discrepancy Matrices by Type (1, 2, 3, 4)

v	q	Total Count	Count by Type	v	q	Total Count	Count by Type
4	1	2	(1,0,1,0)	12	1	11	(1,0,10,0)
	2	1	(1,0,0,0)		2	2	(1,0,1,0)
5	2	1	(1,0,0,0)	3	4	(0,0,4,0)	
	2	1	(1,0,0,0)	4	17	(1,0,16,0)	
6	1	3	(1,0,2,0)	5	56	(7,0,49,0)	
	2	2	(1,0,1,0)	6	1	(1,0,0,0)	
	3	1	(1,0,0,0)	7	480	(15,101,55,309)	
	4	1	(1,0,0,0)	8	1	(1,0,0,0)	
7	2	1	(0,0,1,0)	9	1	(1,0,0,0)	
	4	1	(1,0,0,0)	10	1	(1,0,0,0)	
8	1	5	(1,0,4,0)	13	2	1	(1,0,0,0)
	2	2	(1,0,1,0)		4	1	(0,0,1,0)
	3	6	(2,0,4,0)	6	1	(0,0,1,0)	
	4	1	(1,0,0,0)	8	18	(2,0,16,0)	
	5	1	(1,0,0,0)	10	1	(1,0,0,0)	
	6	1	(1,0,0,0)	14	1	15	(1,0,14,0)
9	2	1	(1,0,0,0)		2	2	(1,0,1,0)
	4	13	(4,0,9,0)		3	1	(0,0,1,0)
	6	1	(1,0,0,0)		4	1	(1,0,0,0)
10	1	7	(1,0,6,0)		5	1	(0,0,1,0)
	2	2	(1,0,1,0)	6	4	(1,0,3,0)	
	3	7	(1,0,6,0)	7	1	(1,0,0,0)	
	4	3	(1,0,2,0)	8	2213	(21,158,198,1836)	
	5	1	(1,0,0,0)	9	24	(2,0,22,0)	
	6	89	(6,26,5,52)	10	11	(6,0,5,0)	
	7	1	(1,0,0,0)	11	1	(1,0,0,0)	
	8	1	(1,0,0,0)	12	1	(1,0,0,0)	
11	2	1	(1,0,0,0)	15	2	1	(1,0,0,0)
	4	1	(0,0,1,0)		4	16	(0,0,16,0)
	6	8	(1,0,7,0)	6	234	(7,0,227,0)	
	8	1	(1,0,0,0)	8	19	(2,0,17,0)	
				10	1	(1,0,0,0)	
				12	1	(1,0,0,0)	

inferior competitor. The design in Example 3.5 is nevertheless E -optimal over $\mathcal{D}(7, 7, 5)$, in part because there is no competing design having discrepancy Δ_{opt} .

The discrepancy matrix for the design in Example 3.2 is Δ_{opt} , proving the design is E -optimal over $\mathcal{D}(7, 28, 5)$ and that it is E -better than the best RGD, found by adding the 21 blocks of a BIBD(7, 21, 5) to the Example 3.5 design. Adding $m - 1$ copies of the BIBD to the design in Example 3.2 produces an E -optimal design in $\mathcal{D}(7, 7 + 21m, 5)$ for every $m \geq 1$.

With the E -best discrepancy matrices in hand, Morgan (2007b) was able to complete the E -optimality design problem for $k > 2$, $v \leq 15$; save a few open cases where the best design is yet to be determined. In every case, optimality was established over the full class $\mathcal{D}(v, b, k)$, using Lemma 3.5 chiefly through Corollary 3.1. At this writing, the designs are available online at www.designtheory.org.

3.5 Small Block Designs

Three senses of what it means for an IBD to be “small” will be explored in this section. All are limiting cases: block size $k = 2$, number of blocks $b = 2$, and small number of units bk . Common to these situations is that at least some of the optimal designs have *very* different combinatorial properties than the designs seen in Section 3.4: they do not “approximate” BIBDs. The auxiliary parameters r , p , λ , and q , defined in (3.11), continue to be used here.

3.5.1 Block Size Two

Block size two would, at first glance, appear to present an optimality problem that, if not simple, would at least be straightforward. This turns out to be decidedly untrue. We begin with an example.

Example 3.11

This design $d \in \mathcal{D}(8, 8, 2)$ is both A -optimal and D -optimal:

$$d : \begin{array}{|c|c|c|c|c|c|c|c|} \hline 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \\ \hline 2 & 3 & 4 & 5 & 6 & 7 & 8 & 1 \\ \hline \end{array}$$

However, the A -optimal design d_1 and the D -optimal design d_2 in $\mathcal{D}(9, 9, 2)$ are quite different from each other:

$$d_1 : \begin{array}{|c|c|c|c|c|c|c|c|c|} \hline 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ \hline 2 & 3 & 4 & 1 & 1 & 1 & 1 & 1 & 1 \\ \hline \end{array}$$

$$d_2 : \begin{array}{|c|c|c|c|c|c|c|c|c|} \hline 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 \\ \hline 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 1 \\ \hline \end{array}$$

Design d_1 in Example 3.11 is especially perplexing due to its grossly unequal replication counts in a setting with no plot excess ($p = 0$). The following theorems will show it is not an isolated example.

The designs in Example 3.11 are all binary, for good reason. One simplification afforded by $k = 2$ is that nonbinary designs are inadmissible, for it is easily demonstrated that a nonbinary block of size two makes no contribution to C_d . Thus, the blocks of a design are a selection, with replacement, of pairs of treatments. Any such design can be cast as an undirected multigraph whose vertices are the treatments and with the number of edges connecting vertices i and i' being the number of blocks consisting of treatments i and i' . These *concurrency graphs* for two of the designs in Example 3.1 are displayed in Figure 3.1.




FIGURE 3.1
 Graphs for the designs d and d_1 in Example 3.11.

If the concurrence graph for a design has the same number of edges joining any two vertices, the design is a BIBD, universally optimal by Theorem 3.1. For a more thorough accounting of design from a graph perspective, see Bailey and Cameron (2013).

A few notions from graph theory are needed for results later in this section. The *degree* of a vertex is the number of edges incident with that vertex. A *leaf* is a vertex of degree one. A graph is *connected* if each vertex can be reached from any other vertex by a path. A *bridge* in a connected graph is an edge that, if removed, creates a disconnected graph. Any edge connected to a leaf is a bridge. A *star* is a collection of leaves incident with a common vertex. A *cycle* is a path with no repeated edges that begins and ends with the same vertex. A *tree* is a connected subgraph with no cycles. A *spanning tree* is a tree that includes every vertex in a graph.

Both graphs in Figure 3.1 are connected. In fact, connectedness of the designs (see Section 3.1) corresponds to connectedness of their concurrence graphs.* The graph for d_1 , which has five leaves and five bridges, is a star attached to a cycle. That for d is a cycle.

Theorem 3.10 (Cheng et al. 1985) *For each given v , there is a b_v such that for all $b \geq b_v$, a D -optimal design in $\mathcal{D}(v, b, 2)$ is*

- (i) A RGD if $p = 0$;
- (ii) A NBBD(1) if $p > 0$ and $\lambda = \lfloor (r + 1)/(v - 1) \rfloor$;
- (iii) A NBBD(2) if $p > 0$ and $\lambda = \lfloor (r + 1)/(v - 1) \rfloor - 1$.

Theorem 3.10 tells us that, for b sufficiently large, D -optimal designs with block size 2 must have the (M, S) -property. In particular, this says that the concurrences are being distributed across the treatment pairs as equally as possible. The graph-theoretic result of the next theorem makes a similar, but much broader, statement: it does not allow “too few” connections between any pair of treatments.

* This is true for the concurrence graph of *any* block design, not just those of block size two.

Theorem 3.11 (Bailey and Cameron 2013) *For any $b \geq v$, the concurrence graph of a D -optimal design in $\mathcal{D}(v, b, 2)$ contains no bridges and thus no leaves.*

Theorem 3.11 becomes especially interesting when compared to Theorem 3.12.

Theorem 3.12 (Bailey 2007) *For $20 \leq v \leq b \leq 5v/4$, any E -optimal design in $\mathcal{D}(v, b, 2)$ has leaves.*

A -optimal designs with b not too much larger than v must also have leaves; see Bailey and Cameron (2013). Thus, many A -optimal and many E -optimal designs with block size two, and with $r = \lfloor 2b/v \rfloor = 2$, have some treatments replicated once! This is in marked contrast to the D -optimality result with large b in Theorem 3.10 and provides myriad examples where optimal designs cannot have (M, S) -property.

The smallest b for $\mathcal{D}(v, b, 2)$ where $p = 0$ (and connected designs exist) is $b = v$. Much is known about this setting. Parts (i) and (ii) of Theorem 3.13 are special cases of results cited at the end of Section 3.5.3; the E -optimality is from Bailey and Cameron (2009). Designs d_1 and d_2 of Example 3.11 illustrate Theorem 3.13.

Theorem 3.13 *For $\mathcal{D}(v, v, 2)$, $v \geq 3$,*

- (i) *The D -optimal designs are the cycles.*
- (ii) *The A -optimal designs are the cycles if $v \leq 8$, $v - 4$ leaves attached to one vertex of a square if $9 \leq v \leq 12$, and $v - 3$ leaves attached to one vertex of a triangle if $v \geq 12$.*
- (iii) *The E -optimal designs are the cycles if $v \leq 6$, and $v - 3$ leaves attached to one vertex of a triangle, or a star with one doubled edge, if $v \geq 6$.*

If b is a multiple of $v(v - 1)/2$, then the optimal designs are precisely the BIBDs. Otherwise, as b grows larger with $p = 0$, RGDs are strong performers for which a number of optimality results are now known. Some of these are instances of results stated earlier in this article, and some have arisen from examination of small v . An E -optimality theorem for RGDs, whose proof relies heavily on Corollary 3.1, is stated next. Recall from Section 3.4.4 that a_{dv} is the largest eigenvalue (other than the concurrence excess q) of the discrepancy matrix $\mathbf{\Delta}_d$.

Theorem 3.14 (Bagchi and Cheng 1993) *For $v \geq 9$ and $1 \leq q \leq (v + 2)/2$, $q \neq \frac{v}{2}$, any RGD with $a_{dv} = 1$ is E -optimal over $\mathcal{D}(v, b, 2)$.*

Bagchi and Cheng (1993) further established E -optimality over all of $\mathcal{D}(v, b, 2)$ for RGDs with $v \leq 8$ and $q = v - 3$. Their results do not claim RGDs are uniquely E -optimal, and as evidenced by Table 3.6, it is not infrequent for certain non-RGDs, falling in the q range of Theorem 3.14, to also be E -optimal. The paper Morgan (2007b), discussed in Section 3.4.4, determined all E -optimal designs in the binary, equireplicate subclass of $\mathcal{D}(v, b, 2)$ for $v \leq 15$.

3.5.2 Two Blocks

Like block size two, with only two blocks, one might initially expect few complications in the optimality problem. Here, this expectation is met, up to a point. For IBDs, we have $k < v$, and if $2k \leq v$, then all designs are disconnected, so take $\frac{v}{2} < k < v$. The parameter r is thus $r = \lfloor 2k/v \rfloor = 1$ from which $p = 2k - v$, that is, $v = 2k - p$. Up to treatment relabeling, there is exactly one binary design (call it d_0 , see Table 3.7), which consequently has (M, S) -property. Design d_0 is a GGDD(3) with groups of size $p, k - p$, and $k - p$ and treatment replication two within the size p group.

An arbitrary design in $d \in \mathcal{D}(v, 2, k)$ must have at least $v - p = 2(k - p)$ treatments replicated exactly once. Partition the treatment set $V = V_1 \cup V_2 \cup V_3$ where V_1 contains those of the singly replicated treatments that appear in the first block, V_2 the singly replicated treatments in the second block, and V_3 the remaining treatments. Averaging C_d within each of the V_i produces \bar{C}_d for a GGDD(3) whose optimality values are lower bounds (see Lemma 3.3) for those of d . Comparing these to the values for d_0 yields Theorem 3.15.

Theorem 3.15 (Morgan and Jin 2007) *The unique binary design d_0 in $\mathcal{D}(v, 2, k)$ is A-optimal and D-optimal.*

Perhaps surprisingly, d_0 need not be the E-optimal design. Depending on k , it can be inferior to the nonbinary design d^* displayed in Table 3.7. The result is stated as Theorem 3.16. The proof in the cited paper is quite lengthy, the averaging technique based on $V_1 \cup V_2 \cup V_3$ being inadequate to the task.

Theorem 3.16 (Morgan and Jin 2007) *In the class $\mathcal{D}(v, 2, k)$, the binary design d_0 is the unique E-optimal design if $\frac{v}{2} < k < \frac{5v}{6}$; the nonbinary design d^* is the unique E-optimal design if $\frac{5v}{6} < k < v$; and both d^* and d_0 , and only these two designs, are E-optimal if $k = \frac{5v}{6}$.*

The E-optimality of the designs in Table 3.5 is a consequence of Theorem 3.16.

TABLE 3.7
Optimal Designs for $\mathcal{D}(v, 2, k)$

d_0 :	1	1	d^* :	1	1
	2	2		2	2
	⋮	⋮		⋮	⋮
	p	p		$2p - k$	$2p - k$
	$p + 1$	$k + 1$		$2p - k + 1$	$p + 1$
	$p + 2$	$k + 2$		$2p - k + 1$	$p + 2$
	⋮	⋮		$2p - k + 2$	$p + 3$
	⋮	⋮		$2p - k + 2$	$p + 4$
	⋮	⋮		⋮	⋮
	⋮	⋮		⋮	⋮
	k	v		p	$v - 1$
				p	v

3.5.3 Minimally Connected Designs

The block design model (3.1) requires at least $1 + (v - 1) + (b - 1) = b + v - 1$ degrees of freedom if all treatment contrasts are to be estimable. A block design produces bk total degrees of freedom, and thus connectedness requires that $bk \geq b + v - 1$. The designs in $\mathcal{D}(v, b, k)$ are *minimally connected* if $bk = b + v - 1$.

Minimal connectedness for block designs is analogous to “saturation” of fractional factorials (see Chapters 7 and 9), and the name *saturated block designs* is also used; in neither case are error degrees of freedom available. Minimally connected classes will be distinguished by the notation $\mathcal{D}(v, b, k)_m$. Lemma 3.6 provides basic properties for these classes.

Lemma 3.6 *For each $d \in \mathcal{D}(v, b, k)_m$, there is exactly one unbiased estimator for each treatment contrast. Consequently, any $d \in \mathcal{D}(v, b, k)_m$ is binary, and no pair of blocks has more than one treatment in common, that is, no pair of treatments occurs in more than one block.*

Proof: Suppose there are two unbiased estimators $p'y$ and $q'y$ for $c'\tau$. Let $a = p - q \neq 0$ and for the block design model (3.2), write $A_d = (\mathbf{1}, X_d, Z)$ and $\theta' = (\mu, \tau', \beta')$. Then $0 = E(a'y) = a'A_d\theta$ for all $\theta \Rightarrow a'A_d = 0' \Rightarrow bk > \text{rank}(A_d) = b + v - 1$, a contradiction. \square

A design in $\mathcal{D}(v, b, k)$ is *D-optimal* if, and only if, its concurrence graph has maximal number of spanning trees over all designs in \mathcal{D} (Cheng 1981). Lemma 3.6 tells us that the concurrence graph for a minimally connected design is a connected, simple graph with no cycles: it is itself a tree. This gives us this curious result:

Theorem 3.17 (Bapat and Dey 1991) *All designs in $\mathcal{D}(v, b, k)_m$ are D-optimal.*

Other criteria do distinguish designs in \mathcal{D}_m . Theorem 3.18 will establish a version of *M-optimality* for a particular member of that class. Let $v_{d1} \geq v_{d2} \geq \dots \geq v_{dm}$ be estimator variances for a set of m contrasts of interest when estimated using design d . Design d^* is *M-optimal* for estimation of the m specified contrasts if d^* minimizes $\sum_{g=1}^m f(v_{dg})$ for every monotonically increasing, convex function f .

Let v_d be the $m \times 1$ vector of the v_{dg} . A necessary and sufficient condition for *M-optimality* of d^* is that for every competing d , v_d *weakly submajorizes* v_{d^*} : $\sum_{g=1}^h v_{dg} \geq \sum_{g=1}^h v_{d^*g}$ for each $h = 1, \dots, m$ (e.g., Bhatia 1997, p. 40). Authors like Bagchi and Bagchi (2001) have pursued *M-optimality* based on the canonical variances e_{di}^{-1} (in which case, type 1 optimality functions are subsumed; compare Definition 3.4). Here, v_d is taken as the vector of variances for the $m = \binom{v}{2}$ pairwise comparisons $\widehat{\tau_i - \tau_{i'}}$.

The designs of interest have one treatment in every block, and all other treatments replicated once; see Table 3.8. The optimality result in Theorem 3.18 is a special case of a more general result proven in the cited paper.

Theorem 3.18 (Jin and Morgan 2008) *The designs in Table 3.8 are uniquely M-optimal for estimation of pairwise contrasts.*

Now the *A-value* is proportional to the average of the variances of the $\widehat{\tau_i - \tau_{i'}}$, and the *MV-value* is the maximum of those variances. Thus, among the results embedded in

TABLE 3.8
M-Optimal Design in $\mathcal{D}(v, b, k)_m$

1	1	1	...	1
2	$k + 1$	$2k$		$v - k + 2$
\vdots	\vdots	\vdots		\vdots
k	$2k - 1$	$3k - 2$		v

Theorem 3.18 are that the designs in Table 3.8 are *A*-optimal and *MV*-optimal over \mathcal{D}_m . These two individual optimalities were first proven by Mandal et al. (1990/1991) and Dey et al. (1995). Bapat and Dey (1991) prove Table 3.8 designs are *E*-optimal.

Table 3.8 designs are maximal trace GGDDs, with $b + 1$ groups and with smallest positive information matrix eigenvalue $1/k$. They maximize, rather than minimize, the replication range ($\max_i r_i - \min_i r_i$) over \mathcal{D}_m and, so like a number of other small designs seen in Sections 3.5.1 and 3.5.2, are well off from having the (*M, S*)-property.

If $bk = b + v$, then one degree of freedom is available for error estimation. Optimal designs are also known for this “saturated plus one” setting; not surprisingly, they are a bit more complicated than those in \mathcal{D}_m . See Balasubramanian and Dey (1996) for *D*-optimality, Krafft and Schaefer (1997) for *A*-optimality, and Dey et al. (1995) for *MV*-optimality. For the *E*-optimality problem, this author is not aware of a general solution, though Theorem 3.13 and Theorem 3.16 with $v = 2(k - 1)$ fall into this category.

3.6 Multiple Blocking Factors

The line-segment experiment described in Section 3.1 motivated variance reduction through blocking by subject. Because subject performance may change with practice, it could also be useful to block on the order of presentation of sheets to subjects. If this is done, the design selection problem will need to consider allocation of treatments relative to two blocking factors: subjects *and* time order. In agricultural experimentation to compare grain varieties, fields of the varieties are typically planted at several locations, and at each location, several harvesting teams may be employed. Now both location and harvesting teams are obvious blocking factors, for both are likely to contribute substantial variation to observed yields. These are just two examples of how more than one blocking factor can be incorporated when designing an experiment, in order to eliminate more than one major source of variability in experimental material.

This section focuses chiefly on the most commonly employed variants of two blocking factors, before briefly exploring some of the other possibilities. Let F_1 and F_2 denote two factors. Two broad design classes can be defined depending on the structural relationship, nested or crossed, between F_1 and F_2 .

Definition 3.6 Blocking factor F_1 is said to nest blocking factor F_2 if any two experimental units in an F_2 -block are in the same F_1 -block. If F_1 nests F_2 , then F_2 is nested in F_1 .

TABLE 3.9
Nested Block Design for 12 Treatments in Blocks of Size Six

1	3	1	2	1	2	1	2
2	7	4	3	3	5	3	4
4	8	8	5	4	9	6	5
5	9	9	6	6	10	10	7
6	11	10	7	7	11	11	8
10	12	12	11	8	12	12	9

Nested block designs are those with two blocking factors, one nesting the other. Members of the two blocking systems in a nested design are descriptively referred to as *superblocks* and *subblocks*. In the agricultural experiment described earlier, locations would be superblocks, and within each location, the harvesting teams would form subblocks. The design in Table 3.9 has four superblocks, each nesting two subblocks (four locations, two teams at each location). Nested block designs are taken up in Sections 3.6.1 and 3.6.2.

The model for observations generated with any nested design is a simple extension of (3.2) for standard block designs. Let b_1 be the number of superblocks and b_2 the number of subblocks per superblock, now using $b (= b_1 b_2)$ for the total number of subblocks. Let k be the number of experimental units per subblock, so that the superblock size is $k_1 = b_2 k$. The class of connected nested designs is denoted by $\mathcal{D}(v, b_1, b_2, k)$. The model for $d \in \mathcal{D}(v, b_1, b_2, k)$ is

$$y = \mu \mathbf{1} + X_d \tau + Z_1 \beta_1 + Z_2 \beta_2 + \epsilon, \tag{3.22}$$

where β_1 and β_2 are the $b_1 \times 1$ and $b_1 b_2 \times 1$ vectors of parameters for the nesting and nested blocks, respectively. With suitable ordering of the observations y and using the Kronecker product \otimes , the corresponding incidence matrices are $Z_1 = I_{b_1} \otimes \mathbf{1}_{k_1}$ and $Z_2 = I_b \otimes \mathbf{1}_k$.

Definition 3.7 *Blocking factors F_1 and F_2 are crossed if each F_1 -block intersects each F_2 -block in the same number of experimental units.*

Designs with two crossed blocking factors are termed *row-column designs*, another descriptive name; see Table 3.10. Members of the two blocking systems are now termed *row blocks* and *column blocks*, and the intersection of a row block with a column block is a *cell*. As is most frequently encountered, the designs in Table 3.10 implement the row-column idea with one unit per cell. More precisely, these are *complete crosses*. Readers interested in incomplete crosses, having either empty cells or unequal numbers of units per cell, are referred to Saharay (1996) as a starting point. Row-column designs are the topic of Section 3.6.3. If the line-segment experiment were to be run as a row-column design, the columns would denote subjects, and the rows would denote order of presentation.

The standard model for row-column designs can also be written in the form (3.22), where now β_1 and β_2 are the $b_1 \times 1$ and $b_2 \times 1$ vectors of parameters for the row blocks and column

TABLE 3.10

Three Row–Column Designs: A 3×12 for Six Treatments and a 4×4 and a 6×6 for Four Treatments

6	2	1	6	5	4	3	3	1	4	5	2
1	6	3	2	6	5	4	1	4	2	3	5
2	1	6	4	3	6	1	5	5	3	2	4

1	2	3	4
2	3	4	1
3	4	1	2
4	3	2	1

1	2	3	4	1	2
2	3	4	1	3	4
3	4	1	2	2	3
4	3	2	1	4	1
1	2	3	4	1	3
3	4	2	1	2	4

Note: Each design has one experimental unit and hence one assigned treatment per cell.

blocks, respectively. The two design models are distinguished by how their blocking structures are reflected in Z_1 and Z_2 . With one unit per cell and with observations in row-major order, they are now $Z_1 = I_{b_1} \otimes \mathbf{1}_{b_2}$ and $Z_2 = \mathbf{1}_{b_1} \otimes I_{b_2}$.

Though substantial, the optimality theory for designs with more than one blocking factor is not as well developed as that in Sections 3.3 through 3.5. The subsections to follow will report the main findings to date. Consistent with earlier sections in this chapter, the number of experimental units in a block is taken to be constant across blocks of the same blocking factor.

3.6.1 Resolvable Block Designs

Various classes of nested designs are defined by various requirements placed on the nesting blocks. A *resolvable block design* is a nested block design for which the superblocks are complete, that is, for which the block design defined by the nesting factor is a RCBD. The design in Table 3.9 is a resolvable block design: the superblocks *resolve* the subblocks into complete replicates of the treatment set. Resolvable designs are, by a wide margin, the most frequently employed nested designs in practice and have been extensively employed in agricultural field trials. The Table 3.9 design, for instance, could be used for a trial comparing 12 crop varieties over four locations (superblocks), using two harvesting teams (subblocks) at each location.

Writing $Z = (Z_1, Z_2)$, the information matrix arising from (3.22) for estimation of τ (see following (3.2)) is $C_d = X'_d(I - P_Z)X_d$. The nesting relationship says the column space of Z_1 is contained in the column space of Z_2 so that $P_Z = P_{Z_2}$ (superblock effects are “washed out” by subblock effects). Consequently, C_d reduces to $C_d = X'_d(I - P_{Z_2})X_d = R_d - \frac{1}{k}N_dN'_d$, where N_d is the $(v \times b)$ treatments \times subblock incidence matrix. It follows that a nested design is connected (i.e., all treatment contrasts are estimable) if and only if the subblock design is connected.

Lemma 3.7 *If a nested block design is Φ -optimal as a member of $\mathcal{D}(v, b, k)$ when ignoring superblocks, then it is Φ -optimal over $\mathcal{D}(v, b_1, b_2, k)$.*

Select any optimal design in $\mathcal{D}(v, b, k)$, then partition its $b = b_1b_2$ blocks into b_1 sets of b_2 blocks each. Lemma 3.7 says this is, an optimal design in $\mathcal{D}(v, b_1, b_2, k)$. If $b_2k = v$

and each treatment appears in each superblock, this is an optimal resolvable design. Simply put, the superblocks are irrelevant to the basic optimality argument. Nested designs would present no new optimality problems were special demands not placed on the nesting blocks. Those demands, though not reflected in the information matrix, are intended to incorporate pragmatic issues with experiment execution. Nesting factors often represent a grouping of subblocks that are similar due to closeness in time or space or that are subject to management by different personnel. A discussion of why nested designs are employed, with several examples, can be found in Section 1 of Morgan (1996).

Superblocks for resolvable designs are complete blocks. Resolvability assures that any problem associated with the nesting factor will impact every treatment equally, with possibilities ranging from model shortcomings at the nesting level to loss of entire superblocks during experimentation. If, in the preceding crop example, disease ruins yields at one location, the remaining locations still give the same amount of data on each variety. Resolvability imparts robustness to otherwise unforeseen difficulties at the superblock level.

The lesson that should be taken from Lemma 3.7 is that although the criterion value for a nested design is precisely that of its subblock design in $\mathcal{D}(v, b, k)$, the superblock requirements restrict the designs in $\mathcal{D}(v, b, k)$ that are even considered. The search for an optimal design in $\mathcal{D}(v, b_1, b_2, k)$ is the search for an optimal design over a subclass of $\mathcal{D}(v, b, k)$. The optimal member of the subclass may, or may not, coincide with an optimal design over all of $\mathcal{D}(v, b, k)$. To give one example, the A -optimal design in $\mathcal{D}(12, 8, 6)$, which is not resolvable, is easily shown to be better than the optimal resolvable design displayed in Table 3.9 with respect to all generalized optimality criteria.

Henceforth, we use $r = bk/v$ in lieu of the number of superblocks b_1 for resolvable designs and $s = v/k$ in lieu of the number of subblocks b_2 per superblock. The class of resolvable IBDs is denoted by $\mathcal{D}_{\mathcal{R}}(v, r, s, k)$. A simple instance where optimality does coincide for $\mathcal{D}_{\mathcal{R}}(v, r, s, k)$ and $\mathcal{D}(v, b, k)$ is when the latter contains a BIBD that is resolvable. For instance, take the ten blocks in Example 3.1, and add ten blocks given by their set complements, to get a BIBD in $\mathcal{D}_{\mathcal{R}}(6, 10, 2, 3)$ that is universally optimal over the larger class $\mathcal{D}(6, 20, 3)$. There is an extensive literature on resolvable BIBDs that will not be covered here; see Chapter II.7 of Colbourn and Dinitz (2007) for an introduction including many existence results.

3.6.1.1 Affine Resolvable Designs

The *affine resolvable designs* provide a very rich class of optimal resolvable designs.

Definition 3.8 *A resolvable design is affine resolvable if any two subblocks in different superblocks intersect in the same number of treatments.*

Denote the common subblock intersection number by ω . Then $k = \omega s$ so that $v = sk = \omega s^2$ is necessarily divisible by a square for an affine resolvable design. This observation points to an interesting combinatorial relationship. Given an affine resolvable design d , create a $v \times r$ array, call it \mathcal{B}_d , as follows. Rows of \mathcal{B}_d correspond to treatments of d . Columns of \mathcal{B}_d correspond to superblocks (replicates) of d . Labeling the subblocks in superblock j by $1, 2, \dots, s$ in any order, the (i, j) entry of \mathcal{B}_d is the label of the subblock in superblock j that contains treatment i . Affineness of d implies that any ordered pair of symbols from $\{1, \dots, s\}$

can be found as rows of any two-columned subarray of B_d exactly ω times. Thus, B_d is an orthogonal array (see Chapter 9) of strength two. Conversely, given a strength-two orthogonal array, the identification can be reversed to create an affine resolvable design. Affine resolvable designs are therefore *combinatorially equivalent* to strength-two orthogonal arrays, a fact that will shortly prove useful.

The importance of affine resolvable designs is made clear in Theorem 3.19.

Theorem 3.19 (Bailey et al. 1995) *Any affine resolvable design in $\mathcal{D}_{\mathcal{R}}(v, r, s, k)$ is M -optimal based on canonical variances. In particular, it is generalized optimal over $\mathcal{D}_{\mathcal{R}}(v, r, s, k)$.*

Also notice that if replicates (superblocks) are lost from an affine resolvable design, the remaining design is still affine resolvable, and so still optimal in all of the senses of Theorem 3.19.

Strength-two orthogonal arrays, and thus affine resolvable designs, have been extensively studied. Relying on that work, Bailey et al. (1995) established existence of affine resolvable designs with up to seven replicates for every $v = \omega s^2$, $\omega \geq 2$. The design in Table 3.9 is affine resolvable. Designs with $\omega = 1$ are also known as square lattices; see Section 3.6.1.2.

MV -optimality is not addressed by Theorem 3.19, and despite the very strong optimality established there, an affine resolvable design need not be an MV -optimal resolvable design (Morgan 2008). Even among nonisomorphic affine resolvable designs in $\mathcal{D}_{\mathcal{R}}(v, r, s, k)$, MV behavior can vary. Morgan (2010) determined MV -best affine resolvable designs based on an aberration criterion for v of the form $\omega 2^k$. There has otherwise been little work on MV -optimality of resolvable designs, affine or not.

3.6.1.2 Other Resolvable Designs

The earliest resolvable designs to be rigorously studied in the statistical literature were the lattices, introduced by Yates (1936b). The m -dimensional lattices for $v = t^m$ treatments have $k = t^u$ and $s = t^{m-u}$. The square lattices are those with $(m, u) = (2, 1)$, corresponding to $\omega = 1$ in Section 3.6.1.1. For $b \leq v$, Cheng and Bailey (1991) proved generalized optimality of square lattices over all binary, equireplicate competitors (not just resolvable designs). Chapter 18 of Hinkelmann and Kempthorne (2005) provides a good introduction to lattice designs. Also covered there are the *rectangular lattices* having $k = s - 1$, known to be highly efficient resolvable designs (Bailey and Speed 1986). Designs in this and the next paragraph are also discussed in Section 1.6.2.

Patterson and Williams (1976a) introduced the *alpha designs* as a flexible method for building resolvable designs having $k < s$. Particular members of this class are known to be highly efficient or optimal, but the technique is too general to admit a comprehensive optimality theory. A readable account of alpha design development and application, with many references, is given by Paterson (1988).

There has been productive work on small resolvable designs. Patterson and Williams (1976b) established a useful link between optimal, binary, equireplicate members of $\mathcal{D}(s, s, k)$ and optimal resolvable designs in $\mathcal{D}_{\mathcal{R}}(v, 2, s, k)$ having two replicates, showing how the former can induce the latter. Efficient resolvable designs having three replicates are provided in Bailey (2005). Morgan and Reck (2007b) established optimality for many classes of resolvable designs having two blocks per replicate.

3.6.2 Other Nested Block Designs

Relative to resolvable designs, there are two directions to explore for the nesting factor: superblocks of size less than v and of size greater than v . Although each of these directions offers multiple possibilities, there has been one main line pursued for each in the statistical literature.

For $b_2k < v$, the emphasis has been on the *nested BIBDs* (NBIBDs). A nested design is a NBIBD if the subblock design is a BIBD in $\mathcal{D}(v, b, k)$, and the superblock design is a BIBD in $\mathcal{D}(v, b_1, b_2k)$. Lemma 3.7 says that a NBIBD is universally optimal over $\mathcal{D}(v, b_1, b_2, k)$. Introduced by Preece (1967), an extensive survey for this topic including a tabling of smaller designs is available in Chapter VI.36 of Colbourn and Dinitz (2007). The BIBD structure at the superblock level provides robustness by assuring a reasonably good design should a superblock be lost, the remaining design being GGDD(2). Moreover, NBIBDs were proven universally optimal by Morgan (1996) when recovering both subblock and superblock information.

For $b_2k > v$, several authors have investigated α -resolvable block designs. The distinguishing characteristic of these nested designs is that the superblocks are multiply complete, with each treatment applied to α -experimental units in each superblock. Section 3.3 of Caliński and Kageyama (1996) contains an introductory treatment of α -resolvability, related extensions, and attendant results, to which the reader is referred for further details. Of special interest are the affine α -resolvable designs, which have any two subblocks intersecting in either ω_1 or ω_2 treatments depending on whether they are in the same, or different, superblocks. The technique employed by Bailey et al. (1995) to prove M -optimality of affine resolvable designs (Theorem 3.19) can be extended to include the affine α -resolvable designs.

3.6.3 Row–Column Designs

As indicated by Definition 3.7 and Table 3.10, row–column designs are represented as rectangular arrays in which rows correspond to levels of a blocking factor F_1 with b_1 (say) levels, columns correspond to levels of another blocking factor F_2 with b_2 levels, cells correspond to experimental units, and numbers in cells are labels for the treatments assigned. Row–column designs are randomized by selecting random permutations of rows, of columns, and of treatment labels, though there are variants on this (e.g., Section 10.2.2 of Hinkelmann and Kempthorne 2008). Row–column designs are widely used for experimentation in many disparate fields including agriculture (where they were first developed), psychology, and industry. Any general, applied textbook on experimental design will offer multiple examples.

Setting $\mathbf{Z} = (\mathbf{Z}_1, \mathbf{Z}_2)$, the information matrix $\mathbf{C}_d = \mathbf{X}'_d(\mathbf{I} - \mathbf{P}_Z)\mathbf{X}_d$ from the row–column version of model (3.22) simplifies to

$$\mathbf{C}_d = \mathbf{R}_d - \frac{1}{b_1}\mathbf{N}_d\mathbf{N}'_d - \frac{1}{b_2}\mathbf{M}_d\mathbf{M}'_d + \frac{1}{b_1b_2}\mathbf{r}_d\mathbf{r}'_d, \quad (3.23)$$

in which $\mathbf{N}_d = ((n_{dij})) = \mathbf{X}'_d\mathbf{Z}_2$ is the treatment/column block incidence matrix and $\mathbf{M}_d = ((m_{dii})) = \mathbf{X}'_d\mathbf{Z}_1$ is the treatment/row block incidence matrix (compare (3.4)).

The class of connected row–column designs is denoted $\mathcal{D}_{\mathcal{RC}}(v, b_1, b_2)$. The row–column design optimality problem amounts to minimizing $\Phi(\mathbf{C}_d)$ over $\mathcal{D}_{\mathcal{RC}}(v, b_1, b_2)$ for one or more of the usual criteria Φ . Any $d \in \mathcal{D}_{\mathcal{RC}}(v, b_1, b_2)$ has two component block designs,

$d_R \in \mathcal{D}(v, b_1, b_2)$ for the row blocks, and $d_C \in \mathcal{D}(v, b_2, b_1)$ for the column blocks. This suggests (compare Lemma 3.7) that the optimality problem might be fruitfully addressed through our considerable knowledge of optimality for the simple block design class $\mathcal{D}(v, b, k)$. This approach can work, provided two conditions can be met.

Theorem 3.20 *Suppose $d_0 \in \mathcal{D}(v, b_2, b_1)$ is Φ -optimal, is equireplicate, and the number of blocks b_2 for d_0 is an integer multiple of v . Then, there is a Φ -optimal row–column design $d^* \in \mathcal{D}_{RC}(v, b_1, b_2)$ for which the column component design d_C^* is d_0 .*

Proof: For any $d \in \mathcal{D}_{RC}(v, b_1, b_2)$, the matrix $\mathbf{M}_d \mathbf{M}'_d - \frac{1}{b_1} \mathbf{r}_d \mathbf{r}'_d$ is nonnegative definite. That is, $\mathbf{C}_{d_C} - \mathbf{C}_d$ is nonnegative definite and thus $\Phi(\mathbf{C}_{d_C}) \leq \Phi(\mathbf{C}_d)$ (see (3.3)(i)).

Write the blocks of d_0 as adjacent columns to give a $b_1 \times b_2$ array. Now permute treatments within each column as necessary to achieve equal replication of treatments within each row. The resulting design, which is d^* of the theorem, is said to be *row regular*. That row regularity can be achieved whenever an equireplicate block design has number of blocks a multiple of v is guaranteed by the basic theory for systems of distinct representatives (e.g., Chapter 5 of van Lint and Wilson 2001). Easily checked is that row regularity implies $\mathbf{M}_d \mathbf{M}'_d - \frac{1}{b_1} \mathbf{r}_d \mathbf{r}'_d = 0$ for any design d , so that \mathbf{C}_{d^*} reduces to \mathbf{C}_{d_0} . Thus,

$$\Phi(\mathbf{C}_{d^*}) = \Phi(\mathbf{C}_{d_0}) \leq \Phi(\mathbf{C}_{d_C}) \leq \Phi(\mathbf{C}_d),$$

for any $d \in \mathcal{D}_{RC}(v, b_1, b_2)$. □

Theorem 3.20 presents a simple, effective method for obtaining optimal row–column designs. The most popular row–column designs in practice fit into its framework. A *Latin square design* $d \in \mathcal{D}_{RC}(v, v, v)$ is a row–column design d for which d_C is a RCBD and d_R is a RCBD. A *Youden design* $d \in \mathcal{D}_{RC}(v, k, v)$ is a row–column design d for which d_C is a BIBD, and d_R is a RCBD. By Theorems 3.1 and 3.20, Latin square and Youden designs are universally optimal in their respective row–column classes. By the same reasoning, all *generalized Latin squares*, for which each of d_C and d_R is either a complete or a generalized complete block design, are universally optimal. A recent introductory review of Latin squares from the statistician’s viewpoint appears in Morgan (2007a). The 4×4 design in Table 3.10 is a Latin square design.

Also included in Theorem 3.20 are row–column designs where the row component is a complete block design and the column component is an optimal IBD, these being optimal cases of *Latin rectangles*. The 3×12 design in Table 3.10 falls into this category, it being a row-regular arrangement of an A - and E -optimal IBD in $\mathcal{D}(6, 12, 3)$. This design would be appropriate for a version of the line-segment experiment as described at the start of Section 3.6, employing twelve subjects (columns) with three presentations of pictures (rows) per subject.

A *generalized Youden design* (GYD) $d \in \mathcal{D}_{RC}(v, b_1, b_2)$ is a row–column design d for which d_C is a BIBD or a KBBB (see Section 3.3) and d_R is either a complete block, a generalized complete block, or a KBB, design. Universal optimality of the row-regular GYDs is immediate. Other GYDs, however, are *not* universally optimal (Das and Dey 1992). Kiefer (1975), employing a very delicate argument, was nonetheless able to prove that all “double KBBB” GYDs, in which both d_R and d_C are KBBBs, are both A -optimal and D -optimal, the latter

being for $v > 4$. A tabling of GYDs is available in Ash (1981). A nonregular 6×6 GYD for four treatments appears in Table 3.10.

Other results employing Theorem 3.20 are easily written. For instance, any GDD in $\mathcal{D}(v, mv, k)$ with $\lambda_2 = \lambda_1 + 1$ can be arranged as a row–column design that is generalized optimal over $\mathcal{D}_{\mathcal{RC}}(v, k, mv)$. However, the row–column design problem becomes quite difficult outside settings where row regularity is possible. Kiefer’s optimality proof for GYDs that lack row regularity is highly nontrivial, even with those designs having the best component designs d_R and d_C possible. The technical difficulties become only more complex for settings where “nice” component designs are not available. The reader is referred to Chapter 4 of Shah and Sinha (1989) for a review of relevant (and limited) results prior to 1989. The remainder of this section will highlight the most attractive of the sparse set of results obtained since. This is an area sorely in need of further development.

Some insight is gained by rewriting the information matrix (3.23) in terms of the component designs. Let C_{d_0} denote the information matrix for a CRD having the same replication numbers as d . Then, $C_{d_0} = \mathbf{R}_d - \frac{1}{b_1 b_2} \mathbf{r}_d \mathbf{r}_d'$ and

$$C_d = C_{d_R} + C_{d_C} - C_{d_0}, \quad (3.24)$$

which tempts one to optimize d_R and d_C separately. Unfortunately, this will not generally work because, for instance, eigenvalues for the individual information matrices may have no specific relationship to those for C_d . Commuting matrices, however, share their eigenspaces, which for equireplicate designs suggests investigating d_R and d_C for which C_{d_R} and C_{d_C} commute. The condition for this turns out to be that $M'_d N_d$ is a multiple of $J_{b_1 \times b_2}$ (Shah and Eccleston 1986), termed *adjusted orthogonality*. The challenge is to show that an adjusted orthogonal design with optimal d_R and optimal d_C is also optimal among a relevant class of designs that is not so restricted. One interesting result of this type is stated in Theorem 3.21.

Theorem 3.21 (Bagchi and Shah 1989) *An adjusted orthogonal row–column design in $\mathcal{D}(v, b_1, b_2)$ for which d_R and d_C are each linked block designs is M -optimal over all equireplicate competitors.*

Linked block design are defined in Section 3.4.4. Bagchi and Shah (1989) also show how to construct designs that satisfy Theorem 3.21. Bagchi and van Berkum (1991) prove E -optimality over all of $\mathcal{D}_{\mathcal{RC}}(v, b_1, b_2)$ for adjusted orthogonal designs with d_C and d_R being a linked block design and an affine resolvable design.

Row–column designs with two rows carry special interest due to their use in two-color microarray experiments (Bailey 2007). The following is a recent result along these lines:

Theorem 3.22 (Chai and Cheng 2011) *If v is even and $2b/v$ is an odd integer, then a design for which d_C is a BIBD, and for which the row replication numbers satisfy $|m_{di1} - m_{di2}| = 1$ for all i , is A -optimal over $\mathcal{D}_{\mathcal{RC}}(v, 2, b)$.*

In the same paper, generalized optimality within the equireplicate subclass of $\mathcal{D}_{\mathcal{RC}}(v, 2, b)$ is proven based on d_C being a GDD(2) (see Section 3.4.2) with $\lambda_2 = \lambda_1 + 1$, with

the row replication difference $m_{di1} - m_{di2}$ depending on group membership. Microarray experiments have also sparked interest in saturated row–column designs; see Chapter 23 and Qu (2010).

On the theme of “small” row–column designs, the optimality question has been thoroughly investigated for $v = 3$ in the pair of papers by Morgan and Parvu (2007) and Parvu and Morgan (2008). The results illustrate well how optimizing d_R and d_C can, depending on b_1 and b_2 , sometimes produce optimal, and sometimes decidedly suboptimal, designs. A little additional terminology is needed before describing the most interesting conclusions.

Let treatment i be replicated r_{di} times. The assignment of treatment i to units is *uniform in rows* if $m_{dil} \in \{\lfloor r_{di}/b_1 \rfloor, \lfloor r_{di}/b_1 \rfloor + 1\}$ for $l = 1, \dots, b_1$, and *uniform in columns* if $n_{dij} \in \{\lfloor r_{di}/b_2 \rfloor, \lfloor r_{di}/b_2 \rfloor + 1\}$ for $j = 1, \dots, b_2$. Row regularity in Theorem 3.20 is the special case of row uniformity for all treatments in an equireplicate design with b_2 an integer multiple of v , while row uniformity for all treatments without row regularity is invoked in Theorem 3.22. A design that is uniform in both rows and columns for all treatments is said to be *uniform*. A row–column design has *maximin replication* if $\min_i r_{di} = \lfloor b_1 b_2 / v \rfloor$. A design is *E_M -optimal* if it is E -optimal, and if among all E -optimal designs, it is M -optimal.

Morgan and Parvu (2007) prove that an A -optimal design in $\mathcal{D}_{RC}(3, b_1, b_2)$ *must* be uniform, but in some cases *cannot* have maximin replication. On the other hand, Parvu and Morgan (2008) prove that, aside from 4×4 layouts, an E_M -optimal design *must* have maximin replication, but in some cases *cannot* be uniform. For instance (see the papers for full results), if $b_1 \equiv 1 \pmod{3}$ and $b_2 \equiv 1 \pmod{3}$, then A -optimal designs are uniform with replications $r_{d1} = b_1(b_2 + 2)/3$ and $r_{d2} = r_{d3} = r_{d1} - b_1$. For the same setting, E_M -optimal designs have $r_{d2} = r_{d3} = r_{d1} - 1 = (b_1 b_2 - 1)/3$, but nonuniformity in all but the smallest cases.

Example 3.12

The two designs $d_1, d_2 \in \mathcal{D}_{RC}(3, 7, 7)$ shown here are, respectively, A -optimal and E_M -optimal. The replication counts (r_1, r_2, r_3) are $(21, 14, 14)$ for d_1 and $(17, 16, 16)$ for d_2 . With their given replications, d_1 is uniform, but d_2 has nonuniformity of treatment 1 in the first column.

$d_1 :$	<table border="1" style="border-collapse: collapse; text-align: center;"> <tr><td>1</td><td>1</td><td>1</td><td>2</td><td>2</td><td>3</td><td>3</td></tr> <tr><td>1</td><td>2</td><td>3</td><td>1</td><td>3</td><td>1</td><td>2</td></tr> <tr><td>1</td><td>2</td><td>3</td><td>2</td><td>3</td><td>1</td><td>1</td></tr> <tr><td>2</td><td>3</td><td>1</td><td>1</td><td>1</td><td>3</td><td>2</td></tr> <tr><td>2</td><td>3</td><td>1</td><td>3</td><td>1</td><td>2</td><td>1</td></tr> <tr><td>3</td><td>1</td><td>2</td><td>1</td><td>2</td><td>1</td><td>3</td></tr> <tr><td>3</td><td>1</td><td>2</td><td>3</td><td>1</td><td>2</td><td>1</td></tr> </table>	1	1	1	2	2	3	3	1	2	3	1	3	1	2	1	2	3	2	3	1	1	2	3	1	1	1	3	2	2	3	1	3	1	2	1	3	1	2	1	2	1	3	3	1	2	3	1	2	1
1	1	1	2	2	3	3																																												
1	2	3	1	3	1	2																																												
1	2	3	2	3	1	1																																												
2	3	1	1	1	3	2																																												
2	3	1	3	1	2	1																																												
3	1	2	1	2	1	3																																												
3	1	2	3	1	2	1																																												

$d_2 :$	<table border="1" style="border-collapse: collapse; text-align: center;"> <tr><td>1</td><td>2</td><td>1</td><td>1</td><td>2</td><td>3</td><td>3</td></tr> <tr><td>2</td><td>1</td><td>3</td><td>3</td><td>1</td><td>2</td><td>3</td></tr> <tr><td>2</td><td>2</td><td>1</td><td>3</td><td>1</td><td>3</td><td>3</td></tr> <tr><td>2</td><td>3</td><td>3</td><td>1</td><td>1</td><td>2</td><td>1</td></tr> <tr><td>3</td><td>1</td><td>2</td><td>2</td><td>2</td><td>1</td><td>1</td></tr> <tr><td>3</td><td>2</td><td>1</td><td>2</td><td>3</td><td>1</td><td>2</td></tr> <tr><td>3</td><td>3</td><td>2</td><td>1</td><td>3</td><td>1</td><td>2</td></tr> </table>	1	2	1	1	2	3	3	2	1	3	3	1	2	3	2	2	1	3	1	3	3	2	3	3	1	1	2	1	3	1	2	2	2	1	1	3	2	1	2	3	1	2	3	3	2	1	3	1	2
1	2	1	1	2	3	3																																												
2	1	3	3	1	2	3																																												
2	2	1	3	1	3	3																																												
2	3	3	1	1	2	1																																												
3	1	2	2	2	1	1																																												
3	2	1	2	3	1	2																																												
3	3	2	1	3	1	2																																												

3.6.4 Designs with Nesting and Crossing

Once one steps away from the basic settings with two crossed, or two nested, blocking factors, the possibilities grow rapidly. This section examines just two of these, each with three blocking factors and each with both nesting and crossing. See Morgan and Bailey (2000) for a systematic approach to optimal design with even more blocking factors.

3.6.4.1 Semi-Latin Squares

Row–column designs are not limited to having one experimental unit for each row / column intersection. With $k > 1$ units per cell, a cell effect can be included in the model so that “cell” becomes a third blocking factor. Because cell blocks are nested within both row blocks and column blocks, a result much like Lemma 3.7 entails: if the cell component block design d_{cell} is optimal over $\mathcal{D}(v, b_1 b_2, k)$, then the row–column design is optimal. Connectedness, too, is determined solely by the cell blocks. There is no new issue here unless, as in Sections 3.6.1 and 3.6.2, there are other demands placed on the nesting factors.

A row–column design is *doubly resolvable* if each treatment is replicated once in each row block and once in each column block. A Latin square is trivially a doubly resolvable design with $b_1 = b_2 = v$ and $k = 1$ unit per cell. The doubly resolvable row–column designs with $b_1 = b_2$ and $k > 1$ are termed *semi-Latin squares*. Writing b for the common number of rows/columns, the class of semi-Latin squares for $v = bk$ treatments is denoted $\mathcal{D}_{SL}(b, k)$. A member of $\mathcal{D}_{SL}(5, 3)$ is displayed in Table 3.11. An excellent introduction to semi-Latin squares with many examples of their use, in areas as widespread as consumer testing, agriculture, and message authentication, is provided by Bailey (1992).

Definition 3.9 *A Trojan square is a semi-Latin square found by superimposing k mutually orthogonal Latin squares (MOLSs) of order b , using a different set of symbols for each square.*

For a thorough accounting of MOLS with a table of squares of small side, see Chapter III.3 of Colbourn and Dinitz (2007). The semi-Latin square in Table 3.11 is a Trojan square. Cheng and Bailey (1991) (using Theorem 3.5 in Section 3.4) proved that Trojan squares are generalized optimal and so A -, D -, and E -optimal, over $\mathcal{D}_{SL}(b, k)$ and indeed over the entire binary, equireplicate subclass of $\mathcal{D}(bk, b^2, k)$. Thus, Trojan squares are preferred whenever they can be found.

There are at most $b - 1$ MOLS of side b , and a full set of $b - 1$ can be found whenever b is a prime or a power of a prime. This gives optimal Trojan squares in $\mathcal{D}_{SL}(4, k \leq 3)$, $\mathcal{D}_{SL}(5, k \leq 4)$, $\mathcal{D}_{SL}(7, k \leq 6)$, $\mathcal{D}_{SL}(8, k \leq 7)$, and $\mathcal{D}_{SL}(9, k \leq 8)$. There is no pair of MOLS of side 6 and so no Trojan square in $\mathcal{D}_{SL}(6, 2)$. Bailey and Royle (1997) determined the A -best, D -best, and E -best designs in the subclass of $\mathcal{D}_{SL}(6, 2)$ having (M, S) -property (all $\lambda_{d_{cell}, i, i'} \in \{0, 1\}$). These designs and all semi-Latin squares with no treatment concurrence exceeding 1 are called *SOMAs* (an acronym for *simple orthogonal multiarray*). Whether or not SOMAs are optimal over all of $\mathcal{D}_{SL}(6, 2)$ remains to be seen. More recently, Soicher (2013) has determined highly efficient squares in $\mathcal{D}_{SL}(6, k)$ for $4 \leq k \leq 10$. A complete

TABLE 3.11

5×5 Semi-Latin Square for 15 Treatments in Blocks of 3

1,6,11	2,7,12	3,8,13	4,9,14	5,10,15
5,9,13	1,10,14	2,6,15	3,7,11	4,8,12
4,7,15	5,8,11	1,9,12	2,10,13	3,6,14
3,10,12	4,6,13	5,7,14	1,8,15	2,9,11
2,8,14	3,9,15	4,10,11	5,6,12	1,7,13

enumeration in the same paper resulted in A -, D -, and E -optimal members of $\mathcal{D}_{SL}(4, k)$ for $4 \leq k \leq 10$. Soicher (2012) proved M -optimality of any semi-Latin square for which every two cells in different rows and columns intersect in the same number of treatments; this may be thought of as an extension of Theorem 3.19 to the doubly resolvable setup.

3.6.4.2 Nested Row–Column Designs

Just as row–column designs are not limited to one unit per cell, they are not limited to a single row–column layout. *Nested row–column designs* are seen in agricultural field trials where there is a separate row–column layout blocking on two sources of variation in each of several fields. They arise whenever a row–column experiment is repeated through time or space and row (column) blocks are unlikely to exert the same effects at each repetition. A nested row–column design can be as simple as two Latin squares, nesting v rows and v columns in each of the squares. The general setup is for b separate $b_1 \times b_2$ layouts, each now called a block, with one experimental unit per cell (a total of bb_1b_2 units). The class of all connected designs is denoted by $\mathcal{D}_{RC}(v, b, b_1 \times b_2)$. Table 3.12 displays two designs in $\mathcal{D}_{RC}(5, 10, 2 \times 2)$.

Theorem 3.23 (Bagchi et al. 1990) *Suppose there is a design $d^* \in \mathcal{D}_{RC}(v, b, b_1 \times b_2)$ for which (i) each block is row regular for the treatments it contains and (ii) the column blocks are a Φ -optimal design in $\mathcal{D}(v, bb_2, b_1)$. Then d^* is Φ -optimal over $\mathcal{D}_{RC}(v, b, b_1 \times b_2)$.*

So, for instance, if the bb_2 columns are the blocks of a BIBD, a row-regular nested row–column design is universally optimal. This is the case with design d_2 in Table 3.12. Likewise, nests of Latin squares, of Youden designs, and of row-regular GYDs are universally optimal. It is known that a nest of GYDs that are not row regular can be suboptimal; see Morgan (1997). Theorem 3.23 is an extension of Theorem 3.20 to the nested setup.

Curiously, if rows, columns, and blocks exert random effects, then an analysis recovering information from the three blocking strata may lead to a design that does not have the row-regular property of Theorem 3.23. The three component designs (blocks, row blocks, column blocks) of design d_1 in Table 3.12 are each BIBDs, together comprising a *BIBD with nested rows and columns* (BIBRC). Morgan and Uddin (1993) give conditions on the variances of the various blocking effects under which a BIBRC will outperform a design that is optimal by Theorem 3.23, in which case d_1 is preferred to d_2 in Table 3.12. Many authors have worked on constructing IBDs of both types, though progress has slowed of late. Especially as regards smaller designs (small b and $b_1b_2 \leq v$), there are few optimality results known, though this is not surprising given the state of knowledge for ordinary ($b = 1$)

TABLE 3.12
Two Nested Row–Column Designs for Five Treatments in Blocks of Size 2×2

d_1 :	<table border="1"><tr><td>1</td><td>5</td></tr><tr><td>3</td><td>2</td></tr></table>	1	5	3	2	<table border="1"><tr><td>1</td><td>4</td></tr><tr><td>5</td><td>3</td></tr></table>	1	4	5	3	<table border="1"><tr><td>1</td><td>3</td></tr><tr><td>2</td><td>4</td></tr></table>	1	3	2	4	<table border="1"><tr><td>1</td><td>2</td></tr><tr><td>4</td><td>5</td></tr></table>	1	2	4	5	<table border="1"><tr><td>2</td><td>3</td></tr><tr><td>4</td><td>5</td></tr></table>	2	3	4	5	<table border="1"><tr><td>3</td><td>1</td></tr><tr><td>2</td><td>5</td></tr></table>	3	1	2	5	<table border="1"><tr><td>5</td><td>1</td></tr><tr><td>3</td><td>4</td></tr></table>	5	1	3	4	<table border="1"><tr><td>2</td><td>1</td></tr><tr><td>4</td><td>3</td></tr></table>	2	1	4	3	<table border="1"><tr><td>4</td><td>1</td></tr><tr><td>5</td><td>2</td></tr></table>	4	1	5	2	<table border="1"><tr><td>4</td><td>2</td></tr><tr><td>5</td><td>3</td></tr></table>	4	2	5	3
1	5																																																	
3	2																																																	
1	4																																																	
5	3																																																	
1	3																																																	
2	4																																																	
1	2																																																	
4	5																																																	
2	3																																																	
4	5																																																	
3	1																																																	
2	5																																																	
5	1																																																	
3	4																																																	
2	1																																																	
4	3																																																	
4	1																																																	
5	2																																																	
4	2																																																	
5	3																																																	
d_2 :	<table border="1"><tr><td>1</td><td>2</td></tr><tr><td>2</td><td>1</td></tr></table>	1	2	2	1	<table border="1"><tr><td>1</td><td>3</td></tr><tr><td>3</td><td>1</td></tr></table>	1	3	3	1	<table border="1"><tr><td>1</td><td>4</td></tr><tr><td>4</td><td>1</td></tr></table>	1	4	4	1	<table border="1"><tr><td>1</td><td>5</td></tr><tr><td>5</td><td>1</td></tr></table>	1	5	5	1	<table border="1"><tr><td>2</td><td>3</td></tr><tr><td>3</td><td>2</td></tr></table>	2	3	3	2	<table border="1"><tr><td>2</td><td>4</td></tr><tr><td>4</td><td>2</td></tr></table>	2	4	4	2	<table border="1"><tr><td>2</td><td>5</td></tr><tr><td>5</td><td>2</td></tr></table>	2	5	5	2	<table border="1"><tr><td>3</td><td>4</td></tr><tr><td>4</td><td>3</td></tr></table>	3	4	4	3	<table border="1"><tr><td>3</td><td>5</td></tr><tr><td>5</td><td>3</td></tr></table>	3	5	5	3	<table border="1"><tr><td>4</td><td>5</td></tr><tr><td>5</td><td>4</td></tr></table>	4	5	5	4
1	2																																																	
2	1																																																	
1	3																																																	
3	1																																																	
1	4																																																	
4	1																																																	
1	5																																																	
5	1																																																	
2	3																																																	
3	2																																																	
2	4																																																	
4	2																																																	
2	5																																																	
5	2																																																	
3	4																																																	
4	3																																																	
3	5																																																	
5	3																																																	
4	5																																																	
5	4																																																	

row–column designs. Bagchi and Bose (2008) developed optimal main effects plans in small nested row–column designs. A number of results are summarized in the survey Morgan (1996).

3.7 Control Treatments and Weighted Optimality

Bailey (2008) described an experiment evaluating impact of several fungicide application regimens for winter wheat. Six treatments are combinations of spray time (early, mid, or late season) and spray amount (full or half spray). A seventh treatment is to do nothing: no spray. This baseline condition can be used to show that fungicide spray does have a strong impact on wheat yield. Differences of the other treatments relative to baseline are expected to be large; the more difficult questions, which are the principal target of the experiment, surround separating out relative effectiveness of the six spray treatments.

Some consumer product trials are aimed at determining if variants of a standard product are comparable to that standard. The primary question is: “Can a product that is ‘equally good’ be produced with differing ingredients?” The goal of these trials is to detect which of the new “test” products differs from the standard and which do not.

The two experimental situations just described earlier share an important commonality: they both incorporate a *control treatment*. They exhibit two of the usual ways in which controls arise, as an established treatment among new variants and as a “do nothing” treatment. The two situations differ, however, on the importance that is placed on comparisons with the control, relative to that on comparisons among the other, test treatments.

Use of a control treatment raises some interesting questions from a design optimality point of view. If, as in the fungicide trial, differences with the control are expected to be large, then it would make sense to invest less replication in the control and more in the test treatments, resulting in improved comparisons among the latter. The consumer product experiment, on the other hand, places high priority on comparisons with the control, indicating likely benefit from its *over*-replication.

Preceding sections have presumed equal interest in all treatments, incorporated in the permutation invariance property (3.3)(ii). By placing neither greater nor lesser emphasis on any treatment, this approach produces optimal designs that in most cases are either equally replicated or as close to equal replication as possible. As seen in the two preceding examples, and as is common for experiments including a control, a different approach is called for when not all treatment comparisons are of equal interest. This section formulates such an approach and presents basic results for its application.

A generic linear model for experiments with treatments and nuisance factors is

$$\mathbf{y} = \mathbf{X}_d \boldsymbol{\tau} + \mathbf{Z} \boldsymbol{\eta} + \boldsymbol{\epsilon}, \quad (3.25)$$

where $\boldsymbol{\eta}$ incorporates the model intercept and all blocking effects. This model can be modified to infuse differential treatment interest. Let $\mathbf{W}_{v \times v}$ be any symmetric, positive definite matrix, scaled so that $\mathbf{1}' \mathbf{W} \mathbf{1} = 1$. $\mathbf{W}^{1/2}$ denotes its symmetric, square root matrix satisfying $\mathbf{W}^{1/2} \mathbf{W}^{1/2} = \mathbf{W}$. Now write $\tilde{\mathbf{X}}_d = \mathbf{X}_d \mathbf{W}^{-1/2}$ and $\tilde{\boldsymbol{\tau}} = \mathbf{W}^{1/2} \boldsymbol{\tau}$. A reexpression of (3.25) is

$$\mathbf{y} = \tilde{\mathbf{X}}_d \tilde{\boldsymbol{\tau}} + \mathbf{Z} \boldsymbol{\eta} + \boldsymbol{\epsilon} \quad (3.26)$$

for which the treatments information matrix, denoted C_{dw} , is

$$C_{dw} = \tilde{X}'_d(\mathbf{I} - P_Z)\tilde{X}_d = W^{-1/2}C_dW^{-1/2}. \tag{3.27}$$

C_{dw} is the *weighted information matrix* for estimation of treatment contrasts. Equations 3.26 and 3.27 tell us that C_{dw} is a proper information matrix for estimation of linear model parameters, though we will not interpret it in that way. Like C_d , it is symmetric, nonnegative definite, and (for any connected design) rank $v - 1$. Unlike C_d , it does not have zero row sums, unless $\mathbf{1}$ is an eigenvector of W . One such case is $W = \frac{1}{v}\mathbf{I}$, giving (aside from the scale factor $\frac{1}{v}$) the ordinary (unweighted) information matrix as a special case.

The matrix W is chosen to change the emphasis an optimality criterion places on different contrasts, particularly those involving specially designated treatments such as controls. Toward that end, it will be profitable (though not necessary, see Stallings and Morgan 2014) to investigate those W that are diagonal. Henceforth,

$$W = \begin{pmatrix} w_1 & 0 & \cdots & 0 \\ 0 & w_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & w_v \end{pmatrix} \tag{3.28}$$

with all $w_i > 0$ and $\sum_i w_i = 1$. Then w_i is the *weight* assigned to the i th treatment. How weights are selected is determined by how they impact optimality evaluations. Conventional (unweighted) optimality theory has $w_i = 1/v$ for all i .

Definition 3.10 Let Φ be any optimality criterion satisfying the properties (3.3). For any given weight matrix W , the weighted criterion Φ_w is

$$\Phi_w(C_d) = \Phi(C_{dw}). \tag{3.29}$$

A design d^* in a class of competitors \mathcal{D} is Φ_w -optimal or weighted optimal with respect to Φ , if $\Phi_w(C_{d^*}) = \min_{d \in \mathcal{D}} \Phi_w(C_d)$.

Let \mathcal{P} be the collection of all $v \times v$ permutation matrices, and let \mathcal{P}_w contain those members of \mathcal{P} that preserve the weight matrix, that is, $\mathcal{P}_w = \{P \in \mathcal{P} : PWP' = W\}$. Weighted criteria inherit these properties from (3.3):

- (i) Φ respects the nonnegative definite ordering: for any $C_1, C_2 \in \mathcal{C}$ such that $C_1 - C_2$ is nonnegative definite, $\Phi_w(C_1) \leq \Phi_w(C_2)$.
 - (ii) Φ_w is invariant to treatment permutation in \mathcal{P}_w : $\Phi_w(PCP') = \Phi_w(C)$ for each $C \in \mathcal{C}$ and every $P \in \mathcal{P}_w$.
 - (iii) Φ_w is convex: $\Phi_w(\alpha C_1 + (1 - \alpha)C_2) \leq \alpha\Phi_w(C_1) + (1 - \alpha)\Phi_w(C_2)$ for all $C_1, C_2 \in \mathcal{C}$.
- (3.30)

Property (i) says weighting cannot change a dominating relationship, (ii) says that weighted criteria do not distinguish among treatments accorded the same weight in (3.28), and (iii) implies that matrix averaging (compare Lemma 3.3) with allowed permutations \mathcal{P}_w

will improve designs from a weighted perspective. The last of these, incidentally, indicates that GGDDs may be a useful design class in the weighted setup.

Definition 3.11 *The weighted variance for contrast $c'\tau$ estimated from design d is*

$$\text{Var}_{dw}(\widehat{c'\tau}) = [c'W^{-1}c]^{-1}\text{Var}_d(\widehat{c'\tau}). \quad (3.31)$$

The multiplier $[c'W^{-1}c]^{-1}$ is the weight of the contrast.

It is through weighted variances that weighted criteria are understood. Many conventional criteria, those that are functions of the eigenvalues e_{di} , are statistically meaningful because all normalized contrast variances are convex combinations of the e_{di}^{-1} (the canonical variances). Denote the positive eigenvalues of the weighted information matrix C_{dw} by $\theta_{d1} \leq \dots \leq \theta_{d,v-1}$. Wang (2009) established that all weighted contrast variances are convex combinations of the θ_{di}^{-1} , showing that eigenvalue-based weighted criteria are summary measures of weighted variance in the same way that their conventional counterparts summarize unweighted variance. Accordingly, the θ_{di}^{-1} are termed *canonical weighted variances*.

Weighted versions of the eigenvalue-based criteria in Table 3.2 are $A_w = \sum_i \theta_{di}^{-1}$ and $E_w = 1/\theta_{d1}$. Morgan and Wang (2010) show that the D criterion is unresponsive to weights (D_w produces the same design ordering regardless of the w_i). They also show that E_w is the largest weighted variance over all contrasts and that $A_{dw} = \sum_i \sum_{j \neq i} \frac{w_i w_j}{v-1} \text{Var}_d(\widehat{\tau_i - \tau_j})$ is a weighted sum of variances for elementary treatment contrasts, providing useful interpretations of each. For brevity, only eigenvalue-based criteria are considered here.

Weight selections w_i are made for the contrast weights $(c'W^{-1}c)^{-1}$ they induce. Though not every conceivable weighting can be captured by the diagonal weight matrix (3.28), situations of practical interest can. For experiments with a control (call it treatment 1) and $v - 1$ equal-interest test treatments, let w_1 be the control weight and $w_2 = \dots = w_v = (1 - w_1)/(v - 1)$ be the common weight for test treatments. With this selection of W , the weight assigned to any elementary control contrast $\tau_1 - \tau_i$ is $(w_1^{-1} + w_2^{-1})^{-1}$, while that for comparing two test treatments is $(2/w_2)^{-1}$. The ratio of these two contrast weights is $\xi = 2(v - 1)w_1/[1 + (v - 2)w_1]$. Selection of w_1 can be determined by the desired value of ξ , which ranges from 0 to 2 as w_1 is varied from 0 to 1. The range of ξ reflects the fact that in an orthogonal design, test–test comparisons can be estimated with variance arbitrarily smaller (subject to sample size) than, though no more than twice that, of test–control comparisons.

The A_w , E_w , and MV_w problems are solved for CRDs and all ξ in Morgan and Wang (2010). Taken up in the following are the implications when blocking in the presence of a control treatment. This topic has been intensively studied for one special case (Section 3.7.1), while general ξ has only of late come under investigation (Section 3.7.2).

A special case of (3.28) is weights proportional to replications ($w_i = r_i/bk$); compare (3.5) and (3.27). With these weights, canonical efficiency factors and canonical weighted variances are proportional. Thus, comparing designs on their average efficiency factor (3.8) is comparing on weighted A -values. From Table 3.3, we can now say that blocking according to the design in Example 3.3 must reduce variance by more than 20% if average weighted variance is to be better than a CRD with the same replication numbers.

If one is not interested in weighted information, then (3.8) should not be used unless restricting to equireplicate designs. It should also be clear that since weighted optimality criteria are based on a fixed choice of weights, (3.8) is only valid for comparing designs d with fixed replication vector r_d .

3.7.1 Design for TvC Experiments

The product-testing experiment described earlier is an example of a *test treatments versus control* (TvC) experiment. TvC experiments seek to evaluate as efficiently as possible comparisons of $v - 1$ test treatments with a control. First rigorously studied by Majumdar and Notz (1983), the relevant information matrix is $C_d^{TvC} = (H' C_d^+ H)^{-1}$, where $H = (1_{v-1}, -I_{v-1})$ is the coefficient matrix for the simple test versus control comparisons.

Lemma 3.8 (Morgan and Wang 2010) *Take $w_2 = \dots = w_v = (1 - w_1)/(v - 1)$ in (3.28) and let control weight $w_1 \rightarrow 1$. The limiting Moore–Penrose inverse of C_{dvw} is*

$$\lim_{\substack{w_2 \rightarrow 0 \\ w_1 \rightarrow 1}} \frac{1}{w_2} C_{dvw}^+ = \begin{pmatrix} 0 & 0 \\ 0 & H' C_d^+ H \end{pmatrix}.$$

Lemma 3.8 places design for TvC experiments in the weighted optimality framework as a limiting case. This relationship has been recently exploited in determining many E -optimal TvC designs. A design is E_2 -optimal if it is E -optimal and if it maximizes the second smallest eigenvalue e_{d2} of C_d among all E -optimal designs.

Theorem 3.24 (Morgan and Wang 2011) *Let $k_0 = \lfloor k/2 \rfloor$. Suppose $d^* \in \mathcal{D}(v, b, k)$ satisfies these conditions:*

- (i) *The control treatment is assigned to $n_{d^*1j} = k_0$ units in block j for $j = 1, 2, \dots, b$.*
- (ii) *d^* is equally replicated and binary in the $v - 1$ test treatments.*
- (iii) *The design $d_1^* \in \mathcal{D}(v - 1, b, k_1 = k - k_0)$ found by deleting all replicates of the control from d^* is connected.*

Then

- *d^* is E -optimal for the TvC problem.*
- *If d_1^* is E -optimal over $\mathcal{D}(v_1, b, k_1)$, then d^* is E_2 -optimal for the TvC problem.*
- *If d_1^* is M -optimal over $\mathcal{D}(v_1, b, k_1)$, then d^* is E_M -optimal for the TvC problem.*

So long as $v_1 = v - 1$ divides bk_1 and $k_1 \leq v_1$ all known M -optimal (this includes BIBDs) and E -optimal (see Section 3.4.4) block designs for $\mathcal{D}(v_1, b, k_1)$ are binary and equireplicate. Theorem 3.24 successfully exploits the considerable knowledge of ordinary block designs in achieving optimality for TvC experiments.

The A -optimality problem for TvC experimentation has a much longer history, having been developed before the general weighted approach was available. Much of the known theory has focused on GGDD(2) designs. Because of the specialized application

to TvC experiments, the designs have been given a distinctive name (for a slightly broader definition, see the original paper Bechhofer and Tamhane 1981).

Definition 3.12 A design $d \in \mathcal{D}(v, b, k)$ is a balanced treatment incomplete block design (BTIBD) if it is a GGDD(2) with groups of sizes $v_1 = 1$ and $v_2 = v - 1$, it is uniform in the control treatment (the treatment in the first group), and it is binary in the $v - 1$ treatments of the second group.

Since a BTIBD is uniform in the control, there are integers t and s for which the total control replication is $tb + s$. That is, the control is assigned to t units in $b - s$ blocks and to $t + 1$ units in s blocks. This is denoted by $\text{BTIBD}(v, b, k; t, s)$. The designs in Example 3.13 below are $\text{BTIBD}(8, 5, 7; 2, 0)$ and $\text{BTIBD}(8, 5, 7; 1, 0)$.

BTIBDs are “maximally averaged” with respect to the allowable permutations \mathcal{P}_w , which is all permutations of the $v - 1$ test treatments. The information matrix C_d^{TvC} for a BTIBD is thus completely symmetric, but Theorem 3.1 does not apply due to this matrix having full rank. The eigenstructure is, nonetheless, very simple, with $v - 2$ of the eigenvalues being identical. A happy consequence is that the A -value for any BTIBD is easily calculated. Indeed, any competing design, regardless of its combinatorial properties, can have its information matrix averaged to produce \bar{C}_d^{TvC} of GGDD(2) form with this same simple eigenstructure, yielding a workable lower bound expression for its A -value. This expression can be minimized over all of $\mathcal{D}(v, b, k)$, and should that minimum occur at a BTIBD, the A -optimal design has been identified. In this way one can prove the following theorem:

Theorem 3.25 (Stufken 1987) If $(k - t - 1)^2 + 1 \leq t^2(v - 1) \leq (k - t)^2$, then a $\text{BTIBD}(v, b, k; t, 0)$ is A -optimal over $\mathcal{D}(v, b, k)$.

Theorem 3.25 is one of the more simply stated of many A -optimality results for BTIBDs. An excellent survey including many examples is available in Majumdar (1996), where variants on the approach outlined earlier, pursued extensively in the 1980s and 1990s, are described. Catalogs of A -optimal TvC designs are available in Hedayat and Majumdar (1984) and Ting and Notz (1988).

Example 3.13

These two designs are both GGDD(2) with overreplication of the control:

$d_1 :$	1	1	1	1	1	1	1	$d_2 :$	1	1	1	1	1	1	1
	1	1	1	1	1	1	1		2	3	4	5	6	7	8
	2	3	4	5	6	7	8		3	4	5	6	7	8	2
	3	4	5	6	7	8	2		6	7	8	2	3	4	5
	5	6	7	8	2	3	4		8	2	3	4	5	6	7

By Theorem 3.25, d_1 is E_M -optimal. By Theorem 3.2 of Majumdar (1996), d_2 is A -optimal.

The optimal designs in Example 3.13 are quite different from one another. Substantial disagreement of criteria as seen here underscores the need for careful criterion choice in

TABLE 3.13
Optimal TwC Designs in $\mathcal{D}(6, 6, 4)$

1	2	2	2	3	1	2	2	2	1	2	2	2	2	3
2	3	2	3	4	2	3	3	3	1	3	3	3	3	4
3	4	3	4	5	3	4	4	4	2	4	4	4	4	5
4	5	4	5	6	4	5	5	5	3	5	5	5	5	6
E: 0.05-0.25, A: 0.05-0.10, MV: 0.05-0.30														
1	2	2	2	3	1	2	2	2	1	2	2	2	2	3
2	3	2	3	4	2	3	3	3	2	3	3	3	3	4
3	4	3	4	5	3	4	4	4	3	4	4	4	4	5
4	5	4	5	6	4	5	5	5	4	5	5	5	5	6
E: 0.30-0.55, A: none, MV: 0.35-0.55														
1	2	2	2	3	1	2	2	2	1	2	2	2	2	3
2	3	2	3	4	2	3	3	3	2	3	3	3	3	4
3	4	3	4	5	3	4	4	4	3	4	4	4	4	5
4	5	4	5	6	4	5	5	5	4	5	5	5	5	6
E: 0.60, A: 0.15-0.40, MV: 0.60														
1	2	2	2	3	1	2	2	2	1	2	2	2	2	3
2	3	2	3	4	2	3	3	3	1	3	3	3	3	4
3	4	3	4	5	3	4	4	4	2	4	4	4	4	5
4	5	4	5	6	4	5	5	5	3	5	5	5	5	6
E: 0.85-1.00, A: 0.75-1.00, MV: 0.85-1.00														
1	2	2	2	3	1	2	2	2	1	2	2	2	2	3
2	3	2	3	4	2	3	3	3	1	3	3	3	3	4
3	4	3	4	5	3	4	4	4	2	4	4	4	4	5
4	5	4	5	6	4	5	5	5	3	5	5	5	5	6

Note: Contrast weight ratios ξ for which each design is optimal are shown at 0.05 increments.

accordance with experimenter goals. Criterion disagreement is more prevalent in weighted situations; see Morgan and Wang (2010).

Gupta et al. (1999) took this line of attack a step further, broadening the class of experiments that fall into the TvC categorization. Though not having the framework for weighted optimality described in (3.27) through (3.29), their work is equivalent to the general A_w problem with W specified by (3.28), allowing all w_1 placing higher priority on control comparisons ($\frac{1}{v} < w_1 < 1$) and so not restricting to the limiting case. Results for the E_w criterion in this same, broader framework were obtained by Wang and Morgan (2010).

3.7.2 Lesser Weight on the Control

At the other end of the weight spectrum are the *test treatments with control* (TwC) experiments. TwC experiments, which include the earlier described fungicide trial, are defined by according lesser ($w_1 < 1/v$) weight to the control. This has been a topic in two recent articles, one of which is briefly described here.

Small w_1 will typically mean small replication for the control. Matrix averaging then becomes a less effective technique (the bounds produced can be further from what is actually attainable) for determining optimal or efficient designs, the difficulties being more acute in small designs. Wang and Morgan (2010) thus resorted to enumeration to find small, optimal TwC designs. For all but the very smallest problems, feasibility restrictions on enumeration are required. These authors conjectured that optimal designs for TwC experimentation in $\mathcal{D}(v, b, k)$ ($k < v$) are binary in all treatments and have range no more than one for the replication numbers of the $v - 1$ test treatments (compare the John–Mitchell conjecture in Section 3.4.1). With these restrictions, they compiled a catalog of optimal TwC designs for all $v \leq 12$ and $b \leq 12$ with $v + b \leq 18$. Table 3.13 displays results of one of these enumerations.

The designs in Table 3.13 hold no surprises. While the criteria do not always agree, replication of the control increases with w_1 (equivalently, as the contrast weight ratio ξ increases) regardless of the criterion. If sufficiently close to the unweighted case ($w_1 = 1/v$; $\xi = 1$), the criteria agree that an equally replicated design (a RGD) is best. The criteria also agree on the same design, having just one replicate of the control, for sufficiently small w_1 .

3.8 Discussion and Directions

Blocking is one of the three fundamental techniques of design propounded by Fisher (1935), growing out of his development of experimental ideas nearly a century ago. Since then, it has been applied in multitudinous forms in a wide variety of experimental settings, with increasing examination from theorists as time has passed. This chapter has attempted to give a reasonably compact overview of the basic blocking implementations and, for them, the mathematical techniques that have been developed to maximize the information that a blocked experiment can provide.

There are many extensions. Covered elsewhere in this volume are blocking of fractions (Chapter 8), and blocking with nonlinear models for responses (Chapter 13). Yet other topics are blocking with nonconstant block sizes (Ghosh et al. 2008), block size $k > v$ (Jacroux and Whittinghill 1988, especially their Example 2.5), blocking with treatments and continuous covariates (Chapter 4 of Liski et al. 2002), blocking in the presence of correlated

errors (Martin and Eccleston 1991) and robustness of blocking schemes to loss of observations (Morgan and Parvu 2008; Godolphin and Warren 2011) or to weight range (Wang and Morgan 2012). The just-cited papers contain references for further exploration.

The simple block designs of Sections 3.3 through 3.5 have been the most intensely studied. It is then not surprising that they are the best understood, this being due in part to greater tractability of the underlying mathematical problems. Yet there is still much unsolved, especially as regards designs where average replication must be small. With small designs, as indicated by Section 3.5, the combinatorial properties that lead to optimality in larger designs do not necessarily prove useful. Paralleling this, optimality bounds from matrix averaging and, from results such as Theorems 3.3 through 3.5, also tend to be unhelpful. In many cases, small designs are simply “too discrete” to closely conform to results found through the smoothness of continuous mathematics. Looking in the other direction, the John–Mitchell conjecture that some RGD must be optimal when a RGD exists has been confirmed for many criteria with arbitrarily large designs by Cheng (1992).

The asymptotics of Cheng (1992) suggest this interesting question: given an optimal design in $\mathcal{D}(v, b, k)$ and a BIBD in $\mathcal{D}(v, b_0, k)$, when is the collection of blocks from both designs an optimal design in $\mathcal{D}(v, b+b_0, k)$? The suggestion is certainly attractive, for adding a BIBD to the IBD d simply adds a constant to the eigenvalues e_{di} . The E -optimality work by Morgan (2007b) described in Section 3.4.4 has established arbitrary *BIBD-extendability* for most E -optimal block designs with $v \leq 15$. Cakiroglu (2013) has shown that extendability holds with respect to A - and D -optimality for many of the regular graph designs in the John–Mitchell catalog. The general question, however, remains open.

In some situations, more than one optimal block design has been identified for the same criterion; see, for instance, Table 3.6. The obvious solution in this case is to bring other criteria to bear, further trimming the competitors by their values on one or more secondary optimality measures. This idea was seen in Theorem 3.24 with E_2 - and E_M -optimality. Alternatively, one might examine the collection of contrast variances that will be produced by each design, seeking simplicity that will aid interpretation of experimental results. It may on occasion be worthwhile to sacrifice a small measure of efficiency in exchange for more readily interpretable results.

When an IBD is needed and an optimal design is not known, the results reported in this chapter give good guidance about what structural properties should be sought. So long as the setting is not too small (cf. Section 3.5), a design incorporating (i) and (ii) of Definition 3.5, and keeping $\text{trace}(C_d^2)$ as small as possible, will be reasonably close to highest efficiency.

References for where optimal block designs can be found are scattered throughout this chapter, but they are by no means complete. This reflects a substantial shortfall of design research: the failure to create easy access to large collections of optimal designs. Practicing statisticians need designs that are readily available, be they from online catalogs or generated by popular software. Unfortunately, an optimal design in a journal article can be only there—it may not exist for much of the experimental world. Two projects to help remedy this situation, *designtheory.org* (Bailey et al. 2006) and *www.iasri.res.in/design*, have appeared in the past decade. While these are very positive steps, there is still much to be done.

For those desiring to generate block designs with specified combinatorial properties, including resolvability restrictions, the freely available software GAP Design (Soicher 2009) is an excellent tool so long as $v + b$ is not too large (typically less than 20, though block size k also plays a role). Numerous techniques for design construction have been developed

through the years, many having roots in the *method of differences* developed by Bose (1939) for construction of BIBDs. Bose's idea was to start with a few, carefully crafted blocks from which all the blocks of a design could be generated through a group action. That general idea is an integral part of GAP Design.

The impact of increasing computational power has not been as great in discrete design as in most other areas of statistics. This is in part because some of the optimality problems, like those for $\mathcal{D}(v, b, k)$ with plot excess $p = 0$, have partially yielded to analytic solution. It is also due in part to combinatorial possibilities becoming too numerous too quickly as design size grows. This could change in the near future if techniques like semidefinite programming, now being brought to bear on optimizing design for regression models (Papp 2012), can be adapted for the design problems considered here. There has certainly been significant, recent progress from reducing problem size through mathematical argument, thus making computational resolution feasible (e.g., Section 3.4.4). However, there will always be a need for, and an appreciation of, complete solutions to design problems like those in Kiefer (1975), Cheng (1978), Bagchi and Bagchi (2001), and others mentioned throughout this chapter. It would be a great loss should search algorithms become replacements for known theory and results for optimal designs. At the least, known optimal block designs should be stored for ready access, not searched at every call with no guarantee of absolute optimality.

References

- Adhikary, B. (1965), On the properties and construction of balanced block designs with variable replications, *Calcutta Stat. Assoc. Bull.*, 14, 36–64.
- Ash, A. (1981), Generalized Youden designs: construction and tables, *J. Stat. Plan. Infer.*, 5, 1–25.
- Bagchi, B. and Bagchi, S. (2001), Optimality of partial geometric designs, *Ann. Stat.*, 29, 577–594.
- Bagchi, S. and Bose, M. (2008), Optimal main effect plans in nested row-column set-up of small size, *Stat. Prob. Lett.*, 78, 2720–2724.
- Bagchi, S. and Cheng, C.-S. (1993), Some optimal designs of block size two, *J. Stat. Plan. Infer.*, 37, 245–253.
- Bagchi, S., Mukhopadhyay, A. C., and Sinha, B. K. (1990), A search for optimal nested row-column designs, *Sankhyā Ser. B*, 52, 93–104.
- Bagchi, S. and Shah, K. R. (1989), On the optimality of a class of row-column designs, *J. Stat. Plan. Infer.*, 23, 397–402.
- Bagchi, S. and van Berkum, E. E. M. (1991), On the optimality of a new class of adjusted orthogonal designs, *J. Stat. Plann. Infer.*, 28, 61–65.
- Bailey, R. A. (1992), Efficient semi-Latin squares, *Stat. Sin.*, 2, 413–437.
- Bailey, R. A. (2005), Six families of efficient resolvable designs in three replicates, *Metrika*, 62, 161–173.
- Bailey, R. A. (2007), Designs for two-colour microarray experiments, *J. R. Stat. Soc. Ser. C*, 56, 365–394.
- Bailey, R. A. (2008), *Design of Comparative Experiments*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge, U.K.: Cambridge University Press.
- Bailey, R. A. and Cameron, P. J. (2009), Combinatorics of optimal designs, in *Surveys in Combinatorics 2009*, Cambridge, U.K.: Cambridge University Press, Vol. 365 of London Mathematics Social Lecture Note Series, pp. 19–73.
- Bailey, R. A. and Cameron, P. J. (2013), Using graphs to find the best block designs, in *Topics in Structural Graph Theory*, Cambridge, U.K.: Cambridge University Press, Vol. 365, pp. 282–317.
- Bailey, R. A., Cameron, P. J., Dobcsányi, P., Morgan, J. P., and Soicher, L. H. (2006), Designs on the web, *Discrete Math.*, 306, 3014–3027.

- Bailey, R. A., Monod, H., and Morgan, J. P. (1995), Construction and optimality of affine-resolvable designs, *Biometrika*, 82, 187–200.
- Bailey, R. A. and Royle, G. (1997), Optimal semi-Latin squares with side six and block size two, *Proc. R. Soc. Lond. Ser. A*, 453, 1903–1914.
- Bailey, R. A. and Speed, T. P. (1986), Rectangular lattice designs: Efficiency factors and analysis, *Ann. Stat.*, 14, 874–895.
- Balasubramanian, K. and Dey, A. (1996), D -optimal designs with minimal and nearly minimal number of units, *J. Stat. Plann. Infer.*, 52, 255–262.
- Bapat, R. B. and Dey, A. (1991), Optimal block designs with minimal number of observations, *Stat. Probab. Lett.*, 11, 399–402.
- Bechhofer, R. E. and Tamhane, A. C. (1981), Incomplete block designs for comparing test treatments with a control: General theory, *Technometrics*, 23, 45–57.
- Bhatia, R. (1997), *Matrix Analysis*, Vol. 169 of *Graduate Texts in Mathematics*, New York: Springer-Verlag.
- Billington, E. J. and Robinson, P. J. (1983), A list of balanced ternary designs with $R \leq 15$, and some necessary existence conditions, *Ars Combin.*, 16, 235–258.
- Bose, R. C. (1939), On the construction of balanced incomplete block designs, *Ann. Eugenics*, 9, 353–399.
- Bose, R. C. and Nair, K. R. (1939), Partially balanced incomplete block designs, *Sankhya*, 4, 337–372.
- Bose, R. C. and Shimamoto, T. (1952), Classification and analysis of partially balanced incomplete block designs with two associate classes, *J. Am. Stat. Assoc.*, 47, 151–184.
- Cakiroglu, S. A. (2013), *A Combinatorial Approach to Optimal Designs*, PhD dissertation, London, U.K.: Queen Mary University.
- Caliński, T. and Kageyama, S. (1996), Block designs: their combinatorial and statistical properties, in *Design and Analysis of Experiments*, Amsterdam, the Netherlands: North-Holland, Vol. 13 of *Handbook of Statistics*, pp. 809–873.
- Chai, F.-S. and Cheng, C.-S. (2011), Some optimal row-column designs, *J. Stat. Theory Pract.*, 5, 59–67.
- Cheng, C.-S. (1978), Optimality of certain asymmetrical experimental designs, *Ann. Stat.*, 6, 1239–1261.
- Cheng, C.-S. (1980), On the E -optimality of some block designs, *J. R. Stat. Soc. Ser. B*, 42, 199–204.
- Cheng, C.-S. (1981), Maximizing the total number of spanning trees in a graph: two related problems in graph theory and optimum design theory, *J. Combin. Theory Ser. B*, 31, 240–248.
- Cheng, C.-S. (1992), On the optimality of (M.S)-optimal designs in large systems, *Sankhyā Ser. A*, 54, 117–125, *Combinatorial Mathematics and Applications* (Calcutta, 1988).
- Cheng, C.-S. (2012), A note on the E -optimality of regular line graph designs, *J. Stat. Theory Pract.*, 6, 162–168.
- Cheng, C.-S. and Bailey, R. A. (1991), Optimality of some two-associate-class partially balanced incomplete-block designs, *Ann. Stat.*, 19, 1667–1671.
- Cheng, C.-S., Masaro, J. C., and Wong, C. S. (1985), Do nearly balanced multigraphs have more spanning trees? *J. Graph Theory*, 9, 342–345.
- Cheng, C.-S. and Wu, C.-F. (1981), Nearly balanced incomplete block designs, *Biometrika*, 68, 493–500.
- Chowla, S. and Ryser, H. J. (1950), Combinatorial problems, *Can. J. Math.*, 2, 93–99.
- Clatworthy, W. H. (1973), *Tables of Two-Associate-Class Partially Balanced Designs*, U. S. Department of Commerce, Washington, DC: National Bureau of Standards, with contributions by Joseph M. Cameron and Janace A. Speckman, An enlarged and revised set of tables based on *Tables of Partially Balanced Designs with Two Associate Classes* by R. C. Bose, Clatworthy and S. S. Shrikhande, Report No. NBS-AMS-63.
- Colbourn, C. J. and Dinitz, J. H. (eds.) (2007), *Handbook of Combinatorial Designs, Discrete Mathematics and Its Applications*, 2nd ed., Boca Raton, FL: Chapman & Hall/CRC Press.
- Constantine, G. M. (1981), Some E -optimal block designs, *Ann. Stat.*, 9, 886–892.
- Constantine, G. M. (1982), On the E -optimality of PBIB designs with a small number of blocks, *Ann. Stat.*, 10, 1027–1031.

- Das, A. and Dey, A. (1992), Universal optimality and nonoptimality of some row-column designs, *J. Stat. Plan. Infer.*, 31, 263–271.
- Dean, A. and Voss, D. (1999), *Design and Analysis of Experiments*, Springer Texts in Statistics, New York: Springer-Verlag.
- Dey, A. (2010), *Incomplete Block Designs*, Hackensack, NJ: World Scientific Publishing Co. Pte. Ltd.
- Dey, A., Shah, K. R., and Das, A. (1995), Optimal block designs with minimal and nearly minimal number of units, *Stat. Sin.*, 5, 547–558.
- Fisher, R. A. (1935), *The Design of Experiments*, Springer Texts in Statistics, Edinburgh, U.K.: Oliver and Boyd.
- Fisher, R. A. (1940), An examination of the different possible solutions of a problem in incomplete blocks, *Ann. Eugenics*, 10, 52–75.
- Ghosh, D. K., Thannippara, A., and Joseph, O. C. (2008), A new class of E -optimal regular graph designs, *Util. Math.*, 76, 249–263.
- Godolphin, J. D. and Warren, H. R. (2011), Improved conditions for the robustness of binary block designs against the loss of whole blocks, *J. Stat. Plan. Infer.*, 141, 3498–3505.
- Gupta, V. K., Ramana, D. V. V., and Parsad, R. (1999), Weighted A -efficiency of block designs for making treatment-control and treatment-treatment comparisons, *J. Stat. Plan. Infer.*, 77, 301–319.
- Hanani, H. (1975), Balanced incomplete block designs and related designs, *Discrete Math.*, 11, 255–369.
- Hedayat, A. S. and Majumdar, D. (1984), A -optimal incomplete block designs for control-test treatment comparisons, *Technometrics*, 26, 363–370.
- Hinkelmann, K. and Kempthorne, O. (2005), *Design and Analysis of Experiments. Vol. 2: Advanced Experimental Design*, Wiley Series in Probability and Statistics, Hoboken, NJ: Wiley-Interscience [John Wiley & Sons].
- Hinkelmann, K. and Kempthorne, O. (2008), *Design and Analysis of Experiments. Vol. 1: Introduction to Experimental Design*, Wiley Series in Probability and Statistics, Hoboken, NJ: Wiley-Interscience [John Wiley & Sons], 2nd ed.
- Jacroux, M. (1980), On the E -optimality of regular graph designs, *J. R. Stat. Soc. Ser. B*, 42, 205–209.
- Jacroux, M. (1985), Some sufficient conditions for the type 1 optimality of block designs, *J. Stat. Plan. Infer.*, 11, 385–398.
- Jacroux, M. (1991), Some new methods for establishing the optimality of block designs having unequally replicated treatments, *Statistics*, 22, 33–48.
- Jacroux, M. and Whittinghill, III, D. C. (1988), On the E - and MV -optimality of block designs having $k \geq v$, *Ann. Inst. Stat. Math.*, 40, 407–418.
- Jin, B. and Morgan, J. P. (2008), Optimal saturated block designs when observations are correlated, *J. Stat. Plan. Infer.*, 138, 3299–3308.
- John, J. A. and Mitchell, T. J. (1977), Optimal incomplete block designs, *J. R. Stat. Soc. Ser. B*, 39, 39–43.
- Kiefer, J. (1975), Construction and optimality of generalized Youden designs, in *A Survey of Statistical Design and Linear Models (Proceedings of the International Symposium, Colorado State University, Ft. Collins, CO 1973)*, Amsterdam, the Netherlands: North-Holland, pp. 333–353.
- Krafft, O. and Schaefer, M. (1997), A -optimal connected block designs with nearly minimal number of observations, *J. Stat. Plan. Infer.*, 65, 375–386.
- Kunert, J. (1985), Optimal repeated measurements designs for correlated observations and analysis by weighted least squares, *Biometrika*, 72, 375–389.
- Liski, E. P., Mandal, N. K., Shah, K. R., and Sinha, B. K. (2002), *Topics in Optimal Design*, Vol. 163 of *Lecture Notes in Statistics*, New York: Springer-Verlag.
- Majumdar, D. (1996), Optimal and efficient treatment-control designs, *Design and Analysis of Experiments*, Vol. 13 of *Handbook of Statistics*, Ghosh, S. and Rao, C. R. (Eds.) Amsterdam, the Netherlands: North-Holland, pp. 1007–1053.
- Majumdar, D. and Notz, W. I. (1983), Optimal incomplete block designs for a comparing treatments with a control, *Ann. Stat.*, 11, 258–266.

- Mandal, N. K., Shah, K. R., and Sinha, B. K. (1990/91), Uncertain resources and optimal designs: problems and perspectives, *Calcutta Stat. Assoc. Bull.*, 40, 267–282.
- Martin, R. J. and Eccleston, B. K. (1991), Optimal incomplete block designs for general dependence structures. *J. Stat. Plan. Infer.*, 28, 67–81.
- Morgan, J. P. (1996), Nested designs, in *Design and Analysis of Experiments*, Vol. 13 of *Handbook of Statistics*, Ghosh, S. and Rao, C. R. (Eds.) Amsterdam, the Netherlands: North-Holland, pp. 939–976.
- Morgan, J. P. (1997), On pairs of Youden designs, *J. Stat. Plan. Infer.*, 60, 367–387.
- Morgan, J. P. (2007a), Latin squares and related designs, in *Encyclopedia of Statistics in Quality and Reliability*, Ruggeiri, F., Kenett, R., and Faltin, F. (Eds.) New York: Wiley, pp. 919–925.
- Morgan, J. P. (2007b), Optimal incomplete block designs, *J. Am. Stat. Assoc.*, 102, 655–663.
- Morgan, J. P. (2008), On MV-optimality of resolvable designs from 2-level orthogonal arrays, *Stat. Appl.*, 6, 87–95.
- Morgan, J. P. (2009), Variance balance and admissibility, *J. Stat. Appl.*, 4, 409–419.
- Morgan, J. P. (2010), Optimal resolvable designs with minimum PV aberration, *Stat. Sin.*, 20, 715–732.
- Morgan, J. P. and Bailey, R. A. (2000), Optimal design with many blocking factors, *Ann. Stat.*, 28, 553–577.
- Morgan, J. P. and Jin, B. (2007), Optimal experimentation in two blocks, *J. Stat. Theory Pract.*, 1, 357–375.
- Morgan, J. P. and Parvu, V. (2007), Optimal row-column design for three treatments, *J. Stat. Plan. Infer.*, 137, 1474–1487.
- Morgan, J. P. and Parvu, V. (2008), Most robust BIBDs, *Stat. Sin.*, 18, 689–707.
- Morgan, J. P. and Reck, B. (2007a), E-optimal design in irregular BIBD settings, *J. Stat. Plan. Infer.*, 137, 1658–1668.
- Morgan, J. P. and Reck, B. H. (2007b), Resolvable designs with large blocks, *Ann. Stat.*, 35, 747–771.
- Morgan, J. P. and Srivastav, S. K. (2000), On the type-1 optimality of nearly balanced incomplete block designs with small concurrence range, *Stat. Sin.*, 10, 1091–1116.
- Morgan, J. P. and Srivastav, S. K. (2002), The completely symmetric designs with blocksize three, *J. Stat. Plann. Infer.*, 106, 21–30, Experimental Design and Related Combinatorics.
- Morgan, J. P. and Stallings, J. W. (2014), The A-criterion of experimental design, *J. Stat. Theory Pract.*, 8, 418–422.
- Morgan, J. P. and Uddin, N. (1993), Optimality and construction of nested row and column designs, *J. Stat. Plan. Infer.*, 37, 81–93.
- Morgan, J. P. and Uddin, N. (1995), Optimal, non-binary, variance balanced designs, *Stat. Sin.*, 5, 535–546.
- Morgan, J. P. and Wang, X. (2010), Weighted optimality in designed experimentation, *J. Am. Stat. Assoc.*, 105, 1566–1580, supplementary materials available online.
- Morgan, J. P. and Wang, X. (2011), E-optimality in treatment versus control experiments, *J. Stat. Theory Pract.*, 5, 99–107.
- Papp, D. (2012), Optimal designs for rational function regression, *J. Am. Stat. Assoc.*, 107, 400–411.
- Parvu, V. and Morgan, J. P. (2008), E-optimal designs for three treatments, *J. Stat. Plan. Infer.*, 138, 642–653.
- Paterson, L. J. (1988), Some recent work on making incomplete-block designs available as a tool for science, *Int. Stat. Rev.*, 56, 129–138.
- Patterson, H. D. and Williams, E. R. (1976a), A new class of resolvable incomplete block designs, *Biometrika*, 63, 83–92.
- Patterson, H. D. and Williams, E. R. (1976b), Some theoretical results on general block designs, in *Proceedings of the Fifth British Combinatorial Conference (University of Aberdeen, Aberdeen, U.K., 1975)*, Winnipeg, Manitoba, Canada: Utilitas Mathematics, pp. 489–496. Congressus Numerantium, No. XV.
- Preece, D. A. (1967), Nested balanced incomplete block designs, *Biometrika*, 54, 479–486.
- Preece, D. A. (1982), Balance and designs: Another terminological tangle, *Util. Math.*, 21, 85–186.

- Qu, X. (2010), Optimal row-column designs in high-throughput screening experiments, *Technometrics*, 52, 409–420.
- Reck, B. and Morgan, J. P. (2005), Optimal design in irregular BIBD settings, *J. Stat. Plan. Infer.*, 129, 59–84.
- Saharay, R. (1996), A class of optimal row-column designs with some empty cells, *Stat. Sin.*, 6, 989–996.
- Seber, G. A. F. (2008), *A Matrix Handbook for Statisticians*, Wiley Series in Probability and Statistics, Hoboken, NJ: Wiley-Interscience [John Wiley & Sons].
- Shah, K. R. and Eccleston, J. A. (1986), On some aspects of row-column designs, *J. Stat. Plan. Infer.*, 15, 87–95.
- Shah, K. R. and Sinha, B. K. (1989), *Theory of Optimal Designs*, Vol. 54 of *Lecture Notes in Statistics*, New York: Springer-Verlag.
- Soicher, L. (2009), The DESIGN package for GAP, Version 1.4, http://designtheory.org/software/gap_design/. Accessed February 7, 2015.
- Soicher, L. H. (2012), Uniform semi-Latin squares and their Schur-optimality, *J. Combin. Des.*, 20, 265–277.
- Soicher, L. H. (2013), Optimal and efficient semi-Latin squares, *J. Stat. Plan. Infer.*, 143, 573–582.
- Srivastav, S. K. and Morgan, J. P. (1998), Optimality of designs with generalized group divisible structure, *J. Stat. Plan. Infer.*, 71, 313–330.
- Stallings, J. W. and Morgan, J. P. (2014), General weighted optimality for designed experiments. Preprint.
- Stufken, J. (1987), A-optimal block designs for comparing test treatments with a control, *Ann. Stat.*, 15, 1629–1638.
- Takeuchi, K. (1961), On the optimality of certain type of PBIB designs, *Rep. Stat. Appl. Res. Un. Jpn. Sci. Eng.*, 8, 140–145.
- Ting, C. P. and Notz, W. (1988), A-optimal complete block designs for treatments-control comparisons, in *Optimal Design and Analysis of Experiments*, Dodge, Y., Fedorov, V. V., and Wynn, H. P. (Eds.) Amsterdam, the Netherlands: North-Holland, pp. 29–37.
- van Lint, J. H. and Wilson, R. M. (2001), *A Course in Combinatorics*, 2nd ed., Cambridge, U.K.: Cambridge University Press.
- Wang, X. (2009), Weighted optimality of block designs, PhD dissertation, Blacksburg, VA: Virginia Polytechnic Institute and State University.
- Wang, X. and Morgan, J. P. (2010), Blocking, efficiency, and weighted optimality, *Biometrika*, 98, 967–978, supplementary materials available online.
- Wang, X. and Morgan, J. P. (2012), RCBDs with a control, *J. Stat. Plan. Infer.*, 142, 1225–1233.
- Yates, F. (1936a), Incomplete randomized blocks, *Ann. Eugen.*, 7, 121–140.
- Yates, F. (1936b), A new method of arranging variety trials involving a large number of varieties, *J. Agric. Sci.*, 26, 424–455.
- Yeh, C.-M. (1988), A class of universally optimal binary block designs, *J. Stat. Plan. Infer.*, 18, 355–361.

4

Crossover Designs

Mausumi Bose and Alope Dey

CONTENTS

4.1	Introduction.....	160
4.1.1	Prologue.....	160
4.1.2	Early History.....	161
4.2	Model for Studying Crossover Designs.....	162
4.2.1	Traditional Model.....	163
4.2.2	Information Matrices.....	163
4.3	Some Families of Crossover Designs.....	166
4.3.1	Designs with $p \geq t$	166
4.3.2	Designs with $p \leq t$	169
4.3.3	Two-Period Designs.....	171
4.3.4	Two-Treatment Designs.....	173
4.4	Constructions of Some Families of Designs.....	173
4.4.1	Balanced Uniform Designs.....	173
4.4.2	Strongly Balanced Uniform Designs.....	175
4.4.3	Patterson Designs.....	176
4.5	Optimality under the Traditional Model.....	177
4.5.1	Balanced Uniform Designs.....	177
4.5.2	Stufken and Patterson Designs.....	178
4.5.3	Strongly Balanced Designs.....	178
4.5.4	Two-Period Designs.....	179
4.5.5	Two-Treatment Designs.....	180
4.6	Some Other Models and Optimal Designs.....	180
4.6.1	Circular Model.....	180
4.6.2	Model with Self- and Mixed Carry-Overs.....	181
4.6.3	Model with Direct-versus-Carry-Over Interactions.....	182
4.6.4	Model with Carry-Over Proportional to Direct Effect.....	183
4.6.5	Mixed Effects Models and Models with Correlated Errors.....	184
4.7	Other Advances.....	187
4.7.1	Crossover Trials for Comparing Treatments versus Control.....	187
4.7.2	Optimal Designs When Some Subjects Drop Out.....	189
4.7.3	Optimal Designs via Approximate Theory.....	191
	Acknowledgment.....	191
	References.....	191

4.1 Introduction

4.1.1 Prologue

In a crossover trial, every *experimental subject* is exposed to a sequence of treatments over time, one treatment being applied to it at each time point. These subjects could be humans, animals, machines, plots of land, etc. The different time points at which the subjects are used are referred to as *periods*. Consider a crossover trial with n experimental subjects, each subject being observed for p periods, resulting in a total of np experimental *units*. We shall assume at the design stage that each such experimental unit yields a single response (which could even possibly be an average or sum over multiple observations). If such a trial aims at drawing inference on a set of t treatments, then any allocation of these t treatments to the np experimental units is called a *crossover design*. In the literature, such designs have also been referred to as *changeover* (see also Chapter 1) or *repeated measurements* designs.

Crossover designs have been extensively applied in a variety of areas including pharmaceutical studies and clinical trials, biological assays, weather modification experiments, sensory evaluation of food products, psychology, bioequivalence studies, and consumer trials. Throughout this chapter, a crossover design with p periods, n subjects, and t treatments will be displayed as a $p \times n$ array, with rows of the array representing the periods, columns representing the subjects, and the numerals $1, 2, \dots, t$ denoting the treatments. The following are two examples of crossover designs:

Example 4.1

(i) $d_1 : t = 2 = p, n = 4.$

1	2	1	2
2	1	2	1

(ii) $d_2 : t = 3 = p, n = 6.$

1	2	3	1	2	3
2	3	1	3	1	2
3	1	2	2	3	1

The design d_1 shown in Example 4.1 is a two-period two-treatment design, also called an AB/BA design, where AB stands for the treatment sequence in which treatment 1 is followed by treatment 2, BA being defined analogously. Such a design is often used in the context of clinical trials where, for example, treatment 1 could be the drug under study and 2 could be a placebo or another drug. An extensive discussion on the analysis of data from such designs is available in Jones and Kenward (2003). An application of this design in the context of pharmacokinetic studies can be found in Jones et al. (1999).

An advantage of a crossover design is that, for the same number of observations, this design requires fewer experimental subjects compared to a traditional design where each subject gives a single observation. This is useful in situations where the subjects are scarce or expensive. However, the very feature of these designs, namely, the repeated use of a subject, also brings in associated problems. For instance, a crossover design results in a longer duration of the experiment compared to a traditional design, and designs with a large number of periods may not be particularly attractive in some areas of application.

More importantly, there is a possibility that the effect of a treatment may continue to linger on in a subject beyond the period in which it is applied. For instance, in design d_2 of Example 4.1, in the second period, the first subject may retain some of the effect of treatment 1 applied to it in the first period, and so the response from the first subject in the second period is affected not only by the *direct* effect of treatment 2 but also possibly by the residual, or *carry-over*, effect of treatment 1. Similarly, the response in the third period of the first subject is influenced by the direct effect of treatment 3, by the carry-over effect of treatment 2, and also possibly by the effect of treatment 1 carrying over across two periods. In particular, an effect carrying over to the immediate next period is referred to as the *first-order* carry-over effect, and extending this idea, there may be *second-order* or even higher-order carry-over effects in subsequent periods.

Thus, there are two types of treatment effects associated with crossover trials, the *direct effects* and the *carry-over effects*, the former effects being usually of primary interest. The presence of carry-over effects complicates the design and analysis of crossover trials. One option of avoiding these is to allow a larger time gap between two successive applications of treatments, with the expectation that the carry-over effect, if any, would washout during this gap. Though this strategy may help in avoiding the carry-over effect, insertion of such gaps, usually called the *rest* (or *washout*) periods, increases the total duration of the trial. Moreover, it can be difficult to determine how long a rest period should be in order to washout the carry-over effect completely. Another reason why such washout periods may make a trial infeasible is apparent in the context of clinical trials where the subjects are patients. In such trials, adopting a washout period is equivalent to denying a patient any treatment during this long gap, and this may be unacceptable on medical or ethical grounds. So, instead of trying to eliminate the carry-over effects by inserting rest periods, if one accepts their possible presence, then the challenge is to come up with an effective design of the trial and its corresponding analysis so that the typical contrasts of interest, namely, the direct effect contrasts, can be estimated efficiently after properly adjusting for these carry-over effects.

Much of the literature on crossover designs deals with solutions to this problem under different assumptions on the nature of the carry-over effects. In the following sections, we provide a survey of some of the important results on the construction and analysis of efficient/optimal crossover designs. Throughout this chapter, we assume that the responses from a crossover trial are quantitative. However, there are situations in practice when such responses may be binary or categorical in nature. We do not elaborate on the analysis of crossover trials with binary or categorical responses and refer the reader to Chapter 6 of Jones and Kenward (2003) for details and additional references.

4.1.2 Early History

Crossover trials have a long history and, apparently, these were first applied in agriculture in 1853. We refer the reader to Jones and Kenward (2003, Section 1.4) for details of a crossover experiment in agriculture conducted by John Bennet Lawes of Rothamsted, England, in 1853. An early use of crossover trials in human nutrition was made by Simpson (1938). These trials were related to experiments on diets for children. In one such trial, four different diets were compared using 24 pairs of children, one male and one female in each pair. Each pair received one of all possible 24 permutations of four diets over four periods in such a way that each treatment was given equally often in each period. Simpson (1938) was aware of carry-over effects and suggested the insertion of a rest (or washout) period between the experimental periods to remove the carry-over effects. He also stated

that the insertion of a washout period to eliminate carry-over effects may not always be the best strategy in all situations, especially when it may be necessary to estimate the carry-over effects themselves, and suggested the use of suitable designs that allow the estimation of both direct and carry-over effects.

Cochran (1939) observed the existence of carry-over effects in long-term agricultural experiments and was one of the first to separate out the direct and carry-over effects while considering an appropriate design for experimentation. In a classic and widely quoted paper, Cochran et al. (1941) considered a crossover trial on Holstein cows for comparing three treatments in three periods. The crossover design used was obtained by using orthogonal Latin squares (see Chapter 1), as in the design d_2 of Example 4.1. Cochran et al. appear to be the first to formally describe the least squares method of estimation of direct and carry-over contrasts. Another early example of an experiment indicating the presence of carry-over effects was quoted by Williams (1949). In this experiment, samples of pulp suspensions were beaten in a Lampen mill, to determine the effect of concentration on the properties of resulting sheets. Observations of the condition of the mill after each beating indicated that certain pulp concentrations had an effect on the mill, which might affect the next beating, indicating the presence of carry-over effect. A design balanced for carry-over effects was therefore used.

An early use of crossover designs was made in biological assays by Fieller (1940). He used a two-period design involving two treatments for comparing the effects of different doses of insulin on rabbits. Finney (1956) also described the design and analysis of several crossover designs for use in biological assay. In subsequent years, the use of these designs in many diverse areas, particularly in clinical trials and pharmaceutical studies, has been extensive. Real-life examples and discussion on various aspects of crossover designs can be found in the books on this topic by Pocock (1983), Ratkowsky et al. (1992), Jones and Kenward (2003), Senn (2003), and Bose and Dey (2009). Over the years, several review papers have been published on these designs, including those by Hedayat and Afsarinejad (1975), Matthews (1988), Stufken (1996), Kenward and Jones (1998), Senn (2000), and Bate and Jones (2008). An early technical report due to Patterson and Lucas (1962) provides tables of useful crossover designs along with detailed steps of their analysis.

4.2 Model for Studying Crossover Designs

Consider a crossover trial in which t treatments are to be compared using n experimental subjects over p time periods. As mentioned earlier, any allocation of the t treatments to the np experimental units is called a crossover design. Let $\Omega_{t,n,p}$ be the collection of all such crossover designs.

For the analysis of data arising from crossover designs, various models have been studied in the literature. We first describe a commonly used model in the following subsection. This model is hereafter called the *traditional model*.

Even though the traditional model has been widely studied, it has also been criticized for being unsuitable for some experimental situations. In order to suit different situations, the traditional model has been variously modified, for example, by making certain assumptions on the form of the carry-over effect, by assuming a certain structure for the correlation of the error terms, or by assuming random subject effects. Some of these modifications will be described in later sections.

4.2.1 Traditional Model

The traditional model described in the following text is an additive linear model, where the expected response from a subject at any given period is the sum of the corresponding subject and period effects, together with the direct effect of the treatment applied at that period and the carry-over effect of the treatment applied in the previous period (if any). For the data from a design $d \in \Omega_{t,n,p}$, the traditional model may be expressed as

$$y_{ij} = \mu + \alpha_i + \beta_j + \tau_{d(i,j)} + \rho_{d(i-1,j)} + \epsilon_{ij},$$

$$1 \leq i \leq p, \quad 1 \leq j \leq n, \quad (4.1)$$

where y_{ij} is the observable random variable corresponding to the observation from the j th subject in the i th period, $d(i, j)$ denotes the treatment allocated to the j th subject in the i th period according to the design d , and μ , α_i , β_j , $\tau_{d(i,j)}$, and $\rho_{d(i-1,j)}$ are, respectively, a general mean, the i th period effect, the j th subject effect, the direct effect of the treatment $d(i, j)$, and the first-order carry-over effect of the treatment $d(i-1, j)$, $1 \leq i \leq p, 1 \leq j \leq n$; the ϵ_{ij} 's are the error components, assumed to be uncorrelated random variables with zero means and constant variance σ^2 . We define $\rho_{d(0,j)} = 0, 1 \leq j \leq n$, to reflect the fact that there are no carry-over effects in the first period. All the parameters in (4.1) are considered as *fixed*, that is, nonrandom. In what follows, the same notation y_{ij} is used for the observation as well as the random variable corresponding to the observation.

4.2.2 Information Matrices

We first express (4.1) in a form that is more convenient to study. Toward this end, let us write the observations from a design d as an ordered vector, where the first p entries are the p observations on subject 1, the next p are the observations on subject 2, ..., and so on. Thus, for any design $d \in \Omega_{t,n,p}$, $\mathbf{y}_d = (y_{11}, \dots, y_{p1}, y_{12}, \dots, y_{p2}, \dots, y_{1n}, \dots, y_{pn})'$ is the $np \times 1$ vector of observations arising out of d with y_{ij} as in (4.1). Here and hereafter, primes denote transposition. Let $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_p)'$, $\boldsymbol{\beta} = (\beta_1, \dots, \beta_n)'$ be, respectively, the $p \times 1$ vector of period effects and the $n \times 1$ vector of subject effects, where α_i and β_j are as in (4.1). Since $d(i, j) \in \{1, 2, \dots, t\}$, for simplicity in notation, we denote the direct (respectively, the first-order carry-over) effect of treatment s by τ_s (respectively, ρ_s), $1 \leq s \leq t$, and write $\boldsymbol{\tau} = (\tau_1, \dots, \tau_t)'$, $\boldsymbol{\rho} = (\rho_1, \dots, \rho_t)'$, $\boldsymbol{\epsilon} = (\epsilon_{11}, \dots, \epsilon_{pn})'$ to denote the $t \times 1$ vector of direct effects, the $t \times 1$ vector of carry-over effects, and the $np \times 1$ vector of error terms, respectively, where ϵ_{ij} is as in (4.1). Also, let $\boldsymbol{\theta} = (\mu, \boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\tau}', \boldsymbol{\rho}')'$ with μ as in (4.1).

Let $\mathbf{1}_a$ and $\mathbf{0}_a$ denote the $a \times 1$ vectors of all ones and all zeros, respectively, and \mathbf{I}_a denote the identity matrix of order a , where a is a positive integer. For positive integers a and b , let \mathbf{O}_{ab} denote the $a \times b$ null matrix and \mathbf{J}_{ab} , the $a \times b$ matrix of all ones; \mathbf{J}_{aa} and \mathbf{O}_{aa} will simply be written as \mathbf{J}_a and \mathbf{O}_a , respectively. A square matrix A of order n is called completely symmetric if $A = a\mathbf{I}_n + b\mathbf{J}_n$ for some scalars a, b .

For a design $d \in \Omega_{t,n,p}$, let \mathbf{T}_{dj} be a $p \times t$ matrix with its (i, s) th entry equal to 1 if subject j receives treatment s in the i th period, and zero otherwise. Similarly, let \mathbf{F}_{dj} be a $p \times t$ matrix with its (i, s) th entry equal to 1 if subject j receives treatment s in the $(i-1)$ th period, and zero otherwise. Since $\rho_{d(0,j)} = 0$ for $1 \leq j \leq n$, the first row of \mathbf{F}_{dj} is zero, and for $2 \leq i \leq p, 1 \leq j \leq n$, the i th row of \mathbf{F}_{dj} is the $(i-1)$ th row of \mathbf{T}_{dj} , that is,

$$F_{dj} = \begin{pmatrix} \mathbf{0}'_{p-1} & 0 \\ I_{p-1} & \mathbf{0}_{p-1} \end{pmatrix} T_{dj}, \quad 1 \leq j \leq n. \quad (4.2)$$

Define $T_d = (T'_{d1}, \dots, T'_{dn})'$, $F_d = (F'_{d1}, \dots, F'_{dn})'$, and let $E(\cdot)$ and $\text{Var}(\cdot)$ denote the expectation and variance operators, respectively.

With the aforementioned notation, model (4.1) can equivalently be written as

$$y_d = X_d \theta + \epsilon, \quad E(\epsilon) = \mathbf{0}_{np}, \quad \text{Var}(\epsilon) = \sigma^2 I_{np}, \quad (4.3)$$

where the design matrix X_d may be written in the following partitioned form:

$$X_d = [\mathbf{1}_{np} \quad P \quad U \quad T_d \quad F_d] = [\mathbf{1}_{np} \quad X_1 \quad X_2], \quad \text{say}, \quad (4.4)$$

P, U, T_d , and F_d being the parts of X_d corresponding to the period, subject, direct, and carry-over effects, respectively, under the design d , and $X_1 = [P \quad U]$, $X_2 = [T_d \quad F_d]$. Furthermore, with the ordering of the observations as in y_d , it is clear that

$$P = \mathbf{1}_n \otimes I_p \quad \text{and} \quad U = I_n \otimes \mathbf{1}_p,$$

where \otimes denotes the Kronecker (tensor) product operator. Hereafter, we will write (4.3) to denote the traditional model.

For a matrix A , let A^- denote an arbitrary generalized inverse (g-inverse) of A (see Chapter 2 for a definition of a g-inverse). Define $P^\perp(A) = I - A(A'A)^-A'$, where I stands for the identity matrix of appropriate order. Then, it can be shown (see e.g., Bose and Dey 2009, Section 1.3) that under model (4.3), after eliminating the nuisance parameters α and β , the information matrix for estimating τ and ρ jointly is of the form

$$\begin{aligned} C_d(\tau, \rho) &= X_2' P^\perp(X_1) X_2 \\ &= X_2' X_2 - X_2' X_1 (X_1' X_1)^- X_1' X_2 \\ &= \begin{bmatrix} T_d' A T_d & T_d' A F_d \\ F_d' A T_d & F_d' A F_d \end{bmatrix}, \end{aligned}$$

with

$$A = (I_n - n^{-1} J_n) \otimes (I_p - p^{-1} J_p). \quad (4.5)$$

We may rewrite $C_d(\tau, \rho)$ as

$$C_d(\tau, \rho) = \begin{bmatrix} C_{d11} & C_{d12} \\ C_{d21} & C_{d22} \end{bmatrix}, \quad (4.6)$$

where

$$\begin{aligned}
 C_{d11} &= \mathbf{R}_d - n^{-1}\mathbf{M}_d\mathbf{M}'_d - p^{-1}\mathbf{N}_d\mathbf{N}'_d + (np)^{-1}\mathbf{r}_d\mathbf{r}'_d, \\
 C_{d12} &= \mathbf{Z}_d - n^{-1}\mathbf{M}_d\bar{\mathbf{M}}'_d - p^{-1}\mathbf{N}_d\bar{\mathbf{N}}'_d + (np)^{-1}\mathbf{r}_d\bar{\mathbf{r}}'_d = \mathbf{C}'_{d21}, \\
 C_{d22} &= \bar{\mathbf{R}}_d - n^{-1}\bar{\mathbf{M}}_d\bar{\mathbf{M}}'_d - p^{-1}\bar{\mathbf{N}}_d\bar{\mathbf{N}}'_d + (np)^{-1}\bar{\mathbf{r}}_d\bar{\mathbf{r}}'_d.
 \end{aligned} \tag{4.7}$$

Here, \mathbf{r}_d (respectively, $\bar{\mathbf{r}}_d$) is the $t \times 1$ replication vector for direct (respectively, carry-over) effects; \mathbf{R}_d (respectively, $\bar{\mathbf{R}}_d$) is the $t \times t$ diagonal matrix with diagonal elements given by the elements of \mathbf{r}_d (respectively, $\bar{\mathbf{r}}_d$); \mathbf{M}_d (respectively, $\bar{\mathbf{M}}_d$) is the $t \times p$ direct (respectively, carry-over) effect versus period incidence matrix; \mathbf{N}_d (respectively, $\bar{\mathbf{N}}_d$) is the $t \times n$ direct (respectively, carry-over) effect versus subject incidence matrix, and \mathbf{Z}_d is the $t \times t$ direct effect versus carry-over effect incidence matrix. It may be verified that these are related to the matrices \mathbf{T}_d and \mathbf{F}_d defined earlier as $\mathbf{r}_d = \mathbf{T}'_d\mathbf{1}_t$, $\bar{\mathbf{r}}_d = \mathbf{F}'_d\mathbf{1}_t$, $\mathbf{R}_d = \mathbf{T}'_d\mathbf{T}_d$, $\bar{\mathbf{R}}_d = \mathbf{F}'_d\mathbf{F}_d$, $\mathbf{Z}_d = \mathbf{T}'_d\mathbf{F}_d$.

Now, let the information matrices of the direct (respectively, carry-over) effects, eliminating the carry-over (respectively, direct) effects, be denoted by \mathbf{C}_d (respectively, $\bar{\mathbf{C}}_d$). Then, it follows from (4.7) that

$$\begin{aligned}
 \mathbf{C}_d &= \mathbf{C}_{d11} - \mathbf{C}_{d12}\mathbf{C}_{d22}^{-1}\mathbf{C}_{d21}, \\
 \bar{\mathbf{C}}_d &= \mathbf{C}_{d22} - \mathbf{C}_{d21}\mathbf{C}_{d11}^{-1}\mathbf{C}_{d12}.
 \end{aligned} \tag{4.8}$$

It can be shown that \mathbf{C}_d and $\bar{\mathbf{C}}_d$ as in (4.8) are invariant with respect to the choice of g-inverses involved. A crossover design is said to be *connected* (cf. Chapter 3) for direct effects if all contrasts among the direct effects are estimable, a necessary and sufficient condition for this being $\text{Rank}(\mathbf{C}_d) = t - 1$. Connectedness for carry-overs is analogously defined.

We now briefly indicate the analysis of the data arising from a crossover design under the model (4.3), assuming that there are no missing observations. The total sum of squares (total SS) with $np - 1$ degrees of freedom (df) can be calculated as usual on the basis of the np individual observations. The sum of squares due to periods (SSP) and that due to subjects (SSS), with $p - 1$ and $n - 1$ df, can also be obtained routinely on the basis of the p period-wise observational totals and the n subject-wise observational totals, respectively (see e.g., Cochran and Cox 1957, Section 4.4). Turning to the direct and carry-over effects, we define the $2t \times 1$ vector of *adjusted* treatment totals as

$$\begin{aligned}
 \mathbf{Q} &= \mathbf{X}'_2\mathbf{y}_d - \mathbf{X}'_2\mathbf{X}_1(\mathbf{X}'_1\mathbf{X}_1)^{-1}\mathbf{X}'_1\mathbf{y}_d \\
 &= \begin{bmatrix} \mathbf{T}'_d\mathbf{A}\mathbf{y}_d \\ \mathbf{F}'_d\mathbf{A}\mathbf{y}_d \end{bmatrix} = \begin{bmatrix} \mathbf{Q}_1 \\ \mathbf{Q}_2 \end{bmatrix},
 \end{aligned} \tag{4.9}$$

where \mathbf{A} is as in (4.5), $\mathbf{Q}_1 = \mathbf{T}'_d\mathbf{A}\mathbf{y}_d$, and $\mathbf{Q}_2 = \mathbf{F}'_d\mathbf{A}\mathbf{y}_d$. Then the SS for the direct and carry-over effects (jointly) after the elimination of the period and subject effects is

$$\text{SS}(\boldsymbol{\tau}, \boldsymbol{\rho}) = \mathbf{Q}'(\mathbf{C}_d(\boldsymbol{\tau}, \boldsymbol{\rho}))^{-1}\mathbf{Q}. \tag{4.10}$$

For a crossover design, which is connected for both direct and carry-over effects, $SS(\boldsymbol{\tau}, \boldsymbol{\rho})$ has $2(t - 1)$ df because $\mathbf{Q}'_1 \mathbf{1}_t = \mathbf{Q}'_2 \mathbf{1}_t = 0$. Thus, the error sum of squares (SSE), with $(n - 1)(p - 1) - 2(t - 1)$ df, can be obtained as

$$SSE = \text{Total SS} - \text{SSP} - \text{SSS} - SS(\boldsymbol{\tau}, \boldsymbol{\rho}).$$

In order to test the significance of the direct effects, one requires the corresponding adjusted SS as given by

$$SS_{\text{adj}}(\boldsymbol{\tau}) = (\mathbf{Q}_1 - \mathbf{C}_{d12} \mathbf{C}_{d22}^- \mathbf{Q}_2)' \mathbf{C}_d^- (\mathbf{Q}_1 - \mathbf{C}_{d12} \mathbf{C}_{d22}^- \mathbf{Q}_2),$$

with $t - 1$ df, where \mathbf{C}_{d11} , \mathbf{C}_{d12} , and \mathbf{C}_{d21} are as in (4.7) and \mathbf{C}_d is as in (4.8). On the basis of $SS_{\text{adj}}(\boldsymbol{\tau})$ and SSE, the F -test can now be employed in a straightforward manner for testing the significance of direct effects. The procedure for testing the significance of carry-over effects is similar.

4.3 Some Families of Crossover Designs

We now introduce a few classes of crossover designs, which have been widely studied in the literature. Apparently, the designs described in Definitions 4.1 and 4.2 were first formally defined and studied by Hedayat and Afsarinejad (1978) and Cheng and Wu (1980), and systematic construction methods for these are available. However, in these designs, the numbers of periods often exceed the numbers of treatments to be compared. In some experiments, it may be difficult to accommodate a large number of periods, and so one may prefer designs with $p < t$. Patterson (1952) was probably the first to give systematic methods of construction for designs with $p \leq t$. Freeman (1959), Patterson and Lucas (1962), Atkinson (1966), Hedayat and Afsarinejad (1975), Constantine and Hedayat (1982), Afsarinejad (1983, 1985), and Stufken (1991) also considered designs with $p \leq t$. Some designs with $p \leq t$ are described in Definitions 4.3 through 4.5. All these designs have nice combinatorial properties and, as a result, they have simple forms of the information matrix for inference on direct effects under model (4.3). Moreover, as will be seen later, they also enjoy excellent optimality properties.

A design is said to be *uniform on periods* if, in each period, it allocates each treatment to the same number of subjects. Similarly, a design is *uniform on subjects* if, for each subject, it allocates each treatment to the same number of periods. A design is simply said to be *uniform* if it is uniform on periods as well as on subjects.

The aforementioned definitions imply that, for a uniform design $d \in \Omega_{t,n,p}$,

$$\mathbf{r}_d = \frac{np}{t} \mathbf{1}_t, \quad \mathbf{R}_d = \frac{np}{t} \mathbf{I}_t, \quad \mathbf{M}_d = \frac{n}{t} \mathbf{J}_{tp}, \quad \mathbf{N}_d = \frac{p}{t} \mathbf{J}_{tn}. \tag{4.11}$$

4.3.1 Designs with $p \geq t$

Definition 4.1 A design $d \in \Omega_{t,n,p}$ is said to be *balanced* if, in the order of application, no treatment precedes itself and each treatment is preceded by every other treatment the same number of times.

Definition 4.2 A design $d \in \Omega_{t,n,p}$ is said to be strongly balanced if, in the order of application, each treatment is preceded by every treatment (including itself) the same number of times.

Clearly, if d is either balanced or strongly balanced, we have

$$\mathbf{Z}_d = \mathbf{T}'_d \mathbf{F}_d = \frac{n(p-1)}{t(t-1)}(\mathbf{J}_t - \mathbf{I}_t) \quad \text{or} \quad \mathbf{Z}_d = \frac{n(p-1)}{t^2} \mathbf{J}_t, \quad (4.12)$$

respectively. We now give some examples of these designs, where we define the positive integers $\lambda_1, \lambda_2, \mu_1$, and μ_2 as $\lambda_1 = n(p-1)/(t(t-1)), \lambda_2 = n(p-1)/t^2, \mu_1 = n/t$, and $\mu_2 = p/t$. Thus, for a balanced design, $\mathbf{Z}_d = \lambda_1(\mathbf{J}_t - \mathbf{I}_t)$ and, for a strongly balanced design, $\mathbf{Z}_d = \lambda_2 \mathbf{J}_t$.

Example 4.2

The designs d_1 and d_2 in Example 4.1 are both balanced uniform designs with $\lambda_1 = 2$. In the following, we give two examples of strongly balanced designs with $t = 3$; the first design has $n = 9, p = 6$, and the second has $n = 6, p = 4$:

$$d_3 \equiv \begin{matrix} 1 & 1 & 1 & 2 & 2 & 2 & 3 & 3 & 3 \\ 1 & 2 & 3 & 1 & 2 & 3 & 1 & 2 & 3 \\ 2 & 2 & 2 & 3 & 3 & 3 & 1 & 1 & 1 \\ 2 & 3 & 1 & 2 & 3 & 1 & 2 & 3 & 1 \\ 3 & 3 & 3 & 1 & 1 & 1 & 2 & 2 & 2 \\ 3 & 1 & 2 & 3 & 1 & 2 & 3 & 1 & 2 \end{matrix}, \quad d_4 \equiv \begin{matrix} 1 & 2 & 3 & 3 & 1 & 2 \\ 2 & 3 & 1 & 2 & 3 & 1 \\ 3 & 1 & 2 & 1 & 2 & 3 \\ 3 & 1 & 2 & 1 & 2 & 3 \end{matrix}.$$

The design $d_3 \in \Omega_{3,9,6}$ is uniform, with $\mu_1 = 3, \mu_2 = 2$, and $\lambda_2 = 5$, while $d_4 \in \Omega_{3,6,4}$ is uniform only on periods with $\mu_1 = 2$ and $\lambda_2 = 2$. Note that d_4 is uniform on subjects in the first $3(= p - 1)$ periods, and its last period is obtained by repeating the allocation in the previous period. Patterson and Lucas (1959) named a design of the form d_4 as an *extraperiod design*.

In view of (4.11) and (4.12), it is easy to see that the properties of uniformity and balance lead to substantial simplifications in the forms of the information matrices. Consider any two designs in $\Omega_{t,n,p}$, say \tilde{d} and d_* , which are balanced uniform and strongly balanced uniform, respectively. Then, on simplification from (4.7), one can show that

$$\begin{aligned} C_{\tilde{d}11} &= \mu_1 p \mathbf{H}_t, & C_{\tilde{d}12} &= -\lambda_1 \mathbf{H}_t, & C_{\tilde{d}22} &= \mu_1(p-1-p^{-1}) \mathbf{H}_t, \\ C_{d_*11} &= \mu_1 p \mathbf{H}_t, & C_{d_*12} &= \mathbf{O}_t, & C_{d_*22} &= \mu_1(p-1-p^{-1}) \mathbf{H}_t, \end{aligned} \quad (4.13)$$

where

$$\mathbf{H}_t = \mathbf{I}_t - t^{-1} \mathbf{J}_t. \quad (4.14)$$

From (4.8), it can be verified that for a balanced uniform design \tilde{d} , the matrices $C_{\tilde{d}}$ and $\bar{C}_{\tilde{d}}$ are completely symmetric, given by

$$C_{\tilde{d}} = \alpha_1 \mathbf{H}_t, \quad \bar{C}_{\tilde{d}} = \alpha_2 \mathbf{H}_t, \quad (4.15)$$

where

$$\begin{aligned} \alpha_1 &= \mu_1 p \left[1 - (p - 1)^2 (t - 1)^{-2} (p^2 - p - 1)^{-1} \right], \\ \alpha_2 &= \mu_1 \left[(p - 1 - p^{-1}) - (p - 1)^2 (t - 1)^{-2} p^{-1} \right]. \end{aligned}$$

Since H_t is an idempotent matrix, a g-inverse of H_t is I_t . Hence, one can consider g-inverses of $C_{\bar{d}}$ and $\bar{C}_{\bar{d}}$ as given, respectively, by

$$(C_{\bar{d}})^- = \alpha_1^{-1} I_t, \quad (\bar{C}_{\bar{d}})^- = \alpha_2^{-1} I_t. \tag{4.16}$$

From (4.16), it is thus clear that the analysis of a balanced uniform crossover design becomes extremely simple and this makes such designs attractive to users. Uniform crossover designs have been used in diverse areas of investigation and, for references to such work, we refer to Bate and Jones (2008).

From (4.8) and (4.13), it is also clear that the analysis of data from a strongly balanced uniform design d_* is further simplified owing to the fact that $C_{d_*,12} = \mathbf{O}_t$. Thus, the direct and carry-over effects are orthogonally estimable under these designs and

$$C_{d_*} = \mu_1 p H_t, \quad \bar{C}_{d_*} = \mu_1 (p - 1 - p^{-1}) H_t. \tag{4.17}$$

Such simple forms of the information matrices make it very convenient to study the statistical properties of these designs.

Again, Definitions 4.1 and 4.2 imply that the parameters t, n , and p need to satisfy certain divisibility requirements for these designs to exist. For instance, a strongly balanced design exists only if t^2 divides $n(p - 1)$. To overcome this problem, Kunert (1983) departed from the requirement of strong balance and introduced *nearly strongly balanced* designs where, instead of requiring that each treatment pair appear equally often in successive periods, he stipulated that each treatment pair appear in successive periods as equally often as possible. Let us write $\mathbf{Z}_d = (z_{dss'})$, that is, $z_{dss'}$ is the number of times treatment s is immediately preceded by treatment s' . A design d is a nearly strongly balanced design if

1. $z_{dss'}$ is equal to either $[n(p - 1)/t^2]$ or $[n(p - 1)/t^2] + 1$, for all $1 \leq s, s' \leq t$;
2. $\mathbf{Z}_d \mathbf{Z}'_d$ is of the form $a \mathbf{I}_t + b \mathbf{J}_t$ for some constants a and b .

In condition 1, $[\cdot]$ is the greatest integer function. Bate and Jones (2006) introduced *nearly balanced designs*, which only require condition 1. We give one example for each of these designs; for further details including their optimality and efficiency properties, we refer to Kunert (1983) and Bate and Jones (2006).

Example 4.3

The design d_5 with $t = 3, n = 6 = p$, is nearly strongly balanced, while d_6 with $t = 5, n = 10, p = 15$ is nearly balanced:

$$\begin{aligned}
 d_5 \equiv & \begin{matrix} 1 & 2 & 3 & 1 & 2 & 3 \\ 2 & 3 & 1 & 1 & 2 & 3 \\ 3 & 1 & 2 & 2 & 3 & 1 \\ 3 & 1 & 2 & 3 & 1 & 2 \\ 2 & 3 & 1 & 3 & 1 & 2 \\ 1 & 2 & 3 & 2 & 3 & 1 \end{matrix} & d_6 \equiv & \begin{matrix} 5 & 1 & 2 & 3 & 4 & 2 & 3 & 4 & 5 & 1 \\ 2 & 3 & 4 & 5 & 1 & 1 & 2 & 3 & 4 & 5 \\ 1 & 2 & 3 & 4 & 5 & 4 & 5 & 1 & 2 & 3 \\ 4 & 5 & 1 & 2 & 3 & 3 & 4 & 5 & 1 & 2 \\ 3 & 4 & 5 & 1 & 2 & 3 & 4 & 5 & 1 & 2 \\ 3 & 4 & 5 & 1 & 2 & 4 & 5 & 1 & 2 & 3 \\ 4 & 5 & 1 & 2 & 3 & 1 & 2 & 3 & 4 & 5 \\ 1 & 2 & 3 & 4 & 5 & 2 & 3 & 4 & 5 & 1 \\ 2 & 3 & 4 & 5 & 1 & 5 & 1 & 2 & 3 & 4 \\ 5 & 1 & 2 & 3 & 4 & 5 & 1 & 2 & 3 & 4 \\ 5 & 1 & 2 & 3 & 4 & 5 & 1 & 2 & 3 & 4 \\ 1 & 2 & 3 & 4 & 5 & 4 & 5 & 1 & 2 & 3 \\ 4 & 5 & 1 & 2 & 3 & 1 & 2 & 3 & 4 & 5 \\ 2 & 3 & 4 & 5 & 1 & 3 & 4 & 5 & 1 & 2 \\ 3 & 4 & 5 & 1 & 2 & 2 & 3 & 4 & 5 & 1 \end{matrix}
 \end{aligned}$$

4.3.2 Designs with $p \leq t$

In this subsection, we consider some designs with $p \leq t$. All these designs enjoy good optimality properties, and we will discuss these properties later in this chapter. We begin with the definition of the class of designs considered by Patterson (1952), which we shall call *Patterson designs*. These designs are very popular among experimenters because they involve a moderate number of subjects for given t and do not involve too many periods. Several families of such designs are known.

Definition 4.3 A design $d \in \Omega_{t,n,p}$, where $p \geq 3$, $t \geq 3$, will be said to be a Patterson design if the following conditions hold:

- (i) d is uniform on periods, so that $n = \mu_1 t$ for some integer $\mu_1 \geq 1$;
- (ii) d is balanced, so that $n(p - 1)/(t(t - 1)) = \lambda_1$ for some integer $\lambda_1 \geq 1$;
- (iii) When the subjects of d are viewed as blocks, they form the blocks of a balanced incomplete block (BIB) design with block size p ;
- (iv) When d is restricted to the first $p - 1$ periods, then again, the subjects of d form the blocks of a BIB design with block size $p - 1$;
- (v) In the set of μ_1 subjects receiving a given treatment in the last period, every other treatment is applied λ_1 times in the first $p - 1$ periods.

Example 4.4

The design d_7 shown in the following is a Patterson design with $t = 4, p = 3, n = 12$:

$$d_7 \equiv \begin{matrix} 1 & 2 & 3 & 4 & 1 & 2 & 3 & 4 & 1 & 2 & 3 & 4 \\ 2 & 1 & 4 & 3 & 3 & 4 & 1 & 2 & 4 & 3 & 2 & 1 \\ 3 & 4 & 1 & 2 & 4 & 3 & 2 & 1 & 2 & 1 & 4 & 3 \end{matrix}$$

For a Patterson design d , it can be shown that

$$C_{d11} = \lambda_1 t \mathbf{H}_t, \quad C_{d12} = -(\lambda_1 t/p) \mathbf{H}_t, \quad C_{d22} = (\lambda_1 (pt - t - 1)/p) \mathbf{H}_t. \tag{4.18}$$

So, analogously to (4.15), the information matrices C_d and \bar{C}_d are both constant multiples of \mathbf{H}_t , leading to considerable simplification in the analysis.

Stufken (1991) introduced a new class of designs and proved that these designs have good optimality properties in certain subclasses of $\Omega_{t,n,p}$. These designs, labeled as *Stufken designs*, are described in the following text. Stronger optimality properties of these designs were established later by Hedayat and Yang (2004).

Definition 4.4 A design $d \in \Omega_{t,n,p}$ will be called a *Stufken design* if it satisfies the following properties:

- (i) d is uniform on periods;
- (ii) The first $p - 1$ periods of d form a BIB design with subjects as blocks;
- (iii) In the last period of d , θ subjects receive a treatment that was not allocated to them in any of the previous periods, while the remaining $n - \theta$ subjects receive the same treatment as in period $p - 1$, where θ is the nearest integer (or one of the nearest integers) to $\frac{n(pt - t - 1)}{(p - 1)t}$;
- (iv) $z_{dss'} - p^{-1} \sum_{j=1}^n n_{dsj} \bar{n}_{ds'j}$ is independent of s and s' , $s \neq s'$, where $Z_d = (z_{dss'})$, $N_d = (n_{dsj})$, $\bar{N}_d = (\bar{n}_{dsj})$ are as defined earlier;
- (v) $\sum_{j=1}^n n_{dsj} n_{ds'j}$ is independent of s and s' , $s \neq s'$.

To construct a Stufken design with $t = 3, n = 36, p = 3$, consider the array $OA_I(6, 3^3, 2)$ shown in its transposed form in Example 4.14. A Stufken design with $t = 3, n = 36, p = 3$, can be constructed by juxtaposing five copies of this transposed array together with one copy of a 3×6 array, which has the first two rows of this transposed array as its first and second rows and the second row of this transposed array repeated as its third row. We shall describe the optimality properties of Stufken designs in Section 4.5.

Kunert and Stufken (2002) studied a general class of designs called *totally balanced* designs, which satisfy more stringent combinatorial conditions and have good statistical properties. These designs are quite general in the sense that, though the number of subjects needs to be a multiple of the number of treatments, there is no restriction on the number of periods, thus allowing $p < t$, $p > t$, or $p = t$. An attractive feature of these designs is that they have good optimality properties, even under models more complicated than the one in (4.3). We define these designs here and will study them again later.

Definition 4.5 A design $d \in \Omega_{t,n,p}$ is called *totally balanced* if

- (i) d is uniform on periods;
- (ii) Each treatment is allocated as equally as possible to each subject in d , that is, each treatment is allocated either $\lfloor p/t \rfloor$ or $\lfloor p/t \rfloor + 1$ times to each subject;
- (iii) The number of subjects where treatments s and s' are both allocated $\lfloor p/t \rfloor + 1$ times in d is the same for every pair $s \neq s', 1 \leq s, s' \leq t$;
- (iv) Each treatment is allocated as equally as possible to each subject in the first $p - 1$ periods of d , that is, each treatment is allocated either $\lfloor (p - 1)/t \rfloor$ or $\lfloor (p - 1)/t \rfloor + 1$ times to each subject over periods $1, \dots, p - 1$;
- (v) The number of subjects where treatments s and s' are both allocated $\lfloor (p - 1)/t \rfloor + 1$ times in the first $p - 1$ periods of d is the same for every pair $s \neq s', 1 \leq s, s' \leq t$;

- (vi) d is balanced;
- (vii) The number of subjects where both treatments s and s' appear $[p/t] + 1$ times in d and the treatment s' does not appear in the last period is the same for every pair $s, s', 1 \leq s, s' \leq t; s \neq s'$.

Interestingly, some of the earlier designs follow as special cases of these designs; for instance, when p is a multiple of t , a totally balanced design is a balanced uniform design. The Patterson design shown in Example 4.4 is also a totally balanced design.

Example 4.5

The following is a totally balanced design with $t = 3, n = 6, p = 4$:

$$d_8 \equiv \begin{matrix} 1 & 2 & 3 & 3 & 1 & 2 \\ 2 & 3 & 1 & 2 & 3 & 1 \\ 3 & 1 & 2 & 1 & 2 & 3 \\ 1 & 2 & 3 & 3 & 1 & 2 \end{matrix}.$$

4.3.3 Two-Period Designs

Clearly, in a crossover design, the number of periods p is at least two. We now review some designs with only two periods, that is, designs with $p = 2$. These designs are of substantial interest in clinical trials and have been studied, among others, by Grizzle (1965), Hills and Armitage (1979), Armitage and Hills (1982), and Willan and Pater (1986).

Hedayat and Zhao (1990) gave an interesting connection between a crossover design with two periods and a block design. We present this result in the following text, where we write \mathcal{C}_d to denote the information matrix for treatments for an arbitrary block design d under the usual additive linear model for block designs (see Chapter 3, Section 3.1).

Theorem 4.1 *Let d be a design in $\Omega_{t,n,2}$, and let there be $b \leq t$ distinct treatments in the first period of d , these treatments being labeled as $1, 2, \dots, b$. Then there exists a block design d_0 with t treatments and b blocks of sizes $\bar{r}_{d1}, \dots, \bar{r}_{db}$, such that the treatment versus block incidence matrix of d_0 equals \mathbf{Z}_d , and the relationship*

$$\mathcal{C}_{d_0} = 2\mathbf{C}_d \tag{4.19}$$

holds, where \mathbf{C}_d is as in (4.8). Conversely, from a block design with t treatments and $b(\leq t)$ blocks, one can obtain a crossover design d in $\Omega_{t,n,2}$, with n equal to the total number of experimental units in the block design, such that (4.19) holds.

This connection is helpful in the study of optimality of two-period crossover designs as one can invoke well-known results on optimality of block designs for this purpose. The following example illustrates Theorem 4.1.

Example 4.6

Let $d \in \Omega_{3,12,2}$ be as follows:

$$d \equiv \begin{matrix} 1 & 2 & 3 & 1 & 2 & 3 & 1 & 2 & 3 & 2 & 3 & 1 \\ 1 & 2 & 3 & 2 & 3 & 1 & 1 & 2 & 3 & 1 & 2 & 3 \end{matrix}.$$

This design has $\bar{r}_{d1} = \bar{r}_{d2} = \bar{r}_{d3} = 4$. Then the corresponding block design d_0 with incidence matrix equal to Z_d is given by the blocks

$$d_0 \equiv \begin{array}{l} \text{Block I: } 1, 1, 2, 3 \\ \text{Block II: } 1, 2, 2, 3. \\ \text{Block III: } 1, 2, 3, 3 \end{array}$$

Conversely, d can be obtained from d_0 , and it can be verified that for these designs, (4.19) holds.

Again, from (4.19), it is clear that a design $d \in \Omega_{t,n,2}$ is connected for direct effects if and only if the corresponding block design d_0 is connected. Because of this fact, it is easy to see that the contrasts among direct effects of treatments cannot be estimated from the design d_1 in Example 4.1. This is because the block design d_0 corresponding to d_1 has the following two blocks:

$$\text{Block I: } 2, 2 \quad \text{and} \quad \text{Block II: } 1, 1.$$

The block design with the aforementioned two blocks is clearly disconnected, implying that d_1 is disconnected too. However, if d_1 is modified to include identical pairs to give a design d^* as

$$d^* \equiv \begin{array}{cccc} 1 & 2 & 1 & 2 \\ 1 & 2 & 2 & 1' \end{array}$$

then the corresponding block design has blocks

$$\text{Block I: } 1, 2 \quad \text{and} \quad \text{Block II: } 2, 1,$$

and is connected, leading to the connectedness of d^* for direct effects.

There is another aspect of two-period designs that makes them interesting and we elaborate on this now. In the context of crossover designs, since the same subject gives multiple responses, it is sometimes reasonable to deviate from the traditional model with uncorrelated errors and, instead, consider a model under which the observations from the same subject are assumed to be correlated, these correlations being the same for all subjects, while those from different subjects remain uncorrelated. Thus, for $p = 2$, the model is the same as (4.3) with the exception that the variance-covariance matrix of the errors is now $\sigma^2 \Sigma$, where $\Sigma = I_n \otimes V$ and $V = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}$, with ρ ($-1 < \rho < 1$) representing the correlation coefficient between the observations arising from the same subject.

An interesting aspect of two-period crossover designs is that the properties of these designs under a model with correlated errors can be studied easily. It turns out that for a design d with $p = 2$, the information matrix for the joint estimation of direct and carry-over effects under a model with correlated errors as specified earlier is proportional to the joint information matrix under (4.3); see Lemma 1.3.1 in Bose and Dey (2009). Hence, for $p = 2$, the optimality properties of a design under the uncorrelated errors model (4.3) remain robust even if the errors are correlated as described earlier.

4.3.4 Two-Treatment Designs

Experiments with only *two treatments* are often used in practice; for example, in medical experiments, one treatment may be a placebo or the standard drug, while the other treatment could be a newly developed drug. The literature on two-treatment designs has been enriched by various authors, including Kershner and Federer (1981), Laska and Meisner (1985), Matthews (1987, 1990), Kunert (1991), Kushner (1997a), Carriere and Reinsel (1992), Carriere and Huang (2000), and Kunert and Stufken (2008). In this context, a class of designs called *dual-balanced designs* are found to have good statistical properties (see Section 4.5.5).

With two treatments, labeled say 1 and 2, consider a treatment sequence of length p , every element of the sequence being either 1 or 2. For any such sequence, its dual is obtained by interchanging the positions of 1 and 2. Then a dual-balanced design is defined as follows.

Definition 4.6 *A design that assigns an equal number of subjects to any treatment sequence and its dual is called a dual-balanced design.*

Example 4.7

The following are examples of dual-balanced designs:

$$d_9 \equiv \begin{matrix} 1 & 2 \\ 2 & 1 \\ 1 & 2 \\ 2 & 1 \end{matrix}, \quad d_{10} \equiv \begin{matrix} 1 & 2 & 1 & 2 \\ 1 & 2 & 2 & 1 \\ 2 & 1 & 2 & 1 \\ 2 & 1 & 1 & 2 \end{matrix}, \quad d_{11} \equiv \begin{matrix} 1 & 2 \\ 1 & 2 \\ 1 & 2 \\ 2 & 1 \end{matrix}.$$

Note that though these designs are called dual balanced, they need not always be balanced in the sense of Definition 4.1. For example, d_9 is balanced, while d_{10} is strongly balanced, and d_{11} is neither balanced nor strongly balanced in the sense of Definitions 4.1 and 4.2.

4.4 Constructions of Some Families of Designs

In this section, we give methods for construction for some selected classes of designs. To see why these methods lead to the designs as claimed, we refer the reader to the related references.

4.4.1 Balanced Uniform Designs

It is clear from Definition 4.1 (Section 4.3.1) and the definition of uniformity (Section 4.3) that $\Omega_{t,n,p}$ can contain a balanced uniform design only if t, n , and p satisfy the following three conditions:

- (i) $n = \mu_1 t$, for some integer $\mu_1 \geq 1$;
- (ii) $p = \mu_2 t$, for some integer $\mu_2 \geq 1$;
- (iii) $n(p - 1) = \lambda_1 t(t - 1)$, for some integer $\lambda_1 \geq 1$.

Case 1: t even. Williams (1949) gave a method for constructing a balanced uniform design $d \in \Omega_{t,t,t}$, where t is any even integer. Starting with the initial $t \times 1$ vector

$$a_0 = \left(1, t, 2, t-1, \dots, \frac{t}{2} - 1, \frac{t}{2} + 2, \frac{t}{2}, \frac{t}{2} + 1 \right)',$$

he obtained $t - 1$ other vectors as $a_u = a_0 + (u, u, \dots, u)'$, $1 \leq u \leq t - 1$, where all entries in a_u are reduced modulo t and every 0 in a_u is replaced by t . Then the $t \times t$ array

$$A_t = [a_0, a_1, \dots, a_{t-1}],$$

is a balanced uniform design in $\Omega_{t,t,t}$. As usual, rows of A_t represent the t periods and the columns represent the t subjects. The design A_t is often called a *Williams square*.

Example 4.8

The following are two examples of Williams squares, or balanced uniform designs, in $\Omega_{t,t,t}$, one with $t = 4$ and $a_0 = (1, 4, 2, 3)'$ and another with $t = 6$ and $a_0 = (1, 6, 2, 5, 3, 4)'$:

					1	2	3	4	5	6
1	2	3	4		6	1	2	3	4	5
4	1	2	3		2	3	4	5	6	1
2	3	4	1	'	5	6	1	2	3	4
3	4	1	2		3	4	5	6	1	2
					4	5	6	1	2	3

For $n = \mu_1 t$ and $p = t$, a balanced uniform design can be obtained by juxtaposing μ_1 copies of a Williams square in $\Omega_{t,t,t}$.

Case 2: t odd. Balanced uniform designs with odd t in $\Omega_{t,t,t}$ are known for only a few values of t , for instance, $t = 9, 15, 21, 27, 39, 55, 57$, while they do not exist for $t = 3, 5, 7$. Higham (1998) proved that a balanced uniform design exists in $\Omega_{t,t,t}$ when t is a composite number. The design for $t = 21$ is shown in Hedayat and Afsarinejad (1975), and designs for $t = 9, 15, 27$ are given in Hedayat and Afsarinejad (1978). The afore mentioned papers may be consulted for more details and references.

However, when $n = 2t$, a balanced uniform design exists in $\Omega_{t,2t,t}$ for all odd t . Williams (1949) gave a construction starting with two initial vectors. Let

$$b_0 = \left(1, t, 2, t-1, \dots, \frac{t+5}{2}, \frac{t-1}{2}, \frac{t+3}{2}, \frac{t+1}{2} \right)',$$

$$c_0 = \left(\frac{t+1}{2}, \frac{t+3}{2}, \frac{t-1}{2}, \frac{t+5}{2}, \dots, t-1, 2, t, 1 \right)'.$$

Note that c_0 is obtained by writing the entries of b_0 in the reverse order. Now, for $1 \leq u \leq t - 1$, let $b_u = b_0 + (u, \dots, u)'$ and $c_u = c_0 + (u, \dots, u)'$, where the elements of b_u and c_u are reduced modulo t and, thereafter, every 0 therein is replaced by t . Then a balanced uniform design in $\Omega_{t,2t,t}$ is given by the $t \times 2t$ array

$$B_t = [b_0 \ b_1 \ \dots \ b_{t-1} \ c_0 \ c_1 \ \dots \ c_{t-1}].$$

The design d_2 with $t = 3$ in Example 4.1 is constructed via this method. A design for $t = 5$ is shown next.

Example 4.9

For $t = 5$, $b_0 = (1, 5, 2, 4, 3)'$ and $c_0 = (3, 4, 2, 5, 1)'$, which lead to the following balanced uniform design in $\Omega_{5,10,5}$:

1	2	3	4	5	3	4	5	1	2
5	1	2	3	4	4	5	1	2	3
2	3	4	5	1	2	3	4	5	1
4	5	1	2	3	5	1	2	3	4
3	4	5	1	2	1	2	3	4	5

There are several simple modifications of the Williams squares that give designs with the same balance properties. For a review of such modifications, we refer the reader to Issac et al. (2001).

4.4.2 Strongly Balanced Uniform Designs

From the definition of uniformity (Section 4.3.1) and Definition 4.2, one may check that $\Omega_{t,n,p}$ can contain a uniform, strongly balanced design only if the following conditions hold:

- (i) $n = \mu_3 t^2$, for some integer $\mu_3 \geq 1$;
- (ii) $p = \mu_2 t$, for some integer $\mu_2 \geq 2$.

Early examples of these designs with $t = 3, n = 18, p = 6$, and $t = 4, n = 16, p = 8$ were given by Quenouille (1953). Later, Berenblut (1964) and Patterson (1973) gave general methods of their construction in $\Omega_{t,n=t^2,p=2t}$. Cheng and Wu (1980) generalized this family to give constructions for situations where t^2 divides n and p is an *even* multiple of t . The design d_3 shown in Example 4.2 is one such design. Starting with designs constructed by Cheng and Wu's method, one may obtain a strongly balanced uniform design in $\Omega_{t,n=\alpha t^2,p=2\beta t}$, where α, β are integers, by juxtaposing copies of this design. Sen and Mukerjee (1987) gave a construction of strongly balanced uniform designs for cases when t^2 divides n and p is an *odd* multiple of t . This, together with the construction of Cheng and Wu (1980), shows that the necessary conditions (i) and (ii) are sufficient as well.

Using orthogonal arrays of strength *two*, Stufken (1996) gave a unified method of construction of strongly balanced uniform designs for general μ_2 , which covers both the odd and even cases. We describe this construction in the following.

An orthogonal array, $OA(n, t^p, 2)$, of strength two is an $n \times p$ array with entries from a set of t symbols, such that any $n \times 2$ subarray contains each ordered pair of symbols equally often as a row, precisely n/t^2 times (see Chapter 9). An $OA(t^2, t^3, 2)$ exists for all $t \geq 2$, and let such an array be denoted by A_0 , its entries being $1, 2, \dots, t$. Let B_0 be an orthogonal array $OA(t^2, t^2, 2)$, obtained from A_0 by deleting its third column. For $1 \leq u \leq t - 1$, let A_u be a $t^2 \times 3$ matrix obtained by adding u to each element of A_0 and similarly, let B_u be a $t^2 \times 2$ matrix obtained by adding u to each element of B_0 , where the elements of A_u and B_u are reduced modulo t , and then every 0 therein is replaced by t .

Finally, let A and B be the $3t \times t^2$ and $2t \times t^2$ matrices defined as

$$A = [A_0 \ A_1 \ \cdots \ A_{t-1}]', \quad B = [B_0 \ B_1 \ \cdots \ B_{t-1}]'$$

Since $\mu_2 \geq 2$, let $\mu_2 = 3\alpha + 2\beta$ for some nonnegative integers α, β . It can then be verified that the $\mu_2 t \times t^2$ array

$$[A' \cdots A' \ B' \cdots B']'$$

consisting of α copies of A and β copies of B is a strongly balanced uniform design in $\Omega_{t,t^2,p=\mu_2 t}$. Now juxtaposing μ_3 copies of this design, we get a strongly balanced uniform design in $\Omega_{t,n=\mu_3 t^2,p=\mu_2 t}$.

Example 4.10

Let $t = 3, n = 9, p = 6$. We start with $A_0 \equiv \text{OA}(9, 3^3, 2)$ as shown in the following in transposed form:

$$A_0' \equiv \begin{bmatrix} 1 & 1 & 1 & 2 & 2 & 2 & 3 & 3 & 3 \\ 1 & 2 & 3 & 1 & 2 & 3 & 1 & 2 & 3 \\ 1 & 2 & 3 & 2 & 3 & 1 & 3 & 1 & 2 \end{bmatrix}.$$

Here, $\mu_2 = 2$ and so we take $\alpha = 0, \beta = 1$. From A_0 , we obtain B_0 and then B_1 and B_2 , leading to the matrix $B = (B_0, B_1, B_2)'$ as indicated earlier. This B will be the design d_3 displayed in Example 4.2.

4.4.3 Patterson Designs

We now describe methods of construction of some families of Patterson designs (see Definition 4.3). Such a family exists in particular when t is a prime or a prime power. Let $u_0 = 0, u_1 = 1, u_2 = x, u_3 = x^2, \dots, u_{t-1} = x^{t-2}$ be the elements of $GF(t)$ (a Galois field of order t), where x is a primitive element. Details on Galois fields may be found for example, in Lidl and Niederreiter (1986). For $1 \leq i \leq t-1$, define a $t \times t$ array L_i whose (α, β) th element equals $u_i u_\alpha + u_\beta, 0 \leq \alpha, \beta \leq t-1$. Then L_1, \dots, L_{t-1} form a complete set of $(t-1)$ mutually orthogonal Latin squares of order t (for more on mutually orthogonal Latin squares, see Chapter 1, Section 1.4.3). Furthermore, L_{i+1} can be obtained by cyclically permuting the last $t-1$ rows of $L_i, 1 \leq i \leq t-2$. The $t \times t(t-1)$ array $L = [L_1 \ L_2 \ \dots \ L_{t-1}]$ is a Patterson design in $\Omega_{t,t(t-1),t}$. On deleting any $t-p$ rows of L , where $t > p \geq 3$, one obtains a Patterson design in $\Omega_{t,t(t-1),p}$. The design in Example 4.4 is obtained by this method, after deleting the last row of the array L for $t = 4$.

Patterson (1952) obtained several families of designs that require fewer subjects than the method described earlier. In particular, the following families of designs were obtained by Patterson (1952):

Family I: $t = 4m + 3, n = t(t-1)/2, p = 3, t$ a prime or a prime power.

Family II: $t = 4u + 3, n = 2t, p = (t+1)/2, t$ a prime.

Example 4.11

The following design is a member of Family I with $t = 11, n = 55, p = 3$:

1	2	3	4	5	6	7	8	9	10	11	1	2	3	4	5	6	7	8	9	10	11
2	3	4	5	6	7	8	9	10	11	1	5	6	7	8	9	10	11	1	2	3	4
4	5	6	7	8	9	10	11	1	2	3	2	3	4	5	6	7	8	9	10	11	1
1	2	3	4	5	6	7	8	9	10	11	1	2	3	4	5	6	7	8	9	10	11
6	7	8	9	10	11	1	2	3	4	5	10	11	1	2	3	4	5	6	7	8	9
5	6	7	8	9	10	11	1	2	3	4	6	7	8	9	10	11	1	2	3	4	5
1	2	3	4	5	6	7	8	9	10	11											
4	5	6	7	8	9	10	11	1	2	3											
10	11	1	2	3	4	5	6	7	8	9											

4.5 Optimality under the Traditional Model

Hedayat and Afsarinejad (1978) initiated the study of optimality of crossover designs. Subsequently, the area of optimal crossover designs has been enriched by the contributions of a number of authors. Many of these results are with respect to the universal optimality criterion of Kiefer (1975) (see Chapter 3). It is well known that universal optimality implies the more common criteria like A -, D -, and E -optimality in the sense that a universally optimal design is also A -, D -, and E -optimal.

Let \mathcal{D} be a class of competing designs in a given context, and let A_d denote the information matrix for a set of relevant parametric functions (e.g., contrasts among the direct or carry-over effects in the setup of this chapter) under a design $d \in \mathcal{D}$ and a given model. Then, we recall from Chapter 3, Theorem 3.1 that a set of sufficient conditions for a design $d^* \in \mathcal{D}$ to be universally optimal over \mathcal{D} is that (1) A_{d^*} is completely symmetric and (2) $\text{trace}(A_{d^*}) \geq \text{trace}(A_d)$ for all $d \in \mathcal{D}$. In this section, we present a selection of results on optimal crossover designs. Throughout this section, we consider the model (4.3).

4.5.1 Balanced Uniform Designs

It is interesting to note how the optimality results on balanced uniform designs have been successively strengthened by various authors. The first result on optimal crossover designs was obtained by Hedayat and Afsarinejad (1978) who proved that a balanced uniform design (Definition 4.1) in $\Omega_{t, \mu_1 t, t}$ is universally optimal for the estimation of both direct and carry-over effects over the class of all uniform designs in $\Omega_{t, \mu_1 t, t}$. This result was strengthened by Cheng and Wu (1980), who removed the restriction of uniformity on the competing designs, but even then, their results are valid only in some subclasses of $\Omega_{t, n, p}, t \geq 3$. For instance, they proved that a balanced uniform design is universally optimal for the estimation of carry-over effects over the class of designs in which (1) $n = \mu_1 t, p = \mu_2 t$ for some integers μ_1, μ_2 , (2) no treatment is assigned to two consecutive periods on the same subject, and (3) each treatment is equally replicated in the first $p - 1$ periods. If in particular $\mu_2 = 1$ also holds (i.e., $p = t$), then restriction (3) is not needed. For the direct effects, they showed that a balanced uniform design is universally optimal for the estimation of direct effects over the class of designs which are uniform on subjects and uniform on the last period. This result on direct effects by Cheng and Wu (1980) was further extended by Kunert (1984a) who removed all restrictions on the competing class and proved that if $t = n = p > 2$, then

a balanced uniform design is universally optimal for the estimation of direct effects over $\Omega_{t,t,t}$ and, if $n = 2t, p = t, t \geq 6$, a balanced uniform design is universally optimal for direct effects over $\Omega_{t,2t,t}$.

A more general result was obtained by Hedayat and Yang (2003) who proved that, for $n = \mu_1 t, t = p > 2$, and $n \leq t(t-1)/2$, a balanced uniform design is universally optimal for direct effects in $\Omega_{t,n,t}$. On recalling that balanced uniform designs have completely symmetric information matrices (see (4.15)), the condition $n \leq t(t-1)/2$ is crucial in their result as this is needed to establish that a balanced uniform design maximizes the trace of the information matrix for direct effects among all designs in the competing class, thereby establishing its universal optimality. It may be noted that when this condition is not satisfied, universal optimality does not hold in general, though it is indeed true for $t = p = 3, n = 6$. Earlier, Street et al. (1990) had shown via a computer search that a balanced uniform design in $\Omega_{3,6,3}$ is A -optimal for direct effects; Hedayat and Yang (2004) extended this result to universal optimality. For larger values of t , they also showed that if $4 \leq p = t \leq 12$ and $n \leq t(t+2)/2$, then a balanced uniform design is universally optimal for the estimation of direct effects over $\Omega_{t,n,t}$.

4.5.2 Stufken and Patterson Designs

When we depart from balanced uniform designs, and focus on designs with uniformity on periods only, several optimality results are again available. For example, the universal optimality of Stufken designs (see Definition 4.4) for direct effects in certain subclasses of $\Omega_{t,n,p}$, was established by Stufken (1991). Kushner (1998) extended these results to show that if $n/(t(p-1))$ is an integer, then the Stufken designs are universally optimal for direct effects in the *entire* class $\Omega_{t,n,p}$. Hedayat and Yang (2004) improved Kushner's result to prove that if a Stufken design exists in $\Omega_{t,n,p}$, then it is universally optimal for direct effects over $\Omega_{t,n,p}$, irrespective of whether or not the aforementioned divisibility condition holds. By this result, a Stufken design, which exists for $t = p = 3, n = 36$, is universally optimal in $\Omega_{3,36,3}$, and thus it dominates the balanced uniform design in this class. Note that here the condition $n \leq t(t-1)/2$ in the result of Hedayat and Yang (2003), mentioned in the earlier subsection, is violated.

The optimality properties of Patterson designs (Definition 4.3) were studied by Shah et al. (2005) who showed that these designs are universally optimal for the estimation of *both direct and carry-over* effects over the subclass of all connected designs in $\Omega_{t,n,p}$ in which no treatment precedes itself.

4.5.3 Strongly Balanced Designs

We now turn to strongly balanced designs. The study of optimality aspects of such designs was initiated by Cheng and Wu (1980). They proved a very general result that shows that a strongly balanced uniform design (Definition 4.2) is universally optimal for the estimation of *both direct and carry-over* effects over the *entire class* $\Omega_{t,n,p}$. However, since such a design exists only if $t^2|n$ and $t|p$, these designs are quite large in size. By relaxing the condition of uniformity on subjects in the class of competing designs, Cheng and Wu (1980) obtained optimal designs that are smaller in size compared to strongly balanced uniform designs. They showed that a strongly balanced design that is uniform on periods and uniform on the subjects in the first $p-1$ periods is also universally optimal for *both direct and carry-over effects* over the *entire class*, $\Omega_{t,n,p}$. The fact that direct and carry-over effects become orthogonally

estimable in a strongly balanced design and the expressions in (4.17) are instrumental in proving the universal optimality over the *entire* class of designs.

4.5.4 Two-Period Designs

Hedayat and Zhao (1990) used the connection established between block designs and two-period crossover designs in Theorem 4.1 (Section 4.3.3) to obtain optimal crossover designs starting from optimal block designs. It follows from Theorem 4.1 that if a block design d_0 is optimal in the class of all proper (equal block size) block designs with t treatments and $b \leq t$ blocks, then the two-period crossover design corresponding to d_0 is also universally optimal for direct effects, in the class of all two-period designs in which b treatments appear in the first period equally often. The treatments need to occur equally often in the first period because the treatments in period 1 of a two-period design correspond to block labels in the block design, while those in period 2 correspond to the treatment labels within these blocks. Consider, for instance, Example 4.6. There, d_0 is a balanced block design and is universally optimal over the entire class of connected block designs with $t = 3$ treatments and $b = 3$ blocks each of size 4. Hence, it follows that the corresponding crossover design d is universally optimal for the estimation of direct effects over the subclass of $\Omega_{3,12,2}$, which are uniform on the first period.

Hedayat and Zhao (1990) also gave a set of necessary and sufficient conditions for a two-period crossover design with t treatments and n subjects to be universally optimal in the *entire* class $\Omega_{t,n,2}$. Their result is as follows:

Theorem 4.2 *Let $n \equiv 0 \pmod{t}$. Then a two-period, t -treatment, n -subject design d^* is universally optimal for direct effects over $\Omega_{t,n,2}$ if and only if*

- (i) $f_{d^*s} \equiv 0 \pmod{t}, 1 \leq s \leq t$, where f_{d^*s} is the number of times treatment s appears in the first period of d^*
- (ii) $z_{d^*s's} = f_{d^*s}/t, 1 \leq s' \leq t$, where $z_{d^*s's}$ is the number of subjects that receive treatment s in the first period and treatment s' in the second period of d^*

Note that the number of distinct treatments in the first period of the design d^* in Theorem 4.2 may be any number $\in \{1, 2, \dots, t\}$. Condition (i) of Theorem 4.2 merely demands that for $1 \leq s \leq t, f_{d^*s} = \mu_s t$, where $\mu_s \geq 0$ is an integer, subject to $\sum_{s=1}^t f_{d^*s} = n$. In particular, one can have an optimal two-period design as given by Theorem 4.2 where the same treatment is used for every subject in the first period, for example, the design d_{12} .

Example 4.12

Using Theorem 4.2, it is easy to see that the designs d_{12} and d_{13} in the following are universally optimal for direct effects over $\Omega_{3,12,2}$ and $\Omega_{6,18,2}$, respectively:

$$d_{12} \equiv \begin{matrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 2 & 3 & 1 & 2 & 3 & 1 & 2 & 3 & 1 & 2 & 3 \end{matrix} \quad ,$$

$$d_{13} \equiv \begin{matrix} 1 & 1 & 1 & 1 & 1 & 1 & 2 & 2 & 2 & 2 & 2 & 2 & 3 & 3 & 3 & 3 & 3 \\ 1 & 2 & 3 & 4 & 5 & 6 & 1 & 2 & 3 & 4 & 5 & 6 & 1 & 2 & 3 & 4 & 5 & 6 \end{matrix} \quad .$$

For the case when n is *not a multiple* of t , Hedayat and Zhao (1990) considered designs with only a single treatment in the first period and all t treatments allocated as equally as possible in the second period. Using the correspondence with block designs, they showed that such crossover designs are A -optimal for direct effects over $\Omega_{t,n,2}$.

4.5.5 Two-Treatment Designs

Matthews (1990) used the approach of approximate design theory (see Kiefer, 1959) to give an easily implementable method for producing optimal dual-balanced designs (Section 4.3.4) with two treatments. He showed that, for even p , any design that is optimal for direct effects is also optimal for carry-over effects, while for odd p , any design that is optimal for carry-over effects is also optimal for direct effects. For example, he showed that the designs d_9 and d_{10} in Example 4.7 are universally optimal designs for both direct and carry-over effects in $\Omega_{2,2,4}$ and $\Omega_{2,4,4}$, respectively. The designs that he obtained as optimal for *both* direct and carry-over effects are identical with the strongly balanced designs, shown to be optimal by Cheng and Wu (1980) and discussed earlier in this section. However, his designs that are optimal *only* for direct effects or *only* for carry-over effects need not be uniform over subjects nor uniform over subjects in the first $p - 1$ periods, and so are not covered by the results of Cheng and Wu (1980).

4.6 Some Other Models and Optimal Designs

Now we consider some models, other than the traditional one, which have been studied in the literature. Recall that the traditional model given in (4.3) makes some implicit assumptions about the carry-over effects. For example, it assumes that only first-order carry-over effects are present, the carry-over effect of a treatment in a period always remains the same no matter which treatment is producing the direct effect in this period, and the carry-over effect of a treatment does not depend on its direct effect.

However, in practice, such assumptions need not hold in all situations and so the validity of the model (4.3) has been questioned, especially in medical applications; see for example, Senn (1992) and Matthews (1994). For examples of situations in nonmedical applications where the simplistic model (4.3) may not be appropriate, see Kempton et al. (2001). For such situations, it becomes necessary to model the carry-over effects differently, and several authors have studied the problem of finding good designs under such modified models.

Moreover, the traditional model (4.3) also assumes that the errors are uncorrelated, an assumption that may not be met in typical crossover trials where it might be more realistic to expect that the observations arising from the same subject over time are correlated. In view of this, models with correlated errors have been studied too. In the following subsections, we review some of these models and describe a selection of optimality results under these models.

4.6.1 Circular Model

Recall that in (4.3), it was assumed that there are no carry-over effects in the first period. Models with carry-over effects in the first period too have been studied by some authors. For this, they proposed the inclusion of a preperiod or baseline period (0th period) when

each subject receives the same treatment as the one allocated to it in the p th period. Even though no observation is taken during this preperiod, the treatments applied in this period cause a carry-over effect to be present in the first period, thereby creating a *circular* pattern of carry-overs across the p periods. A model for studying these experiments is termed a *circular* model, its only change from (4.3) being that instead of $\rho_{d(0,j)} = 0$ as in (4.3), it now has $\rho_{d(0,j)} = \rho_{d(p,j)}$, $1 \leq j \leq n$. Consequently, the matrix F_{dj} defined in (4.2) in the context of model (4.3) now takes the form

$$F_{dj} = \begin{pmatrix} \mathbf{0}'_{p-1} & 1 \\ \mathbf{I}_{p-1} & \mathbf{0}_{p-1} \end{pmatrix} T_{dj}, \quad 1 \leq j \leq n.$$

This leads to a considerable simplification in the analysis, but this simplicity comes at the expense of having a preperiod of experimentation. This model has been sporadically used in the literature, the traditional *noncircular* model being far more popular. For some results on optimality under the circular model, one may refer to Magda (1980) and Kunert (1984b).

4.6.2 Model with Self- and Mixed Carry-Overs

In the models described so far, when a treatment is applied to a subject in a period, the carry-over effect of this treatment in the following period is always the same, irrespective of which treatment follows it. However, in some crossover trials, for example, in medical applications, such a constant form of the carry-over may not be realistic and the carry-over effect of a treatment may depend on whether it is being followed by itself, or by a different treatment. For such situations, Afsarinejad and Hedayat (2002) introduced the *self- and mixed carry-over* model where they assumed that the carry-over effect of a treatment is of two types: if a treatment is followed by itself on a subject, then the carry-over effect of the former treatment is called a *self*-carry-over effect, while if it is followed by a different treatment, then it is called a *mixed* carry-over effect. Afsarinejad and Hedayat (2002) studied only two-period designs under this model, while a study of designs in $\Omega_{t,n,p}$ was developed by Kunert and Stufken (2002). More recently, Kunert and Stufken (2008) studied the optimality of two-treatment designs under a model with self- and mixed carry-over effects. The model with self- and mixed carryovers is as follows:

$$y_{ij} = \begin{cases} \alpha_i + \beta_j + \tau_{d(i,j)} + \nu_{d(i-1,j)} + \epsilon_{ij}, & \text{if } d(i,j) \neq d(i-1,j) \\ \alpha_i + \beta_j + \tau_{d(i,j)} + \chi_{d(i-1,j)} + \epsilon_{ij}, & \text{if } d(i,j) = d(i-1,j) \end{cases}, \quad (4.20)$$

where $\chi_{d(i-1,j)}$ is the self-carry-over effect and $\nu_{d(i-1,j)}$ is the mixed carry-over effect of treatment $d(i-1,j)$, $\nu_{d(0,j)} = \chi_{d(0,j)} = 0$, $1 \leq i \leq p$, $1 \leq j \leq n$, and all other terms in (4.20) are as in (4.1). Analogous to (4.3) and remembering that $d(i,j) \in \{1, \dots, t\}$, the model in (4.20) can equivalently be written as

$$\mathbf{y}_d = \mathbf{P}\boldsymbol{\alpha} + \mathbf{U}\boldsymbol{\beta} + \mathbf{T}_d\boldsymbol{\tau} + \mathbf{G}_d\boldsymbol{\nu} + \mathbf{S}_d\boldsymbol{\chi} + \boldsymbol{\epsilon}, \quad (4.21)$$

where $\boldsymbol{\chi} = (\chi_1, \dots, \chi_t)'$, $\boldsymbol{\nu} = (\nu_1, \dots, \nu_t)'$ and \mathbf{G}_d and \mathbf{S}_d are the design matrices for the mixed carry-over and self-carry-over effects, respectively, all other terms in (4.21) being as in (4.3). Model (4.20) is presented here in the same form as was considered by Hedayat and Afsarinejad (2002). This model remains unaffected if, as in (4.1), a general mean term is included in it. This is because the column space of \mathbf{P} or \mathbf{U} includes the vector $\mathbf{1}_{np}$.

For the simpler case $p = 2$, Afsarinejad and Hedayat (2002) obtained optimal two-period designs under (4.21) by invoking the connection between a block design and a two-period crossover design as given in Theorem 4.1. They proved that a symmetric BIB design with t treatments and block size k can be used to obtain a design that is optimal for direct effects under (4.21) over the subclass of designs in $\Omega_{t,tk,2}$, which are uniform on the first period.

The case $p > 2$ for model (4.21) was studied by Kunert and Stufken (2002). To identify the optimal design in this class, they first found an upper bound on the information matrix for direct effects under (4.21) (in the Loewner sense) and then showed that this bound is attained by a totally balanced design (see Definition 4.5). Next, they maximized the trace of the upper bound and showed that a totally balanced design again attains this maximum. Thus, they established the following result.

Theorem 4.3 For $t \geq 3$ and $3 \leq p \leq 2t$, if a totally balanced design $d^* \in \Omega_{t,n,p}$ exists, then d^* is universally optimal for the estimation of direct effects over $\Omega_{t,n,p}$ under (4.21).

By this theorem, the design shown in Example 4.5 is universally optimal for direct effects over $\Omega_{3,6,4}$ under (4.21).

4.6.3 Model with Direct-versus-Carry-Over Interactions

In the earlier model, the carry-over effect of a treatment was only of two types, the mixed carry-over effect being the same no matter which treatment followed it. One can extend this idea to postulate a model where the treatments allocated to the same subject in two successive periods may have an interaction effect, in addition to the usual direct and carry-over effects. Such a model was proposed by Sen and Mukerjee (1987) and is given by

$$\begin{aligned} y_{1j} &= \mu + \alpha_1 + \beta_j + \tau_{d(1,j)} + \epsilon_{1j}, \quad 1 \leq j \leq n, \\ y_{ij} &= \mu + \alpha_i + \beta_j + \tau_{d(i,j)} + \rho_{d(i-1,j)} + \gamma_{d(i,j),d(i-1,j)} + \epsilon_{ij}, \\ & \quad 2 \leq i \leq p, \quad 1 \leq j \leq n, \end{aligned} \tag{4.22}$$

where $\gamma_{d(i,j),d(i-1,j)}$ is the interaction effect between treatments $d(i,j)$ and $d(i-1,j)$, $d(i,j) \in \{1, \dots, t\}$, and all other terms are as in (4.1). For $t=2$, this model is equivalent to model (4.20).

Sen and Mukerjee (1987) showed that strongly balanced uniform designs are universally optimal for *direct* effects under the nonadditive model (4.22). However, this result does not have an exact counterpart for the estimation of carry-over effects. Sen and Mukerjee (1987) proved that a strongly balanced uniform design satisfying certain extra combinatorial conditions is universally optimal for estimation of carry-over effects under the model (4.22).

Further results on optimal crossover designs under (4.22) were obtained recently by Park et al. (2011). They considered a particular class of strongly balanced designs with $n = t^2$ subjects, which are uniform on the periods, and obtained a lower bound for the A -efficiency

of the designs for estimating the direct effects. They then showed that such designs are highly efficient for any number of periods p , $2 \leq p \leq 2t$.

4.6.4 Model with Carry-Over Proportional to Direct Effect

In some crossover experiments, it is believed that the carry-over effect of a treatment is proportional to its direct effect, thus requiring yet another modification of the traditional model. The constant of proportionality may be either positive or negative, but it is generally unknown. Cross (1973) gave an example of such a situation where subjects were asked to rate the loudness levels of different sound stimuli, and it was found that subjects generally gave a higher rating to a stimulus when it was preceded by a loud sound and gave a lower rating to the same sound when it was preceded by a soft sound. Thus, here, the constant of proportionality is positive. Schifferstein and Oudejans (1996) described another experiment where subjects were asked to rate the saltiness of several saline solutions. It was observed that the subjects rated a solution to be less salty if it was immediately preceded by a solution with high salt concentration, while they rated the same solution to be more salty when preceded by one with low salt concentration. Here, the constant of proportionality is negative.

The model where carry-over effects are assumed to be proportional to direct effects is given by

$$y_{ij} = \mu + \alpha_i + \beta_j + \tau_{d(i,j)} + \gamma\tau_{d(i-1,j)} + \epsilon_{ij}, \quad 1 \leq i \leq p, \quad 1 \leq j \leq n, \quad (4.23)$$

where γ is the constant of proportionality and all other terms are as in (4.1). We also assume that $t \geq 3$, because for $t = 2$, the model reduces to (4.1).

Even though this model has fewer parameters than (4.1), it is technically much harder to analyze as it is nonlinear in τ and γ , both being unknown. In order to linearize (4.23), Kempton et al. (2001) used a Taylor series expansion about τ_0 and γ_0 , the true values of τ and γ , respectively. Then the information matrix for direct effects can be obtained under the linearized model. However, this information matrix now depends on the unknowns τ_0 and γ_0 . Bailey and Kunert (2006) and Bose and Stufken (2007) also studied this model, and we refer to them for expressions of the information matrices as obtained by them.

To overcome the aforesaid difficulty, Kempton et al. (2001) considered the distribution of possible values of τ_0 and studied the performance of a design based on (a) the \bar{A} -criterion, which is the averaged version of the usual A -criterion, the average being taken over a multivariate normal distribution of τ_0 with zero mean vector and dispersion matrix $I_t - t^{-1}J_t$, and (b) the IA -criterion, where the A -criterion is averaged over both τ_0 and γ_0 , with τ_0 distributed as in (a), γ_0 having the uniform distribution on $[-1,1]$, and τ_0 and γ_0 being independent. Note that the \bar{A} -criterion is a local optimality criterion as it depends on γ_0 . They proved the following results:

1. Let $d \in \Omega_{t,n,p}$, where $n = \mu_1 t$, $\mu_1 \geq 1$, $p = t \geq 3$, be a balanced uniform design. Then, under model (4.23), d is \bar{A} -optimal for the estimation of direct effects over the class of all uniform designs in $\Omega_{t,n,p}$ for all γ_0 .
2. Let $d \in \Omega_{t,n,p}$, where $n = \mu_1 t$, $\mu_1 \geq 1$, $p = t + 1$, $t \geq 3$, be a strongly balanced design that is uniform on periods and uniform on subjects in the first $p - 1$ periods. Then, under model (4.23), d is IA -optimal for the estimation of direct effects over $\Omega_{t,n,p}$.

Additional optimality results under (4.23) were obtained by Bailey and Kunert (2006). Among other things, they showed that if d^* is a totally balanced design with $t \geq p \geq 3$ or $t \geq 3, p = 2$, then for all $\gamma_0 \in [-1, 1]$, d^* is \bar{A} -optimal for direct effects over all designs in $\Omega_{t,n,p}$, which do not assign the same treatment to successive periods in any subject.

Bose and Stufken (2007) obtained optimal designs under the model (4.23) when γ is known and not necessarily restricted to the interval $[-1, 1]$. Under this assumption, the model (4.23) becomes linear, and the more stringent universal optimality criterion can be used for obtaining optimal designs for given γ . While we refer the reader to the original source for details, we give in the following some examples of universally optimal designs under (4.23).

Example 4.13

Let $t = 3 = n = p$. The following designs are universally optimal for direct effects over $\Omega_{3,3,3}$ for any γ in the intervals $(0.52, 11.48)$, $(-4.73, -1.27)$, and $(-1.27, 0.52)$, respectively:

$$\begin{matrix} 1 & 2 & 3 & , & 1 & 2 & 3 & , & 1 & 2 & 3 \\ 2 & 3 & 1 & , & 2 & 3 & 1 & , & 2 & 3 & 1 \\ 2 & 3 & 1 & , & 1 & 2 & 3 & , & 3 & 1 & 2 \end{matrix}$$

4.6.5 Mixed Effects Models and Models with Correlated Errors

Several authors considered models with random subject effects. This leads to a mixed effects version of model (4.1). With $p = 2$ periods, Carriere and Reinsel (1993) considered the situation where the t^2 possible treatment sequences are assigned to the subjects at random, the l th sequence being assigned to n_l subjects, $1 \leq l \leq t^2$. Accordingly, they modified model (4.1) to the following form:

$$y_{ijl} = \mu + \alpha_i + \tau_{d(i,l)} + \rho_{d(i-1,l)} + \beta_{jl} + \epsilon_{ijl}, \tag{4.24}$$

where $d(i, l)$ is the treatment in the period i in the sequence l , y_{ijl} is the response obtained in period i from the j th subject assigned to the sequence l , β_{jl} is the random subject effect of the j th subject assigned to sequence l , and $\mu, \alpha_i, \tau_s, \rho_s$ are as in model (4.1). The random subject effects and the errors ϵ_{ijl} are assumed to be mutually uncorrelated random variables with means zero and variances σ_β^2 and σ^2 , respectively, $1 \leq i \leq 2, 1 \leq j \leq n_l, 1 \leq l \leq t^2$. Thus, (4.24) is a mixed effects model.

Then for any $d \in \Omega_{t,n,2}$, the information matrix for the direct effects, C_d , under the model (4.24) is given by

$$\sigma^2 C_d = (1 + v)^{-1} (R_d - v^2 \bar{R}_d - n^{-1} (1 - v^2) \bar{r}_d \bar{r}'_d - Z_d \bar{R}_d^- Z'_d), \tag{4.25}$$

where $v = \sigma_\beta^2 / (\sigma^2 + \sigma_\beta^2)$ and $R_d, \bar{R}_d, \bar{r}_d$, and Z_d are as defined in Section 4.2.2. Using (4.25), Carriere and Reinsel (1993) proved, among other things, that a strongly balanced two-period design that is uniform on periods is universally optimal for direct effects over $\Omega_{t,n,2}$.

A mixed effects model for general $p(\geq 2)$ was studied by Mukhopadhyay and Saha (1983), who assumed that the β_j 's in model (4.1) are mutually uncorrelated random variables with means zero and constant variances, these being also uncorrelated with the

error variable. Under this mixed effects model, they studied the optimality of balanced and strongly balanced uniform designs, when the variances are known. This mixed effects model has been considered more recently by Hedayat et al. (2006) who obtained optimal and efficient crossover designs under such a model.

As has been mentioned in Section 4.3.3, in the context of crossover designs, the assumption of independent errors may not be realistic in some situations and a model with *correlated errors* may seem more appealing. To incorporate the correlations among observations within subjects, we may modify the traditional model (4.3) to

$$y_d = X_d\theta + \epsilon, E(\epsilon) = \mathbf{0}, \text{Var}(\epsilon) = I_n \otimes V, \tag{4.26}$$

where V is a positive definite matrix of order p , representing the variance-covariance matrix of the errors corresponding to observations from the same subject, and all other terms are as in (4.3). So now the responses from different subjects are uncorrelated, while those from the same subject can be correlated, these correlations being the same for all subjects. We may take various forms of V to reflect the actual error structure in different situations; for $V = I_p$, we are back to the traditional model. After some algebra, it can be seen that the information matrix for direct effects under model (4.26) for a design d is given by

$$C_d = T'_d(I_n \otimes V^{-1/2})P^\perp \left((I_n \otimes V^{-1/2}) (P U F_d) \right) (I_n \otimes V^{-1/2}) T_d. \tag{4.27}$$

As in Section 4.3.3, for the special case $p = 2$, the problem of finding an optimal design under a correlated model with $V = (1 - \rho)I + \rho J$ and given ρ is equivalent to that of finding an optimal design under (4.3), where ρ is the correlation coefficient between a pair of observations from the same subject. For $p > 2$, the task of finding optimal/efficient designs under the model (4.26) becomes simplified in the particular case $t = 2$. In this case, using the approximate theory approach, Kushner (1997a) gave a set of necessary and sufficient conditions for a dual-balanced design (Section 4.3.4) to be universally optimal for direct effects. He also gave an expression for computing the efficiency of a dual-balanced design for direct effects. Some authors have assumed that the errors within each subject follow a stationary first-order autoregressive process. This reflects the belief that the correlation between the observations from different periods on the same subject diminishes with time. So V is of the form

$$V = \begin{pmatrix} 1 & \rho & \rho^2 & \dots & \rho^{p-1} \\ \rho & 1 & \rho & \dots & \rho^{p-2} \\ \rho^2 & \rho & 1 & \dots & \rho^{p-3} \\ \dots & \dots & \dots & \dots & \dots \\ \rho^{p-1} & \dots & \rho^2 & \rho & 1 \end{pmatrix}, \tag{4.28}$$

or a multiple of this. For model (4.26) with such a V , several authors (see e.g., Matthews, 1987; Kunert, 1991; and Kushner, 1997a) have studied the problem of obtaining efficient or optimal designs, mainly using the approximate theory.

For example, if we consider the following two pairs of dual sequences,

$$\begin{matrix} 1 & 2 & 1 & 2 \\ 1 & 2 & 2 & 1 \\ 2 & 1 & 2 & 1 \end{matrix},$$

then a design that allocates a proportion of $\rho/(3\rho - 1)$ subjects to each of the first two sequences and a proportion of $(\rho - 1)/(2(3\rho - 1))$ to each of the last two sequences is universally optimal for direct effects for all $\rho \in (-1, 0]$. Similarly, if we consider the following two pairs of dual sequences $(1, 2, 1)', (2, 1, 2)', (1, 2, 2)', (2, 1, 1)'$, then a design that allocates a proportion of $\rho^2/(3 + \rho)^2$ subjects to each of the first two sequences and a proportion of $(6\rho - \rho^2 + 9)/(2(3 + \rho^2))$ to each of the last two sequences is optimal for direct effects for all $\rho \in (0, \rho_1)$, where $\rho_1 \approx 0.464$ (Matthews 1987). In practice, however, ρ is unknown. Taking cognizance of this fact, Matthews (1987) showed that if we simply consider the four sequences displayed and allocate each to one subject, then the resulting design in $\Omega_{2,4,3}$ has an efficiency of at least 90% when $-0.8 \leq \rho \leq 0.8$. Thus, for a large range of possible ρ values, this dual-balanced design is efficient, even when errors are correlated. He also gave tables of efficiencies of several dual-balanced designs with three and four periods and for several values of ρ . These tables can be used to choose an efficient design for experimentation in situations where the parameter combinations are such that an optimal design is not available. Kunert (1991), too, identified efficient designs and gave a method for constructing such efficient designs for any given value of ρ and p . Again, these efficient designs are dual-balanced designs.

For arbitrary t and p , Kushner (1997b, 1998) gave a general approximate theory approach for identifying optimal designs under the model (4.26). Kunert (1985), Gill (1992), Donev (1998), among others, studied optimality under the exact theory approach and under various assumptions. Martin and Eccleston (1998) studied variance-balanced crossover designs, which allow all elementary contrasts (Chapter 3) of direct effects to be estimated with a constant variance and also ensure the same for all elementary contrasts of carry-over effects under the model (4.26). They showed that such a design in $\Omega_{t,n,p}$ can be constructed from an orthogonal array of type I of strength two. Note that a type I orthogonal array, $OA_I(n, t^p, 2)$, of strength two is an $n \times p$ array with entries from a set of t symbols, such that any $n \times 2$ subarray contains each ordered pair of distinct symbols equally often as a row, precisely $n/(t(t - 1))$ times (see Chapter 9). An example is given in the following.

Example 4.14

A type I orthogonal array $OA_I(6, 3^3, 2)$ is shown as follows in transposed form:

1	2	3	1	2	3
2	3	1	3	1	2
3	1	2	2	3	1

Kunert and Martin (2000) obtained a general result on the optimality of designs given by an $OA_I(n, t^p, 2)$ under model (4.26). Their result is as follows:

Theorem 4.4 For $t \geq p > 2$, let $d^* \in \Omega_{t,n,p}$ be a crossover design given by a type I orthogonal array, $OA_I(n, t^p, 2)$, with columns of the array representing the periods and the rows representing the subjects of d^* . Then under (4.26) where V is any known positive definite matrix, d^* is universally optimal for the estimation of direct effects over the class of designs which are binary on subjects in the sense that a treatment is applied to a subject at most once.

Kunert and Martin (2000) also showed that these designs are quite efficient even over the general class $\Omega_{t,n,p}$. Moreover, these designs often require fewer subjects than that required

by the optimum design obtained through an approximate theory approach for the same number of treatments and periods.

Hedayat and Yan (2008) extended the self- and mixed carry-over model (4.21) to one with correlated errors as in model (4.26). They considered two forms of V , one where the errors within each subject follow a stationary first-order autoregressive process as in (4.28) and another where they follow a stationary first-order moving average process. In the second case,

$$V = I_p + \rho W, \quad \rho \in (-1/2, 1/2), \quad (4.29)$$

with $W = ((w_{ij}))$, $w_{ij} = 1$ if $|i - j| = 1$ and $= 0$; otherwise, $1 \leq i, j \leq p$. They studied the performance of the designs in $\Omega_{t,n,p}$ given by an $OA_I(n, t^p, 2)$ and proved the following theorem:

Theorem 4.5 For $p = 3$ and $t \geq 3$, let $d^* \in \Omega_{t,n,3}$ be a design given by a type I orthogonal array, $OA_I(n, t^p, 2)$. Then d^* is universally optimal for direct effects over $\Omega_{t,n,3}$ under a model with self- and mixed carry-over effects and with variance-covariance structure of errors within each subject given by either

- (i) V as in (4.28) and any $\rho \in (-1, 1)$, or
- (ii) V as in (4.29) and any $\rho \in (-1/2, 1/2)$.

Example 4.15

The array given in Example 4.14 may be looked upon as a design in $\Omega_{3,6,3}$. By Theorem 4.5, this design is universally optimal for direct effects in $\Omega_{3,6,3}$ under the model (4.26) with correlated errors, the error structure being given by either (4.28) or (4.29) and the correlations as specified in (i) and (ii).

4.7 Other Advances

4.7.1 Crossover Trials for Comparing Treatments versus Control

A problem that arises often in practice concerns the evaluation of the performance of a number of *test* treatments vis-a-vis a standard treatment, called *control*. The test treatments, for instance, could be a number of new drugs, whose efficacy has to be evaluated relative to an established drug, which acts as the control treatment. For direct effects, the parametric function of interest in the present context is the contrast vector

$$\begin{pmatrix} \tau_0 - \tau_1 \\ \dots \\ \tau_0 - \tau_t \end{pmatrix} = \begin{pmatrix} 1 & -1 & 0 & \dots & 0 \\ 1 & 0 & -1 & \dots & 0 \\ \vdots & & & & \\ 1 & 0 & 0 & \dots & -1 \end{pmatrix} \boldsymbol{\tau} = \mathbf{B}\boldsymbol{\tau} \text{ (say),}$$

where τ_0 is the direct effect of the control treatment and, for $1 \leq i \leq t$ and τ_i is the direct effect of the i th test treatment. The contrasts of interest for carry-over effects can be defined similarly. Let $\Omega_{t+1,n,p}$ be the class of all crossover designs involving t test treatments and a control. For a design $d \in \Omega_{t+1,n,p}$, if \mathcal{I}_d is the information matrix for $B\tau$ under the model (4.3), then it can be shown that

$$\mathcal{I}_d = \bar{I}' C_d \bar{I},$$

where $\bar{I} = (0_t, I_t)'$ and C_d is as given by (4.8) under a design involving $t + 1$ treatments.

For the test-control experiments, the A - and MV -optimality (see Chapter 3) criteria seem to be natural and are frequently used. A design $d^* \in \Omega_{t+1,n,p}$ is A -optimal for $B\tau$ if it minimizes $\text{trace}(\mathcal{I}_d^{-1})$ over $\Omega_{t+1,n,p}$ and is MV -optimal if it minimizes $\max \{\text{Var}(\hat{\tau}_0 - \hat{\tau}_i) : 1 \leq i \leq t\}$ where, for $1 \leq i \leq t$, $\hat{\tau}_0 - \hat{\tau}_i$ is the best linear unbiased estimator of $\tau_0 - \tau_i$. The problem of finding A - and MV -optimal crossover designs has been addressed by several authors and we describe in the following some of their results.

Majumdar (1988) showed that if $t = c^2$ for some positive integer c and $d_0 \in \Omega_{c^2+c,n,p}$ is a strongly balanced uniform crossover design, then a design d^* obtained from d_0 by replacing each of the treatment labels $c^2 + 1, c^2 + 2, \dots, c^2 + c$, by the control treatment label 0, keeping everything else unchanged, is an A - and MV -optimal design for direct effects for comparing c^2 test treatments with a control, under the model (4.3). Following the approach of Majumdar (1988), Ting (2002) also obtained additional optimal/efficient crossover designs for comparing several test treatments with a control. Hedayat and Zhao (1990) considered two-period designs for the problem, and starting from designs as given by Theorem 4.2 (Section 4.5.4), they obtained A - and MV -optimal designs when the number of test treatments is c^2 for some integer c . Following Hedayat and Zhao (1990), for example, the design in Example 4.16 with four test treatments and a control is A - and MV -optimal for direct effects over $\Omega_{4+1,18,2}$, where the control treatment is labeled 0. This design is obtained from design d_{13} in Example 4.12 by replacing the treatment labels 5 and 6 by the control treatment label 0, keeping everything else unchanged.

Example 4.16

The following design is A - and MV -optimal for direct effects for comparing control 0 to $t = 4$ test treatments:

1	1	1	1	1	1	2	2	2	2	2	2	3	3	3	3	3	3
1	2	3	4	0	0	1	2	3	4	0	0	1	2	3	4	0	0

For some other results on two-period designs for comparing test treatments with a control, see Koch et al. (1989) and Hedayat and Zhao (1990).

The problem of finding optimal designs when $3 \leq p \leq t + 1$ has been addressed by Hedayat and Yang (2005, 2006), Yang and Park (2007), and Yang and Stufken (2008). Hedayat and Yang (2005) defined a class of designs called *totally balanced test-control incomplete crossover designs* and showed that if such a design satisfies a certain additional condition, then it is A - and MV -optimal over the subclass of designs for which (1) the control treatment appears equally often in the p periods and (2) no treatment precedes itself. Hedayat and Yang (2005) also gave construction procedures of these optimal designs. Two such designs are given in the next examples.

Example 4.17

Suppose $t = 4$ test treatments are to be compared with a control. Then, the following design, obtained on replacing the treatment symbol 5 by the control treatment label 0 in the uniform design shown in Example 4.9, is A - and MV -optimal over the subclass of $\Omega_{4+1,10,5}$ satisfying (1) and (2) earlier:

1	2	3	4	0	3	4	0	1	2
0	1	2	3	4	4	0	1	2	3
2	3	4	0	1	2	3	4	0	1
4	0	1	2	3	0	1	2	3	4
3	4	0	1	2	1	2	3	4	0

Example 4.18

Suppose $t = 3$ test treatments are to be compared with a control. Then, the following design is A - and MV -optimal over the subclass of $\Omega_{3+1,10,5}$ satisfying (1) and (2) earlier:

0	0	0	2	3	1	2	3	1
1	2	3	0	0	0	1	2	3
2	3	1	1	2	3	0	0	0

For some more results on optimal crossover designs for test–control comparisons, see Hedayat and Yang (2005, 2006). Yang and Park (2007) obtained designs that are optimal or efficient over a wider class of competing designs, but their results are only for three-period designs. Yang and Stufken (2008) obtained further efficient and optimal crossover designs under a wide variety of models including (4.3), (4.20), and some variants of these. They also gave two algorithms for generating highly efficient designs under several models.

Extending the results of Hedayat and Yang (2005) to the case of random subject effects, Yan and Locke (2010) showed that under the model (4.3), totally balanced test–control designs with $p = 3, 4, 5$ periods are highly efficient in the class of designs in which the control treatment appears equally often in all periods and no treatment is immediately preceded by itself.

Before concluding this section, we comment on the replication numbers of the control vis-a-vis test treatments in the optimal designs studied here. Commonly, in the absence of carry-over effects, optimal designs for control versus test treatments have the control replicated more often than the test treatments. However, this is not necessarily true for crossover designs. For instance, in Example 4.16 with only two periods, the overall replication of the control is less than that of each test treatment, but in the second period, the control is replicated twice as many times as each test treatment. Again, in Example 4.17, the control treatment is replicated as often as each test treatment, while in Example 4.18, the control is replicated more often. This may be attributed to the versatility of the approaches underlying the identification and construction of optimal designs. Nevertheless, for large n , the replication number of the control could possibly exceed that of the test treatments in some optimal crossover designs.

4.7.2 Optimal Designs When Some Subjects Drop Out

In practice, it may happen that a crossover trial cannot be performed for the initially planned p periods for all the subjects. Such a situation arises, for example, in clinical trials where certain patients drop out from the study before the entire sequence of p treatments

assigned to them can be completed. When subjects drop out before the trial is completed, the final *implemented* design is a truncated version of the originally *planned* design. If these two designs are denoted by d_{imp} and d_{plan} , respectively, then d_{imp} may not remain optimal/efficient even if d_{plan} is an optimal design and, in certain extreme cases, d_{imp} may not even remain connected. Therefore, while choosing d_{plan} , the possibility of subjects dropping out has to be taken into consideration.

Low et al. (1999, 2003) suggested a computer-intensive method to ascertain the robustness of d_{plan} with $p > 2$ when the subjects drop out at random. They used the means and standard deviations of certain performance measures based on the A - and MV -optimality criteria and used these to assess the performance of d_{plan} . This line of work was further enriched by Majumdar et al. (2008) who started with a d_{plan} , which is a balanced uniform design in $\Omega_{t,\mu_1 t,t}$, and explored the situation where all subjects remain in the experiment in the first $t - u$ periods, and then start dropping out at random, $1 \leq u \leq t - 1$. The design consisting of the first $t - u$ periods was called by them as minimal and denoted by d_{min} . They obtained a sufficient condition for d_{min} , and hence d_{imp} , to remain connected and also gave an upper bound on the maximum loss of efficiency due to subject dropouts in the last u periods when d_{plan} is a balanced uniform design in $\Omega_{t,\mu_1 t,t}$. The following example illustrates some of their findings.

Example 4.19

Consider the following balanced uniform designs for $t = 3, 4, 5$:

$$d_{14} \equiv \begin{matrix} 1 & 2 & 3 & 2 & 3 & 1 \\ 3 & 1 & 2 & 3 & 1 & 2 \\ 2 & 3 & 1 & 1 & 2 & 3 \end{matrix}, \quad d_{15} \equiv \begin{matrix} 4 & 1 & 2 & 3 \\ 1 & 2 & 3 & 4 \\ 3 & 4 & 1 & 2 \\ 2 & 3 & 4 & 1 \end{matrix}, \quad d_{16} \equiv \begin{matrix} 1 & 2 & 3 & 4 & 5 & 3 & 4 & 5 & 1 & 2 \\ 5 & 1 & 2 & 3 & 4 & 4 & 5 & 1 & 2 & 3 \\ 2 & 3 & 4 & 5 & 1 & 2 & 3 & 4 & 5 & 1 \\ 4 & 5 & 1 & 2 & 3 & 5 & 1 & 2 & 3 & 4 \\ 3 & 4 & 5 & 1 & 2 & 1 & 2 & 3 & 4 & 5 \end{matrix}.$$

If these designs are used as d_{plan} , then Majumdar et al. (2008) showed that if no observation is taken in the last period, the d_{imp} designs arising out of d_{14} and d_{16} remain connected but that arising out of d_{15} becomes disconnected. A similar observation was made by Low et al. (1999) too.

Bose and Bagchi (2008) studied the optimality aspects of the designs d_{min} when d_{plan} belongs to a class of designs, say, \mathcal{D}_1 , consisting of *locally balanced crossover* designs, introduced by Anderson and Preece (2002). Among other things, they showed that a design $d_{plan} \in \mathcal{D}_1$ is itself universally optimal for direct and carry-over effects over the binary subclass in $\Omega_{t,n,t}$, and furthermore, d_{min} obtained from a member of \mathcal{D}_1 remains optimal over the binary subclass in $\Omega_{t,n,t-u}$ when d_{min} consists of $t - u \geq 3$ periods.

Example 4.20

For example, the design d_{min} obtained from the following design with $t = 5 = p$, $n = 20$ is optimal for $t - u \geq 3$:

$$\begin{matrix} 1 & 2 & 3 & 4 & 5 & 2 & 3 & 4 & 5 & 1 & 1 & 2 & 3 & 4 & 5 & 2 & 3 & 4 & 5 & 1 \\ 2 & 3 & 4 & 5 & 1 & 4 & 5 & 1 & 2 & 3 & 5 & 1 & 2 & 3 & 4 & 5 & 1 & 2 & 3 & 4 \\ 5 & 1 & 2 & 3 & 4 & 5 & 1 & 2 & 3 & 4 & 2 & 3 & 4 & 5 & 1 & 4 & 5 & 1 & 2 & 3 \\ 3 & 4 & 5 & 1 & 2 & 1 & 2 & 3 & 4 & 5 & 4 & 5 & 1 & 2 & 3 & 3 & 4 & 5 & 1 & 2 \\ 4 & 5 & 1 & 2 & 3 & 3 & 4 & 5 & 1 & 2 & 3 & 4 & 5 & 1 & 2 & 1 & 2 & 3 & 4 & 5 \end{matrix}.$$

For some more results on efficient crossover designs when subjects drop out at random, see Zhao and Majumdar (2012). In a recent paper, Zheng (2013) obtained necessary and sufficient conditions for a crossover design to be universally optimal in approximate design theory in the presence of subject dropout. He also provided an algorithm to derive efficient exact designs.

4.7.3 Optimal Designs via Approximate Theory

Most of the results on optimal crossover designs described earlier concern exact designs where each subject is allocated a sequence of treatments over the p periods and, thus, the number of subjects assigned to a treatment sequence is a nonnegative integer. It follows then that the proportion of subjects receiving a treatment sequence is of the form u/n where $0 \leq u \leq n$ and n is the total number of subjects. Because of the discreteness of u , this approach does not allow the use of techniques based on calculus, and one has to employ combinatorial arguments to arrive at optimal designs. In contrast to the exact theory, one can often achieve considerable simplicity by allowing the aforementioned proportions to vary continuously over $[0, 1]$, such that the sum of these proportions over all treatment sequences is unity. As a result, one now has a continuous design framework, which allows the development of an approximate design theory and methods based on calculus can be employed to determine these in an optimal fashion.

For $t = 2$ treatments, the approximate design theory was used by Laska et al. (1983), Matthews (1987, 1990), and Kushner (1997a), and we have already described some of the results obtained by these authors earlier in this chapter. A more detailed study of optimal crossover designs using the approximate theory was made by Kushner (1997b, 1998), who obtained optimality results for direct effects with arbitrary number of treatments under a correlated errors model as given by (4.26). Note that the approach of Kushner (1997b) can be employed to arrive at optimal designs under the uncorrelated errors model as well. Based on the methods of Kushner (1997b), Bose and Shah (2005) obtained optimal designs for the estimation of carry-over effects under (4.3). A detailed exposition of these methods based on approximate theory is beyond the scope of this chapter, and an interested reader may refer to the aforementioned references for more details.

Acknowledgment

The work of A. Dey was supported by the National Academy of Sciences, India.

References

- Afsarinejad, K. (1983). Balanced repeated measurements designs. *Biometrika* **70**, 199–204.
- Afsarinejad, K. (1985). Optimal repeated measurements designs. *Statistics* **16**, 563–568.
- Afsarinejad, K. and A. S. Hedayat (2002). Repeated measurements designs for a model with self and carryover effects. *J. Stat. Plan. Infer.* **106**, 449–459.
- Anderson, I. and D. A. Preece (2002). Locally balanced change-over designs. *Util. Math.* **62**, 33–59.

- Armitage, P. and M. Hills (1982). The two-period crossover trial. *Statistician* **31**, 119–131.
- Atkinson, G. F. (1966). Designs for sequences of treatments with carry-over effects. *Biometrics* **22**, 292–309.
- Bailey, R. A. and J. Kunert (2006). On optimal crossover designs when carryover effects are proportional to direct effects. *Biometrika* **93**, 613–625.
- Bate, S. T. and B. Jones (2006). The construction of nearly balanced and nearly strongly balanced uniform cross-over designs. *J. Stat. Plan. Infer.* **136**, 3248–3267.
- Bate, S. T. and B. Jones (2008). A review of uniform cross-over designs. *J. Stat. Plan. Infer.* **138**, 336–351.
- Berenblut, I. I. (1964). Change-over designs with complete balance for first residual effects. *Biometrics* **20**, 707–712.
- Bose, M. and S. Bagchi (2008). Optimal crossover designs under premature stopping. *Util. Math.* **75**, 273–285.
- Bose, M. and A. Dey (2009). *Optimal Crossover Designs*. Hackensack, NJ: World Scientific.
- Bose, M. and K. R. Shah (2005). Estimation of residual effects in repeated measurements designs. *Calcutta Stat. Assoc. Bull. Special Volume* **56**, 125–143.
- Bose, M. and J. Stufken (2007). Optimal crossover designs when carryover effects are proportional to direct effects. *J. Stat. Plan. Infer.* **137**, 3291–3302.
- Carriere, K. C. and R. Huang (2000). Crossover designs for two-treatment clinical trials. *J. Stat. Plan. Infer.* **87**, 125–134.
- Carriere, K. C. and G. C. Reinsel (1992). Investigation of dual-balanced crossover designs for two treatments. *Biometrics* **48**, 1157–1164.
- Carriere, K. C. and G. C. Reinsel (1993). Optimal two-period repeated measurements designs with two or more treatments. *Biometrika* **80**, 924–929.
- Cheng, C.-S. and C. F. Wu (1980). Balanced repeated measurements designs. *Ann. Stat.* **8**, 1272–1283. Corrigendum: *ibid.* **11** (1983), 349.
- Cochran, W. G. (1939). Long-term agricultural experiments. *J. R. Stat. Soc.* **6B**, 104–148.
- Cochran, W. G., K. M. Autrey and C. Y. Cannon (1941). A double change-over design for dairy cattle feeding experiments. *J. Dairy Sci.* **24**, 937–951.
- Cochran, W. G. and G. M. Cox (1957). *Experimental Designs*, 2nd ed. New York: Wiley.
- Constantine, G. and A. Hedayat (1982). A construction of repeated measurements designs with balance for residual effects. *J. Stat. Plan. Infer.* **6**, 153–164.
- Cross, D. V. (1973). Sequential dependencies and regression in psychophysical judgments. *Percept. Psychophys.* **14**, 547–552.
- Donev, A. (1998). Crossover designs with correlated observations. *Stat. Med.* **8**, 249–262.
- Fieller, E. C. (1940). The biological standardisation of insulin. *J. R. Stat. Soc., Suppl.* **1**, 1–64.
- Finney, D. J. (1956). Cross-over designs in bioassay. *Proc. R. Soc.* **B 145**, 42–61.
- Freeman, G. H. (1959). The use of the same experimental material for more than one set of treatments. *Appl. Stat.* **8**, 13–20.
- Gill, P. S. (1992). Balanced change-over designs for autocorrelated observations. *Aust. J. Stat.* **34**, 415–420.
- Grizzle, J. E. (1965). The two-period change-over design and its use in clinical trials. *Biometrics* **21**, 469–480. Corrigendum: *ibid.* **30** (1974), 727.
- Hedayat, A. and K. Afsarinejad (1975). Repeated measurements designs, I. In *A Survey of Statistical Designs and Linear Models* (J. N. Srivastava, Ed.). Amsterdam, the Netherlands: North-Holland, pp. 229–242.
- Hedayat, A. and K. Afsarinejad (1978). Repeated measurements designs, II. *Ann. Stat.* **6**, 619–628.
- Hedayat, A. S., J. Stufken, and M. Yang (2006). Optimal and efficient crossover designs when subject effects are random. *J. Am. Stat. Assoc.* **101**, 1031–1038.
- Hedayat, A. S. and Z. Yan (2008). Crossover designs based on type I orthogonal arrays for a self and mixed carryover effects model with correlated errors. *J. Stat. Plan. Infer.* **138**, 2201–2213.
- Hedayat, A. S. and M. Yang (2003). Universal optimality of balanced uniform crossover designs. *Ann. Stat.* **31**, 978–983.

- Hedayat, A. S. and M. Yang (2004). Universal optimality for selected crossover designs. *J. Am. Stat. Assoc.* **99**, 461–466.
- Hedayat, A. S. and M. Yang (2005). Optimal and efficient crossover designs for comparing test treatments with a control treatment. *Ann. Stat.* **33**, 915–943.
- Hedayat, A. S. and M. Yang (2006). Efficient crossover designs for comparing test treatments with a control treatment. *Discrete Math.* **306**, 3112–3124.
- Hedayat, A. and W. Zhao (1990). Optimal two-period repeated measurements designs. *Ann. Stat.* **18**, 1805–1816. Corrigendum: *ibid.* **20** (1992), 619.
- Higham, J. (1998). Row-complete Latin squares of every composite order exist. *J. Combin. Des.* **6**, 63–77.
- Hills, M. and P. Armitage (1979). The two-period crossover clinical trial. *Br. J. Clin. Pharmacol.* **8**, 7–20.
- Issac, P. D., A. M. Dean and T. Ostrom (2001). Generating pairwise balanced Latin squares. *Stat. Appl.* **3**, 25–46.
- Jones, B. and M. G. Kenward (2003). *Design and Analysis of Cross-Over Trials*, 2nd ed. London, U.K.: Chapman and Hall.
- Jones, B., J. Wang, P. Jarvis and W. Byrom (1999). Design of cross-over trials for pharmacokinetic studies. *J. Stat. Plan. Infer.* **78**, 307–316.
- Kempton, R. A., S. J. Ferris and O. David (2001). Optimal change-over designs when carryover effects are proportional to direct effects of treatments. *Biometrika* **88**, 391–399.
- Kenward, M. G. and B. Jones (1998). Crossover trials. In *Encyclopedia of Statistical Sciences*, Update Vol. 2 (S. Kotz, C. B. Read and D. L. Banks, Eds.). New York: Wiley, pp. 167–175.
- Kershner, R. P. and W. T. Federer (1981). Two-treatment crossover designs for estimating a variety of effects. *J. Am. Stat. Assoc.* **76**, 612–619.
- Kiefer, J. (1959). Optimum experimental designs (with discussion). *J. R. Stat. Soc. Ser. B* **21**, 272–319.
- Kiefer, J. (1975). Construction and optimality of generalized Youden designs. In *A Survey of Statistical Designs and Linear Models* (J. N. Srivastava, Ed.). Amsterdam, the Netherlands: North-Holland, pp. 333–353.
- Koch, G. G., I. A. Amara, B. W. Brown, T. Colton and D. B. Gillings (1989). A two-period crossover design for the comparison of two active treatments and placebo. *Stat. Med.* **8**, 487–504.
- Kunert, J. (1983). Optimal design and refinement of the linear model with applications to repeated measurements designs. *Ann. Stat.* **11**, 247–257.
- Kunert, J. (1984a). Optimality of balanced uniform repeated measurements designs. *Ann. Stat.* **12**, 1006–1017.
- Kunert, J. (1984b). Designs balanced for circular carry-over effects. *Commun. Stat. Theory Methods.* **13**, 2665–2671.
- Kunert, J. (1985). Optimal repeated measurements designs for correlated observations and analysis by weighted least squares. *Biometrika* **72**, 375–389.
- Kunert, J. (1991). Cross-over designs for two treatments and correlated errors. *Biometrika* **78**, 315–324.
- Kunert, J. and R. J. Martin (2000). Optimality of type I orthogonal arrays for cross-over models with correlated errors. *J. Stat. Plan. Infer.* **87**, 119–124.
- Kunert, J. and J. Stufken (2002). Optimal crossover designs in a model with self and mixed carryover effects. *J. Am. Stat. Assoc.* **97**, 898–906.
- Kunert, J. and J. Stufken (2008). Optimal crossover designs for two treatments in the presence of mixed and self carryover effects. *J. Am. Stat. Assoc.* **103**, 1641–1647.
- Kushner, H. B. (1997a). Optimality and efficiency of two-treatment repeated measurements designs. *Biometrika* **84**, 455–468. Corrigendum: *ibid.* **86** (1999), 234.
- Kushner, H. B. (1997b). Optimal repeated measurements designs: The linear optimality equations. *Ann. Stat.* **25**, 2328–2344. Corrigendum: *ibid.* **27** (1999), 2081.
- Kushner, H. B. (1998). Optimal and efficient repeated measurements designs for uncorrelated observations. *J. Am. Stat. Assoc.* **93**, 1176–1187.
- Laska, E. and M. Meisner (1985). A variational approach to optimal two-treatment crossover designs: Application to carryover-effect models. *J. Am. Stat. Assoc.* **80**, 704–710.
- Laska, E., M. Meisner and H. B. Kushner (1983). Optimal crossover designs in the presence of crossover effects. *Biometrics* **39**, 1087–1091.

- Lidl, R. and H. Niederreiter (1986). *Introduction to Finite Fields and Their Applications*. Cambridge, U.K.: Cambridge University Press.
- Low, J. L., S. M. Lewis and P. Prescott (1999). Assessing robustness of crossover designs to subjects dropping out. *Stat. Comput.* **9**, 219–227.
- Low, J. L., S. M. Lewis and P. Prescott (2003). An application of Pólya theory to cross-over designs with dropout. *Util. Math.* **63**, 129–142.
- Magda, C. G. (1980). Circular balanced repeated measurements designs. *Commun. Stat. Theory Methods* **9**, 1901–1918.
- Majumdar, D. (1988). Optimal repeated measurements designs for comparing test treatments with a control. *Commun. Stat. Theory Methods* **17**, 3687–3703.
- Majumdar, D., A. M. Dean and S. M. Lewis (2008). Uniformly balanced repeated measurements designs in the presence of subject dropout. *Stat. Sin.* **18**, 235–253.
- Martin, R. J. and J. A. Eccleston (1998). Variance-balanced change-over designs for dependent observations. *Biometrika* **85**, 883–892.
- Matthews, J. N. S. (1987). Optimal crossover designs for the comparison of two treatments in the presence of carryover effects and autocorrelated errors. *Biometrika* **74**, 311–320. Corrigendum: *ibid.* **75** (1988), 396.
- Matthews, J. N. S. (1988). Recent developments in crossover designs. *Int. Stat. Rev.* **56**, 117–127.
- Matthews, J. N. S. (1990). Optimal dual-balanced two treatment crossover designs. *Sankhyā Ser. B* **52**, 332–337.
- Matthews, J. N. S. (1994). Modelling and optimality in the design of crossover studies for medical applications. *J. Stat. Plan. Infer.* **42**, 89–108.
- Mukhopadhyay, A. C. and R. Saha (1983). Repeated measurements designs. *Calcutta Stat. Assoc. Bull.* **32**, 153–168.
- Park, D. K., M. Bose, W. I. Notz and A. M. Dean (2011). Efficient crossover designs in the presence of interactions between direct and carry-over treatment effects. *J. Stat. Plan. Infer.* **141**, 846–860.
- Patterson, H. D. (1952). The construction of balanced designs for experiments involving sequences of treatments. *Biometrika* **39**, 32–48.
- Patterson, H. D. (1973). Quenouille's changeover designs. *Biometrika* **60**, 33–45.
- Patterson, H. D. and H. L. Lucas (1959). Extra-period change-over designs. *Biometrics* **15**, 116–132.
- Patterson, H. D. and H. L. Lucas (1962). *Change-Over Designs*. Raleigh, NC: North Carolina Agricultural Experiment Station Technical Bulletin No. 147.
- Pocock, S. (1983). *Clinical Trials*. New York: Wiley.
- Quenouille, M. H. (1953). *The Design and Analysis of Experiment*. London, U.K.: Griffin.
- Ratkowsky, D. A., M. A. Evans and J. R. Alldredge (1992). *Cross-Over Experiments: Design, Analysis, and Application*. New York: Marcel Dekker.
- Schiffstein, H. N. V. and I. M. Oudejans (1996). Determination of cumulative successive contrasts in saltiness intensity judgements. *Percept. Psychophys.* **58**, 713–724.
- Sen, M. and R. Mukerjee (1987). Optimal repeated measurements designs under interaction. *J. Stat. Plan. Infer.* **17**, 81–91.
- Senn, S. (1992). Is the simple carryover model useful? *Stat. Med.* **11**, 715–726.
- Senn, S. (2000). Crossover design. In *Encyclopedia of Biopharmaceutical Statistics 2* (S. Chow, Ed.). New York: Marcel Dekker, pp. 142–149.
- Senn, S. (2003). *Cross-over Trials in Clinical Research*, 2nd ed. Chichester, U.K.: Wiley.
- Shah, K. R., M. Bose and D. Raghavarao (2005). Universal optimality of Patterson's designs. *Ann. Stat.* **33**, 2854–2872.
- Simpson, T. W. (1938). Experimental methods and human nutrition (with discussion). *J. R. Stat. Soc. Suppl.* **5**, 46–69.
- Street, D. J., J. A. Eccleston and W. H. Wilson (1990). Tables of small optimal repeated measurements designs. *Aust. J. Stat.* **32**, 345–359.
- Stufken, J. (1991). Some families of optimal and efficient repeated measurements designs. *J. Stat. Plan. Infer.* **27**, 75–83.

- Stufken, J. (1996). Optimal crossover designs. In *Handbook of Statistics* 13 (S. Ghosh and C. R. Rao, Eds.). Amsterdam, the Netherlands: North-Holland, pp. 63–90.
- Ting, C. P. (2002). Optimal and efficient repeated measurements designs for comparing test treatments with a control. *Metrika* **56**, 229–238.
- Willan, A. R. and J. L. Pater (1986). Carryover and the two-period crossover clinical trial. *Biometrics* **42**, 593–599.
- Williams, E. J. (1949). Experimental designs balanced for the estimation of residual effects of treatments. *Aust. J. Sci. Res. A* **2**, 149–168.
- Yan, Z. and C. Locke (2010). Crossover designs for comparing test treatments to a control treatment when subject effects are random. *J. Stat. Plan. Infer.* **140**, 1214–1224.
- Yang, M. and M. Park (2007). Efficient crossover designs for comparing test treatments with a control treatment when $p = 3$. *J. Stat. Plan. Infer.* **137**, 2056–2067.
- Yang, M. and J. Stufken (2008). Optimal and efficient crossover designs for comparing test treatments to a control under various models. *J. Stat. Plan. Infer.* **138**, 278–285.
- Zhao, S. and D. Majumdar (2012). On uniformly balanced crossover designs efficient under subject dropout. *J. Stat. Theory Pract.* **6**, 178–189.
- Zheng, W. (2013). Universally optimal crossover designs under subject dropout. *Ann. Stat.* **41**, 63–90.

5

Response Surface Experiments and Designs

André I. Khuri and Siuli Mukhopadhyay

CONTENTS

5.1	Introduction.....	198
5.2	Traditional Response Surface Designs.....	199
5.2.1	Designs for First-Degree Models.....	202
5.2.1.1	2^k Factorial Design.....	202
5.2.1.2	Plackett–Burman Design/Nonregular Factorial Design.....	203
5.2.1.3	Simplex Design.....	203
5.2.2	Designs for Second-Degree Models.....	203
5.2.2.1	3^k Factorial Design.....	203
5.2.2.2	Central Composite Design.....	204
5.2.2.3	Box–Behnken Design.....	205
5.2.3	Importance of Model Bias in the Choice of Design.....	207
5.3	Response Surface Designs with Some Difficult-to-Change Factors.....	208
5.4	Response Surface Models with Mixed Effects.....	209
5.5	Robust Parameter Design.....	211
5.5.1	Taguchi Approach.....	211
5.5.2	Response Surface Approach.....	212
5.6	GLMs.....	214
5.6.1	Choice of Designs for GLMs.....	214
5.6.1.1	Locally Optimal Designs.....	215
5.6.1.2	Two-Stage Sequential Designs.....	215
5.6.1.3	Bayesian Designs.....	215
5.6.1.4	Quantile Dispersion Graphs (QDGs) Approach.....	216
5.6.2	Model Misspecification in GLMs.....	216
5.6.3	Further Extensions.....	217
5.7	Multiresponse Experiments.....	217
5.7.1	Designs for Multiresponse Experiments.....	219
5.7.2	Multiresponse Optimization.....	219
5.8	Graphical Procedures for the Evaluation and Comparison of Response Surface Designs.....	220
5.8.1	Numerical Example.....	222
5.9	Mixture Designs.....	224
5.9.1	Constraints on the Mixture Components.....	226
5.9.2	Slack-Variable Models.....	227

5.10 Applications of RSM.....228
 5.10.1 Applications of RSM in Agricultural Sciences.....228
 5.10.2 Applications of RSM in Food Sciences.....229
5.11 New Research Directions in RSM.....230
References.....230

5.1 Introduction

Response surface methodology (RSM) is concerned with the development of an empirical relationship between a response variable y and a set of control variables, x_1, x_2, \dots, x_k , that represent levels of quantitative factors believed to affect the response values. Such a relationship can be approximately represented by a polynomial model, typically of the first degree or the second degree in x_1, x_2, \dots, x_k . The model is then fitted to a data set generated by observing y at certain values of the control variables, referred to as locations, within an experimental region denoted by \mathcal{R} . If the fitted model is determined to be an adequate representation of the response, then it can be used to

1. Determine, through hypothesis testing, significance of the factors whose levels are represented by x_1, x_2, \dots, x_k .
2. Predict the response at locations within \mathcal{R} , that is, for given settings of the control variables.
3. Determine optimum operating conditions on x_1, x_2, \dots, x_k that maximize (or minimize) the expected response over a certain region of interest within \mathcal{R} .

The settings of the control variables used to obtain the response values for fitting the assumed model constitute the so-called response surface design. The choice of this design is very important since it can have a significant impact on the reliability of the predicted response values, the adequacy of the fitted model, and the ensuing optimization results. The actual choice of design is based on certain criteria that pertain, for example, to the prediction variance and the protection against a possible bias that can result from fitting the wrong model.

One of the characteristics of RSM is its sequential nature. By this, we mean that an experiment is performed in stages whereby information acquired in one stage is used to plan the strategy in the next stage. For example, an exploratory experiment is initially run to determine which factors are important in terms of influencing a response of interest. Data collected in this experiment can then be used to gain more information about the response and its influential factors. This process of sequential experimentation is repeated several times until a clear understanding of the overall system is realized. The fitted model in the last stage can then be used to determine optimum operating conditions on the factors that maximize (or minimize) the response.

Factorial experiments play an important role in RSM. As will be seen later in Section 5.2.1, two-level factorial designs are used to fit simple models in the initial stages of the experiment. Three-level factorial designs (see Section 5.2.2) are used to fit second-degree models that are useful in determining optimal response values. Fractions of factorial designs can be used if the number of factors is large. In a factorial experiment, the response

is observed under all or a fraction of the possible combinations of the levels of the factors considered in the experiment. Inferences can then be made with regard to the main effects of the factors as well as their interactions. This investigation also provides information regarding the intermediate levels of the factors. In other words, it is possible to interpolate the response values at the combinations of levels not tried in the experiment. To explain this, let us consider, for example, two quantitative factors, A and B , having three levels each. In this case, we use the method of orthogonal polynomials to set up the polynomial effects, namely, the linear and quadratic effects that partition the main effects of each factor that carries two degrees of freedom (see Montgomery, 2013, Section 5.5). These polynomial effects have one degree of freedom each. The interaction effect, $A * B$, with its four degrees of freedom is partitioned into four single-degree-of-freedom polynomial effects, namely, $A_L * B_L$, $A_L * B_Q$, $A_Q * B_L$, and $A_Q * B_Q$, where the subscripts L and Q denote the linear and quadratic effects of the corresponding factor. With commonly used designs, all the aforementioned effects have independent estimates. An analysis of variance (ANOVA) can then lead to the determination of the polynomial effects that are significant. This result in a regression model that when fitted to the data can, in general, provide an adequate description of the response data. Such a model is then used to predict the response at intermediate levels of A and B . A similar analysis can be conducted if one factor is quantitative and the other is qualitative, but only the main effect of the quantitative factor is partitioned into polynomial effects. Furthermore, if a quantitative factor has k levels, then its effect can be partitioned into $k - 1$ independent polynomial effects each having one degree of freedom.

The original work in the area of RSM, which dates from the 1950s, has been widely used in the chemical industry giving rise to many interesting applications. A review of such applications was given by Hill and Hunter (1966). Since the 1970s, the scope of RSM's applicability has widened considerably with applications in the physical and engineering sciences, food science, and the biological and chemical sciences, to name just a few (see the review article by Myers et al. 1989). The more recent use in RSM of generalized linear models (GLMs), developed by Nelder and Wedderburn (1972), extended the applicability of RSM to several new areas, including the medical sciences and the pharmaceutical industry (see the review article by Myers et al. 2004). Section 5.10 highlights the use of RSM in agricultural and food sciences.

5.2 Traditional Response Surface Designs

By a *traditional response surface design* (TRSD), we mean a design typically used to fit a first-degree or a second-degree polynomial model. Such models provide low-degree approximations to the mean response in a given experimental situation. They are simple and do not require the availability of a large number of experimental runs, especially in the initial stages of the experiment where, for example, a first-degree model is often used to explore the nature of the response surface. Most of the designs for these models were developed in the 1950s and early 1960s. See Box and Wilson (1951), Box (1952), Box and Draper (1959), Box and Hunter (1957), Box and Behnken (1960), and Plackett and Burman (1946).

In a TRSD setup, a response variable y is represented by a linear model of the form

$$y = f'(x)\beta + \epsilon, \quad (5.1)$$

where $x = (x_1, x_2, \dots, x_k)'$ is a vector of k control variables, $f(x)$ is a vector function of p elements that consist of a constant and powers and cross products of x_1, x_2, \dots, x_k up to a certain degree denoted by d (≥ 1), β is a vector of unknown parameters, and ϵ is a random experimental error assumed to have a zero mean. Hence, the mean of y , $\eta(x)$, is given by

$$\eta(x) = f'(x)\beta, \quad (5.2)$$

which is a polynomial model of degree d . Two commonly used models are the first-degree and second-degree models, namely,

$$\eta(x) = \beta_0 + \sum_{i=1}^k \beta_i x_i, \quad (5.3)$$

and

$$\eta(x) = \beta_0 + \sum_{i=1}^k \beta_i x_i + \sum_{i < j} \beta_{ij} x_i x_j + \sum_{i=1}^k \beta_{ii} x_i^2, \quad (5.4)$$

where β_0 , β_i , β_{ij} , and β_{ii} are unknown parameters that make up, in the respective models, the elements of β .

Given a series of n ($n \geq p$) experimental runs resulting in responses y_1, \dots, y_n , we get from Model (5.1) the following representation for the u th response value:

$$y_u = f'(x_u)\beta + \epsilon_u, \quad u = 1, 2, \dots, n, \quad (5.5)$$

where $x_u = (x_{u1}, x_{u2}, \dots, x_{uk})'$ is the vector of design settings of x_1, x_2, \dots, x_k at the u th run, and ϵ_u is the corresponding error term. The $n \times k$ matrix, D , whose rows consist of the values of x'_u , $u = 1, 2, \dots, n$, is called the design matrix. Thus,

$$D = \begin{bmatrix} x_{11} & x_{12} & \dots & x_{1k} \\ x_{21} & x_{22} & \dots & x_{2k} \\ \dots & \dots & \dots & \dots \\ x_{n1} & x_{n2} & \dots & x_{nk} \end{bmatrix}. \quad (5.6)$$

Model (2.5) can then be expressed as

$$y = X\beta + \epsilon, \quad (5.7)$$

where $y = (y_1, y_2, \dots, y_n)'$, X is the so-called model matrix whose rows are the values of $f'(x_u)$, $u = 1, 2, \dots, n$, and $\epsilon = (\epsilon_1, \epsilon_2, \dots, \epsilon_n)'$. In a typical response surface experiment, X is of full column rank; hence, $X'X$ is a nonsingular matrix.

Estimation of β is usually done by using the method of ordinary least squares. If ϵ is assumed to have a zero mean vector and a variance–covariance matrix $\sigma^2 I_n$, then the least-squares estimator of β is

$$\hat{\beta} = (X'X)^{-1}X'y. \quad (5.8)$$

The variance–covariance matrix of $\hat{\beta}$ is given by

$$\text{Var}(\hat{\beta}) = \sigma^2(X'X)^{-1}. \quad (5.9)$$

Using $\hat{\beta}$ in place of β in (5.2) gives an estimate of the response surface at a point x that is represented as

$$\hat{y}(x) = f'(x)\hat{\beta}. \quad (5.10)$$

This is called the prediction equation, and $\hat{y}(x)$ is referred to as the predicted response at x , assuming that the form of the fitted model is correct. Its mean is $f'(x)\beta$, the same as $\eta(x)$ in (5.2), and its variance, which is called the prediction variance, is given by

$$\text{Var}[\hat{y}(x)] = \sigma^2 f'(x)(X'X)^{-1}f(x). \quad (5.11)$$

The size of the prediction variance in a response surface experiment is of major concern. Large prediction variances indicate unreliable predictions within \mathcal{R} . Since the prediction variance depends on D , the choice of the design matrix is of paramount importance in getting good quality predictions. This is of course contingent on (5.2) being an adequate model for the true mean response. Certain choices of D have desirable properties. For example, a design is said to be *orthogonal* if the matrix $X'X$ is diagonal resulting in the elements of $\hat{\beta}$ being uncorrelated as can be seen from (5.9). Another property of design is that of *rotatability*. A design is *rotatable* if the prediction variance is constant at all points that are equidistant from the design center, normally chosen to be the point at the origin of the coordinates system, but note that this depends on the scaling of x . This is an invariance property for the prediction variance that allows equal-quality predictions in any direction provided that the distance from the origin is unchanged. This can be helpful in the comparison of the predicted response values on the surfaces of concentric hyperspheres centered at the origin. Note that properties such as orthogonality and rotatability depend on the design used and the form of the fitted model since the matrix X depends on both. A rotatable design has the additional *uniform precision* property if the prediction variances at the origin and at a distance of one from the origin are equal. This can be helpful in keeping changes in the prediction variance minimal in the vicinity of the origin. Further details concerning all three design properties can be found in, for example, Khuri and Cornell (1996).

It is very common in RSM to have coded variables used in place of the original control variables. This is done in order to facilitate the construction of the response surface design. Coding removes the units of measurement of the control variables so that distances measured along the axes of the coded variables are standardized, that is, converted to unitless terms. Two of the most obvious advantages of coding are

1. Computational ease and increased accuracy in estimating the model parameters,
2. Enhanced interpretability of the parameter estimates.

There are several ways to code variables. Quite often, the following coding convention for a model with k control variables is given by

$$x_{ui} = \frac{X_{ui} - \bar{X}_i}{s_{xi}},$$

where X_{ui} and x_{ui} are the original and coded u th levels of the i th variable, $u = 1, 2, \dots, n$, $i = 1, 2, \dots, k$ with n being the number of experimental runs, \bar{X}_i is the average of the X_{ui} 's, and s_{xi} is a measure of spread given by:

$$s_{xi} = \left[\sum_{u=1}^n \frac{(X_{ui} - \bar{X}_i)^2}{n} \right]^{\frac{1}{2}}.$$

Under this coding convention, $\sum_{u=1}^n x_{ui} = 0$, and $\sum_{u=1}^n x_{ui}^2 = n$.

5.2.1 Designs for First-Degree Models

We begin by describing designs for models of the form given in (5.3). Such models are of exploratory nature used in the initial stages of experimentation where little is known about the response. As a result, the corresponding designs are typically simple and, in general, of low cost. They are referred to as *first-order designs*. Of these designs, we will mention the 2^k factorial designs (k is the number of control variables considered in an experiment), Plackett–Burman designs, and simplex designs. All three designs have the orthogonality property.

5.2.1.1 2^k Factorial Design

In this design, each factor, or control variable, is measured at two levels, and all possible combinations of these levels are considered. Such combinations are called treatments. Thus the design matrix D consists of k columns and 2^k unreplicated rows whose (u, i) th element is the value of the i th control variable used in the u th experimental run ($u = 1, 2, \dots, 2^k$; $i = 1, 2, \dots, k$). If the two levels of each factor are coded as 1 and -1 , then each element of D is equal to 1 or -1 . This design is very easy to apply and allows adequate estimation of not only the parameters in model (5.3) but also all possible interactions among the factors. If only model (5.3) is fitted, then the design affords $2^k - k - 1$ degrees of freedom for the error term in the corresponding ANOVA table, which is a larger number of degrees of freedom than can be generated by many RSM designs. The only drawback is the cost of running a complete factorial experiment when k is large (≥ 6). In this case, fractions of 2^k can be considered. The construction of such fractions can be found in, for example, Chapter 3 in Khuri and Cornell (1996) and in Chapter 7.

5.2.1.2 Plackett–Burman Design/Nonregular Factorial Design (Plackett and Burman 1946)

This design is a fraction of a 2^k design whose number of points, n , is only equal to $k + 1$ and is therefore quite economical. Unlike regular fractions, which require n to be a power of 2, nonregular designs only require n to be a multiple of 4. Since n is equal to the number of parameters in model (5.3), it is considered to be a *saturated* design. In Plackett and Burman (1946), design arrangements can be found for $k = 3, 7, 11, \dots, 99$ factors. Such designs can be used with qualitative and quantitative factors. For more details, refer to Chapter 9.

5.2.1.3 Simplex Design

This design is less frequently used than the first two designs. It is a saturated design with $n = k + 1$ points located at the vertices of a k -dimensional simplex. For example, for $k = 2$, the simplex is an equilateral triangle centered at the origin, and for $k = 3$, the simplex is a tetrahedron. A method for constructing a simplex design in $k = n - 1$ dimensions is described in Box (1952).

The levels of some of the factors of a simplex may not be practicable from the experimental point of view. For example, under a simplex design, one factor has a coded range extending from $-\sqrt{\frac{3}{2}}$ to $\sqrt{\frac{3}{2}}$, rather than the traditional $-1, 1$ range as in 2^k factorial and Plackett–Burman designs. The former settings may be harder to apply and are therefore not favorable to the experimenter.

5.2.2 Designs for Second-Degree Models

Having determined what factors to account for using a first-degree model, the experimenter often proceeds to the next stage where a higher-degree model is fitted to acquire more information about the nature of the response surface under investigation. The second-degree model in (5.4) is a natural extension to model (5.3). The number of parameters in (5.4) is $p = 1 + 2k + \frac{1}{2}k(k - 1)$; hence, any second-order design (i.e., a design for fitting model (5.4)) must have at least this number of distinct points in order to support estimation of the parameters.

The use of second-degree models becomes necessary when curvature is suspected in the response surface. In this case, first-degree models will suffer from lack of fit and should therefore be expanded by the addition of higher-order terms to their fitted equations. The most commonly used designs for fitting second-degree models are the 3^k factorial, the central composite design (CCD), and the Box–Behnken design.

5.2.2.1 3^k Factorial Design

This design is constructed by taking all possible combinations of the levels of k factors that have three levels each. Thus, the number of points in this design is 3^k . Running a complete 3^k factorial design can therefore be very costly if k is large.

To reduce the cost of the experiment, a fraction of a 3^k design can be considered. For example, we can have a one-third fraction of a full 3^k factorial containing 3^{k-1} runs. In general, if $m < k$, then a 3^{-m} th fraction of the 3^k design will have 3^{k-m} runs. Such a fraction is called a 3^{k-m} fractional factorial design. Coverage of these designs can be found in, for example, McLean and Anderson (1984).

Orthogonality for second-order designs is possible, but only when the variables in model (5.4) are expressed in terms of orthogonal polynomials as shown in Box and Hunter (1957) (see pp. 200–201) and in Khuri and Cornell (1996) (see Section 4.3). This is necessary because the corresponding $X'X$ matrix for model (5.4) is not diagonal since some of the off-diagonal elements of $X'X$, namely, $\sum_{u=1}^n x_{ui}^2, \sum_{u=1}^n x_{ui}^2 x_{uj}^2$ ($i \neq j$), are positive, where x_{ui} is the u th setting of the i th control variable, $u = 1, 2, \dots, n; i = 1, 2, \dots, k$, and n is the number of rows in the X matrix. However, if model (5.4) is expressed in terms of orthogonal polynomials, then orthogonality can be achieved for some designs. In particular, the 3^k design is orthogonal if the three levels of each factor are equally spaced and if these levels are coded so that $\sum_{u=1}^n x_{ui} = 0$ and $\sum_{u=1}^n x_{ui}^2 = 1$ (see Khuri and Cornell 1996, Sections 4.3 and 4.5.1).

As was pointed out in Section 5.2, orthogonality is a desirable property since it results in a fitted model whose estimated parameters are uncorrelated. This facilitates making simultaneous inferences about the unknown parameters of the true model. Furthermore, the variance of the predicted response at any point x in a region of interest is expressible as a weighted sum of the variances of the parameter estimates.

5.2.2.2 Central Composite Design

Box and Wilson (1951) introduced the CCD as an alternative to the 3^k factorial design. The makeup of the CCD is as follows:

1. A factorial portion consisting of a complete or fractional 2^k factorial design whose settings are coded using the usual ± 1 values. Let F denote the number of points in this portion. If a fractional factorial is to be used, it should be chosen so that all main effects and two-factor interactions can be estimated.
2. An axial portion consisting of k pairs of points, the i th pair of which comprises of two points on the i th coordinate axis that are symmetric with respect to the origin at a distance of α .
3. n_0 center-point replications.

Thus the number of points in this design is $n = F + 2k + n_0$. The values of α and n_0 are chosen so that the CCD can attain desirable properties. For example, choosing $\alpha = F^{\frac{1}{4}}$ causes the design to be rotatable. Furthermore, n_0 can be chosen so that the design is either orthogonal or has the uniform precision property.

An example of a design matrix for a CCD when $k = 2, \alpha = \sqrt{2}, n_0 = 1$ is

$$\begin{matrix} & x_1 & x_2 \\ \begin{bmatrix} -1 & -1 \\ 1 & -1 \\ -1 & 1 \\ 1 & 1 \\ \sqrt{2} & 0 \\ -\sqrt{2} & 0 \\ 0 & \sqrt{2} \\ 0 & -\sqrt{2} \\ 0 & 0 \end{bmatrix} & . \end{matrix}$$

A table listing desirable values of n_0 for a rotatable CCD to be orthogonal or have the uniform precision property is given in Box and Hunter (1957). Other criteria for the choice of n_0 , including good detectability of lack of fit of the second-degree model, can be found in Draper (1982).

The CCD is perhaps the most popular second-order design due to its flexibility in choosing the values of α and n_0 , and for requiring a considerably smaller number of experimental runs than the 3^k design, especially when $k \geq 4$. Furthermore, since it consists of a 2^k complete or fractional factorial design plus added points, it can be divided to specify a response surface investigation in two stages. For example, suppose that an experimenter starts out a study using a 2^k factorial experiment to fit a first-degree model. Upon testing its adequacy, it was determined that it suffered from a significant lack of fit, indicating inadequacy of the first-degree model. The experimenter can then augment, in a second stage, the initial 2^k design with the additional axial runs and center points in order to fit a second-degree model. This process clearly explains the sequential nature of RSM as was pointed out in Section 5.1.

A variant of the CCD, known as the San Cristobal design (SCD), was introduced by Rojas (1962) for use in experiments in sugar farming. It was named after the sugar-milling town of that name in Mexico. This design is used to fit a second-degree model in situations in which the levels of the k control variables, for example, levels of fertilizers, are restricted to be positive or zero. As described by Haines (2006), it consists of 2^k factorial points together with center and axial points, all contained within the positive orthant. The main difference with the CCD is that the SCD includes a control, that is, the point $(0, 0, \dots, 0)$ corresponding to no application of any fertilizer. More specifically, the SCD in k control variables consists of the following:

1. A full 2^k factorial or a 2^{k-p} fractional factorial design replicated r times with the control variables taking the coded values 0 or 2. This gives a total of $F = r2^k$ or $r2^{k-p}$ such points.
2. n_0 center points of the form $(1, 1, \dots, 1)$.
3. c axial points of the form $(1, 1, \dots, 1 - \alpha, \dots, 1)$ for each of the k variables where $0 < \alpha \leq 1$.
4. One axial point of the form $(1, 1, \dots, c\alpha + 1, \dots, 1)$ for each of the k variables.

Note that the total number of points in this design is $n = F + n_0 + k(c + 1)$. The constraint $\alpha \leq 1$ is needed in order to ensure nonnegative applications of fertilizer. The parameters r , n_0 , c , and α can be chosen to achieve certain design properties such as orthogonality. The SCD offers agriculturists interesting alternatives to factorial and fractional factorial designs in that they can accommodate unusual design settings and often require fewer runs. Haines (2006) also evaluated the performance of this design and reviewed some of its basic properties.

5.2.2.3 Box–Behnken Design

This design, which was introduced in Box and Behnken (1960), consists of a particular subset of the factorial combinations from the 3^k factorial design where each factor has three levels. More specifically, it is formed by combining a two-level factorial design with a balanced incomplete block design (BIBD). To illustrate this combination, consider,

for example, a BIBD with three treatments and three blocks of the form shown in the following text.

Blocks	Treatments		
	1	2	3
1	*	*	
2	*		*
3		*	*

Let us combine this design with a 2² factorial design (with levels at ±1) as follows: The two asterisks in every block of the BIBD are replaced by the two columns of the 2² design. In places where no asterisk is marked, the corresponding variable is set at the center level of the factor (i.e., 0). Such a combination is augmented with a certain number of center points. This results in a Box–Behnken design in three variables as shown in Table 5.1. The last row in this table contains vectors of center points.

The Box–Behnken design is popular among industrial research workers because it is economical and provides the settings -1, 0, and 1, which are easy to attain in an experiment for each control variable. For three factors, the number of runs required by a Box–Behnken design is low. However, as the number of factors increases, this advantage is no longer there. This design, however, is not always rotatable. Box and Behnken (1960) list a number of arrangements of this design for k = 3, 4, 5, 6, 7, 9, 10, 11, 12, 16 control variables, some of which use two-level subdesigns in more than two factors or partially BIBDs in place of BIBDs.

TABLE 5.1
Box–Behnken Design in Three Variables

x ₁	x ₂	x ₃
-1	-1	0
1	-1	0
-1	1	0
1	1	0
—	—	—
-1	0	-1
1	0	-1
-1	0	1
1	0	1
—	—	—
0	-1	-1
0	1	-1
0	-1	1
0	1	1
—	—	—
0	0	0

More details concerning the aforementioned three designs and other lesser-used second-order designs are given in standard texts on RSM, such as Khuri and Cornell (1996), Myers et al. (2009), and Box and Draper (2007). Chapter 3 discusses BIBDs in detail.

5.2.3 Importance of Model Bias in the Choice of Design

The main design criteria in classical RSM are (1) minimization of the prediction variance and (2) minimization of the bias caused by fitting the wrong polynomial model. While, in (1), the fitted model is assumed to be the “true” model, the approach in (2) regards that the fitted model does not necessarily represent the true model, and that model bias should be accounted for when choosing a design. Box and Draper (1959) (see also Box and Draper 1963) emphasized the importance of bias contribution in the choice of design and advocated the consideration of the so-called integrated mean squared error (IMSE) that incorporates prediction variance and model bias. More specifically, the IMSE is expressible as (see Khuri and Cornell 1996)

$$\text{IMSE} = V + B,$$

where

$$V = \frac{n\Omega}{\sigma^2} \int_{\mathcal{R}} \text{Var}[\hat{y}(x)] dx, \quad (5.12)$$

and

$$B = \frac{n\Omega}{\sigma^2} \int_{\mathcal{R}} \{E[\hat{y}(x)] - \zeta(x)\}^2 dx. \quad (5.13)$$

In (5.12) and (5.13), $\Omega^{-1} = \int_{\mathcal{R}} dx$ is the volume of the region of interest, \mathcal{R} , n is the total number of observations used for fitting the model, σ^2 is the error variance, $dx = dx_1 dx_2 \cdots dx_k$, and $\zeta(x)$ is the “true” mean response at a point x in \mathcal{R} . Although it is not necessary, for design purposes, $\zeta(x)$ is often represented by a polynomial model of degree higher than the fitted model, and the quantity $E[\hat{y}(x)] - \zeta(x)$ is the bias of response prediction due to model misspecification. It follows that a reasonable choice of design is based on the minimization of the IMSE. However, this is not possible without the specification of the form of $\zeta(x)$ that is unknown. Designs are usually chosen by minimizing either V or B , where in case of the latter, a particular form of the true model is usually assumed. Box and Draper (1959) argued that unless V is considerably larger than B , a good choice of design is to have one that minimizes B since such a design has characteristics similar to those of a design that minimizes IMSE.

More specifically, if we assume that $\zeta(x)$ is of the form

$$\zeta(x) = f'_1(x)\beta_1 + f'_2(x)\beta_2, \quad (5.14)$$

where $f'_1(x)\beta_1$ is the portion associated with the fitted model and $f'_2(x)\beta_2$ is the portion that was not fitted (both β_1 and β_2 are unknown), then it can be shown (see, e.g., Khuri and Cornell 1996, Chapter 6) that B can be expressed as

$$B = \frac{n}{\sigma^2} \beta_2' [\mu_{22} - \mu_{12}' \mu_{11}^{-1} \mu_{12} + (A - \mu_{11}^{-1} \mu_{12})' \mu_{11} (A - \mu_{11}^{-1} \mu_{12})] \beta_2, \quad (5.15)$$

where $A = (X_1' X_1)^{-1} X_1' X_2$ and μ_{ij} are the region moment matrices,

$$\mu_{ij} = \Omega \int_{\mathcal{R}} f_i(x) f_j'(x) dx, \quad i, j = 1, 2. \quad (5.16)$$

Note that in the expression for A , X_1 is the model matrix with n rows for the fitted model, and X_2 corresponds to the unfitted portion of the model. Box and Draper (1959) state that, in particular, B is minimized when $(1/n)X_1' X_1 = \mu_{11}$ and $(1/n)X_1' X_2 = \mu_{12}$. In this case, B takes the value

$$B = \frac{n}{\sigma^2} \beta_2' (\mu_{22} - \mu_{12}' \mu_{11}^{-1} \mu_{12}) \beta_2. \quad (5.17)$$

This value of B depends on the unknown value of β_2 . Note that arriving at such an expression for B was based on assuming a particular form for the true model. In general, such an assumption cannot be fully justified as the true form of the model may be unknown, but designs may still be compared for their potential bias properties under realistic alternative model forms.

Additional details concerning the role of bias in the choice of design can be found in Chapter 20.

5.3 Response Surface Designs with Some Difficult-to-Change Factors

In some response surface experiments, particularly those carried out in certain industrial settings, the levels of some factors are easy to change, while the levels of other factors may be harder to change. This can be due to operational difficulties associated with frequent changes of the levels in the latter factors or to the cost of these changes. For example, in the process of increasing the strength of steel, two factors were considered influential, namely, the temperature of heating the steel and the heat treatment time. Two temperature settings, 1600°F and 1700°F, and three times, 10, 20, and 30 min, were selected. The process began by heating a particular furnace to one of the two temperatures, for example, 1600°F, and then, three specimens of steel were inserted. One specimen was removed after 10 min and then soaked and air cooled. After 10 more minutes, another specimen was removed and subjected to the same treatment. The same process was applied to the last remaining specimen after 10 min of removing the second specimen. Thus, the times of exposure to the first selected temperature were 10, 20, and 30 min for the first, second, and third specimens, respectively. Next, the temperature of the furnace was changed to the second setting, 1700°F, and the process was repeated as before using three new specimens.

In this experiment, the furnace temperature is a difficult-to-change factor since it takes a great deal of time to heat the furnace to the desired temperature, especially if the furnace

is large. On the other hand, heating time for a steel specimen is much easier to change. This situation can be handled by running a factorial experiment with certain restrictions imposed on randomization. These restrictions have an impact on how the analysis is to be conducted. This topic is covered in detail in Chapter 8.

Several authors considered designs for experiments where the levels of some factors can be difficult to change. For example, Box (1996) indicated that certain designs can be more efficient and easier to run than completely randomized designs in industrial situations. Draper and John (1998) recommended certain modifications to the CCD and Box–Behnken designs to produce designs with a high degree of rotatability that accommodate difficult-to-change factors. Mee and Bates (1998) considered designs for two- and three-level factorial experiments involving silicon wafers. Kowalsky et al. (2006) also considered modifying the CCD to allow the estimation of separate models for the process mean and process variance. Their article was an extension of the one by Vining et al. (2005). More references concerning designs for experiments with difficult-to-change factors can be found in Myers et al. (2004).

5.4 Response Surface Models with Mixed Effects

The traditional model in classical RSM includes only fixed continuous polynomial effects. This excludes other effects of discrete nature that are not measured on a continuous scale such as block effects. For example, in the semiconductor industry, measurements on the resistance in computer chips may be taken using several silicon wafers drawn from a large lot. In order to account for any possible variation among the wafers, effects due to blocks representing the wafers should be considered in addition to the effects of continuous factors that influence the resistance response. In the early days of RSM, block effects were considered, but only under the special case of designs that *block orthogonally* as in Box and Hunter (1957) who introduced the concept of orthogonal blocking for designs for second-degree models. By that, we mean designs that allow estimation of the mean response in a manner that is invariant to the block effects as if they do not exist. Orthogonal blocking requires certain conditions on the design used. For example, for second-degree models, a design blocks orthogonally if (1) each block consists of a first-order orthogonal design and, (2) for each block, the sum of squared values for the i th control variable, $i = 1, 2, \dots, k$, is proportional to the number of runs in the block.

The CCD can be blocked orthogonally where each of the factorial and axial portions of the design forms a first-order orthogonal design. The appropriate number of center-point replications will have to be determined to satisfy condition (2). For more details, see Khuri and Cornell (1996), Section 8.2. A general condition for orthogonal blocking for models of degree d (≥ 2) was given in Khuri (1992).

Suppose that the mean response in model (5.2) is estimated using a design divided into b blocks of sizes n_1, n_2, \dots, n_b . Modifying model (5.5) to include a block effect results in the following model:

$$y_u = f'(x_u)\beta + z'_u\gamma + \Phi'(x_u)\Lambda z_u + \epsilon_u, \quad u = 1, 2, \dots, n, \quad (5.18)$$

where $n = \sum_{i=1}^b n_i$ z_u is a vector of b elements whose i th element takes the value 1 if the u th trial is in the i th block and is zero otherwise; $\gamma = (\gamma_1, \gamma_2, \dots, \gamma_b)'$, where γ_i denotes the

effect of the i th block; $\phi(x_u)$ is the same as $f(x_u)$ except that it is missing the first element of $f(x_u)$ that is equal to 1; Λ is a matrix that contains interaction coefficients between the fixed polynomial terms and the blocks; and ϵ_u is a random experimental error. The model can be expressed in vector form as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{Z}_b\boldsymbol{\gamma} + \sum_{j=2}^p \mathbf{U}_j\boldsymbol{\delta}_j + \boldsymbol{\epsilon}, \quad (5.19)$$

where $\mathbf{Z}_b = \text{diag}(\mathbf{1}_{n_1}, \mathbf{1}_{n_2}, \dots, \mathbf{1}_{n_b})$, \mathbf{U}_j is a matrix of order $n \times b$ whose i th column is obtained by multiplying the elements of the j th column of \mathbf{X} with the corresponding elements of the i th column of \mathbf{Z}_b ($i = 1, 2, \dots, b; j = 2, 3, \dots, p$), $\boldsymbol{\delta}_j$ is a vector of interaction coefficients that consists of the elements of the $(j - 1)$ st row of Λ , p is the number of columns of \mathbf{X} , and $\boldsymbol{\epsilon}$ is the vector consisting of all the ϵ_u 's ($u = 1, 2, \dots, n$). For more details, see Chapter 8 in Khuri and Cornell (1996).

It is assumed that $\boldsymbol{\gamma}$, $\boldsymbol{\delta}_2, \boldsymbol{\delta}_3, \dots, \boldsymbol{\delta}_p$ are normally and independently distributed with zero means and variance–covariance matrices $\sigma_\gamma^2 \mathbf{I}_b, \sigma_2^2 \mathbf{I}_b, \dots, \sigma_p^2 \mathbf{I}_b$, respectively. The random error vector $\boldsymbol{\epsilon}$ is independent of the previous random effects and is assumed to have the normal distribution with mean $\mathbf{0}$ and a variance–covariance matrix Δ_b given by

$$\Delta_b = \text{diag}(\tau_1^2 \mathbf{I}_{n_1}, \tau_2^2 \mathbf{I}_{n_2}, \dots, \tau_b^2 \mathbf{I}_{n_b}), \quad (5.20)$$

where the τ_i^2 's are unknown variance components. Note that the variance–covariance structure for Δ_b indicates that the error variances are different for the different blocks, but observations obtained within a block have the same variance. Using model (5.19), the mean and variance-covariance matrix of \mathbf{y} can then be written, respectively, as

$$E(\mathbf{y}) = \mathbf{X}\boldsymbol{\beta}, \quad (5.21)$$

$$\begin{aligned} \text{Var}(\mathbf{y}) &= \boldsymbol{\Sigma}_b \\ &= \sigma_\gamma^2 \mathbf{Z}_b \mathbf{Z}_b' + \sum_{j=2}^p \sigma_j^2 \mathbf{U}_j \mathbf{U}_j' + \Delta_b. \end{aligned} \quad (5.22)$$

Therefore, the estimated generalized least-squares estimator of $\boldsymbol{\beta}$ is of the form

$$\hat{\boldsymbol{\beta}}^* = (\mathbf{X}' \hat{\boldsymbol{\Sigma}}_b^{-1} \mathbf{X})^{-1} \mathbf{X}' \hat{\boldsymbol{\Sigma}}_b^{-1} \mathbf{y}, \quad (5.23)$$

where $\hat{\boldsymbol{\Sigma}}_b$ is the same as $\boldsymbol{\Sigma}_b$ except that appropriate estimates of the variance components in $\boldsymbol{\Sigma}_b$, such as maximum likelihood or restricted maximum likelihood estimates, are used. The corresponding estimated variance–covariance matrix of $\hat{\boldsymbol{\beta}}^*$ is approximately given by

$$\widehat{\text{Var}}(\hat{\boldsymbol{\beta}}^*) \approx (\mathbf{X}' \hat{\boldsymbol{\Sigma}}_b^{-1} \mathbf{X})^{-1}. \quad (5.24)$$

Consequently, the predicted response at a point x in a region \mathcal{R} is

$$\hat{y}(x) = f'(x) \hat{\boldsymbol{\beta}}^*, \quad (5.25)$$

and the corresponding prediction variance is approximately of the form

$$\widehat{\text{Var}}[\hat{y}(x)] \approx f'(x)(X' \hat{\Sigma}_b^{-1} X)^{-1} f(x). \quad (5.26)$$

Tests concerning the fixed effects (i.e., the elements of β) and the random effects (i.e., the variance components, $\sigma_\gamma^2, \sigma_2^2, \sigma_3^2, \dots, \sigma_p^2$) were outlined in Khuri (2006).

The special case of homogeneous error variances, that is, when the diagonal elements of Δ_b in formula (5.20) are equal, was considered by Khuri (1996b). In this case, more simplified tests concerning the fixed and random effects in model (5.19) can be derived than in the case of heterogeneous error variances. More recently, Saha and Khuri (2009) used a graphical technique for the comparison of designs for response surface models with random block effects.

5.5 Robust Parameter Design

The area of robust parameter design (RPD) is a well-developed methodology introduced by Genichi Taguchi for product quality improvement. Taguchi identified two types of inputs in a process: control factors (x), which are easy to manipulate, and noise factors (z), which are difficult to control when the product is in actual use, but can be controlled under laboratory conditions. The noise factors are the sources of variations in the process response when the system is used in practice. The main aim of an RPD is to determine the settings of the control factors that achieve a desired mean response and to make the process robust, or less sensitive, to the effects of the noise variables.

5.5.1 Taguchi Approach

The Taguchi method has been discussed in details by several authors, for example, Kacker (1985), Taguchi (1986), Khuri and Cornell (1996, Chapter 11), and Wu and Hamada (2009). More references are given in Khuri and Mukhopadhyay (2010). To solve the RPD problem, Taguchi proposed the use of two design matrices, one for the control variables called the control array and the other for the noise variables called the noise array. The arrays used for the control and noise variables are usually orthogonal arrays (refer to Chapter 9). Each of the level combinations of the control array is then crossed with all the level combinations of the noise array to give a crossed array design. Let us consider an example on epitaxial layer growth from Wu and Hamada (2009). There were eight control factors and two noise factors. A 16-run fraction of a 2^8 factorial design was used as the control array and a 2×4 complete factorial design as the noise array. Thus, at each control setting, there are eight observations, say, y_i , $i = 1, \dots, 8$, one for each of the different noise settings. Taguchi classified the RPD problem into three categories based on the following three goals of the experimenter: the smaller the better, the larger the better, and target is the best. For each of these classifications, a performance criterion or a signal-to-noise ratio (SNR) was defined in terms of the y_i 's. As an example, when the goal was to maximize the response, the SNR to be used is $-10 \log \left[\frac{1}{8} \sum_{i=1}^8 \frac{1}{y_i^2} \right]$. Rather than performing an analysis on the actual observations, an analysis would be performed with SNR as the response variable.

Even though the Taguchi method was an important step toward process improvement, it was widely criticized by several authors. See, for example, Box (1985, 1988), Easterling (1985), Pignatiello and Ramberg (1985), Nair (1992), and Myers et al. (2009), to name just a few. Some of the main criticisms of the Taguchi method are the following:

1. Interactions among the control factors are ignored.
2. The required number of experimental runs in the crossed array scheme is much larger than is needed.
3. The SNRs do not provide adequate means to identify the control factors that affect only the process mean as distinguished from those that affect only the process variance, as well as those that affect both the process mean and process variance.

For these reasons, other approaches were proposed as alternative solutions to the RPD problem. One of the main alternative solutions is the adoption of the response surface approach. More details concerning the Taguchi approach can be found in, for example, the texts by Khuri and Cornell (1996) and Myers et al. (2009).

5.5.2 Response Surface Approach

The response surface approach to robust designs can further be categorized into the dual-model and the single-model approaches.

1. *Dual-model approach:* Vining and Myers (1990) proposed that Taguchi's aim of keeping the mean on target while simultaneously minimizing the variance could also be achieved in a response surface framework. They introduced a dual-model approach by fitting two separate response surface models in only the control settings, x , to the process mean (μ) and the process variance (σ^2) whose estimates are, respectively,

$$\hat{\mu}(x) = \hat{\beta}_{0\mu} + x' \hat{\beta}_{\mu} + x' \hat{B}_{\mu} x,$$

$$\hat{\sigma}^2(x) = \hat{\gamma}_{0\sigma} + x' \hat{\gamma}_{\sigma} + x' \hat{C}_{\sigma} x,$$

where $\hat{\beta}_{0\mu}$, $\hat{\beta}_{\mu}$, \hat{B}_{μ} , $\hat{\gamma}_{0\sigma}$, $\hat{\gamma}_{\sigma}$, and \hat{C}_{σ} are the least-squares estimates. Any TRSD with an outer array set up at each point is used to fit the mean and variance models. The values of the sample means and variances computed from the outer array are used as data points to fit the aforementioned two models. All the model parameters are estimated using the least-squares method (for details, see Vining and Myers 1990, p. 40). Following the variance modeling of Bartlett and Kendall (1946), several authors (Box and Meyer 1986 and Nair and Pregibon 1988) have also suggested modeling the logarithm of $\hat{\sigma}^2$. However, Bartlett and Kendall (1946) advises against the use of the natural logarithm of $\hat{\sigma}^2$ if the number of points is less than 5. Vining and Myers (1990) proposed the use of any variance stabilizing transformation of $\hat{\sigma}^2$ for cases where the number of outer array points is < 10 .

The variance function, $\hat{\sigma}^2(x)$, is then minimized with respect to the control settings, x , under the constraint that the mean response, $\hat{\mu}(x)$, is equal to some prespecified target quantity, using the dual response optimization method of Myers and Carter (1973). Del Castillo and Montgomery (1993) showed how the

same goals as in Vining and Myers (1990) could be realized by using nonlinear programming techniques.

2. *Single-model approach:* Welch et al. (1990) showed that using a single experimental design for the control and noise variables greatly reduced the number of runs needed in an experiment as compared to the crossed array design of Taguchi. This single design for both types of factors is known as a combined array. Using the combined array, a single model in the control variables, x , and the noise variables, z , was fitted to the response y ,

$$y(x, z) = \beta_{0s} + g'(x)\beta_s + z'\delta + g'(x)\Delta z + \epsilon,$$

where $g(x)$ is a known function of the control variables and β_s , δ , Δ are the corresponding regression coefficients of $g(x)$, z , and the interaction coefficients between the polynomial terms involving control variables and z . The error term, ϵ , is usually assumed to be normally distributed with mean 0 and variance σ^2 .

Note that the combined array allowed estimation of interactions among the control variables as well as interactions between the control and noise variables, if the design is constructed to support this. Other authors who used the single-model approach were Shoemaker et al. (1991), Myers et al. (1992a), and Borkowski and Lucas (1997) among several others. Shoemaker et al. (1991) illustrated, through an example, the superiority of the combined array approach over the crossed arrays from the standpoint of allowing the estimation of control-by-control, control-by-noise, and noise-by-noise two-factor interactions. The initial motivation for using the combined array approach is the substantial reduction in the size of the experiment. In spite of this reduction in size, a carefully selected combined array will nonetheless provide better information about the effects of control and noise variables on the response than can be obtained from a larger combined array. In contrast, the crossed array approach facilitates estimation of many effects that are unlikely to be important. Under the combined array approach, the size of the experiment can be made smaller without sacrificing important information.

Myers et al. (1992a) showed that the single response model of Welch et al. (1990) fitted to the response could be used to model the dual response surfaces for the mean and the variance,

$$E[y(x, z)] = \beta_{0s} + g'(x)\beta_s,$$

and

$$\text{Var}[y(x, z)] = (\delta' + g'(x)\Delta)\text{Var}(z)(\delta' + g'(x)\Delta)' + \sigma^2,$$

where $\text{Var}(z)$ denotes the variance–covariance matrix of the vector of the noise variables z . To obtain the control settings for which the variance is minimized and the mean response remains close to the prespecified target, say, T , Myers et al. (1992a) minimized the criterion, $E[\hat{y}(x, z) - T]^2 = (E[\hat{y}(x, z)] - T)^2 + \text{Var}[\hat{y}(x, z)]$, with respect to x , where $\hat{y}(x, z)$ is the predicted response. In their work, Myers et al. (1992a) assumed that effects of both control and noise variables were fixed in the experiment. However, in reality the effects of the noise variable are not fixed but random in the process. Later, to incorporate the random nature of the noise variables, Khuri (1996b) proposed a mixed-effects model where the control factor

was fixed but the noise factor was assumed to be random. For illustration, he presented an example of wafers used in the semiconductor industry for studying the effects of several factors on resistance in computer chips. Since the wafers were drawn from a large lot, their effects in the model should be considered as random.

5.6 GLMs

Most traditional response surface techniques were developed within the framework of linear models under the strong assumption of continuous responses (quite often, normally distributed) with uncorrelated errors and equal error variances. However, very often, data from clinical/epidemiological trials do not satisfy these assumptions. In clinical/epidemiological trials, we come across data that are more variable, and the assumption of homogeneous error variances is not satisfied. In dose–response experiments, we come across responses that are binary in nature, for example, in Phase I clinical trials, we get data in the 1/0 format where 1 represents a toxic side effect in the subject, while 0 says that there was no toxic effect. For such data, statistical analysis using GLMs would be more appropriate. GLMs were introduced by Nelder and Wedderburn (1972) as an extension of the class of linear models. Under the framework of GLMs, discrete as well as continuous responses can be accommodated, and normality and constant variances are no longer a requirement for the response. A classic reference for GLMs is the book by McCullagh and Nelder (1989) (see also Lindsey 1997, Dobson 2001, McCulloch and Searle 2001, and Myers et al. 2002). Chapter 13 provides a detailed discussion on GLMs.

GLMs are usually specified by the following three components: (1) *distributional component*, a sample of n independent random variables, y_1, \dots, y_n , from an exponential family; (2) *linear predictor*, a linear regression function, η , called the linear predictor, in k control variables, x_1, \dots, x_k , of the form $\eta(x) = f'(x)\beta$, where $f(x)$ is a known vector function of x and β is a vector of unknown parameters; and (3) *link function*, $\eta(x) = g(\mu(x))$, where $\mu(x)$ is the mean response and g is a known monotone differentiable function whose inverse exists and is denoted by h . Estimation of β is usually based on the method of maximum likelihood, which is usually carried out using an iterative weighted least-squares procedure (see McCullagh and Nelder 1989, pp. 40–44). Using the resulting estimate of β , an estimate of the mean response and approximate estimates of its variance and bias can be obtained (for details, see Robinson and Khuri 2003).

5.6.1 Choice of Designs for GLMs

Unlike linear models, the prediction variances and the mean squared error of predictions (MSEPs) of the mean response both depend on the unknown parameter vector β . Thus, to choose a good design, the experimenter requires some prior knowledge on β . This dependence causes a great difficulty in the construction and choice of designs for GLMs. By the choice of a good design, we mean the determination of the settings of the control variables and the number of experimental units to be assigned at these particular settings so that the prediction variance or the MSEP are low. A detailed review of design issues for GLMs is given in Khuri et al. (2006). Chapter 13 contains more details on selecting optimal designs in GLMs.

We next discuss some common approaches to address that the assessment of designs depends on the unknown parameters for most GLM models.

5.6.1.1 *Locally Optimal Designs*

In this approach, the experimenter assumes certain initial values for the unknown parameters. Using these initial values, designs are found based on various optimality criteria such as D -optimality or A -optimality. Since the optimal designs depend on the initial values for the parameters, they are called locally optimal. Research on locally optimal designs may be categorized into a study of binary data using logistic regression and a study of Poisson count data. Abdelbasit and Plackett (1983), Khan and Yazdi (1988), Sitter and Wu (1993), and Mathew and Sinha (2001), among others, are some of the early references on locally optimal designs for logistic regression using a single response. More recently, Biedermann et al. (2006, 2007) found locally optimal designs for logistic regression models. Dror and Steinberg (2006) used clustering techniques to find locally D -optimal designs with several predictor variables, but only a single binary response. Yang et al. (2011) extended the algebraic approach of Yang and Stufken (2009) to construct optimal designs for GLMs with multiple covariates. A recent work by Biedermann and Woods (2011) proposed optimal designs for generalized nonlinear models, where the linear predictor is replaced by a nonlinear function in the parameters. For multiresponse binary data, very little work can be found in the literature. For bivariate responses, Heise and Myers (1996) studied optimal designs, while Zocchi and Atkinson's (1999) work was based on multinomial logistic models. For Poisson count data, optimal results were discussed by Liski et al. (2002) and Wang et al. (2006). Russell et al. (2009) considered optimal designs for a Poisson regression model with a log link and multiple predictor variables.

5.6.1.2 *Two-Stage Sequential Designs*

If "good" initial values of the parameters are not available, then the optimal design selected for those parameter values may not be very efficient. Keeping this in mind, Wu (1985), Sitter and Forbes (1997), and Sitter and Wu (1999) proposed that instead of relying solely on the initial values of the parameters and selecting an optimal design, the experiment should be divided into two stages: In the first stage, estimates of the parameters are computed. These estimates are then used to find additional design points in the second stage so that the combined design is "good" (Sitter and Forbes 1997). More recently, Dror and Steinberg (2008) proposed sequential optimal designs for GLMs that can be applied to cases with multiple linear predictors and can also be used with any GLMs and not just those with binary responses.

5.6.1.3 *Bayesian Designs*

The optimality criteria for Bayesian designs are usually the integrated versions of the classical optimality criteria like D - or A -optimality, where the integration is with respect to the prior distribution of the parameter vector. Due to the computational difficulties in finding the exact posterior distribution of β , most of the Bayesian criteria are based on normal approximations to the posterior distribution. The most common form of the normal approximation used is that the parameter vector β follows the normal distribution

with mean $\hat{\beta}$ and variance $[nI(\hat{\beta}, \zeta)]^{-1}$, where $\hat{\beta}$ is the MLE of β , ζ is the design measure, and I is the expected information matrix. Chaloner and Larntz (1989) proposed the following Bayesian optimality criterion analogous to D -optimality, $E_{\beta}[\log \det I(\beta, \zeta)]$, where the expectation is taken according to the prior distribution of β . Some of the papers illustrating the Bayesian approach in designs for GLMs include Zacks (1977), Chaloner and Larntz (1989, 1992), Atkinson et al. (1993), Dette and Sperlich (1994), Chaloner and Verdinelli (1995), and Mukhopadhyay and Haines (1995). More details can be found in Chapter 13.

5.6.1.4 Quantile Dispersion Graphs (QDGs) Approach

In this approach, several designs for GLMs are compared based on the distribution of the MSEF corresponding to each design. A description of the distribution of the MSEF is given in terms of its quantiles. Since the MSEF depends on the unknown parameters, the quantiles of the MSEF are also functions of the unknown parameters. The dispersion of these quantiles over a certain parameter space is determined and then depicted by the so-called quantile dispersion graphs (QDGs). If an initial data set is available, it is used to find the confidence region of the parameter vector. Some values are then randomly selected from the confidence region to address the issue of dependence of the MSEF on the parameters. Robinson and Khuri (2003) were the first to implement the QDGs to find designs for a single binary response using a logistic regression model. Their work was extended by Khuri and Mukhopadhyay (2006) to Poisson count data with a log-linear link function. Both the papers by Robinson and Khuri (2003) and Khuri and Mukhopadhyay (2006) discussed situations with several control variables. Mukhopadhyay and Khuri (2008a) applied QDGs to find designs for multivariate GLMs. Their methodology was illustrated by a multinomial response in a combination drug therapy study on male mice.

5.6.2 Model Misspecification in GLMs

The functional form of the true relationship between the response and the control variables affecting it is usually unknown in practice. The experimenter assumes a model of a certain form to approximate the unknown relationship and bases all derived conclusions on such a model. However, there always remains a possibility that the assumed form of the model is incorrect and might be vastly different from the true relationship. In most cases, inferences drawn are affected by the assumed model form and are therefore incorrect. This is known as model misspecification. Chapter 20 covers robust designs.

In the case of GLMs, the form of the true linear predictor is usually unknown. Suppose that the experimenter assumes a linear predictor of the form

$$\eta(\mathbf{x}) = \mathbf{f}'(\mathbf{x})\beta. \quad (5.27)$$

An estimate of β is obtained under this model form and then used to estimate the mean response. Suppose that the true linear predictor is of the form

$$\eta_T(\mathbf{x}) = \mathbf{f}'(\mathbf{x})\beta + v(\mathbf{x}), \quad (5.28)$$

where $v(\mathbf{x})$ is an unknown function, possibly highly nonlinear. Thus assuming a model of the form (5.27) when in fact (5.28) is the true model gives a misspecified linear predictor. There may also be cases when the link function and/or the linear predictor together

are misspecified. For example, the experimenter may assume a logistic link for a binary response when the true link function is of the probit type.

In the context of linear models, several authors considered the effect of model misspecification on design selection. However, not much work has been done on model misspecification in nonlinear models or GLMs. Sinha and Wiens (2002) studied model misspecification for nonlinear models. In recent years, selecting designs robust to model misspecification have been extended to GLMs by Abdelbasit and Butler (2006), Woods et al. (2006), and Dror and Steinberg (2006). For logistic regression models, Adewale and Wiens (2009) developed an average mean squared error criterion that generates designs insensitive to possible misspecifications in the linear predictors. Their work was later extended by Adewale and Xu (2010) where misspecifications due to both linear predictors and link functions were also studied. More recently, Mukhopadhyay and Khuri (2012) applied the QDGs approach to compare designs for GLMs in the presence of model misspecification of the type given by model (5.28). Additional details concerning misspecifications in GLMs can be found in Chapter 20.

5.6.3 Further Extensions

Very recently, Woods and van De Ven (2011) proposed methods for constructing D-optimal designs for GLMs with blocks. They considered experiments with an exponentially distributed response represented by a model fitted by using the method of generalized estimating equations.

Although several researchers have studied designs for GLMs, very little work has been done in the context of generalized linear mixed models (GLMMs). These models are used in analyzing clustered correlated data in binary and count response studies. To define a GLMM, we again need to specify three components (Sinha and Xu 2011): Elements of the response vector $\mathbf{y} = (y_1, y_2, \dots, y_n)'$, conditional on the random vector \mathbf{v} , are independent and follow a distribution in the exponential family. The linear predictor is $\eta(x) = \mathbf{f}'(x)\boldsymbol{\beta} + \mathbf{z}'\mathbf{v}$, where $\mathbf{f}(x)$ is a row of the model matrix \mathbf{X} and \mathbf{z} is a row of the matrix \mathbf{Z} for the random effects. The linear predictor is related to the mean response through a link function g , where the inverse of g exists. Sinha and Xu (2011) constructed D-optimal sequential designs for the analysis of longitudinal or repeated measurements data using GLMMs. The performance of their proposed designs was studied using various simulation studies.

5.7 Multiresponse Experiments

Response surface methods are quite often concerned with experiments in which several responses are measured for each setting of a group of control variables. Such experiments are known as *multiresponse experiments*. Some popular examples of multiresponse experiments involve taking measurement of the efficacious and the toxic effects of a drug when administered to a subject or measurements of yield and cost of a product in a manufacturing experiment. Schmidt et al. (1979) investigated the effects of cysteine and calcium chloride on the textural and water-holding characteristics of dialyzed whey protein concentrates gel systems. This experiment involved four response variables and is therefore a multiresponse experiment. Another example concerning a multiresponse experiment was

described in Evans et al. (1982) who considered data of seed-germination percentages after 4 weeks of incubation of four plant species in response to 55 alternating and constant-temperature regimes in dark laboratory germinators. Many other multiresponse scenarios were cited in the works of Hill and Hunter (1966), Myers et al. (1989), and Khuri (1996a). In a multiresponse experiment, the measured response values are usually correlated within a single run, but are otherwise assumed to be independent. Specifically, responses taken from a single experimental unit or subject are modeled to allow for correlations, while those taken from different units are regarded as independent. Thus, the correlation structure existing among the responses should be taken into account in any analysis concerning the multiresponse data. This includes the selection of an experimental design that depends on the variance-covariance matrix of all the responses, as will be seen in Section 5.7.1.

Suppose there are q responses measured for each setting of k control variables, $\mathbf{x} = (x_1, \dots, x_k)'$, and n is the total number of experimental runs. Using a polynomial regression model, we can write

$$\mathbf{Y}_i = \mathbf{X}_i \boldsymbol{\beta}_i + \boldsymbol{\epsilon}_i, \quad i = 1, \dots, q,$$

where \mathbf{Y}_i is an $n \times 1$ vector corresponding to the i th response, \mathbf{X}_i is an $n \times p_i$ matrix of rank p_i of known functions of the setting of the control variables, $\boldsymbol{\beta}_i$ is a $p_i \times 1$ vector of unknown parameters, and $\boldsymbol{\epsilon}_i$ is a random experimental error vector such that $E(\boldsymbol{\epsilon}_i) = 0$, $\text{Var}(\boldsymbol{\epsilon}_i) = \sigma_{ii} \mathbf{I}_n$ and $\text{Cov}(\boldsymbol{\epsilon}_i, \boldsymbol{\epsilon}_j) = \sigma_{ij} \mathbf{I}_n$, $i \neq j = 1, \dots, q$. Combining the aforementioned q equations, we get the multiresponse model

$$\tilde{\mathbf{Y}} = \mathbf{X} \boldsymbol{\beta} + \boldsymbol{\epsilon},$$

where $\tilde{\mathbf{Y}} = [\mathbf{Y}'_1, \mathbf{Y}'_2, \dots, \mathbf{Y}'_q]'$, $\boldsymbol{\beta} = [\boldsymbol{\beta}'_1, \boldsymbol{\beta}'_2, \dots, \boldsymbol{\beta}'_q]'$, $\boldsymbol{\epsilon} = [\boldsymbol{\epsilon}'_1, \boldsymbol{\epsilon}'_2, \dots, \boldsymbol{\epsilon}'_q]'$, and $\mathbf{X} = \text{diag}(\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_q)$ is a block diagonal matrix. The best linear unbiased estimator of $\boldsymbol{\beta}$ is given by $\hat{\boldsymbol{\beta}} = (\mathbf{X}' \Delta^{-1} \mathbf{X})^{-1} \mathbf{X}' \Delta^{-1} \tilde{\mathbf{Y}}$, where $\Delta = (\boldsymbol{\Sigma}^{-1} \otimes \mathbf{I}_n)^{-1}$ and $\boldsymbol{\Sigma}$ is the variance-covariance matrix whose (i, j) th element is σ_{ij} ($i, j = 1, 2, \dots, q$). If $\boldsymbol{\Sigma}$ is unknown, an estimate of it can be used. Zellner (1962) provided the estimate $\hat{\boldsymbol{\Sigma}}$ whose (i, j) th element is given by

$$\hat{\sigma}_{ij} = \frac{1}{n} \left\{ \mathbf{Y}'_i \left[\mathbf{I}_n - \mathbf{X}_i (\mathbf{X}'_i \mathbf{X}_i)^{-1} \mathbf{X}'_i \right] \left[\mathbf{I}_n - \mathbf{X}_j (\mathbf{X}'_j \mathbf{X}_j)^{-1} \mathbf{X}'_j \right] \mathbf{Y}_j \right\}, \quad i, j = 1, 2, \dots, q.$$

More details concerning the use of this estimator can be found in Khuri and Cornell (1996, p. 254).

For a general multiresponse model, Box and Draper (1965) used Bayesian techniques to estimate the parameter vector $\boldsymbol{\beta}$. They assumed normal responses with the variance-covariance matrix $\boldsymbol{\Sigma}$ and noninformative prior distributions for $\boldsymbol{\beta}$ and $\boldsymbol{\Sigma}$. Box et al. (1973) later showed that Box and Draper's estimation criterion failed to give meaningful results in the presence of exact linear dependencies among the responses. For example, in a chemical mechanism, an exact linear relationship exists among the amounts of the substituents in order to maintain the carbon balance (Khuri and Cornell 1996, Section 7.2).

5.7.1 Designs for Multiresponse Experiments

The design selection problem in a multiresponse experiment is more complex than in the case of a single response. In choosing a design for a multiresponse experiment, one should consider all the responses simultaneously; otherwise, the selected design may be efficient for one response but inefficient for the others. Draper and Hunter (1966) used Bayesian methods with a noninformative prior distribution for β to obtain a design criterion for the estimation of the parameters. Fedorov (1972) introduced an algorithm for the sequential construction of a D -optimal design for multiresponse models. However, his method required knowledge of the variance–covariance matrix Σ . Wijesinha and Khuri (1987a) suggested a solution to the problem of unknown Σ by developing a sequential procedure for the construction of the design where an updated estimate of Σ was obtained at each stage of the sequential procedure. Wijesinha and Khuri (1987b) also introduced a design criterion for constructing optimal designs to improve the power of the multivariate lack of fit test for a linear multiresponse model. In clinical trials, it is of interest to find the dose levels/combinations of dose levels that optimize the toxicity and safety responses simultaneously. Recently, several papers (viz., Dragalin and Fedorov 2008, Dragalin et al. 2008 and Fedorov et al. 2011) have been published that propose adaptive/sequential procedures based on optimal design methods as proposed in Fedorov (1972) and Cook and Fedorov (1995) for finding the dose level when two responses are involved.

5.7.2 Multiresponse Optimization

A primary objective in RSM is to determine the settings of the control variables that give rise to an optimum response. However, when there is more than one response, the problem of finding common settings of the control variables that will optimize all the responses is usually not possible since the settings that optimize one response will generally not optimize the others. Consider a situation where a researcher is studying the effect of various fertilizing agents on the yield and cost of a certain crop product. He or she is interested in determining the values of the fertilizers that maximize the crop yield while minimizing the cost of production. It is very likely in this situation for the yield to be maximized at certain settings of the fertilizers, while the cost is minimized for a different set of values of the fertilizers. In this case, determining the settings of the fertilizers that are favorable to both responses, yield and cost, is not possible.

Lind et al. (1960) considered superimposing contours of all the responses to determine the region where all the responses are “near” optimal. However, this method becomes very cumbersome and difficult to use as the numbers of control and response variables increase. A more general approach is to set up the multiresponse optimization problem as a constrained optimization problem. Suppose that there are q responses from which one response is selected, say, Y_1 , identified as the primary response. Then the constrained optimization problem may be framed as in Myers et al. (2004), by finding the settings of $x = (x_1, x_2, \dots, x_k)' \in \mathcal{R}$ that optimize the predicted response corresponding to Y_1 such that the other $(q - 1)$ responses lie in the intervals $[l_i, u_i]$, $i = 2, \dots, q$. Myers and Carter (1973) studied the constrained optimization problem where there are two responses, labeled as the primary and the secondary responses. They fitted appropriate models to the two responses and then determined the settings of the control variables that optimized the estimated primary response while keeping the secondary response within a certain interval. Some extensions for optimization of nondegenerate and degenerate dual response systems within a spherical experimental region were suggested by Del Castillo et al. (1997, 1999) and Fan (2000).

Harrington (1965) introduced the desirability approach for solving the multiresponse optimization problem. Each predicted response was transformed into a desirability function using an exponential transformation that increased as the desirability of the corresponding property increased. Derringer and Suich (1980) proposed more general transformations that provided the user with greater flexibility. The desirability functions for all the responses were then incorporated into a single function, usually the geometric mean, which was then maximized with respect to the control variables. Del Castillo (1996) modified the desirability approach so that the desirability functions were differentiable and more efficient gradient-based optimization methods could be used. Kim and Lin (2000) used the desirability approach for multiresponse optimization; however, they did not assume any form or degree of the estimated response models. Later, Wu (2005) proposed a double-exponential transformation for the desirability functions accounting for correlations among the responses.

Khuri and Conlon (1981) argued that the desirability approach does not take into account correlations that may exist among the responses and instead suggested the *generalized distance approach*. Since the individual optima for the responses are not usually attained at the same settings of the control variables, the generalized distance approach determined a set of conditions on the control variables that were “favorable” to all the responses. The distance of the vector of individual optima of the predicted responses from the vector of estimated mean responses, weighted by the inverse of the variance–covariance matrix of the predicted responses, was determined. Minimizing this distance function with respect to the control variables, Khuri and Conlon arrived at a set of conditions for a so-called compromise optimum. Other approaches for multiresponse optimization based on a loss function, usually the mean squared error loss, taking into account the heterogeneity in responses, were discussed by Pignatiello (1993), Ames et al. (1997), Vining (1998), and Ko et al. (2005). Multiresponse optimization for GLMs using the distance approach was more recently discussed by Mukhopadhyay and Khuri (2008c).

5.8 Graphical Procedures for the Evaluation and Comparison of Response Surface Designs

In this section, we discuss graphical procedures for design comparison. Unlike single-valued criteria like D -optimality that only measure the performance of a design at a single point in the experimental region, these techniques enable comparison of designs over the entire region. A design is said to be “good” if it has low and stable values of prediction variance or MSEF throughout the experimental region. Graphical procedures give a more complete picture of the design’s prediction capability since it may vary very well when measured at different parts of the experimental region. However, these graphical procedures are not used to develop better designs but only to compare several designs and choose the best among them. In recent years, graphical techniques for evaluating the prediction capability of designs throughout the experimental region have become popular. We next discuss three such graphical techniques.

Giovannitti-Jensen and Myers (1989) and Myers et al. (1992b) proposed 2D plots known as *variance dispersion graphs* (VDGs), which display the minimum, maximum, and average prediction variances for a given design over spheres of varying radii centered at the origin.

By changing the radius of a sphere, these graphs are able to cover the entire region, \mathcal{R} , from the center to the boundary. Later, various authors, including Vining (1993) and Borkowski (1995), worked on construction of VDGs. These graphs were also used for studying process robustness by Borrer et al. (2002). Trinca and Gilmour (1998) extended VDGs to be used in response surface designs with blocks. It should be noted that the VDGs depend on the form of the fitted model as well as on the coding of the control variables used in the model. This follows from the dependence of the prediction variance on the model matrix that is influenced by the coding convention.

Zahran et al. (2003) argued that VDGs were unable to provide complete information about the distribution of the scaled prediction variance (SPV) on a given sphere (SPV is the prediction variance multiplied by the number of experimental runs and divided by the experimental error variance). They suggested that to get all the useful information, the volume associated with the SPV on a given sphere should be considered and introduced the *fraction of design space* (FDS) criterion. The FDS values for each design were plotted against SPV to obtain the FDS plots. FDS plots have been used by Goldfarb et al. (2004) in mixture experiments and by Liang et al. (2006) for split-plot designs. For examining design robustness related to SPV values, Ozol-Godfrey et al. (2005) used FDS plots. These plots were also used to assess different designs for GLMs under parameter misspecification by Ozol-Godfrey et al. (2008).

Khuri et al. (1996) proposed the comparison of several response surface designs using the distribution of the SPV on a given sphere based on its quantiles. Instead of comparing average or extreme values of the SPV, they reasoned that the entire distribution of the SPV should be taken into account. Using several examples, they showed that two designs may have similar VDG patterns, but the distributions of their SPV in terms of quantiles are different. They determined the quantiles of the SPV using a set of points on a sphere of radius r inside the experimental region. Using different values of r , they were able to cover the entire experimental region. For each r , the quantiles for a given design were then plotted against their corresponding probabilities to obtain the quantile plots. Khuri et al. (1999) used quantile plots to compare designs for a constrained mixture region. More recently, Mukhopadhyay and Khuri (2008b) used quantile plots to compare the robustness of several response surface designs to model misspecification on the basis of the MSEF. The proposed approach uses four criterion functions derived from the MSEF.

QDGs were first proposed by Khuri (1997) for comparing several designs for estimating variance components in an ANOVA situation, where the factors are discrete rather than continuous. Later, Lee and Khuri (1999, 2000) extended the use of QDGs to unbalanced random one-way and two-way models, respectively. Their QDGs were based on both ANOVA and maximum likelihood estimators. QDGs were also used by Khuri and Lee (1998) in the comparison of designs for nonlinear models and for three-fold nested designs by Jung et al. (2008). More recently, Saha and Khuri (2009) compared designs for response surface models with random block effects based on QDGs.

Robinson and Khuri (2003) extended the use of QDGs to compare designs for logistic regression models. QDGs were later used for count data by Khuri and Mukhopadhyay (2006). In general, designs for GLMs were compared on the basis of the distribution of their MSEF values over the entire experimental region. The distribution of the MSEF was considered in terms of its quantiles. Since in the case of GLMs the MSEF depends on unknown parameters, then so do the quantiles. For a given design, the quantiles of the MSEF were obtained within the experimental region, \mathcal{R} . This region was partitioned into several concentric regions, denoted by \mathcal{R}_λ , using a shrinkage parameter λ (where by shrinking we

mean reducing the size of the region). Suppose we denote the p th quantile corresponding to a design D on the surface of \mathcal{R}_λ by $Q_D(p, \beta, \lambda)$, where β is the unknown parameter vector of the linear predictor. To address the dependence of the quantiles on the parameters, a parameter space was considered from which a subset of parameter values, denoted by C , was chosen. Values of $Q_D(p, \beta, \lambda)$ were computed for $\beta \in C$. The minimum and the maximum values, $Q_D^{\min}(p, \beta, \lambda)$ and $Q_D^{\max}(p, \beta, \lambda)$, of the p th quantile with respect to the parameter values in C were computed. The QDGs for each design were then obtained by plotting $Q_D^{\min}(p, \beta, \lambda)$ and $Q_D^{\max}(p, \beta, \lambda)$ against p . Using the QDGs approach, a design is said to be good if it has small and close values of $Q_D^{\min}(p, \beta, \lambda)$ and $Q_D^{\max}(p, \beta, \lambda)$ over the range of p . Small values of the minimum and maximum quantiles imply small MSEF, while close values indicate robustness to the changes in the parameter values. Mukhopadhyay and Khuri (2008a) also used QDGs to evaluate and compare designs for multivariate GLMs. Since in the multivariate setup the MSEF is a matrix, they used a scalar-valued function of the MSEF, namely, the largest eigenvalue of the MSEF matrix (EMSEF), as their comparison criterion. As before, EMSEF also depends on the unknown parameter vector. The dispersion of the quantiles of the EMSEF over the space of the unknown parameters was determined and then depicted by the QDGs.

5.8.1 Numerical Example

We consider a data set from a manufacturing process studying the effect of temperature and processing time on product quality. This data set was also considered by Ozol-Godfrey et al. (2008) for comparing designs using FDS plots. The response variable is binary in nature taking the value 1 if the product is acceptable and zero otherwise. The initial design, D_1 , is a 2^2 factorial design with five replications at each of the corner points. The factors are both recoded to lie in the interval $[-1, 1]$. Ozol-Godfrey et al. (2008) arrived at the following fitted model:

$$\log \left(\frac{\hat{\pi}(x)}{1 - \hat{\pi}(x)} \right) = 2 + 0.5x_1 - 0.8x_2 + 0.3x_1x_2.$$

Design D_1 is compared with another design, denoted by D_2 , which is a replicated (five times) 2^2 factorial design with a center point (five replications). Through FDS plots, the authors conclude that design D_1 performs (in terms of SPV) better than D_2 in most of the design space. Adding center runs to a design helps the performance of the design (lower prediction variances) at the center. Using graphical techniques, Vining and Myers (1991) considered the impact of adding center runs to a 2^3 factorial design using a first-degree model with a full second-degree model as an alternative model. They demonstrated that the addition of center runs increases the MSEF near the perimeter of the region.

Here, we use QDGs to compare the two designs and try to see if we can obtain any additional information about the performances of the two designs in the design space. The designs are compared on the basis of the quantiles of their MSEF as well as the SPV. The quantiles are computed on the surfaces of concentric regions \mathcal{R}_λ to cover the entire design space. To address the dependence of the quantiles of the MSEF/SPV on the unknown parameter values, we consider for each of the four parameters in the model a set of five points consisting of the point estimate of the parameter value in addition to two pairs of points that are symmetric with respect to the point estimate. The points in the first pair are at a distance of one standard error from the estimate; in the second pair, the points are two standard errors away from the estimate. These 625 points constructed by selecting the

aforementioned five points for each parameter is denoted by C. The QDGs for the SPV are presented in Figure 5.1; those for the MSEP are shown in Figure 5.2.

From Figure 5.1, we note that the prediction capability of design D_2 is better than D_1 for small values of λ , that is, near the center of the region, while for higher values of λ , D_1 performs better than D_2 . However, from Figure 5.2 (QDGs based on MSEP), we get a different message. We see that design D_2 performs better than D_1 over the entire design space, when compared on the basis of MSEP. We also observe that the performances of both designs deteriorate as we move out toward the boundary. From the gap between the minimum and the maximum quantiles of both designs, we can say that the designs are sensitive to the changes in the parameter values. So we see that basing our comparison




FIGURE 5.1 QDGs comparing designs D_1 (2^2 factorial) and D_2 (2^2 factorial with center runs) based on SPV.

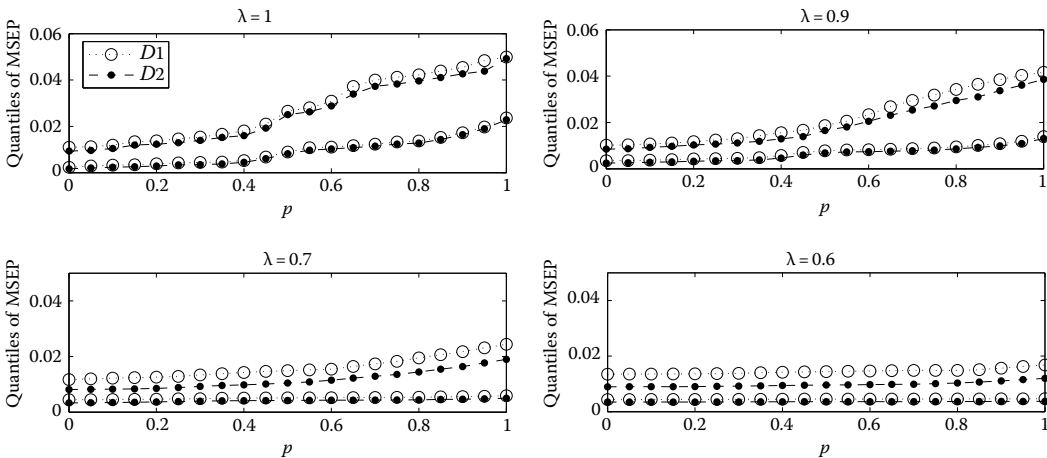


FIGURE 5.2 QDGs comparing designs D_1 (2^2 factorial) and D_2 (2^2 factorial with center runs) based on MSEP.

on MSEP instead on SPV, we get a different result. This is due to the fact that the MSEP takes into account not only the variance but also the bias in estimation. Also, unlike the FDS plots, the QDGs could tell exactly where in the design region the performance of D_2 is better than that of D_1 . Additionally, the QDGs assessed the robustness of the designs to the parameter values.

5.9 Mixture Designs

A mixture experiment is one in which the response depends only on the relative proportions of the ingredients, or components, present in a mixture. For example, the tint strength of a house paint is affected by the proportions of two known pigments and a certain solvent in the paint, but does not depend on the amount of the mixture. Also, the tensile strength of stainless steel depends on the proportions of iron, copper, nickel, and chromium in the alloy.

If x_i denotes the proportion of the i th component of k components, then $x_i \geq 0$ for $i = 1, 2, \dots, k$, and

$$\sum_{i=1}^k x_i = 1. \tag{5.29}$$

The constraint in (5.29) on the values of x_i causes mixture experiments to be different from general response surface experiments where the values of the control variables do not necessarily add up to 1. Because of this constraint, the factor space of interest, or experimental region \mathcal{R} , is a subset of a $k - 1$ dimensional space (or a $(k - 1)$ simplex). For example, for $k = 2$, \mathcal{R} is a line segment; for $k = 3$, it is an equilateral triangle; and for $k = 4$, it is a regular (equal sided) tetrahedron.

In a mixture experiment, polynomial models, typically used in RSM, are expressed in light of the restriction in (5.29). For example, for $k = 2$, the second-degree model,

$$\eta = \beta_0 + \sum_{i=1}^2 \beta_i x_i + \beta_{12} x_1 x_2 + \sum_{i=1}^2 \beta_{ii} x_i^2,$$

can be rewritten, after replacing x_1^2 by $x_1(1 - x_2)$ and x_2^2 by $x_2(1 - x_1)$, as

$$\eta = \beta_0(x_1 + x_2) + \sum_{i=1}^2 \beta_i x_i + \beta_{12} x_1 x_2 + \beta_{11} x_1(1 - x_2) + \beta_{22} x_2(1 - x_1),$$

which can be expressed as

$$\eta = \beta'_1 x_1 + \beta'_2 x_2 + \beta'_{12} x_1 x_2,$$

where $\beta'_1 = \beta_0 + \beta_1 + \beta_{11}$, $\beta'_2 = \beta_0 + \beta_2 + \beta_{22}$, and $\beta'_{12} = \beta_{12} - \beta_{11} - \beta_{22}$. For the sake of simplicity, we can drop the primes from the betas in this model.

Using the restriction in (5.29), Scheffé (1958) introduced the following so-called canonical forms of mixture polynomial models:

Linear

$$\eta = \sum_{i=1}^k \beta_i x_i. \quad (5.30)$$

Quadratic

$$\eta = \sum_{i=1}^k \beta_i x_i + \sum_{i<j} \beta_{ij} x_i x_j. \quad (5.31)$$

Full cubic

$$\eta = \sum_{i=1}^k \beta_i x_i + \sum_{i<j} \beta_{ij} x_i x_j + \sum_{i<j} \delta_{ij} x_i x_j (x_i - x_j) + \sum_{i<j<k} \beta_{ijk} x_i x_j x_k. \quad (5.32)$$

Special cubic

$$\eta = \sum_{i=1}^k \beta_i x_i + \sum_{i<j} \beta_{ij} x_i x_j + \sum_{i<j<k} \beta_{ijk} x_i x_j x_k. \quad (5.33)$$

Since the mixture factor space is a simplex, design points for a mixture experiment must be at the vertices, on the edges or faces, or in the interior of the simplex. Scheffé (1958) suggested some designs for fitting the polynomial models he had already introduced. His choice of points was initially motivated by the need to simplify the derivation of the least-squares estimates of the models' parameters by solving simple equations. Of course, such a need is no longer necessary. Still, however, his designs are quite adequate and suitable for practical experimentation.

For fitting Scheffé's canonical polynomial models, Scheffé (1958) suggested *simplex-lattice* designs. By definition, a $\{k, m\}$ simplex-lattice design for k components consists of all design points that can be formed by using for each of the k components, x_1, x_2, \dots, x_k , the $(m + 1)$ equally spaced levels

$$x_i = 0, \frac{1}{m}, \frac{2}{m}, \dots, 1, \quad i = 1, 2, \dots, k$$

provided that $\sum_{i=1}^k x_i = 1$. For example, the $\{3, 2\}$ simplex lattice consists of the following six points on the boundary of the triangular factor space:

$$(x_1, x_2, x_3) = (1, 0, 0), (0, 1, 0), (0, 0, 1), \left(\frac{1}{2}, \frac{1}{2}, 0\right), \left(\frac{1}{2}, 0, \frac{1}{2}\right), \left(0, \frac{1}{2}, \frac{1}{2}\right).$$

In general, the number of design points in the $\{k, m\}$ simplex lattice is

$$\binom{k+m-1}{m} = \frac{(k+m-1)!}{m!(k-1)!}.$$

An alternative design to the $\{k, m\}$ simplex lattice is the simplex-centroid design suggested by Scheffé (1963). It contains $2^k - 1$ points consisting of the k permutations of $(1, 0, 0, \dots, 0)$, the $\binom{k}{2}$ permutations of $(\frac{1}{2}, \frac{1}{2}, 0, 0, \dots, 0)$, the $\binom{k}{3}$ permutations of $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}, 0, 0, \dots, 0)$, and so on, down to the centroid point $(\frac{1}{k}, \frac{1}{k}, \dots, \frac{1}{k})$.

The $\{k, m\}$ simplex-lattice and simplex-centroid designs are boundary-point designs in that, with the exception of the overall centroid, all the design points are located on the boundaries (vertices, edges, faces, etc.) of the simplex factor space. Quite often, however, it may be desirable to have more experimental runs in the interior of the simplex where information about complete mixtures can be obtained. It is therefore recommended to augment such designs with axial runs that are positioned inside the simplex factor space. A discussion concerning these axial runs and how they can be selected is given in (Cornell 1990, Section 2.15).

5.9.1 Constraints on the Mixture Components

Quite often in some mixture experiments, values of the mixture components, x_1, x_2, \dots, x_k , may be restricted by certain lower and/or upper bounds. For example, x_i may be required to be positive inside the factor space, that is, the i th components must be present in all the blends, as in the case of a fruit drink that is required to contain at least a certain proportion of orange juice. If lower bounds are placed on all mixture components, then we have

$$x_i \geq L_i \geq 0, \quad i = 1, 2, \dots, k, \tag{5.34}$$

where the L_i 's are given values such that $L = \sum_{i=1}^k L_i < 1$. The resulting factor space is reduced to a subspace of the original simplex region that is also a simplex. In order to maintain summability to one of the mixture components inside the reduced space, the following transformation is needed:

$$x'_i = \frac{x_i - L_i}{1 - L}, \quad i = 1, 2, \dots, k, \tag{5.35}$$

where x'_1, x'_2, \dots, x'_k are called *pseudocomponents*. Note that $x'_i \geq 0$ for all i and $\sum_{i=1}^k x'_i = 1$. Thus, under the transformation in (5.35), the reduced mixture space becomes an ordinary mixture space when described in the pseudocomponents.

In case of both lower and upper bounds on the mixture components, we get the following double inequalities:

$$0 \leq L_i \leq x_i \leq U_i \leq 1, \quad i = 1, 2, \dots, k, \tag{5.36}$$

where L_i and U_i are given constants. In this case, the factor space is a convex polyhedron that is more complicated in shape than the simplex-shaped region defined by having only lower bounds. In such cases, computer-generated designs, based, for example, on the D -optimality criterion, can be used to find designs to fit Scheffé's mixture models. In this

respect, the vertices of the constrained region, as well as convex combinations of some of the vertices, are candidates for design points. This gives rise to the so-called extreme vertices designs. See the work by McLean and Anderson (1966). Crosier (1984) derived a formula for calculating the number of vertices of any constrained region.

5.9.2 Slack-Variable Models

Scheffé's polynomial models can alternatively be expressed in terms of $k - 1$ independent variables. This can be done by selecting one mixture component, say, x_i for some i , and expressing it in terms of the remaining mixture components (using the restriction in [5.29]), then substituting it into Scheffé's model. This mixture component is called a *slack variable*, and the resulting model, which now depends only on the remaining mixture components, is called a *slack-variable model*. If a complete Scheffé's model is used, then, in terms of prediction, it makes no difference which component is designated as the slack variable. By a complete model, we mean a model that contains all polynomial terms, up to a certain degree, along with their lower-order terms and cross products. If such a model is used, then it will be equivalent to all its corresponding slack-variable models (see Khuri 2005). By definition, two models are equivalent if the column spaces of their corresponding model matrices (i.e., the X matrix as in (5.7)) are the same. For such models, the predicted response values and corresponding regression sums of squares are identical.

The concept of model equivalence is related to the definition of a well-formulated model. By this, we mean a model with the property that for every one of its polynomial terms, all polynomial terms inferior to it are included in the model (a polynomial term, such as t_1 , is said to be inferior to another polynomial term, say, t_2 , if $t_2 = t_1 t_3$, where t_3 is a polynomial term different from 1). We can then see that a complete Scheffé's model is well formulated.

Well formulation of a model becomes an important consideration when the model is subjected to any variable selection procedure. This is true because such a procedure may select as the "best submodel" of a given size one that is not well formulated. This is undesirable since for such a submodel, the predicted response and the regression sum of squares can change depending on how the control variables in the submodel are coded. Thus, any variable selection procedure should be restricted to submodels that are well formulated.

In particular, if Scheffé's complete model and one of its equivalent slack-variable models are both reduced by deleting some of their terms using the same variable selection procedure, then the best t -variable submodel obtained from Scheffé's model may not be equivalent to the best t -variable submodel obtained from the slack-variable model (t is the number of parameters in the submodel). This is true because the t -dimensional column spaces of the matrices for the two submodels may be different. Note that all submodels considered must be well formulated. Khuri (2005) demonstrated that a Scheffé's submodel may or may not be equivalent to a given slack-variable submodel. Furthermore, the slack-variable submodels may not be equivalent to one another. Thus, when making a comparison among the t -variable Scheffé's and slack-variable submodels using, for example, R^2 , the coefficient of determination; MS_E , the error mean square; and $PRESS$, the prediction error sum of squares, the choice of the best fitting submodel may turn out to be a submodel of the Scheffé type or one of the slack-variable submodels. Thus, this choice depends on which mixture component is selected as a slack variable. If all the aforementioned submodels were equivalent, then they would have the same summary statistics

values (i.e., R^2 , MS_E , and $PRESS$). Khuri (2005) also studied the effect of collinearity on Scheffé's and its slack-variable submodels (a model is said to suffer from collinearity if there is a near linear dependency among the columns of its model matrix). He used certain diagnostic measures of collinearity to discriminate among the submodels (in mixture experiments, collinearity can occur when one or more of the mixture components have very small ranges, which cause the mixture space to be highly constrained). It was concluded that collinearity affected Scheffé's and the slack-variable submodels differently, depending on the size of the fitted submodel.

5.10 Applications of RSM

Mead and Pike (1975) were among the first to explore the use of RSM in biological research. They provided a survey of a large number of papers in biological journals to determine the extent of using RSM ideas. They reported that "not much awareness of current RSM methods was shown" and proposed a "joint development by biologists and statisticians of particular biologically reasonable models for particular practical research problems." This is a good advice since the practical research worker will be more interested in methods that pertain to his or her particular field of application rather than pursuing general results.

Fortunately, RSM has since become more applicable to a wide spectrum of research areas, including those with biological and agricultural applications. The development of new statistical software and the introduction of fast computers made it a lot easier for practitioners to attempt more advanced RSM techniques than was possible before. The food industry, in particular, has been a prime user of RSM since the early 1970s. Myers et al. (1989) devoted two sections to review various applications of RSM in the food and biological sciences.

Since RSM's original work was widely used in the physical and engineering sciences as well as industrial research, more attention was initially given to these areas. This was reflected in the review articles by Hill and Hunter (1966) and Myers et al. (1989). This trend continued in the 1980s with the flurry of interest in the Taguchi method and its response surface alternatives. In the remainder of this section, we shall bring attention to applications in other areas that have not received as much attention, as in the case of the agricultural and food sciences. The cited papers represent only a small sample from a much larger collection.

5.10.1 Applications of RSM in Agricultural Sciences

Edmondson (1991) provided an interesting application of RSM to greenhouse experiments and presented some valuable insights into the use of RSM in an agricultural setting versus an industrial one. Response surface techniques were utilized by Keisling et al. (1984) to predict weed age and future weed size from weed height. The objectives of their study were to (1) utilize response models to generate data for describing weed interference in soybeans, (2) present strategies for estimating multispecies interference, and (3) project yield loss from existing data. The study was designed to produce information to assist soybean producers in recognizing economically detrimental threshold levels of weed infestations that require the initiation of control measures. Broudiscou et al. (1999) investigated the effects of several mineral compounds on feed degradation and microbial growth in a continuous culture system using RSM. The models considered were of the second degree fitted to data generated by a nonstandard design that consisted of 16 points giving seven levels to each of the four

factors in the experiment. The design had good characteristics (by comparison to a CCD with 25 experimental runs, as shown in their Table II on p. 257), was close to being orthogonal, and almost rotatable. The authors used the measure of rotatability by Khuri (1988) to assess the percent rotatability of their design, which turned out to be 99.6 as compared to 89.2 for the CCD. Furthermore, the design was also more G efficient than the CCD.

RSM has received attention for modeling the performance of agronomic experiments. For example, Gilmour and Trinca (2005) showed, using turnip yield data, that due to the complexity in biological systems, second-degree polynomial models may not provide an adequate fit. In these situations, researchers should instead use fractional polynomial response models. They considered a data set from studying the spacing and density effects on turnip yield given by Mead (1988) and fitted a fractional polynomial model. Salawu et al. (2007) used inverse polynomials to model the yield of maize against three control variables, namely, levels of nitrogen, phosphorous, and potassium. A 3^3 factorial design was used, and the experiment was conducted in a randomized complete block layout with two replications per treatment combination. The inverse polynomial model (introduced by Nelder 1966) provided a better fit than the traditional second-degree model. The latter model may produce negative estimates of the yield response, which, of course, must be positive. This shows that taking into account any physical knowledge about the response can be very beneficial when choosing an appropriate model.

5.10.2 Applications of RSM in Food Sciences

Food science has also benefited greatly from the application of RSM to its various areas of research. Diniz and Martin (1996) used RSM to study the effects of pH, temperature, and enzyme–substrate ratio (E/S) on the degree of hydrolysis of dogfish muscle protein. The effects of the hydrolysis variables were estimated using a Box–Behnken design. This design was also utilized by Jiang et al. (2009) to investigate the effects of x_1 = ultrasonic temperature (30°C–70°C), x_2 = power (120–300 W), and x_3 = time (10–50 min) on ultrasonic-assisted extraction for oligosaccharides from longan fruit pericarp (OLFP). Their fitted second-degree model was then used to obtain optimum conditions on x_1 , x_2 , x_3 that maximize the OLFP response. Optimization was also the goal of a study conducted by Cao et al. (2009) to study the effects of temperature, pH, and E/S ratio on the response, degree of hydrolysis (DH) for a marine shrimp called *Acetes chinensis* that was harvested in China. The design used was a CCD for three control variables with $n_0 = 6$ center-point replications and an axial parameter $\alpha = 1.682$. This causes the design to be rotatable. Also, since $n_0 = 6$, the design has the additional uniform precision property (see Table 4.3 in Khuri and Cornell 1996). The results of Cao et al.'s study indicated that hydrolysis of shrimp (*A. chinensis*) resulted in a maximum DH value of about 26.33% under the optimal conditions on temperature, pH, and E/S ratio.

Another optimization experiment was carried out by Zhang et al. (2007) in a study concerning pyridoxine (PN), which is one of the three members of the vitamin B_6 group. It has broad applications in the food industry, cosmetics, and medical supplies. RSM was successfully applied to determine optimum operating conditions for maximum conversion of PN. The control variables were reaction temperature, reaction time, enzyme loading, molar substrate ratio, and water activity. The design used was a CCD whose factorial portion consisted of a one-half fraction of a 2^5 factorial, its axial portion contained 10 points with an axial parameter $\alpha = 2$ and $n_0 = 6$ center-point replications. This design is rotatable since $\alpha = F^{1/4}$ where $F = 16$ is the number of factorial points. It also has the uniform precision

property since $n_0 = 6$ (see Table 4.3 in Khuri and Cornell 1996). A listing of several applications of RSM in the optimization of chemical and biochemical processes was given by Bas and Boyaci (2007). In addition to their review of the recent literature on RSM applications to the aforementioned areas, they also provided a critique concerning the misuses of RSM in some of the reviewed articles.

5.11 New Research Directions in RSM

The area of GLMs is a relatively new area in RSM. We believe that the use of GLMs in RSM will continue to grow and expand. For example, there is a need to develop computational methods to evaluate optimal designs when the dimension of the design space is high. Finding robust and efficient designs in high-dimensional problems will involve formidable computational challenges.

An emerging area is the use of graphical techniques to compare response surface designs. It is known that the minimization of the MSEP criterion depends on the unknown parameters of the linear predictor for a GLM. This makes the choice of designs for GLMs dependent on unknown parameters. The method of QDGs was developed to deal with the design dependence problem, as was described in Khuri and Mukhopadhyay (2006). Mukhopadhyay and Khuri (2012) used QDGs to compare designs for GLMs under a misspecified linear predictor. There is also a need to study the effects of misspecifications of the link function and/or the parent distribution of the data on the shape of the quantile plots of the QDG approach. Furthermore, there is a need to explore the design dependence problem in multiresponse situations involving several response variables that may be correlated.

The use of QDGs has thus far been limited to just comparing several given designs. However, it is important to be able to start with an initial design and then construct a more efficient design by augmenting the initial design with several points selected in a sequential manner from the design space through the use of QDGs. This will be consistent with the sequential nature of RSM. Solving this problem can lead to the development of optimal designs for GLMs using QDGs.

References

- Abdelbasit, K. M., Butler, N. A., 2006. Minimum bias designs for generalized linear models. *Sankhya: The Indian Journal of Statistics* 68, 587–599.
- Abdelbasit, K. M., Plackett, R. L., 1983. Experimental design for binary data. *Journal of the American Statistical Association* 78, 90–98.
- Adewale, A. J., Wiens, D. P., 2009. Robust designs for misspecified logistic models. *Journal of Statistical Planning and Inference* 139, 3–15.
- Adewale, A. J., Xu, X., 2010. Robust designs for generalized linear models with possible overdispersion and misspecified link functions. *Computational Statistics and Data Analysis* 54, 875–890.
- Ames, A. E., Mattucci, N., McDonald, S., Szonyi, G., Hawkins, D. M., 1997. Quality loss function for optimization across multiple response surfaces. *Journal of Quality Technology* 29, 339–346.
- Atkinson, A. C., Chaloner, K., Herzberg, A. M., Juritz, J., 1993. Optimum experimental designs for properties of a compartmental model. *Biometrics* 49, 325–337.

- Bartlett, M. S., Kendall, D. G., 1946. The statistical analysis of variance—Heterogeneity and the logarithmic transformation. *Journal of Royal Statistical Society* 8, 128–133.
- Bas, D., Boyaci, I. H., 2007. Modeling and optimization I: Usability of response surface methodology. *Journal of Food Engineering* 78, 836–845.
- Biedermann, S., Dette, H., Zhu, W., 2006. Optimal designs for dose-response models with restricted design spaces. *Journal of the American Statistical Association* 101, 747–759.
- Biedermann, S., Dette, H., Zhu, W., 2007. Compound optimal designs for percentile estimation in dose-response models with restricted design intervals. *Journal of Statistical Planning and Inference* 137, 3838–3847.
- Biedermann, S., Woods, D., 2011. Optimal designs for generalized non-linear models with application to second-harmonic generation experiments. *Applied Statistics* 60, 281–299.
- Borkowski, J. J., 1995. Spherical prediction-variance properties of central composite and Box-Behnken designs. *Technometrics* 37, 399–410.
- Borkowski, J. J., Lucas, J. M., 1997. Designs of mixed resolution for process robustness studies. *Technometrics* 39, 63–70.
- Borror, C. M., Montgomery, D. C., Myers, R. H., 2002. Evaluation of statistical designs for experiments involving noise variables. *Journal of Quality Technology* 34, 54–70.
- Box, G. E. P., 1952. Multi-factor designs of first order. *Biometrika* 39, 49–57.
- Box, G. E. P., 1985. Discussion of off-line quality control, parameter design, and the Taguchi methods. *Journal of Quality Technology* 17, 198–206.
- Box, G. E. P., 1988. Signal-to-noise ratios, performance criteria, and transformations. *Technometrics* 30, 1–17.
- Box, G. E. P., 1996. Split-plot experiments. *Quality Engineering* 8, 515–520.
- Box, G. E. P., Behnken, D. W., 1960. Some new three-level designs for the study of quantitative variables. *Technometrics* 2, 455–475.
- Box, G. E. P., Draper, N. R., 1959. A basis for the selection of a response surface design. *Journal of the American Statistical Association* 54, 622–654.
- Box, G. E. P., Draper, N. R., 1963. The choice of a second order rotatable design. *Biometrika* 50, 335–352.
- Box, G. E. P., Draper, N. R., 1965. The Bayesian estimation of common parameters from several responses. *Biometrika* 52, 355–365.
- Box, G. E. P., Draper, N. R., 2007. *Response Surfaces, Mixtures, and Ridge Analyses*, 2nd edn. Wiley, Hoboken, NJ.
- Box, G. E. P., Hunter, J. S., 1957. Multifactor experimental designs for exploring response surfaces. *Annals of Mathematical Statistics* 28, 195–241.
- Box, G. E. P., Hunter, W. G., MacGregor, J. F., Erjavec, J., 1973. Some problems associated with the analysis of multiresponse data. *Technometrics* 15, 33–51.
- Box, G. E. P., Meyer, R. D., 1986. Dispersion effects from fractional designs. *Technometrics* 28, 19–27.
- Box, G. E. P., Wilson, K. B., 1951. On the experimental attainment of optimum conditions (with discussion). *Journal of the Royal Statistical Society B* 13, 1–45.
- Broudiscou, L. P., Papon, Y., Broudiscou, A. F., 1999. Effects of minerals on feed degradation and protein synthesis by rumen micro-organisms in a dual effluent fermenter. *Reproduction Nutrition Development* 39, 255–268.
- Cao, W., Zhang, C., Hong, P., Ji, H., 2009. Optimising the free radical scavenging activity of shrimp protein hydrolysate produced with alcalase using response surface methodology. *International Journal of Food Science and Technology* 44, 1602–1608.
- Chaloner, K., Larntz, K., 1989. Optimal Bayesian designs applied to logistic regression experiments. *Journal of Statistical Planning and Inference* 21, 191–208.
- Chaloner, K., Larntz, K., 1992. Bayesian design for accelerated life testing. *Journal of Statistical Planning and Inference* 33, 245–259.
- Chaloner, K., Verdinelli, I., 1995. Bayesian experimental design: A review. *Statistical Science* 10, 273–304.
- Cook, D., Fedorov, V., 1995. Constrained optimization of experimental design. *Statistics* 26, 129–178.
- Cornell, J. A., 1990. *Experiments with Mixtures*. Wiley, New York.

- Crosier, R. B., 1984. Mixture experiments: Geometry and pseudocomponents. *Technometrics* 26, 209–216.
- Del Castillo, E., 1996. Multiresponse process optimization via constrained confidence regions. *Journal of Quality Technology* 28, 61–70.
- Del Castillo, E., Fan, S. K., Semple, J., 1997. Computation of global optima in dual response systems. *Journal of Quality Technology* 29, 347–353.
- Del Castillo, E., Fan, S. K., Semple, J., 1999. Optimization of dual response systems: a comprehensive procedure for degenerate and nondegenerate problems. *European Journal of Operational Research* 112, 174–186.
- Del Castillo, E., Montgomery, D. C., 1993. A nonlinear programming solution to the dual response problem. *Journal of Quality Technology* 25, 199–204.
- Derringer, G., Suich, R., 1980. Simultaneous optimization of several response variables. *Journal of Quality Technology* 12, 214–219.
- Detle, H., Sperlich, S., 1994. A note on Bayesian D-optimal designs for generalization of the simple exponential growth model. *South African Statistical Journal* 28, 103–117.
- Diniz, F. M., Martin, A. M., 1996. Use of response surface methodology to describe the combined effects of pH, temperature and E/S ratio on the hydrolysis of dogfish (*Squalus acanthias*) muscle. *International Journal of Food Science and Technology* 31, 419–426.
- Dobson, A. J., 2001. *An Introduction to Generalized Linear Models*, 2nd edn. Chapman and Hall/CRC, London, U.K.
- Dragalin, V., Fedorov, V., 2008. Adaptive designs for selecting drug combinations based on efficacy-toxicity response. *Journal of Statistical Planning and Inference* 138, 352–373.
- Dragalin, V. V., Fedorov, V., Wu Y., 2008. Two-stage design for dose-finding that accounts for both efficacy and toxicity. *Statistics in Medicine* 27, 5156–5176.
- Draper, N. R., 1982. Center points in second-order response surface designs. *Technometrics* 24, 127–133.
- Draper, N. R., Hunter, W. G., 1966. Design of experiments for parameter estimation in multiresponse situations. *Biometrika* 53, 525–533.
- Draper, N. R., John, J. A., 1998. Response surface designs where levels of some factors are difficult to change. *Australian and New Zealand Journal of Statistics* 40, 487–495.
- Dror, H. A., Steinberg, D. M., 2006. Robust experimental design for multivariate generalized linear models. *Technometrics* 48, 520–529.
- Dror, H. A., Steinberg, D. M., 2008. Sequential experimental designs for generalized linear models. *Journal of the American Statistical Association* 103, 288–298.
- Easterling, R. B., 1985. Discussion of off-line quality control, parameter design, and the Taguchi methods. *Journal of Quality Technology* 17, 198–206.
- Edmondson, R. N., 1991. Agricultural response surface experiments based on four-level factorial designs. *Biometrics* 47, 1435–1448.
- Evans, R. A., Easi, D. A., Book, D. N., Young, J. A., 1982. Quadratic response surface analysis of seed-germination trials. *Weed Science* 30, 411–416.
- Fan, S. K. S., 2000. A generalized global optimization algorithm for dual response systems. *Journal of Quality Technology* 32, 444–456.
- Fedorov, V. V., 1972. *Theory of Optimal Experiments*. Academic Press, New York.
- Fedorov, V., Wu, Y., Zhang R., 2011. Optimal dose finding designs with correlated continuous and discrete responses. *Statistics in Medicine* 31, 217–234.
- Gilmour, S. G., Trinca, L. A., 2005. Fractional polynomial response surface models. *Journal of Agricultural, Biological, and Environmental Statistics* 10, 50–60.
- Giovannitti-Jensen, A., Myers, R. H., 1989. Graphical assessment of the prediction capability of response surface designs. *Technometrics* 31, 159–171.
- Goldfarb, H. B., Borror, C. M., Montgomery, D. C., Anderson-Cook, C. M., 2004. Three-dimensional variance dispersion graphs for mixture-process experiments. *Journal of Quality Technology* 36, 109–124.

- Haines, L. M., 2006. Evaluating the performance of non-standard designs: The San Cristobal design. In: Khuri, A. I. (Ed.), *Response Surface Methodology and Related Topics*. World Scientific, Singapore, pp. 251–281.
- Harrington, E. C., Jr., 1965. The desirability functions. *Industrial Quality Control* 12, 494–498.
- Heise, M. A., Myers, R. H., 1996. Optimal designs for bivariate logistic regression. *Biometrics* 52, 613–624.
- Hill, W. J., Hunter, W. G., 1966. A review of response surface methodology: A literature review. *Technometrics* 8, 571–590.
- Jiang, G., Jiang, Y., Yang, B., Yu, C., Tsao, R., Zhang, H., Chen, F., 2009. Structural characteristics and antioxidant activities of oligosaccharides from longan fruit pericarp. *Journal of Agricultural and Food Chemistry* 57, 9293–9298.
- Jung, B. C., Khuri, A. I., Lee, J., 2008. Comparison of designs for the three-fold nested random model. *Journal of Applied Statistics* 35, 701–715.
- Kackar, R. N., 1985. Off-line quality control, parameter design, and the Taguchi method. *Journal of Quality Technology* 17, 176–188.
- Keisling, T. C., Oliver, L. R., Crowley, R. H., Baldwin, F. L., 1984. Potential use of response surface analyses for weed management in soybeans (glycine max). *Weed Science* 32, 552–557.
- Khan, M. K., Yazdi, A. A., 1988. On D-optimal designs for binary data. *Journal of Statistical Planning and Inference* 18, 83–91.
- Khuri, A. I., 1988. A measure of rotatability for response surface designs. *Technometrics* 30, 95–104.
- Khuri, A. I., 1992. Response surface models with random block effects. *Technometrics* 34, 26–37.
- Khuri, A. I., 1996a. Multiresponse surface methodology. In: Ghosh, S., Rao, C. R. (Eds.), *Handbook of Statistics*, Vol. 13. Elsevier Science B.V., Amsterdam, the Netherlands, pp. 377–406.
- Khuri, A. I., 1996b. Response surface models with mixed effects. *Journal of Quality Technology* 28, 177–186.
- Khuri, A. I., 1997. Quantile dispersion graphs for analysis of variance estimates of variance components. *Journal of Applied Statistics* 24, 711–722.
- Khuri, A. I., 2005. Slack-variable models versus Scheffé's mixture models. *Journal of Applied Statistics* 32, 887–908.
- Khuri, A. I., 2006. Mixed response surface models with heterogeneous within-block error variances. *Technometrics* 48, 206–218.
- Khuri, A. I., Conlon, M., 1981. Simultaneous optimization of multiple responses represented by polynomial regression functions. *Technometrics* 23, 363–375.
- Khuri, A. I., Cornell, J. A., 1996. *Response Surfaces*, 2nd edn. Dekker, New York.
- Khuri, A. I., Harrison, J. M., Cornell, J. A., 1999. Using quantile plots of the prediction variance for comparing designs for a constrained mixture region: An application involving a fertilizer experiment. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 48, 521–532.
- Khuri, A. I., Kim, H. J., Um, Y., 1996. Quantile plots of the prediction variance for response surface designs. *Computational Statistics and Data Analysis* 22, 395–407.
- Khuri, A. I., Lee, J., 1998. A graphical approach for evaluating and comparing designs for nonlinear models. *Computational Statistics and Data Analysis* 27, 433–443.
- Khuri, A. I., Mukherjee, B., Sinha, B. K., Ghosh, M., 2006. Design issues for generalized linear models: A review. *Statistical Science* 21, 376–399.
- Khuri, A. I., Mukhopadhyay, S., 2006. GLM designs: The dependence on unknown parameters dilemma. In: Khuri, A. I. (Ed.), *Response Surface Methodology and Related Topics*. World Scientific, Singapore, pp. 203–223.
- Khuri, A. I., Mukhopadhyay, S., 2010. Response surface methodology. In: Wegman, E. J., Said, Y. H., and Scott, D. W. (Eds.), *Wiley Interdisciplinary Reviews—Computational Statistics*, Vol. 2, No. 2. Wiley, Hoboken, NJ, pp. 128–149.
- Kim, K. J., Lin, D. K. J., 2000. Simultaneous optimization of mechanical properties of steel by maximizing exponential desirability functions. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 49, 311–325.

- Ko, Y. H., Kim, K. J., Jun, C. H., 2005. A new loss function-based method for multiresponse optimization. *Journal of Quality Technology* 37, 50–59.
- Kowalsky, S. M., Vining, G. G., Montgomery, D. C., Borrer, C. M., 2006. Modifying central composite design to model the process mean and variance when there are hard-to-change factors. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 55, 615–630.
- Lee, J., Khuri, A. I., 1999. Graphical technique for comparing designs for random models. *Journal of Applied Statistics* 26, 933–947.
- Lee, J., Khuri, A. I., 2000. Quantile dispersion graphs for the comparison of designs for a random two-way model. *Journal of Statistical Planning and Inference* 91, 123–137.
- Liang, L., Anderson-Cook, C. M., Robinson, T. J., Myers, R. H., 2006. Three-dimensional variance dispersion graphs for split-plot designs. *Journal of Computational and Graphical Statistics* 15, 757–778.
- Lind, E. E., Goldin, J., Hickman, J. B., 1960. Fitting yield and cost response surfaces. *Chemical Engineering Progress* 56, 62–68.
- Lindsey, J. K., 1997. *Applying Generalized Linear Models*. Springer, New York.
- Liski, E. P., Mandal, N. K., Shah, K. R., Sinha, B. K., 2002. Topics in Optimal Design. Lecture Notes in Statistics 163. Springer, New York.
- Mathew, T., Sinha, B. K., 2001. Optimal designs for binary data under logistic regression. *Journal of Statistical Planning and Inference* 93, 295–307.
- McCullagh, P., Nelder, J. A., 1989. *Generalized Linear Models*, 2nd edn. Chapman and Hall, London, U.K.
- McCulloch, C. E., Searle, S. R., 2001. *Generalized, Linear, and Mixed Models*. Wiley, New York.
- McLean, R. A., Anderson, V. L., 1966. Extreme vertices design of mixture experiments. *Technometrics* 8, 447–454.
- McLean, R. A., Anderson, V. L., 1984. *Applied Factorial and Fractional Designs*. Dekker, New York.
- Mead, R., 1988. *The Design of Experiments*. Cambridge University Press, Cambridge, U.K.
- Mead, R., Pike, D. J., 1975. A review of response surface methodology from a biometric viewpoint. *Biometrics* 31, 803–851.
- Mee, R. W., Bates, R. L., 1998. Split-lot designs: Experiments for multistage batch processes. *Technometrics* 40, 127–140.
- Montgomery, D. C., 2013. *Design and Analysis of Experiments*, 8th edn. Wiley, New York.
- Mukhopadhyay, S., Haines, L. M., 1995. Bayesian D-optimal designs for the exponential growth model. *Journal of Statistical Planning and Inference* 44, 385–397.
- Mukhopadhyay, S., Khuri, A. I., 2008a. Comparison of designs for multivariate generalized linear models. *Journal of Statistical Planning and Inference* 138, 169–183.
- Mukhopadhyay, S., Khuri, A. I., 2008b. A new graphical approach for comparing response surface designs on the basis of the mean squared error of prediction criterion. *Statistics and Applications* 6, 293–324.
- Mukhopadhyay, S., Khuri, A. I., 2008c. Optimization in a multivariate generalized linear model situation. *Computational Statistics and Data Analysis* 52, 4625–4634.
- Mukhopadhyay, S., Khuri, A. I., 2012. Comparison of designs for generalized linear models under model misspecification. *Statistical Methodology* 9, 285–304.
- Myers, R. H., Carter, W. H., 1973. Response surface techniques for dual response systems. *Technometrics* 15, 301–317.
- Myers, R. H., Khuri, A. I., Carter, W. H., Jr., 1989. Response surface methodology: 1966–1988. *Technometrics* 31, 137–157.
- Myers, R. H., Khuri, A. I., Vining, G. G., 1992a. Response surface alternatives to the Taguchi robust parameter design approach. *The American Statistician* 46, 131–139.
- Myers, R. H., Montgomery, D. C., Anderson-Cook, C. M., 2009. *Response Surface Methodology: Process and Product Optimization Using Designed Experiments*, 3rd edn. Wiley, New York.
- Myers, R. H., Montgomery, D. C., Vining, G. G., 2002. *Generalized Linear Models with Applications in Engineering and the Sciences*. Wiley, New York.

- Myers, R. H., Montgomery, D. C., Vining, G. G., Borror, C. M., Kowalski, S. M., 2004. Response surface methodology: A retrospective and literature survey. *Journal of Quality Technology* 36, 53–77.
- Myers, R. H., Vining, G. G., Giovannitti-Jensen, A., Myers, S. L., 1992b. Variance dispersion properties of second order response surface designs. *Journal of Quality Technology* 24, 1–11.
- Nair, V. N., 1992. Taguchi's parameter design: A panel discussion. *Technometrics* 34, 127–161.
- Nair, V. N., Pregibon, D., 1988. Analyzing dispersion effects from replicated factorial experiments. *Technometrics* 30, 247–257.
- Nelder, J. A., 1966. Inverse polynomials, a useful group of multi-factor response functions. *Biometrics* 22, 128–141.
- Nelder, J. A., Wedderburn, R. W. M., 1972. Generalized linear models. *Journal of the Royal Statistical Society A* 135, 370–384.
- Ozol-Godfrey, A., Anderson-Cook, C. M., Montgomery, D. C., 2005. Fraction of design space plots for examining model robustness. *Journal of Quality Technology* 37, 223–235.
- Ozol-Godfrey, A., Anderson-Cook, C. M., Robinson, T. J., 2008. Fraction of design space plots for generalized linear models. *Journal of Statistical Planning and Inference* 138, 203–219.
- Pignatiello, J. J., Jr., 1993. Strategies for robust multiresponse quality engineering. *IIE Transactions* 25, 5–15.
- Pignatiello, J. J., Jr., Ramberg, J. S., 1985. Off-line quality control, parameter design, and the Taguchi methods—A discussion. *Journal of Quality Technology* 17, 198–206.
- Plackett, R. L., Burman, J. P., 1946. The design of optimum multifactor experiments. *Biometrika* 33, 305–325.
- Robinson, K. S., Khuri, A. I., 2003. Quantile dispersion graphs for evaluating and comparing designs for logistic regression models. *Computational Statistics and Data Analysis* 43, 47–62.
- Rojas, B. A., 1962. The San Cristobal design for fertilizer experiments. In: *Proceedings of the 11th Congress of the International Society for Sugar Cane Technology*. Reduit, Mauritius, pp. 197–203.
- Russell, K. G., Woods, D. C., Lewis, S. M., Eccleston, J. A., 2009. D-optimal designs for Poisson regression models. *Statistica Sinica* 19, 721–730.
- Saha, S., Khuri, A. I., 2009. Comparison of designs for response surface models with random block effects. *Quality Technology and Quantitative Management* 6, 219–234.
- Salawu, I. S., Adeyemi, R. A., Aremu, T. A., 2007. Modified inverse polynomial and ordinary polynomial as a response surface model: A case study of nitrogen, phosphate and potassium levels on the yield of maize. *International Journal of Pure and Applied Sciences* 1, 18–24.
- Scheffé, H., 1958. Experiments with mixtures. *Journal of the Royal Statistical Society B* 20, 344–360.
- Scheffé, H., 1963. The simplex-centroid design for experiments with mixtures. *Journal of the Royal Statistical Society B* 25, 235–263.
- Schmidt, R. H., Illingworth, B. L., Deng, J. C., Cornell, J. A., 1979. Multiple regression and response surface analysis of the effects of calcium chloride and cysteine on heat-induced whey protein gelation. *Journal of Agricultural and Food Chemistry* 27, 529–532.
- Shoemaker, A. C., Tsui, K. L., Wu, C. F. J., 1991. Economical experimentation methods for robust design. *Technometrics* 33, 415–427.
- Sinha, S., Wiens, D. P., 2002. Robust sequential designs for nonlinear regression. *Canadian Journal of Statistics* 30, 601–618.
- Sinha, S. K., Xu, X., 2011. Sequential D-optimal designs for generalized linear mixed models. *Journal of Statistical Planning and Inference* 141, 1394–1402.
- Sitter, R. R., Forbes, B. E., 1997. Optimal two-stage designs for binary response experiments. *Statistica Sinica* 7, 941–956.
- Sitter, R. R., Wu, C. F. J., 1993. Optimal designs for binary response experiments: Fieller, D, and A criteria. *Scandinavian Journal of Statistics* 20, 329–341.
- Sitter, R. R., Wu, C. F. J., 1999. Two-stage design of quantal response studies. *Biometrics* 55, 396–402.
- Taguchi, G., 1986. *Introduction to Quality Engineering*. UNIPUB/Kraus International, White Plains, NY.

- Trinca, L. A., Gilmour, S. G., 1998. Variance dispersion graphs for comparing blocked response surface designs. *Journal of Quality Technology* 30, 314–327.
- Vining, G. G., 1993. A computer program for generating variance dispersion graphs. *Journal of Quality Technology* 25, 45–58.
- Vining, G., 1998. A compromise approach to multiresponse optimization. *Journal of Quality Technology* 30, 309–313.
- Vining, G. G., Kowalsky, S. M., Montgomery, D. C., 2005. Response surface designs within a split-plot structure. *Journal of Quality Technology* 37, 115–129.
- Vining, G. G., Myers, R. H., 1990. Combining Taguchi and response surface philosophies: A dual response approach. *Journal of Quality Technology* 22, 38–45.
- Vining, G. G., Myers, R. H., 1991. A graphical approach for evaluating response surface designs in terms of the mean squared error of prediction. *Technometrics* 33, 315–326.
- Wang, Y., Myers, R. H., Smith, E. P., Ye, K., 2006. D-optimal designs for Poisson regression models. *Journal of Statistical Planning and Inference* 136, 2831–2845.
- Welch, W. J., Yu, T. K., Kang, S. M., Sacks, J., 1990. Computer experiments for quality control by parameter design. *Journal of Quality Technology* 22, 15–22.
- Wijesinha, M., Khuri, A. I., 1987a. The sequential generation of multiresponse D-optimal designs when the variance-covariance matrix is not known. *Communications in Statistics–Simulation and Computation* 16, 239–259.
- Wijesinha, M. C., Khuri, A. I., 1987b. Construction of optimal designs to increase the power of the multiresponse lack of fit test. *Journal of Statistical Planning and Inference* 16, 179–192.
- Woods, D. C., Lewis, S. M., Eccleston, J. A., Russell, K. G., 2006. Designs for generalized linear models with several variables and model uncertainty. *Technometrics* 48, 284–292.
- Woods, D. C., van De Ven, P., 2011. Blocked designs for experiments with correlated non-normal response. *Technometrics* 53, 173–182.
- Wu, C. F. J., 1985. Efficient sequential designs with binary data. *Journal of the American Statistical Association* 80, 974–984.
- Wu, C. F. J., Hamada, M., 2009. *Experiments: Planning, Analysis, and Parameter Design Optimization*. Wiley-Interscience, New York.
- Wu, F. C., 2005. Optimization of correlated multiple quality characteristics using desirability function. *Quality Engineering* 17, 119–126.
- Yang, M., Stufken, J., 2009. Support points of locally optimal designs for nonlinear models with two parameters. *The Annals of Statistics* 37, 518–541.
- Yang, M., Zhang, B., Huang, S., 2011. Optimal designs for binary response experiments with multiple variables. *Statistica Sinica*, 21, 1415–1430.
- Zacks, S., 1977. Problems and approaches in design of experiments for estimation and testing in non-linear models. In: Krishnaiah, P. R. (Ed.), *Multivariate Analysis*, Vol. 4. North-Holland, Amsterdam, the Netherlands, pp. 209–223.
- Zahrn, A., Anderson-Cook, C. M., Myers, R. H., 2003. Fraction of design space to assess prediction capability of response surface designs. *Journal of Quality Technology* 35, 377–386.
- Zellner, A., 1962. An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias. *Journal of the American Statistical Association* 57, 348–368.
- Zhang, D., Bai, S., Dong, X., Sun, Y., 2007. Optimization of lipase-catalyzed regioselective acylation of pyridoxine (vitamin B6). *Journal of Agricultural and Food Chemistry* 55, 4526–4531.
- Zocchi, S. S., Atkinson, A. C., 1999. Optimum experimental designs for multinomial logistics models. *Biometrics* 55, 437–444.

6

Design for Linear Regression Models with Correlated Errors

Holger Dette, Andrey Pepelyshev, and Anatoly Zhigljavsky

CONTENTS

6.1	Estimation and Design for Correlated Errors: The Main Approaches.....	238
6.1.1	Introduction.....	238
6.1.2	Different Versions of the Optimal Design Problem.....	240
6.1.2.1	General Regression Problem.....	240
6.1.2.2	Design Space \mathcal{X}	240
6.1.2.3	Vector of Regression Functions $f(x)$	240
6.1.2.4	Covariance Kernels K	241
6.1.2.5	Designs.....	241
6.1.2.6	Interpretation of Approximate Designs.....	242
6.1.2.7	Methods of Estimation and Covariance Matrices.....	243
6.1.2.8	Optimality Criteria.....	243
6.1.2.9	Information Matrix.....	244
6.1.2.10	Possible Generalizations of the Model (6.2).....	244
6.1.3	Information Contained in Design Points.....	245
6.1.3.1	Points Providing Zero Information.....	245
6.1.3.2	Smit's Paradox.....	246
6.1.3.3	Estimates with Zero Variance.....	246
6.2	Designs for One-Parameter Models.....	246
6.2.1	Designs for the BLUE.....	247
6.2.1.1	Designs for a Continuous Observation.....	247
6.2.1.2	Results of Sacks and Ylvisaker.....	247
6.2.2	Optimal Design for OLS Estimation.....	249
6.2.2.1	Location-Scale Model.....	249
6.2.2.2	Hajek Result.....	251
6.2.3	Optimal Design for SLS.....	252
6.3	Optimal Designs for the BLUE.....	253
6.3.1	BLUE for Continuous Observations.....	254
6.3.2	Results of Sacks and Ylvisaker.....	255
6.3.3	Exchange-Type Algorithms.....	256
6.3.4	Constructing Optimal Designs by Expansion of the Covariance Kernel.....	257
6.3.5	Method of Virtual Noise.....	258
6.3.6	Design for Prediction in the Quadratic Model.....	259
6.4	Optimal Designs for OLS.....	260
6.4.1	OLS for Approximate Designs.....	260
6.4.2	Results of Bickel and Herzberg.....	261

6.4.3	Results for Long-Range Dependence Error Structure.....	263
6.4.3.1	Linear Regression.....	267
6.4.4	Optimal Designs for OLS.....	268
	Acknowledgments.....	271
	Appendix 6.A Proofs.....	271
	Appendix 6.B Common Correlation Functions.....	273
	References.....	273

The present chapter provides a survey of results on experimental design for linear regression models with correlated responses.

6.1 Estimation and Design for Correlated Errors: The Main Approaches

6.1.1 Introduction

The common linear regression model is given by

$$y(\mathbf{x}) = \theta_1 f_1(\mathbf{x}) + \cdots + \theta_m f_m(\mathbf{x}) + \varepsilon(\mathbf{x}), \quad (6.1)$$

where $f_1(\mathbf{x}), \dots, f_m(\mathbf{x})$ are given linearly independent functions, $\varepsilon(\mathbf{x})$ denotes a random error process or field with $E[\varepsilon(\mathbf{x})] = 0$, \mathbf{x} is the explanatory variable, which varies in the design space $\mathcal{X} \subset \mathbb{R}^d$, and $\theta_1, \dots, \theta_m$ are unknown parameters.

We assume that N observations can be taken at experimental conditions $\mathbf{x}_1, \dots, \mathbf{x}_N$ to estimate the parameters in the linear regression model (6.1). If an estimate of $\boldsymbol{\theta} = (\theta_1, \dots, \theta_m)^T$ has been chosen, the quality of the statistical analysis can be further improved by choosing an appropriate design for the experiment (in this chapter, we use the sign T to denote the transposition). In particular, an optimal design minimizes a functional of the variance-covariance matrix of the estimate, where the functional should reflect certain aspects of the goal of the experiment. In contrast to the case of uncorrelated errors, where numerous results and a rather complete theory are available (see e.g., Chapters 2 and 3 and the monograph of Pukelsheim (2006), the construction of optimal designs for dependent observations is intrinsically more difficult because the information matrix cannot be decomposed into the sum of information matrices for the individual design points. However, this problem is of particular practical interest as, in many applications, the observations are correlated. Typical examples include models where the explanatory variable x represents time and where all observations correspond to one subject.

Because explicit solutions of optimal design problems for correlated observations are rarely available, several authors have proposed to determine optimal designs based on asymptotic arguments, see for example, Sacks and Ylvisaker (1966, 1968), Bickel and Herzberg (1979), N  ther (1985a), and Zhigljavsky et al. (2010). Roughly speaking, there exist three approaches to embed the optimal design problem for regression models with correlated observations in an asymptotic optimal design problem.

The first one is due to Sacks and Ylvisaker (1966, 1968), who assumed that the covariance structure of the error process $\varepsilon(\mathbf{x})$ is fixed and that the number of design points tends to

infinity. As a result of this assumption, the design points become very close to each other and the corresponding asymptotic optimal designs depend only on the behavior of the correlation function in a neighborhood of the point 0.

Alternatively, Bickel and Herzberg (1979) and Bickel et al. (1981) considered a different model, where the ordinary least squares (OLS) estimate is used and the correlation function depends on the number of observations N . The covariance matrix of the estimate is of order $O(1)$ in the model considered by Sacks and Ylvisaker (1966) and of order $O(1/N)$ in the model discussed by Bickel and Herzberg (1979). Therefore, the approach of Bickel and Herzberg (1979) makes the optimal designs derived for the dependent and independent cases more comparable. These authors assumed that the observation errors $\varepsilon(t)$ in the model (6.1) is a stationary process with short-range dependence, where the correlation function of this process $\rho(t) = E[\varepsilon(0)\varepsilon(t)]$ satisfies $\rho(t) = o(1/t)$ if $t \rightarrow \infty$. Dette et al. (2009) extended results of Bickel and Herzberg (1979) to the case where the error process has long-range dependence.

Recently, Zhigljavsky et al. (2010) modified the Bickel–Herzberg approach and allowed the variance (in addition to the correlation function) to vary as the number of observations changes. As a result, the asymptotic covariance matrices may contain a kernel with a singularity at the diagonal.

Significant research has been devoted to constructing exact optimal designs for the best linear unbiased estimator (BLUE); see Section 6.3. The simplest algorithm was proposed in Brimkulov et al. (1980). Other algorithms are based on the method of virtual noise developed by Pázman and Müller (2001) and the method of the expansion of the covariance kernel suggested by Fedorov and Müller (2007). Note that the BLUE can only be used if the correlation structure of errors is known, and its misspecification can lead to a severe loss of efficiency. On the other hand, the OLS estimate does not employ the correlation structure. Obviously, the OLS estimate can be less efficient than the BLUE, but in many cases, the loss of efficiency is either small or negligible.

The structure of this chapter is as follows: In Section 6.1.2, we introduce different variations of the optimal design problem for the linear regression model (6.1). In particular, we consider various assumptions about the design space \mathcal{X} , the vector-function $f(x) = (f_1(x), \dots, f_m(x))^T$, the covariance kernel $K(x, z) = E[\varepsilon(x)\varepsilon(z)]$, and also different sets of designs, three different estimates of the unknown parameter θ , and corresponding covariance matrices. In Section 6.1.3, we briefly discuss the concept of information contained in design points and discuss some well-known paradoxes.

Section 6.2 is devoted to the problem of designing experiments for one-parameter models. This problem is often easier than similar problems for the multi-parameter case, and in some cases, it can be solved explicitly, see, for example, Theorem 6.5. The easiest one-parameter model is the so-called location-scale model where the variance of the OLS estimate leads to a convex design optimality functional and makes many tools of the convex optimization theory (see Section 6.2) applicable. Moreover, if the correlation function ρ of the stationary error process in a location-scale model is convex for $x > 0$, then the OLS estimate coincides with the BLUE; see Theorem 6.4.

Section 6.3 is devoted to the problem of optimal design for the BLUE. We review classical Sacks–Ylvisaker results (Sacks and Ylvisaker 1966, 1968) and also more recent results of Fedorov–Müller and Pázman–Müller (Pázman and Müller 2001; Fedorov and Müller 2007). Moreover, we also consider the well-known exchange algorithm for the construction of N -point optimal designs. Additionally, we review one of Harman’s results on optimal design for prediction in the case of the quadratic model and Wiener process (Harman and Stulajter 2010).

Section 6.4 reviews some results concerning characterization and construction of optimal designs for the OLS estimate. We explain the classical Bickel–Herzberg approach and its extension for the case of a long-range dependent error process. We also review some of the recent results of the authors concerning the explicit construction of optimal designs for models with $m \geq 2$ parameters and some particular covariance kernels.

Appendices 6.A and 6.B contain selected proofs and a table of common correlation functions that appear in discussions in the literature on optimal design for correlated observations.

The authors are aware of the following three substantial surveys devoted to the theory of optimal designs for correlated observations; see also a short survey in Müller's book (Müller 2007): the first one is an excellent book by Näther (1985a), and the other two are the surveys by Cambanis (1985) and Fedorov (1996). As much research has been done in recent years, we feel that there is a need for a new survey on the subject.

6.1.2 Different Versions of the Optimal Design Problem

6.1.2.1 General Regression Problem

The general multiparameter linear regression model (6.1) can be written as

$$y(x) = \boldsymbol{\theta}^T \mathbf{f}(x) + \varepsilon(x), \quad (6.2)$$

where the explanatory variable x belongs to a design space \mathcal{X} , $\mathbf{f}(x) = (f_1(x), \dots, f_m(x))^T$ is a vector of linearly independent regression functions, and $\varepsilon(x)$ denotes random process or field with $E[\varepsilon(x)] = 0$, and $E[\varepsilon(x)\varepsilon(z)] = K(x, z)$. The vector of parameters $\boldsymbol{\theta} = (\theta_1, \dots, \theta_m)^T$ is unknown and has to be estimated on the basis of observations taken from one realization of a stochastic process (or field) $y(x)$. The function $K(\cdot, \cdot)$ will be called a covariance kernel.

Throughout this article, we consider different variations and specifications of the model (6.2) in relation to the problem of optimal design.

6.1.2.2 Design Space \mathcal{X}

We make a distinction between the following forms of the design space:

- (a) \mathcal{X} is a finite set.
- (b) \mathcal{X} is an interval $[-1, 1]$.
- (c) $\mathcal{X} \subseteq \mathbb{R}$.
- (d) $\mathcal{X} \subseteq \mathbb{R}^d$ with $d \geq 1$.

In the case $d = 2$, the model (6.2) is called spatial model. If $d \geq 2$, then $\varepsilon(x)$ is called a random field. If nothing is stated about \mathcal{X} , then the most general case, which is (d), is considered.

6.1.2.3 Vector of Regression Functions $\mathbf{f}(x)$

We will distinguish the following forms of the vector of functions $\mathbf{f}(x)$:

- (a) The general case with $m \geq 1$.
- (b) The case of $m = 1$ where the model (6.2) is called a one-parameter model; see Section 6.2.

- (c) The case of $m = 1$ and $f(x) = 1$ where the model (6.2) is called a location-scale model; see Section 6.2.2.

To avoid technical difficulties, we always assume that all the components of the vector $f(x)$ are continuous functions.

6.1.2.4 Covariance Kernels K

We will distinguish the following cases for the covariance kernel $K(\cdot, \cdot)$:

1. The general positive definite function $K(\cdot, \cdot)$.
2. Kernels of stationary processes or fields, which have the form $K(x, z) = \sigma^2 \rho(x - z)$, where $\rho(\cdot)$ is a correlation function. Examples of commonly used correlation functions are given in Table 6.1 and discussed in Appendix 6.B.
3. Kernels with nugget term, that is, $K(x, z) = \gamma K_0(x, z) + (1 - \gamma) \delta_{x,z}$, where $0 < \gamma < 1$, the kernel K_0 is continuous on the diagonal, and δ denotes Kronecker's symbol.
4. Kernels with singularity at the diagonal. These kernels possess the property that $K(x, z) \rightarrow \infty$ as $x \rightarrow z$ and $K(x, x)$ is not defined for some $x \in \mathcal{X}$ (see Section 6.4.4 for examples).

By the definition, a kernel $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is called positive definite if $K(x, z) = K(z, x)$ for all $x, z \in \mathcal{X}$, and for any set of distinct points x_1, \dots, x_N in \mathcal{X} , the matrix

$$\Sigma = (K(x_i, x_j))_{i,j=1}^N,$$

is nonnegative definite. We shall call the kernel K strictly positive definite if the inequality

$$\iint K(x, z) \zeta(dx) \zeta(dz) > 0,$$

holds for any signed measure ζ on \mathcal{X} such that $0 < |\zeta|(\mathcal{X}) < \infty$.

6.1.2.5 Designs

The following three types of designs will be considered in this article:

1. An exact N -point design $\xi_N = \{x_1, \dots, x_N\}$, where $x_i \in \mathcal{X}$.
2. An approximate design $\xi(dx)$ corresponding to a probability measure on the design space \mathcal{X} (see interpretation in the following text).
3. A signed design defined as a signed measure $\xi(dx)$ on the design space \mathcal{X} with

$$|\xi|(\mathcal{X}) = \xi^+(\mathcal{X}) + \xi^-(\mathcal{X}) < \infty.$$

We denote the spaces of exact N -point designs on \mathcal{X} , approximate designs on \mathcal{X} , and signed design measures on \mathcal{X} by Ξ_N , Ξ , and $\Xi^{(S)}$, respectively.

TABLE 6.1

Commonly Used Correlation Functions, $d = 1, \lambda > 0, R > 0$

Name	$\rho(x), x \geq 0$
Exponential	$e^{-\lambda x}$
Gaussian	$e^{-\lambda x^2}$
Rational quadratic	$\frac{1}{(1 + \lambda x^2)^\nu}, \nu > 0$
Spherical, triangular	$(1 - \frac{x}{R}) \mathbf{1}_{[0,R]}(x)$
Spherical, circular	$\left(1 - \frac{2}{\pi} \left(\frac{x}{R} \sqrt{1 - \frac{x^2}{R^2}} + \arcsin\left(\frac{x}{R}\right)\right)\right) \mathbf{1}_{[0,R]}(x)$
Spherical	$\left(1 - \frac{3x}{R} + \frac{x^3}{2R^3}\right) \mathbf{1}_{[0,R]}(x)$
Penta-spherical	$\left(1 - \frac{15x}{8R} + \frac{5x^3}{4R^3} - \frac{3x^5}{R^5}\right) \mathbf{1}_{[0,R]}(x)$
Cubic	$\left(1 - \frac{7x^2}{R^2} + \frac{35x^3}{4R^3} - \frac{7x^5}{2R^5} + \frac{3x^7}{4R^7}\right) \mathbf{1}_{[0,R]}(x)$
Stable	$\exp(-\lambda x^\nu), 0 < \nu \leq 2$
Oscillating, damped cosine	$e^{-\lambda x} \cos(\omega x), \omega > 0$
Oscillating, hole effect	$\frac{1}{\omega x} \sin(\omega x), \omega > 0$
Oscillating, Bessel	$\Gamma(\nu + 1) 2^\nu (\lambda x)^{-\nu} J_\nu(\lambda x), \nu \geq (d + 1)/2$
Poisson	$\frac{1 - \beta^2}{1 - 2\beta \cos(2\pi x) + \beta^2}, 0 < \beta < 1$
Cauchy family	$\frac{1}{(1 + x ^\beta)^\alpha/\beta}, \beta > 0, 0 < \alpha \leq 1$
Mittag-Leffler family	$E_{\nu,\beta}(- x ^\alpha), 0 < \alpha \leq 1, 0 < \nu \leq 1, \beta \geq \nu$ $E_{\nu,\beta}(-x) = \Gamma(\beta) \sum_{k=0}^{\infty} \frac{(-x)^k}{\Gamma(\nu k + \beta)}$ $E_{1,1}(-x) = e^{-x}, E_{1,2}(-x) = (1 - e^{-x})/x$ $E_{1,3}(-x) = 2(e^{-x} - 1 + x)/x^2$ $E_{1/2,1}(-x) = e^{x^2} \left(1 - \frac{2}{\sqrt{\pi}} \int_0^x e^{-u^2} du\right)$
Scaling	$(1 - x^2)e^{-x^2/2}$
Singular logarithmic	$-\ln x^2$
Singular rational	$\frac{1}{ x ^\alpha}, 0 < \alpha < 1$

6.1.2.6 Interpretation of Approximate Designs

Consider an approximate design ξ , which is a probability measure on the design space \mathcal{X} . In asymptotic investigations for the case $\mathcal{X} \subseteq \mathbb{R}$, it is usually assumed that a sequence of exact designs $\xi_N = \{x_{1,N}, \dots, x_{N,N}\}$ is generated for increasing N using a continuous nondecreasing function $a : [0, 1] \rightarrow \mathcal{X}$ by

$$x_{i,N} = a\left(\frac{i-1}{N-1}\right), \quad i = 1, \dots, N,$$

where the function a is the inverse of a distribution function corresponding to ξ . For the multidimensional space, we can suppose that exact designs are generated as centers of partitions of the design space assuming that the diameters of partitions tend to zero but the volumes of the cells $C_{i,N}$ in these partitions are proportional to $\xi(C_{i,N})$.

6.1.2.7 Methods of Estimation and Covariance Matrices

Assume that an exact design $\xi_N = \{x_1, \dots, x_N\}$ with a corresponding vector of observations $\mathbf{y} = (y_1, \dots, y_N)^T$ is given. We consider the following three estimates of the unknown parameters θ : the BLUE, OLS, and signed least squares (SLS). These estimates are, respectively, defined by,

$$\text{BLUE} \quad \hat{\theta} = (\mathbf{X}^T \Sigma^{-1} \mathbf{X})^{-1} \mathbf{X}^T \Sigma^{-1} \mathbf{y}, \tag{6.3}$$

$$\text{OLS} \quad \tilde{\theta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}, \tag{6.4}$$

$$\text{SLS} \quad \tilde{\theta}_S = (\mathbf{X}^T \mathbf{S} \mathbf{X})^{-1} \mathbf{X}^T \mathbf{S} \mathbf{y}, \tag{6.5}$$

where $\mathbf{X} = (f_i(x_j))_{j=1, \dots, N}^{i=1, \dots, m}$ is an $N \times m$ matrix, $\Sigma = (K(x_i, x_j))_{i,j=1, \dots, N}$ is an $N \times N$ matrix, and \mathbf{S} is an $N \times N$ diagonal matrix with $+1$ and -1 on the diagonal. The covariance matrices of the estimates (6.3), (6.4), and (6.5) are given by

$$\begin{aligned} \text{Var}(\hat{\theta}) &= (\mathbf{X}^T \Sigma^{-1} \mathbf{X})^{-1}, \\ \text{Var}(\tilde{\theta}) &= (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \Sigma \mathbf{X} (\mathbf{X}^T \mathbf{X})^{-1}, \\ \text{Var}(\tilde{\theta}_S) &= (\mathbf{X}^T \mathbf{S} \mathbf{X})^{-1} \mathbf{X}^T \mathbf{S} \Sigma \mathbf{S} \mathbf{X} (\mathbf{X}^T \mathbf{S} \mathbf{X})^{-1}, \end{aligned} \tag{6.6}$$

respectively. Note that for the BLUE, there exists a nontrivial optimal design problem in the space Ξ_N , but the corresponding problem in Ξ is trivial: we simply observe the whole process $y(x)$; see Section 6.3. On the other hand, for the OLS estimate, the optimal design problems in the spaces Ξ_N and Ξ are meaningful; see Section 6.4. For the SLS estimate, we consider the optimal design problem in the space $\Xi^{(S)}$; see Section 6.2.3.

6.1.2.8 Optimality Criteria

The following criteria of design optimality (to be minimized) will be considered:

- (a) Minimization of the variance of an estimate in a one-parameter model.
- (b) Universal optimality of the covariance matrix (see Section 6.4.4).
- (c) Minimization of a functional of the covariance matrix, for example, the determinant or the trace (see Sections 6.3 and 6.4).
- (d) Minimization of the mean squared error (MSE) of the best linear predictor at some point (see Section 6.3.6).

6.1.2.9 Information Matrix

To express information obtained from a design ξ_N , we use the Fisher information matrix based on the assumption of normality of the observations

$$\mathbf{IM}(\xi_N) = -E \left[\frac{\partial^2 \ln p(\mathbf{y}|\xi_N)}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}^T} \right],$$

where $p(\mathbf{y}|\xi_N)$ is the normal density with mean $(\boldsymbol{\theta}^T \mathbf{f}(x_1), \dots, \boldsymbol{\theta}^T \mathbf{f}(x_N))$ and covariance matrix $\Sigma = (K(x_i, x_j))_{i,j=1}^N$. Standard calculus gives that the matrix $\mathbf{IM}(\xi_N)$ equals

$$\mathbf{IM}(\xi_N) = \mathbf{X}^T \Sigma^{-1} \mathbf{X}.$$

That is, the Fisher information matrix $\mathbf{IM}(\xi_N)$ is the inverse of the covariance matrix of the BLUE of $\boldsymbol{\theta}$.

A very important observation concerning the information matrix $\mathbf{IM}(\xi_N)$ is the fact, unlike the independent error situation, that in general, $\mathbf{IM}(\xi_N)$ cannot be decomposed into the sum of the information measures for the individual design points. This makes the problem of designing experiments for correlated observations much more difficult than in the case of uncorrelated errors.

6.1.2.10 Possible Generalizations of the Model (6.2)

Possible generalizations of the model (6.2) that are not considered in this survey include the following:

- (a) Regression models having general expected response $E[y(x)] = \eta(x, \boldsymbol{\theta})$, which is non-linear in its parameters; see Atkinson (2008), Dette et al. (2010), and Fedorov et al. (2012).
- (b) Models with covariance kernel K depending on unknown parameters; see Müller and Stehlík (2004), Pázman (2010), and Zimmerman (2006).
- (c) Random-effect and mixed-effect models including population models; see Atkinson (2008), Fedorov (1996), Schmelter (2007), Dette and Holland-Letz (2009), Dette et al. (2010), and Holland-Letz et al. (2011, 2012).
- (d) Treatment models where only block designs are of interest; see Cutler (1993a,b), Kiefer and Wynn (1981, 1984), and Kunert et al. (2010).
- (e) Models where observations are split into groups corresponding to independent realizations of a stochastic process, resulting the block-diagonal structure of the matrix Σ ; see Dette et al. (2010), Holland-Letz et al. (2011), and Schmelter (2007).
- (f) Models with observational noise, that is, $cov(y(x_i), y(x_j)) = K(x_i, x_j) + \tau^2 \delta_{ij}$, so that we need to repeat observations at the same points to estimate τ^2 ; see Bates et al. (1996), and Bettinger et al. (2008).

6.1.3 Information Contained in Design Points

6.1.3.1 Points Providing Zero Information

To express information obtained from an exact design ξ_N , we use the information matrix

$$\mathbf{IM}(\xi_N) = \mathbf{X}^T \Sigma^{-1} \mathbf{X} = \sum_{i,j \in \mathcal{N}} f(x_i) g_{ij}^{\mathcal{N}} f^T(x_j),$$

where $g_{ij}^{\mathcal{N}}$ stands for the i, j th element of the matrix Σ^{-1} , $(g_{ij}^{\mathcal{N}})_{i,j=1}^{\mathcal{N}} = \Sigma^{-1}$, and $\mathcal{N} = \{1, \dots, N\}$. Despite the fact that the information matrix $\mathbf{IM}(\xi_N)$ cannot be decomposed into the sum of information measures for each point, we can characterize the points that provide zero information.

Lemma 6.1 (Pázman 2010, Lemma 1) *Let $\xi_N = \{x_1, \dots, x_N\}$ be an exact design and \mathcal{A} be a subset of \mathcal{N} . If there exist vectors $\mathbf{a}_i \in \mathbb{R}^m$, $i \in \mathcal{A}$, such that the vector of regression functions $\mathbf{f}(x)$ can be represented as*

$$\mathbf{f}(x) = \sum_{i \in \mathcal{A}} K(x, x_i) \mathbf{a}_i,$$

for all $x \in \{x_1, \dots, x_N\}$, then all points in the set $\{x_i \mid i \notin \mathcal{A}\}$ provide zero information; that is,

$$\mathbf{IM}(\{x_i\}_{i \in \mathcal{A}}) = \mathbf{IM}(\xi_N).$$

The proof of Lemma 6.1 is given in Appendix 6.A.

Corollary 6.1 (Pázman 2010, Corollary 1) *Let $\xi_N = \{x_1, \dots, x_N\}$ be an exact design. Define the vector*

$$\mathbf{a}_i = \sum_{j=1}^N g_{ij}^{\mathcal{N}} \mathbf{f}(x_j).$$

If $\mathbf{a}_i = 0$ for some i , then the point x_i provides zero information.

Example 6.1

Consider the linear regression model, that is, $\mathbf{f}(x) = (1, x)^T$, with covariance kernel $K(x, z) = \max\{0, 1 - |x - z|\}$. Then the three-point design $\{-1, 0, 1\}$ gives the same information as the whole process observed on $[-1, 1]$ and the same as any N -point design that includes points $-1, 0, 1$ in its support.

Let us now present two paradoxes that are specific for the case of correlated observations; see Näther (1985a).

6.1.3.2 Smit's Paradox

Consider the location-scale model $y(x) = \theta + \varepsilon(x)$ with the correlation function $\rho(t) = e^{-|t|}$ and the design interval $\mathcal{X} = [-1, 1]$. Let us compare two estimates of the parameter θ :

- The mean for an exact design

$$\bar{\theta}_{\xi_N} = N^{-1} \sum_{i=1}^N y(x_i).$$

- The mean for the continuous design

$$\bar{\theta}_c = \int_{-1}^1 y(x) dx / 2.$$

For the design $\xi_5 = \{-1, -0.5, 0, 0.5, 1\}$, straightforward calculation shows that $\text{Var}(\bar{\theta}_{\xi_5}) = 0.529$ while $\text{Var}(\bar{\theta}_c) = 0.568$. We can see that $\text{Var}(\bar{\theta}_{\xi_5}) < \text{Var}(\bar{\theta}_c)$; that is, the variance of the mean $\bar{\theta}_{\xi_N}$ using five observations is smaller than the variance of the mean for the continuous design; this is called Smit's paradox (see Smit 1961).

Consider now the design $\xi_9 = \{-1, -\frac{3}{4}, -\frac{1}{2}, \dots, \frac{1}{2}, \frac{3}{4}, 1\}$ with nine support points. Note that ξ_9 is obtained from ξ_5 by adding four points. Calculus gives that $\text{Var}(\bar{\theta}_{\xi_9}) = 0.542$ and we can observe that $\text{Var}(\bar{\theta}_{\xi_9}) > \text{Var}(\bar{\theta}_{\xi_5})$. This means that for correlated observations, the variance of the mean $\bar{\theta}_{\xi_N}$ can be increased by additional observations, which never happens in the case of independent errors.

It is worth noting that the variance of the continuous BLUE of θ in the location-scale model is 0.5, that is, slightly smaller than $\text{Var}(\bar{\theta}_{\xi_5})$. This means that the observation of a process at five points gives almost the same information as the continuous observation of the process.

6.1.3.3 Estimates with Zero Variance

Another interesting effect happens if we consider the location-scale model with a correlation function such that $\rho(2) = -1$. Then for the two-point design $\xi_2 = \{-1, 1\}$, we have $\text{Var}(\bar{\theta}_{\xi_2}) = 0$. This means that the BLUE yields exactly the true value of the parameter θ , which can never be in the case of uncorrelated observations. In general, estimation with zero variance is possible only if the correlation function is positive semidefinite but not positive definite.

6.2 Designs for One-Parameter Models

Throughout this section, we consider the one-parameter model

$$y(x) = \theta f(x) + \varepsilon(x), \tag{6.7}$$

where θ is a scalar parameter, $f(x)$ is a continuous function on \mathcal{X} , and $E[\varepsilon(x)] = 0$, and $E[\varepsilon(x)\varepsilon(z)] = K(x, z)$.

6.2.1 Designs for the BLUE

6.2.1.1 Designs for a Continuous Observation

Suppose that $\mathcal{X} \subset \mathbb{R}$ and that an observation of the whole process $\{y(x)\}_{x \in \mathcal{X}}$ is available. The estimate $\hat{\theta}$ is called BLUE if $\hat{\theta}$ admits the representation

$$\hat{\theta} = \int y(x) dG(x),$$

where G is a function of bounded variation, that is, $G \in BV(\mathcal{X})$, $E[\hat{\theta}] = \theta$, and

$$E(\hat{\theta} - \theta)^2 = \inf \left\{ E \left(\int y(x) dG(x) - \theta \right)^2 : G \in BV(\mathcal{X}), E \left(\int y(x) dG(x) \right) = \theta \right\}.$$

Note that the condition of unbiasedness in terms of G has the form

$$\int f(x) dG(x) = 1.$$

The following result is proved in N  ther (1985a), p. 19. This result gives a necessary condition for an estimator to be the BLUE.

Theorem 6.1 (N  ther 1985a, Theorem 2.3) *If*

$$\int K(x, z) dG(x) = Cf(z),$$

for all $z \in \mathcal{X}$ and $\int f(x) dG(x) = 1$, then the estimate

$$\hat{\theta}(G) = \int y(x) dG(x),$$

is the BLUE. Moreover,

$$\text{Var}(\hat{\theta}) = C.$$

The existence of the solution of the Wiener–Hopf integral equation

$$\int K(x, z) dG(x) = f(z),$$

in the general case is a very hard and often ill-posed problem. Some analytic results have been obtained in the case of stationary processes having the spectral density in the form of the ratio of polynomials (see Pisarenko and Rozanov 1963 and N  ther 1985a, Sec. 2.3).

6.2.1.2 Results of Sacks and Ylvisaker

Let $\{\xi_N\}_{N \in \mathbb{N}}$ be a sequence of designs that converges to a continuous design μ , where for each $N \in \mathbb{N}$, ξ_N is an N -point design. Then the design problem can be viewed as how the discrete BLUE $\hat{\theta}(\xi_N)$ approximates the continuous BLUE $\hat{\theta}(\mu)$. The best known asymptotic

result was obtained by Sacks and Ylvisaker (1966, 1968), who studied a sequence of exact designs, which is asymptotically optimal in the sense that the convergence

$$\lim_{N \rightarrow \infty} |\hat{\theta}(\xi_N) - \hat{\theta}(\mu)|^2 = 0,$$

holds with the best possible convergence rate.

Suppose that the design space \mathcal{X} is an interval $[a, b]$. The sequence $\{\xi_N\}$ is called asymptotic optimal for the BLUE if

$$\lim_{N \rightarrow \infty} \frac{\text{Var}(\hat{\theta}(\xi_N)) - \text{Var}(\hat{\theta}(\mu))}{\inf_{\tilde{\xi}_N} \text{Var}(\hat{\theta}(\tilde{\xi}_N)) - \text{Var}(\hat{\theta}(\mu))} = 1.$$

To formulate the main result of Sacks and Ylvisaker (1966), we first define $\alpha(x)$ (assuming its existence) as a difference of the two limits as z approaches x from below and above

$$\alpha(x) = \lim_{z \nearrow x} \frac{\partial K(x, z)}{\partial z} - \lim_{z \searrow x} \frac{\partial K(x, z)}{\partial z}, \quad (6.8)$$

where the covariance kernel K can correspond to a stationary or non-stationary process.

Theorem 6.2 (Sacks and Ylvisaker 1966) *Assume that $\alpha(x) > 0$ and the function $f(x)$ in the model (6.7) enables the representation*

$$f(x) = \int K(z, x)h(z)dz,$$

where $h(z)$ is continuous and $x, z \in \mathcal{X} = [a, b]$. Then the sequence $\{\xi_N\}_{N \in \mathbb{N}}$ defined by $\xi_N = \{x_{11}, \dots, x_{NN}\}$ where x_{iN} is such that

$$\int_a^{x_{iN}} |\alpha(x)h^2(x)|^{1/3} dx = \frac{i-1}{N-1} \int_a^b |\alpha(x)h^2(x)|^{1/3} dx,$$

$i = 1, \dots, N$, is asymptotically optimal. Moreover,

$$\text{Var}(\hat{\theta}(\xi_N)) \rightarrow \text{Var}(\hat{\theta}(\mu))$$

with the rate $O(N^{-2})$.

Example 6.2

Consider the one-parameter model

$$y(x) = \theta x^2 + \sigma W(x),$$

where $x \in [a, b]$ and $\{W(x)\}_{x \in [a, b]}$ is a Wiener process. The equidistant design (i.e., the design with equally spaced design points) on $[a, b]$ is asymptotically optimal for estimating θ ; see Sacks and Ylvisaker (1966). Moreover, for the model

$$y(x) = \theta x^\gamma + \sigma W(x), \quad \gamma < 1/2, \quad x \in [0, 1],$$

the design with points $x_i = (i/N)^{3/(2\gamma-1)}, i = 1, \dots, N$, is asymptotically optimal.

Example 6.3

Consider the location-scale model with correlation structure given by the stationary Ornstein–Uhlenbeck process, that is, $\rho(t) = e^{-\lambda|t|}$. Then equidistant designs are optimal; see Kiselak and Stehlík (2008) and Zagoraïou and Baldi-Antognini (2009).

6.2.2 Optimal Design for OLS Estimation

In view of (6.6), for an N -point design $\xi_N \in \Xi_N$, the variance of the OLS estimate of θ in the model (6.7) is given by

$$\text{Var}(\tilde{\theta}) = \frac{\sum_{i=1}^N \sum_{j=1}^N K(x_i, x_j) f(x_i) f(x_j)}{\left(\sum_{i=1}^N f(x_i)\right)^2}.$$

Consequently, for an approximate design $\xi \in \Xi$, we consider the functional

$$D(\xi) = \left[\int f^2(x) \xi(dx) \right]^{-2} \iint K(x, z) f(x) f(z) \xi(dx) \xi(dz) \tag{6.9}$$

as the design optimality functional. In general, the optimality functional (6.9) is not convex, and therefore the problem of finding the optimal design is hard. The situation is much simpler in the case of the location-scale model where $f(x) = 1$ for all $x \in \mathcal{X}$.

6.2.2.1 Location-Scale Model

For the function $f(x) = 1$, the design optimality functional (6.9) becomes

$$D(\xi) = \iint K(x, z) \xi(dx) \xi(dz). \tag{6.10}$$

Lemma 6.2 *The functional D defined in (6.10) is convex. Moreover, if the covariance kernel K is strictly positive definite, then D is strictly convex. That is,*

$$D((1 - \alpha)\xi + \alpha\xi_0) < (1 - \alpha)D(\xi) + \alpha D(\xi_0),$$

for all $0 < \alpha < 1$ and any two measures ξ , and ξ_0 on \mathcal{X} such that $\xi - \xi_0$ is a non zero (signed) measure.

The proof of Lemma 6.2 is given in Appendix.

The following result uses Lemma 6.2 and serves as an equivalence theorem (see Chapter 2, Theorem 2.1), which can be used to verify the optimality of a given design.

Theorem 6.3 (Zhigljavsky et al. 2010)

(i) An approximate design ξ^* minimizes the functional D defined in (6.10) if and only if

$$\min_{x \in \mathcal{X}} b(x, \xi^*) \geq D(\xi^*), \quad (6.11)$$

where the function b is given by

$$b(x, \xi) = \int K(x, z) \xi(dz).$$

(ii) In particular, a design ξ^* is optimal if the function $b(\cdot, \xi^*)$ is constant, that is,

$$b(x, \xi^*) = D(\xi^*),$$

for all $x \in \mathcal{X}$.

In the following examples, we present cases where analytical expressions for optimal designs can be found and verified using Theorem 6.3. In these examples, we suppose that $\mathcal{X} = [-1, 1]$. Details can be found in Zhigljavsky et al. (2010).

Example 6.4

For the location-scale model with exponential correlation function $\rho(t) = e^{-\lambda|t|}$, the optimal design ξ^* is a mixture of the continuous uniform measure on the interval $[-1, 1]$ and a two-point discrete measure supported on $\{-1, 1\}$, that is, the design ξ^* has the density

$$p^*(x) = \omega^* \left(\frac{1}{2} \delta_1(x) + \frac{1}{2} \delta_{-1}(x) \right) + (1 - \omega^*) \frac{1}{2} \mathbf{1}_{[-1,1]}(x),$$

where $\omega^* = 1/(1 + \lambda)$, $\delta_x(\cdot)$ denotes the Dirac measure concentrated at the point x and $\mathbf{1}_A(\cdot)$ is the indicator function of a set A . Note that the function $b(\cdot, \xi^*)$ is constant and given by $D(\xi^*) = 1/(1 + \lambda)$.

Example 6.5

For the location-scale model with triangular correlation function $\rho(t) = \max\{0, 1 - \lambda|t|\}$, we have the following on optimal designs minimizing (6.10):

- For $\lambda \in \mathbb{N}$, the optimal design is a discrete uniform measure supported at the $1 + 2\lambda$ equidistant points, $t_j = j/\lambda - 1$, $j = 0, 1, \dots, 2\lambda$. For this design, $D(\xi^*) = 1/(1 + 2\lambda)$.
- For any $\lambda > 0$, the optimal design ξ^* is a discrete symmetric measure supported at $2n$ points $\pm t_1, \pm t_2, \dots, \pm t_n$ with weights w_1, \dots, w_n at t_1, \dots, t_n , where $n = \lceil 2\lambda \rceil$,

$$(w_1, \dots, w_n) = \frac{1}{n(n+1)} (\lceil n/2 \rceil, \dots, 3, n-2, 2, n-1, 1, n).$$




FIGURE 6.1 Support points of the optimal designs in the location-scale model with triangular correlation function $\rho(t) = \max\{0, 1 - \lambda|t|\}$ for different values of λ .

The symbol $[z]$ stands for the smallest integer that is larger or equal to z . Here, t_1, \dots, t_n denote the ordered quantities $|u_1|, \dots, |u_n|$, where $u_j = -1 + j/\lambda, j = 1, \dots, n-1, u_n = 1$. Moreover, $D(\xi^*) = 2\lambda/(n(n+1))$.

The support points for various values of λ are depicted in Figure 6.1, where a vertical line at λ intersects the curves.

Example 6.6

Let $\rho(t) = -\ln(t^2)$. Then the asymptotic optimal design minimizing the functional (6.10) is the arcsine density on the interval $[-1, 1]$ with density

$$p^*(x) = \frac{1}{\pi\sqrt{1-x^2}}.$$

Example 6.7

Let $\rho(t) = 1/|t|^\alpha$ with $0 < \alpha < 1$. Then the asymptotic optimal design minimizing the functional (6.10) is a beta distribution on the interval $[-1, 1]$ with density

$$p^*(x) = \frac{2^{-\alpha}}{B\left(\frac{1+\alpha}{2}, \frac{1+\alpha}{2}\right)} (1+x)^{\frac{\alpha-1}{2}} (1-x)^{\frac{\alpha-1}{2}}.$$

6.2.2.2 Hajek Result

The following result is proved by Hajek (1956).

Theorem 6.4 (Hajek 1956) Consider the location-scale model with stationary error process $\{\varepsilon(x)\}_{x \in \mathcal{X}}$. Suppose that the correlation function ρ is convex on $(0, \infty)$. Let G^* be a function such that the estimate

$$\hat{\theta} = \int y(x) dG^*(x),$$

is the BLUE. Then the function G^* is monotonically increasing.

Thus, for the model (6.7) with $f(x) = 1$ and a stationary error process having a convex correlation function, G^* is a proper distribution function corresponding to a probability measure. This means that the pair OLS plus optimal design for OLS is the best possible pair and coincides with the BLUE for a continuous observation of the process.

Note that the exponential and triangular correlation functions are convex.

Example 6.8

For the location-scale model with exponential correlation function

$$\rho(t) = e^{-\lambda|t|}, \quad \mathcal{X} = [a, b],$$

we have that

$$dG^*(x) = \frac{1}{2 + \lambda(b-a)} [\delta_a(x) + \delta_b(x) + 1] dx,$$

and

$$\hat{\theta} = \frac{1}{2 + \lambda(b-a)} \left[y(a) + y(b) + \int_a^b y(x) dx \right],$$

is the BLUE; see Näther (1985a, p. 57) for details.

6.2.3 Optimal Design for SLS

Consider the one-parameter model (6.7) with the design space as a discrete set $\mathcal{X} = \{x_1, \dots, x_n\}$ and $f(x_i) \neq 0$ for all $i = 1, \dots, n$ (if $f(x_j) = 0$ for some j , then the point x_j can be removed from \mathcal{X} without changing the value of the estimates). Assume also that we are using SLS and the matrix S with signs on the diagonal can be chosen along with the weights p_1, \dots, p_n assigned to the points x_1, \dots, x_n in \mathcal{X} . The design optimality functional $D(\xi)$ becomes

$$D(\xi) = \frac{\sum_{i=1}^n \sum_{j=1}^n K(x_i, x_j) f(x_i) f(x_j) s_i s_j p_i p_j}{\left(\sum_{i=1}^n f(x_i) s_i p_i \right)^2} \quad (6.12)$$

with $s_i \in \{-1, 1\}$ for all $i = 1, \dots, n$ and $\sum_{i=1}^n p_i = 1$.

Denote $w_i = s_i p_i$ and call it signed weight of a point x_i in the design ξ . Since $\sum_{i=1}^n |w_i| = \sum_{i=1}^n p_i = 1$, the signed measure ξ , which assigns weights w_i to points x_i , $i = 1, \dots, n$, belongs to the space $\Xi^{(S)}$ of signed measures.

The problem of finding an optimal design and an optimal SLS estimate simultaneously in the linear regression model with one parameter consists therefore in optimizing the

functional

$$\frac{\sum_{i=1}^n \sum_{j=1}^n K(x_i, x_j) f(x_i) f(x_j) w_i w_j}{(\sum_{i=1}^n f(x_i) w_i)^2} \tag{6.13}$$

with respect to the signed weights $\{w_1, \dots, w_n\}$ such that $\sum_{i=1}^n |w_i| > 0$. Note that the value of the functional (6.13) does not change if we change all the weights $w_i \rightarrow cw_i$ ($i = 1, \dots, n$) for arbitrary $c \neq 0$.

Despite the fact that the functional is not convex, the problem of determining the optimal design can be easily solved by applying the Cauchy–Schwarz inequality.

Theorem 6.5 *Assume that the matrix $\Sigma = (K(x_i, x_j))_{i,j=1,\dots,n}$ is positive definite and $f(x_i) \neq 0$ for all $i = 1, \dots, n$. Then the optimal weights w_1^*, \dots, w_n^* minimizing (6.13) subject to $\sum_{i=1}^n |w_i| = 1$ are given by*

$$w_i^* = c \frac{(\Sigma^{-1} \mathbf{f})_i}{f(x_i)} \quad i = 1, \dots, n,$$

where $c = (\sum_{i=1}^n (\Sigma^{-1} \mathbf{f})_i / f(x_i))^{-1}$ and $\mathbf{f} = (f(x_1), \dots, f(x_n))^T$.

The proof of Theorem 6.5 is given in Appendix 6.A.

For the design $\xi^* = \{x_1, \dots, x_n; w_1^*, \dots, w_n^*\}$, we have $D(\xi^*) = (\mathbf{f}^T \Sigma^{-1} \mathbf{f})^{-1}$, that is, the variance of the SLS estimate coincides with the variance of the BLUE constructed using all observations. This means that the pair {SLS estimate, design ξ^* } provides the optimal pair {estimate, design} for the problem (6.7). This result (in a slightly different form) is obtained in (Näther 1985a, Theorem 5.3).

Example 6.9

Consider the location-scale model on $\mathcal{X} = \{-1, 0, 1\}$ and the Gaussian correlation function $\rho(t) = e^{-t^2/2}$. The optimal signed measure ξ^* assigns the weights 0.455, -0.09 , and 0.455, respectively, to the points $-1, 0$, and 1 . It gives $\text{Var}(\hat{\theta}_{\text{BLUE}}) = \text{Var}(\tilde{\theta}_S | \xi^*) = 0.563$. Note that for OLS the optimal design ξ_{OLS}^* is concentrated at only two points -1 and 1 giving them weights 0.5 each. It gives $\text{Var}(\tilde{\theta} | \xi_{\text{OLS}}) = 0.568$, which is slightly larger than the variance of the BLUE.

6.3 Optimal Designs for the BLUE

Consider the general model (6.2) with $\theta \in \mathbb{R}^m$. Recall from Section 6.1.2 that having N observations the BLUE of the parameter θ has the form

$$\hat{\theta} = (\mathbf{X}^T \Sigma^{-1} \mathbf{X})^{-1} \mathbf{X}^T \Sigma^{-1} \mathbf{y},$$

which is the solution of the weighted least squares problem, and the covariance matrix of the BLUE is given by

$$\text{Var}(\hat{\boldsymbol{\theta}}) = (\mathbf{X}^T \boldsymbol{\Sigma}^{-1} \mathbf{X})^{-1}.$$

The results in this section are the multi-parameters counterparts of the one-parameter results of Section 6.2.

6.3.1 BLUE for Continuous Observations

Similar to the one-parameter case of Section 6.2.1, suppose that an observation of the whole process $\{y(x)\}_{x \in \mathcal{X}}$ is available. The estimate $\hat{\boldsymbol{\theta}}$ is called BLUE if $\hat{\boldsymbol{\theta}}$ admits the representation

$$\hat{\boldsymbol{\theta}} = \int y(x) d\mathbf{G}(x),$$

where $\mathbf{G}(x) = (G_1(x), \dots, G_m(x))^T$ is a vector of functions with bounded variation, $E[\hat{\boldsymbol{\theta}}] = \boldsymbol{\theta}$ and

$$E\|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|^2 = \inf \left\{ E \left(\left\| \int y(x) d\mathbf{G}(x) - \boldsymbol{\theta} \right\|^2 \right) : \mathbf{G} \text{ such that } E \left(\int y(x) d\mathbf{G}(x) \right) = \boldsymbol{\theta} \right\}.$$

Note that the condition of unbiasedness in terms of \mathbf{G} has the form

$$\int \mathbf{f}(x) d\mathbf{G}^T(x) = \mathbf{I}_m,$$

where \mathbf{I}_m is the $m \times m$ identity matrix.

The following result is proved in N  ther (1985a, p. 19), and is a generalization of Theorem 6.1.

Theorem 6.6 (N  ther 1985a) *If*

$$\int K(x, z) d\mathbf{G}(x) = \mathbf{C}\mathbf{f}(z),$$

for all $z \in \mathcal{X}$ *and*

$$\int \mathbf{f}(x) d\mathbf{G}^T(x) = \mathbf{I}_m,$$

then the estimate

$$\hat{\boldsymbol{\theta}}(G) = \int y(x) d\mathbf{G}(x),$$

is the BLUE. Moreover, $\text{Var}(\hat{\boldsymbol{\theta}}) = \mathbf{C}$.

Note that numerical computation of the continuous BLUE can be done as follows; see Näther (1985a) for more details: First, one has to find a solution H of the integral equation

$$\int K(x, z)dH(x) = f(z),$$

and then the matrix

$$C = \left(\int f(x)dH^T(x) \right)^{-1}$$

is defined. If the matrix C is non-singular, then the BLUE is given by

$$\hat{\theta} = C \int y(x)dH(x).$$

Remark 6.1 Note that there does not exist a design problem in the case of continuous observation since measurements are performed at all points. However, observation of a process by exact designs can approximate continuous observation in many ways. Suppose that a sequence of exact designs ξ_N converges to a continuous measure μ and another sequence of exact designs $\tilde{\xi}_N$ converges to a continuous measure $\tilde{\mu}$. Then it follows that variances of the BLUE for designs ξ_N and $\tilde{\xi}_N$ converge to the same value if the supports of the measures μ and $\tilde{\mu}$ are the same.

6.3.2 Results of Sacks and Ylvisaker

Let $\{\xi_N\}_{N \in \mathbb{N}}$ be a sequence of designs that converges to a continuous design μ and $\mathcal{X} = [a, b]$. Then the problem is how the discrete BLUE $\hat{\theta}(\xi_N)$ approximates the continuous BLUE $\hat{\theta}(\mu)$.

In the multi parameter case, the covariance matrix $D(\xi_N) = (X^T \Sigma^{-1} X)^{-1}$ converges to a limiting matrix as $N \rightarrow \infty$. To compare matrices corresponding to different designs, we have to use a suitable functional, for example, the L -optimality functional $\Psi(M) = \text{tr}(LM)$ with a given positive definite matrix L and M is any matrix. The sequence $\{\xi_N\}_{N \in \mathbb{N}}$ is called asymptotically L -optimal for the BLUE if

$$\lim_{N \rightarrow \infty} \inf_{\tilde{\xi}_N} \frac{\Psi(D(\xi_N) - D(\mu))}{\Psi(D(\tilde{\xi}_N) - D(\mu))} = 1.$$

In addition to the L -criterion, some other ways to define criteria of the asymptotic optimality are considered in Sacks and Ylvisaker (1968). The following result is a generalization of Theorem 6.2.

Theorem 6.7 (Sacks and Ylvisaker 1968) *Suppose that the function α defined in (6.8) is positive and the components of the vector $f(x) = (f_1(x), \dots, f_m(x))^T$ satisfy the representation*

$$f_j(x) = \int K(z, x)h_j(z) dz,$$

where the functions $h_j(z)$ are continuous, $j = 1, \dots, m$. Consider the design $\xi_N = \{x_{11}, \dots, x_{NN}\}$ where the points x_{iN} are defined by

$$\int_a^{x_{iN}} |\alpha(x)\mathbf{h}^T(x)\mathbf{Lh}(x)|^{1/3} dx = \frac{i-1}{N-1} \int_a^b |\alpha(x)\mathbf{h}^T(x)\mathbf{Lh}(x)|^{1/3} dx,$$

$i = 1, \dots, N$, $\mathbf{h} = (h_1, \dots, h_m)^T$, and $\alpha(x)$ as in (6.8), then the sequence $\{\xi_N\}_{N \in \mathbb{N}}$ is asymptotically L -optimal.

Remark 6.2 In the linear and quadratic regression model with exponential covariance function $e^{-\lambda|t|}$, the exact N -point D -optimal design converges to the equally spaced design as $\lambda \rightarrow 0$; see Dette et al. (2008).

In the following sections, we present three methods of numerical construction of optimal designs: the method of exchange of points, the method of virtual noise, and the method using the expansion of the covariance kernel.

6.3.3 Exchange-Type Algorithms

The steps of the exchange algorithm for computing an exact Ψ -optimal N -point design are as follows. Note that the exchange algorithm for models with uncorrelated observations is given in Chapter 21.

First, we have to choose a starting design $\xi_N^{(0)}$ such that the covariance matrix $\mathbf{D}(\xi_N^{(0)})$ is nonsingular. At iteration j , one point from the design $\xi_N^{(j)}$ is replaced by another point from the design space, where we need to find a replacement giving a decrease of the Ψ -criterion. The algorithm has to be stopped if the decrease is smaller than a given tolerance bound and proceeds to the next iteration otherwise.

For the D -optimality criterion, Brimkulov et al. (1980) proposed the procedure where simultaneously a new point is introduced, which is defined by

$$\mathbf{x}^+ = \arg \max_{\mathbf{x} \in \mathcal{X} \setminus \xi_N^{(j)}} \phi(\mathbf{x}, \xi_N^{(j)}),$$

and a point is removed, which is defined by

$$\mathbf{x}^- = \arg \min_{\mathbf{x} \in \xi_N^{(j)}} \phi(\mathbf{x}, \xi_N^{(j)}).$$

Here, the function ϕ is given by

$$\phi(\mathbf{x}, \xi) = \frac{\psi^2(\mathbf{x}, \xi) + \tilde{\mathbf{f}}^T(\mathbf{x}, \xi) \mathbf{I} \mathbf{M}^{-1}(\xi) \tilde{\mathbf{f}}(\mathbf{x}, \xi)}{\psi^2(\mathbf{x}, \xi)},$$

where

$$\begin{aligned}\psi^2(\mathbf{x}, \xi) &= K(\mathbf{x}, \mathbf{x}) - \mathbf{k}^T(\mathbf{x}, \xi) \Sigma^{-1}(\xi) \mathbf{k}(\mathbf{x}, \xi), \\ \mathbf{k}(\mathbf{x}, \xi) &= (K(\mathbf{x}, x_1), \dots, K(\mathbf{x}, x_N)), \quad \mathbf{IM}(\xi) = \mathbf{X}^T(\xi) \Sigma^{-1}(\xi) \mathbf{X}(\xi), \\ \tilde{\mathbf{f}}(\mathbf{x}, \xi) &= \mathbf{f}(\mathbf{x}) - \mathbf{X}^T(\xi) \Sigma^{-1}(\xi) \mathbf{k}(\mathbf{x}, \xi), \quad \mathbf{X}(\xi) = (\mathbf{f}(x_1), \dots, \mathbf{f}(x_N))^T.\end{aligned}$$

Note that the procedure proposed in Brimkulov et al. (1980) is based on ideas of using an analog of the sensitivity function from the equivalence theorem for the case of uncorrelated observations. Uciniski and Atkinson (2004) have developed formulas for the straightforward exchange algorithm that provides the best decrease of the optimality functional as follows:

1. Select an initial design $\xi^{(0)} = \{x_1^{(0)}, \dots, x_N^{(0)}\}$ such that $x_i^{(0)} \neq x_j^{(0)}$ for $i \neq j$ and $\det \mathbf{IM}(\xi^{(0)}) \neq 0$. Define the matrices $\mathbf{X}^{(0)} = (\mathbf{f}(x_1^{(0)}), \dots, \mathbf{f}(x_N^{(0)}))^T$, $\Sigma^{(0)} = (K(x_i^{(0)}, x_j^{(0)}))_{i,j=1}^N$, and $\mathbf{IM}^{(0)} = \mathbf{X}^{(0)T} (\Sigma^{(0)})^{-1} \mathbf{X}^{(0)}$.
2. Set $j = 0$.
3. Determine

$$(i^*, z^*) = \arg \max_{(i,z) \in \{1, \dots, N\} \times \mathcal{X}} \Delta(\mathbf{x}_i, \mathbf{z}),$$

where

$$\Delta(\mathbf{x}_i, \mathbf{z}) = (\det \mathbf{IM}(\xi_{x_i \leftrightarrow z}^{(j)}) - \det \mathbf{IM}(\xi^{(j)})) / \det \mathbf{IM}(\xi^{(j)}),$$

and $\xi_{x_i \leftrightarrow z}^{(j)}$ denotes the design obtained from $\xi^{(j)}$ if the points x_i and z are interchanged.

4. If $\Delta(\mathbf{x}_i, \mathbf{z}) < \delta$, where δ is some given positive tolerance, then terminate. Otherwise, set $\xi^{(j+1)} = \xi_{x_{i^*} \leftrightarrow z^*}^{(j)}$ and determine $\mathbf{X}^{(j+1)}$, $\Sigma^{(j+1)}$, and $\mathbf{IM}^{(j+1)}$ corresponding to $\xi^{(j+1)}$ (expressions simplifying the numerical computation are given in Uciniski and Atkinson 2004). Set $j \leftarrow j + 1$ and go to step 3.

This method has been used in a number of practical examples; see Glatzer and Müller (1999), Müller (2005), Müller (2007), Müller and Stehlík (2010), and Stehlík et al. (2008).

6.3.4 Constructing Optimal Designs by Expansion of the Covariance Kernel

Fedorov and Müller (2004) proposed to approximate the model (6.2) by the following mixed-effect model

$$y(\mathbf{x}) = \boldsymbol{\theta}^T \mathbf{f}(\mathbf{x}) + \sum_{j=1}^q \beta_j \psi_j(\mathbf{x}) + \varepsilon_o(\mathbf{x}),$$

where $\boldsymbol{\theta}$ is the vector fixed-effect parameters, β_j are random-effect parameters, $\psi_j(\mathbf{x})$ are eigenfunctions of Mercer's expansion (Kanwal 1997) of the covariance kernel, and $\varepsilon_o(\mathbf{x})$

is an error process with no correlation. Then an optimal design is determined using the truncated D -criterion, namely, the minimization of the determinant of the covariance matrix for the parameter θ , while both parameters θ and β are considered as unknown parameters. In general, the computation of optimal designs requires the knowledge of eigenfunctions ψ_j . Fedorov and Müller (2004) developed an approximation of the sensitivity function $\phi(x, \xi)$, which is used in the exchange algorithm for computing discrete optimal designs for models with uncorrelated observations. Specifically, for a discrete design $\xi = \{x_1, \dots, x_n; w_1, \dots, w_n\}$, the function $\phi(x, \xi)$ has the form

$$\phi(x, \xi) = \tilde{f}^T(x, \xi)M^{-1}(\xi)\tilde{f}(x, \xi),$$

where

$$\begin{aligned}\tilde{f}^T(x, \xi) &= f^T(x) + k^T(x, \xi)\Sigma_{S_\xi}^{-1} \left(W + \Sigma_{S_\xi}^{-1} \right)^{-1} WX, \\ M &= X^T(W - W \left(W + \Sigma_{S_\xi}^{-1} \right)^{-1} W)X,\end{aligned}$$

and the matrices X and W are defined by $X = (f(x_1), \dots, f(x_n))^T$, $W = \frac{N}{s^2} \text{diag}\{w_1, \dots, w_n\}$, $k(x, \xi) = (K(x, x_1), \dots, K(x, x_n))$, $S_\xi = \text{supp}(\xi)$, $\Sigma_{S_\xi} = (K(x_i, x_j))_{i,j=1}^n$, and s^2 is a tuning parameter that should be close to zero, for example, $s^2 = 10^{-6}$.

This method has been used in a number of practical examples (see Fedorov and Flanagan 1997; Müller 2005, 2007).

6.3.5 Method of Virtual Noise

Pázman and Müller (2001) have proposed the method of virtual noise to determine optimal designs. This method considers the following extended model:

$$\tilde{y}(x) = \theta^T f(x) + \varepsilon(x) + \epsilon(x),$$

where $\varepsilon(x)$ is the original stochastic process such that $\text{Cov}(\varepsilon(x_1), \varepsilon(x_2)) = K(x_1, x_2)$ and $\epsilon(x)$ is an additional heteroscedastic white noise depending on a design. The two processes $\varepsilon(x)$ and $\epsilon(x)$ are assumed to be uncorrelated. For a given design $\xi = \{x_1, \dots, x_n; w_1, \dots, w_n\}$, the variance of white noise at design points is given by $\text{Var}(\epsilon(x_j)) = \gamma\sigma^2 \ln(\max_i w_i/w_j)$, where γ is a tuning parameter, which should be small, for example, $\gamma = 10^{-6}$.

The information matrix for the model with virtual noise is given by

$$\mathbf{M}_\epsilon(\xi) = X^T \left(\Sigma_{S_\xi} + \gamma \text{diag} \left(\ln \left[\frac{\max_i w_i}{w_1} \right], \dots, \ln \left[\frac{\max_i w_i}{w_n} \right] \right) \right)^{-1} X,$$

where $X = (f(x_1), \dots, f(x_n))^T$. Note that despite the presence of the weights w_i , the designs ξ in this approach can only be considered as exact N -point designs (see Pázman and Müller 1998).

To deal with the optimal design problem of finding the exact N -point optimal design $\xi_N \in \Xi_N$, which maximizes $\Phi(\mathbf{M}_\epsilon(\xi_N))$, several algorithms have been proposed (see Müller and Pázman 1999; Müller and Pázman 2003). These algorithms consist of a step-wise one-point correction of the design. For example, in Müller and Pázman (2003) the

design $\xi^{(j)}$ is updated as

$$\xi^{(j+1)} = \frac{j}{j+1} \xi^{(j)} + \frac{1}{j} \delta_{x^*},$$

where $\delta_x = \{x; 1\}$ is the one-point measure supported at $x \in \mathcal{X}$. The point x^* minimizes the directional derivative of $\Phi(\mathbf{M}_\epsilon(\xi^{(j)}))$ in the direction of δ_x where $\xi^{(j)} = \{x_1, \dots, x_n; w_1^{(j)}, \dots, w_n^{(j)}\}$. Thus $x^* = x_{i^*}$, where

$$i^* = \arg \min_{i=1, \dots, n} \frac{1}{w_i^{(j)}} \left(d(x_i) - \frac{\mathbf{1}_{B_j}(i)}{N_{B_j}} \sum_{k=1}^n d(x_k) \right),$$

$d(x) = \mathbf{a}^T(x) \nabla \Phi(\mathbf{M}) \mathbf{a}(x)$, $\mathbf{a}(x) = \sum_{k=1}^n g_{ik} f(x_k)$, $\Sigma^{-1} = (g_{ik})_{i,k=1}^n$, $\mathbf{M} = \mathbf{X}^T \Sigma^{-1} \mathbf{X}$, $\nabla \Phi(\mathbf{M}) = \partial \Phi(\mathbf{M}) / \partial \mathbf{M}$, $\mathbf{X} = (f_i(x_j))_{j=1, \dots, n}^{i=1, \dots, m}$, N_{B_j} is the cardinality of the set $B_j = \{i \in \{1, \dots, n\} : w_i = \max_k w_k^{(j)}\}$, and $\mathbf{1}_{B_j}(i)$ is the indicator function of the set B_j . An initial design $\xi^{(0)}$ can be chosen as the uniform discrete design supported at points forming a discretization of the design space.

This method has been used in a number of practical examples (see Müller 2005, 2007). A relationship between the method of virtual noise and the method of the expansion of the covariance kernel is discussed in Pázman (2010) and Pázman and Müller (2010).

6.3.6 Design for Prediction in the Quadratic Model

Consider the quadratic model

$$y(x) = \theta_1 + \theta_2 x + \theta_3 x^2 + \sigma W(x), \tag{6.14}$$

where $x \geq 0$, $f(x) = (1, x, x^2)^T$, and $W(x)$ is the standardized Wiener process with

$$K(u, v) = \text{cov}(W(u), W(v)) = \min(u, v).$$

Note that the Wiener process is nonstationary. In addition to estimation of the vector of parameters $\theta = (\theta_1, \theta_2, \theta_3)^T$, one can be interested in prediction of the process $\{y(u)\}_{u \in [a, b]}$ at a point x , where $x > b$ and $[a, b]$ is a design interval. The best linear unbiased predictor is given by

$$\hat{y}(x) = \hat{\theta}^T f(x) + \mathbf{k}^T(x, \xi) \Sigma^{-1} (\mathbf{y} - \mathbf{X} \hat{\theta}),$$

where $\mathbf{k}(x, \xi) = (K(x, x_1), \dots, K(x, x_N))^T$, $\mathbf{y} = (y(x_1), \dots, y(x_N))^T$ is a vector of observations at design points x_1, \dots, x_N , and $\hat{\theta}$ is the BLUE of θ . The MSE of $\hat{y}(x)$ can easily be calculated as

$$\text{MSE}(\hat{y}(x)) = K(x, x) - \mathbf{k}^T(x, \xi) \Sigma^{-1} \mathbf{k}(x, \xi) + \mathbf{c}_x^T (\mathbf{X}^T \Sigma^{-1} \mathbf{X})^{-1} \mathbf{c}_x,$$

where $\mathbf{X} = (f_i(x_j))_{j=1, \dots, N}^{i=1, \dots, m}$ and $\mathbf{c}_x = f(x) - \mathbf{X}^T \Sigma^{-1} \mathbf{k}(x, \xi)$.

Theorem 6.8 (Harman and Stulajter 2011) *Consider the process (6.14) with design interval $[a, b]$. Then the N -point design with points $x_i = a + (b - a)(i - 1)/(N - 1)$, $i = 1, \dots, N$, is optimal for estimating the unknown parameters $\theta_1, \theta_2, \theta_3$ with respect to any continuous Loewner isotonic criterion as well as for the MSE of the best linear unbiased predictor.*

Further results on optimal designs for prediction of processes can be found in Harman and Stulajter (2010, 2011) and Zimmerman (2006). Note that the design problem for prediction is a major problem in computer experiments (see Bates et al. 1996; Dette and Pepelyshev 2010; Pronzato and Müller 2012).

6.4 Optimal Designs for OLS

Recall from (6.2) and (6.4) that the OLS estimate is given by $\tilde{\theta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}$ with covariance matrix

$$\text{Var}(\tilde{\theta}) = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \Sigma \mathbf{X} (\mathbf{X}^T \mathbf{X})^{-1}. \quad (6.15)$$

Note that the BLUE can only be used if the correlation structure of the errors is known, and its misspecification can lead to a considerable loss of efficiency. At the same time, the OLS estimate does not employ the structure of the correlation. Obviously, the OLS estimates can be less efficient than the BLUE, but in many cases, the loss of efficiency is small. For example, consider the location-scale model with a stationary error process, the Gaussian correlation function $\rho(t) = e^{-\lambda^2 t^2}$, and the exact design $\xi = \{-1, -2/3, -1/3, 1/3, 2/3, 1\}$. Suppose that the specified value of λ equals 1 while the true value is 2. Then the variance of the BLUE is 0.528, while the variance of the OLS estimate is 0.433. If the specified value of λ equals the true value, then the variance of the BLUE is 0.382. A similar relationship between the variances holds if the location-scale model and the Gaussian correlation function are replaced by a polynomial model and a triangular or exponential correlation function, respectively. For a more detailed discussion concerning advantages of the OLS against the weighted least squares estimate, see Bickel and Herzberg (1979) and Section 6.1 in Näther (1985a).

Some results on the efficiency of the OLS estimation compared to the BLUE estimation are obtained in Kiefer and Wynn (1981), Bischoff (1995a,b), and Puntanen et al. (2011).

6.4.1 OLS for Approximate Designs

Consider the model (6.2) and the case when a continuous observation of a process is available, and let ξ be an approximate design with nonsingular matrix $\mathbf{M}(\xi) = \int \mathbf{f}(x) \mathbf{f}^T(x) \xi(dx)$. Hence, the least squares problem has the form

$$\inf_{\theta} \int (y(x) - \theta^T \mathbf{f}(x))^2 \xi(dx)$$

with the solution

$$\tilde{\theta} = \mathbf{M}^{-1}(\xi) \int f(x)y(x)\xi(dx),$$

which is called the continuous LSE. The covariance matrix of the estimate $\tilde{\theta}$ has the form

$$\text{Var}(\tilde{\theta}) = \mathbf{M}(\xi)^{-1} \mathbf{B}(\xi, \xi) \mathbf{M}(\xi)^{-1},$$

where

$$\mathbf{B}(\xi, \nu) = \iint K(x, z) f(x) f^T(z) \xi(dx) \nu(dz).$$

The general approximate design problem is therefore given by

$$\min_{\xi \in \Xi} \Phi(\mathbf{D}(\xi)),$$

where Φ is some functional on the set of $m \times m$ matrices and the matrix \mathbf{D} is defined by

$$\mathbf{D}(\xi) = \mathbf{M}(\xi)^{-1} \mathbf{B}(\xi, \xi) \mathbf{M}(\xi)^{-1}.$$

6.4.2 Results of Bickel and Herzberg

Consider the general model (6.2) with stationary error process and $\mathcal{X} = [-T, T]$. Suppose that for N observations, the correlation function is given by

$$\rho_N(t) = \rho_0(Nt), \quad (6.16)$$

where $\rho_0(t) = \gamma\rho(t) + (1 - \gamma)\delta_t$ and $\rho(t) \rightarrow 0$ as $t \rightarrow \infty$, $\gamma \in (0, 1]$. The following regularity conditions are needed to present a main result of Bickel and Herzberg (see Bickel and Herzberg 1979) described in Theorem 6.9.

(C1) *The regression functions $f_1(t), \dots, f_m(t)$ are linearly independent and bounded on the interval $[-T, T]$ and satisfy a first order Lipschitz condition, that is,*

$$|f_i(t) - f_i(s)| \leq \kappa|t - s|$$

and

$$|f_i(t)| \leq \kappa$$

for all $t, s \in [-T, T]$, $i = 1, \dots, m$.

(C2) *There exists a twice differentiable quantile function $a : (0, 1) \rightarrow \mathbb{R}$ and a positive constant $\kappa < \infty$ such that for all $u \in (0, 1)$,*

$$\frac{1}{\kappa} \leq a'(u) \leq \kappa, \quad |a''(u)| \leq \kappa. \quad (6.17)$$

The quantile function a is used to generate exact N -point designs $\xi_N = \{t_{1N}, \dots, t_{NN}\}$, that is,

$$t_{iN} = a\left(\frac{i-1}{N-1}\right), \quad i = 1, \dots, N. \quad (6.18)$$

(C3) The correlation function $\rho(t)$ is differentiable with bounded derivative, that is, $|\rho'(t)| \leq \kappa$, $t \in (0, \infty)$, and satisfies $\rho'(t) \leq 0$ for sufficiently large t . This assumption implies that $\rho(t)$ is nonnegative for sufficiently large t .

(C4) The correlation function ρ is integrable, that is, $\int |\rho(t)| dt < \infty$. As a consequence, the function

$$Q(t) = \sum_{j=1}^{\infty} \rho(jt), \quad (6.19)$$

is well defined and finite for all $t > 0$.

Theorem 6.9 (Bickel and Herzberg 1979) Consider the model (6.2) with correlation function ρ_N defined in (13.6) for observations at N points t_{1N}, \dots, t_{NN} defined in (6.18). Assume that the correlation function ρ , the quantile function a , and the regression functions f_1, \dots, f_m satisfy the regularity assumptions (C1)–(C4), and suppose that the elements of the matrix

$$\mathbf{R}(a) = \left(\int_0^1 f_i(a(u))f_j(a(u))Q(a'(u)) du \right)_{i,j=1}^m,$$

exist and are finite. Then the variance–covariance matrix (6.15) of the least squares estimate is well defined and

$$\lim_{N \rightarrow \infty} \sigma^{-2} N \text{Var}(\tilde{\boldsymbol{\theta}}) = \mathbf{W}^{-1}(a) + 2\gamma \mathbf{W}^{-1}(a) \mathbf{R}(a) \mathbf{W}^{-1}(a),$$

where γ is the parameter of the correlation function $\rho_\theta(t)$ and the matrix $\mathbf{W}(a)$ is given by

$$\mathbf{W}(a) = \left(\int_0^1 f_i(a(u))f_j(a(u)) du \right)_{i,j=1}^m.$$

The conditions on the quantile function a imply that the corresponding design ξ has a continuous density, say, $p : [0, 1] \rightarrow \mathbb{R}$. Therefore, the matrices $\mathbf{W}(a)$ and $\mathbf{R}(a)$ can be expressed in terms of p as follows:

$$\mathbf{R}(\xi) = \left(\int_0^1 f_i(t)f_j(t)Q(1/p(t))p(t) dt \right)_{i,j=1}^m,$$

and

$$W(\xi) = \left(\int_0^1 f_i(t)f_j(t)p(t) dt \right)_{i,j=1}^m .$$

In the remaining part of this section, we consider the one-parameter case when $f = f_1$. Define $H(t) = Q(t) - tQ'(t)$, where $Q(t)$ is given by (6.19), and the function $q(x, \mu, \tau)$ by

$$q(t, \mu, \tau) = \begin{cases} \frac{1}{H^{-1}(\mu(1 - \tau/f^2(t)))}, & \mu(1 - \tau/f^2(t)) \geq 0, \\ 0, & \text{otherwise.} \end{cases}$$

Theorem 6.10 (Bickel and Herzberg 1979) *Assume that the regression function f in the one-parameter linear model (6.7) is continuous and Q defined in (6.19) is strictly convex. Then the optimal design exists, and its density is of the form $q(t, \mu^*, \tau^*)$, where the parameters μ^* and τ^* satisfy the equations*

$$\int q(t, \mu^*, \tau^*) dt = 1$$

and

$$\frac{2 \int Q(1/q(t))f^2(t)q(t)dt}{\int f^2(t)q(t)dt} = \mu^* - \frac{1}{2\gamma}.$$

6.4.3 Results for Long-Range Dependence Error Structure

As in the previous section, we consider the general model (6.2) with stationary error process having short-range dependence. Suppose that for N observations, the correlation function is given by

$$\rho_N(t) = \rho_o(Nt),$$

where $\rho_o(t) = \gamma\rho(t) + (1-\gamma)\delta_t$ and $\rho(t) \rightarrow 0$ as $t \rightarrow \infty, \gamma \in (0, 1]$. As in the previous section, we assume that regularity conditions (C1)–(C3) are satisfied, and instead of assumption (C4) of Bickel and Herzberg (1979), we now assume that

$$\int_0^\infty |\rho(t)| dt = \infty. \tag{6.20}$$

The condition (6.20) means the long-range dependence of the observations. Note that in this case, it follows that

$$\int_0^\infty |\rho(t)| dt = \sum_{k=0}^\infty |\rho(k)| = \infty,$$

where $\rho(k) = \text{cov}(\varepsilon(u), \varepsilon(u+k))$ for any u . The correlation function of a stationary process with long-range dependence can be written as

$$\rho_\alpha(t) = \frac{L(t)}{|t|^\alpha}, \quad |t| \rightarrow \infty, \tag{6.21}$$

where $0 < \alpha \leq 1$ and $L(t)$ is a slowly varying function (SVF) for large t (see Doukhan et al. 2003). In particular, ρ_α satisfies

$$\rho_\alpha(t) = O(1/|t|^\alpha), \quad |t| \rightarrow \infty,$$

and we will say that $\rho_\alpha(t)$ belongs to SVF family.

First, we introduce two parametric families of correlation functions that are important in applications. The correlation function ρ_α belongs to the Cauchy family if it is defined by

$$\rho_\alpha(t) = \frac{1}{(1 + |t|^\beta)^{\alpha/\beta}}, \tag{6.22}$$

where $\beta > 0, 0 < \alpha \leq 1$. The correlation function ρ_α belongs to the Mittag-Leffler family if it is defined by

$$\rho_\alpha(t) = E_{\nu, \beta}(-|t|^\alpha), \quad E_{\nu, \beta}(-t) = \Gamma(\beta) \sum_{k=0}^{\infty} \frac{(-t)^k}{\Gamma(\nu k + \beta)}, \tag{6.23}$$

where $0 < \alpha \leq 1, 0 < \nu \leq 1, \beta \geq \nu$ (for more details, see Dette et al. 2009).

In the following, we present optimal designs for the three families of correlation functions, which are given by (6.21) through (6.23). The function $Q(t) = \sum_{j=1}^{\infty} \rho(jt)$ plays an important role in the asymptotic analysis by Bickel and Herzberg (1979), but in the case of long-range dependence, this function is infinite. For an asymptotic analysis under long-range dependence, we introduce the function

$$Q_\alpha(t) = \lim_{N \rightarrow \infty} \frac{1}{d_\alpha(N)} \sum_{j=1}^N \rho_\alpha(jt), \tag{6.24}$$

where the normalizing sequence is given by

$$d_\alpha(N) = \begin{cases} N^{1-\alpha} & \text{if } \alpha < 1 \text{ and } \rho_\alpha \text{ has the form (6.22) or (6.23),} \\ \ln N & \text{if } \alpha = 1 \text{ and } \rho_\alpha \text{ has the form (6.22) or (6.23),} \\ L(N)N^{1-\alpha} & \text{if } \alpha < 1 \text{ and } \rho_\alpha \text{ has the form (6.21),} \\ L(N) \ln N & \text{if } \alpha = 1 \text{ and } \rho_\alpha \text{ has the form (6.21),} \end{cases}$$

and shown in Lemma 6.3, which the function $Q_\alpha(t)$ is well defined.

Lemma 6.3 (Dette et al. 2009) *If the correlation function $\rho_\alpha(t)$ belongs either to the Cauchy, Mittag-Leffler, or SVF family, then the limit in (6.24) exists and is given by*

$$Q_\alpha(t) = \begin{cases} \frac{c}{(1-\alpha)|t|^\alpha}, & 0 < \alpha < 1, \\ \frac{c}{|t|}, & \alpha = 1, \end{cases}$$

where

$$c = \begin{cases} \frac{\Gamma(\beta)}{\Gamma(\beta-\nu)}, & \text{if } \rho_\alpha(t) \text{ belongs to the Mittag-Leffler family,} \\ 1, & \text{otherwise.} \end{cases}$$

The following result describes the asymptotic behavior of the OLS for the case of the long-range dependence.

Theorem 6.11 (Dette et al. 2009) *Consider the model (6.2) with correlation function ρ_N defined in (6.16) for observations at N points t_{1N}, \dots, t_{NN} defined by (6.18). Assume that the correlation function ρ_α is either an element of the Cauchy, Mittag-Leffler, or SVF family. If $\int_0^1 Q_\alpha(a'(t)) dt < \infty$ and the regularity assumptions (C1)–(C3) stated in the previous subsection are satisfied, then we obtain for the variance-covariance matrix of the least squares estimate defined in (6.15):*

$$\sigma^{-2} \frac{N}{d_\alpha(N)} \text{Var}(\tilde{\theta}) = 2\gamma W^{-1}(a)R_\alpha(a)W^{-1}(a) + O(1/d_\alpha(N)),$$

where the matrices W and R_α are given by

$$W(a) = \left(\int_0^1 f_i(a(u))f_j(a(u)) du \right)_{i,j=1}^m,$$

$$R_\alpha(a) = \left(\int_0^1 f_i(a(u))f_j(a(u))Q_\alpha(a'(u)) du \right)_{i,j=1}^m.$$

Note that the constant γ only appears as a factor in the asymptotic variance-covariance matrices of the least squares estimate. Because most optimality criteria are positively homogeneous (see, e.g., Pukelsheim 1993), it is reasonable to consider the matrix

$$W^{-1}(a)R_\alpha(a)W^{-1}(a),$$

which is proportional to the asymptotic variance-covariance matrix of the least squares estimate. Moreover, if the function a corresponds to a continuous distribution with a density, say ϕ , then $a'(t) = 1/\phi(t)$ and the asymptotic variance-covariance matrix of the least squares estimate is proportional to the matrix

$$D_\alpha(\phi) = W^{-1}(\phi)R_\alpha(\phi)W^{-1}(\phi),$$

where the matrices $W(\phi)$ and $R_\alpha(\phi)$ are given by

$$W(\phi) = \left(\int_{-T}^T f_i(t)f_j(t)\phi(t) dt \right)_{i,j=1,\dots,m},$$

$$R_\alpha(\phi) = \left(\int_{-T}^T f_i(t)f_j(t)Q_\alpha(1/\phi(t))\phi(t) dt \right)_{i,j=1,\dots,m}$$

$$= \frac{c}{1-\alpha} \left(\int_{-T}^T f_i(t)f_j(t)\phi^{1+\alpha}(t) dt \right)_{i,j=1,\dots,m},$$

respectively, and we have used the representation $Q_\alpha(t) = c/((1-\alpha)|t|^\alpha)$ for the last identity. An (asymptotic) optimal design for classical least squares estimation minimizes an appropriate function of the matrix $D_\alpha(\phi)$.

Note that under long-range dependence, the variance–covariance matrix of the least squares estimate converges more slowly to zero than in the case of independent or short-range dependent errors. In the case of short-range dependence, no normalization is necessary other than normalizing the variance–covariance matrix. Under long-range dependence, an additional factor $d_\alpha(N)/N$ is needed. Moreover, it is worthwhile to note that under long-range dependence, the asymptotic variance–covariance matrix is fully determined by the function $Q_\alpha(t)$ and does not depend on the particular correlation function ρ_α . In the following section, we discuss several examples in order to illustrate the concept.

In most cases, the asymptotic optimal designs for the regression model (6.2) have to be found numerically; explicit solutions are only possible for very simple models. The following result established optimal designs for linear models with one parameter.

Theorem 6.12 (Dette et al. 2009) *Assume that the correlation function ρ_α is either an element of the Cauchy, Mittag-Leffler, or SVF family. Then, for the one-parameter linear regression model (6.7), the asymptotic optimal design exists; it is absolute continuous with respect to the Lebesgue measure and has the density*

$$p^*(t) = \begin{cases} \frac{1}{H_\alpha^{-1}(\mu - \tau/f^2(t))} = \left(\frac{1-\alpha}{1+\alpha}(\mu - \tau/f^2(t)) \right)^{\frac{1}{\alpha}}, & \mu - \tau/f^2(t) \geq 0, \\ 0, & \text{otherwise,} \end{cases} \quad (6.25)$$

where the constants μ and τ are given by

$$\mu = 2 \frac{\int f^2(t)Q_\alpha(1/p^*(t))p^*(t) dt}{\int f^2(t)p^*(t) dt},$$

$$\tau = \int f^2(t)Q_\alpha(1/p^*(t))p^*(t) dt + \int f^2(t)Q'_\alpha(1/p^*(t)) dt.$$

We now consider two special cases, which are of particular importance. If the number of parameters m equals 1 and $f(t) \equiv 1$, we obtain the location-scale model, and the asymptotic optimal density is the uniform density, that is,

$$p^*(t) = \begin{cases} \frac{1}{2T}, & |t| \leq T, \\ 0, & \text{otherwise.} \end{cases} \tag{6.26}$$

Similarly, in the linear model through the origin, we have $m = 1, f(t) \equiv t$, and the asymptotic optimal density is given by

$$p(t) = \begin{cases} 0, & |t| \leq \sqrt{\tau/\mu}, \\ \left(\frac{1-\alpha}{1+\alpha}(\mu - \tau/t^2)\right)^{1/\alpha}, & \sqrt{\tau/\mu} \leq |t| \leq T, \\ 0, & \text{otherwise,} \end{cases}$$

where

$$\mu = 2 \frac{\int t^2 p^{1+\alpha}(t) dt}{(1 - \alpha) \int t^2 p(t) dt}, \quad \tau = \int t^2 p^{1+\alpha}(t) dt,$$

and α is the parameter of the correlation function. These formulas are given for $0 < \alpha < 1$. For $\alpha = 1$ and $f(t) = t$, the asymptotic optimal density is the uniform density (6.26). The optimal densities for the parameters $\alpha = 1/4, 1/2, 3/4, 0.95$ and $T = 1$ are displayed in Figure 6.2.

6.4.3.1 Linear Regression

Consider the model (6.2) with $m = 2, f_1(t) = 1, f_2(t) = t$, which corresponds to the linear regression model. In this case, the optimal design for estimating the slope (i.e., the




FIGURE 6.2 Asymptotic optimal design densities on the interval $[-1, 1]$ for the linear regression model through the origin.




FIGURE 6.3
Asymptotic D -optimal design densities for the linear regression model on the interval $[-1, 1]$.

e_2 -optimal design) has the density (6.25), while the D -optimal designs have to be determined numerically in all cases. Some D -optimal design densities on the interval $[-1, 1]$ corresponding to the parameters $\alpha = 1/4, 1/2, 3/4, 0.95$ are displayed in Figure 6.3.

6.4.4 Optimal Designs for OLS

In this section, we consider the general model (6.2) with arbitrary covariance kernel. For an exact N -point design ξ_N , the covariance matrix of the least squares estimate $\tilde{\theta} = \tilde{\theta}_{\xi_N}$ given in (6.15) can be written as

$$\text{Var}(\tilde{\theta}) = D(\xi_N) = M^{-1}(\xi_N)B(\xi_N, \xi_N)M^{-1}(\xi_N), \tag{6.27}$$

where the matrices M and B are given by

$$M(\xi_N) = \int f(x)f^T(x)\xi_N(dx), \tag{6.28}$$

$$B(\xi_N, \xi_N) = \iint K(x, z)f(x)f^T(z)\xi_N(dx)\xi_N(dz). \tag{6.29}$$

The definition of the matrices $M(\xi)$ and $B(\xi, \xi)$ can be extended to an approximate design ξ , provided that the corresponding integrals exist. The matrix

$$D(\xi) = M^{-1}(\xi)B(\xi, \xi)M^{-1}(\xi), \tag{6.30}$$

is called the covariance matrix for the design ξ (as in Section 6.4.1) if the matrices $B(\xi, \xi)$ and $M^{-1}(\xi)$ are well defined. An (approximate) optimal design minimizes a functional of the covariance matrix $D(\xi)$ over the set Ξ .

Note that in general the function $D(\xi)$ is not convex (with respect to the Loewner ordering) on the space of all approximate designs. This implies that even if one determines

optimal designs by minimizing a convex functional, say Φ , of the matrix $D(\xi)$, the corresponding functional $\xi \rightarrow \Phi(D(\xi))$ is in general not convex on the space of designs Ξ . Consider, for example, the case $m = 1$ where $D(\xi)$ is given by

$$D(\xi) = \left[\int f^2(x) \xi(dx) \right]^{-2} \int \int K(x, z) f(x) f(z) \xi(dx) \xi(dz); \tag{6.31}$$

then it is obvious that this functional is not necessarily convex. On the other hand, for the location-scale model, we have $m = 1, f(x) = 1$ for all $x \in \mathcal{X}$, and this expression reduces to

$$D(\xi) = \int \int K(x, z) \xi(dx) \xi(dz).$$

In the case, $K(x, z) = \sigma^2 \rho(x - z)$, where ρ is a correlation function; this functional is convex on the set of probability measures on the domain \mathcal{X} ; see Lemma 6.1 in Zhigljavsky et al. (2010) and Lemma 4.3 in N  ther (1985a). For this reason (namely, the convexity of the functional $D(\xi)$), most of the literature on (asymptotic) optimal design problems for least squares estimation in the presence of correlated observations considers the location-scale model. This corresponds to the estimation of the mean of a stationary process; see, for example, Boltze and N  ther (1982) and N  ther (1985a,b).

Consider an optimality functional Φ on the space of the non-negative definite matrices, and define

$$\varphi(\mathbf{x}, \xi) = \mathbf{f}^T(\mathbf{x}) D(\xi) \mathbf{C}(\xi) M^{-1}(\xi) \mathbf{f}(\mathbf{x}), \tag{6.32}$$

$$b(\mathbf{x}, \xi) = \text{tr} \left[\mathbf{C}(\xi) M^{-1}(\xi) \mathbf{B}(\xi, \xi_{\mathbf{x}}) M^{-1}(\xi) \right], \tag{6.33}$$

where

$$\mathbf{C}(\xi) = \left. \frac{\partial \Phi(D)}{\partial D} \right|_{D=D(\xi)}$$

(here, we assume differentiability of the optimality functional). The following result is a reformulation of the necessary condition of design optimality.

Theorem 6.13 (Dette et al. 2013) *Let ξ^* be any design minimizing the functional $\Phi(D(\xi))$. Then the inequality*

$$\varphi(\mathbf{x}, \xi^*) \leq b(\mathbf{x}, \xi^*) \tag{6.34}$$

holds for all $\mathbf{x} \in \mathcal{X}$, where the functions $\varphi(\mathbf{x}, \xi)$ and $b(\mathbf{x}, \xi)$ are defined in (6.32) and (6.33), respectively. Moreover, there is equality in (6.34) for ξ^ -almost all \mathbf{x} , that is, $\xi^*(\mathcal{A}) = 0$ where*

$$\mathcal{A} = \mathcal{A}(\xi^*) = \{ \mathbf{x} \in \mathcal{X} \mid \varphi(\mathbf{x}, \xi^*) < b(\mathbf{x}, \xi^*) \}$$

is the set of $\mathbf{x} \in \mathcal{X}$ such that the inequality (6.34) is strict.

In this section, we consider the matrix $D(\xi)$ defined in (6.30) as the matrix optimality functional, which we are going to minimize on the set Ξ of all designs, such that the matrices $B(\xi, \xi)$ and $M^{-1}(\xi)$ (and therefore the matrix $D(\xi)$) are well defined. We say that a design ξ^* is universally optimal if

$$D(\xi^*) \leq D(\xi)$$

in the sense of the Loewner ordering for any design $\xi \in \Xi$; see Pukelsheim (1993). Note that a design ξ^* is universally optimal if and only if ξ^* is c -optimal for any nonzero vector $c \in \mathbb{R}^m$; that is, $c^T D(\xi^*) c \leq c^T D(\xi) c$ for any $\xi \in \Xi$ and any nonzero $c \in \mathbb{R}^m$.

For a given design $\xi \in \Xi$ with nonsingular matrix $M(\xi)$, introduce the vector-valued function

$$g(x) = \int K(x, z) f(z) \xi(dz) - \Lambda f(x), \quad x \in \mathcal{X}, \quad (6.35)$$

where $\Lambda = B(\xi, \xi) M^{-1}(\xi)$. This function satisfies the equality

$$\int g(x) f^T(x) \xi(dx) = 0. \quad (6.36)$$

Additionally, as the vector of regression functions $f(\cdot)$ is continuous on \mathcal{X} , the function $g(\cdot)$ is continuous too.

The following theorem gives a necessary conditions for the universal optimality of a design.

Theorem 6.14 (Dette et al. 2013) *Consider the regression model (6.1) with covariance kernel K , a design $\xi \in \Xi$, and the corresponding vector-function $g(x)$ defined in (6.35). If the design ξ is universally optimal, then the function $g(x)$ can be represented in the form $g(x) = \gamma(x)f(x)$, where $\gamma(x)$ is a nonnegative function defined on \mathcal{X} such that $\gamma(x) = 0$ for all x in the support of the design ξ .*

Let us now discuss the case when the regression functions are proportional to eigenfunctions of the integral operator induced by the covariance kernel. To be precise, let \mathcal{X} denote a compact subset of a metric space, and let ν denote a measure on the corresponding Borel field with positive density. Consider the integral operator

$$T_K(f)(\cdot) = \int K(\cdot, z) f(z) \nu(dz) \quad (6.37)$$

on $L_2(\nu)$. Under certain assumptions on the kernel (e.g., if $K(x, z)$ is symmetric, continuous, and positive definite), T_K defines a symmetric, compact self-adjoint operator. In this case, Mercer's theorem (see, e.g., Kanwal 1997) shows that there exists a countable number of eigenfunctions $\varphi_1, \varphi_2, \dots$ with positive eigenvalues $\lambda_1, \lambda_2, \dots$ of the operator K , that is,

$$T_K(\varphi_\ell) = \lambda_\ell \varphi_\ell, \quad \ell = 1, 2, \dots \quad (6.38)$$

The next statement follows directly from Theorem 6.14.

Theorem 6.15 (Dette et al. 2013) *Let \mathcal{X} be a compact subset of a metric space and assume that the covariance kernel K defines an integral operator T_K of the form (6.37), where the eigenfunctions satisfy (6.38). Consider the regression model (6.1) with $f(x) = \mathbf{L}(\varphi_{i_1}(x), \dots, \varphi_{i_m}(x))^T$ and the covariance kernel K , where $\mathbf{L} \in \mathbb{R}^{m \times m}$ is a non-singular matrix. Then the design ν is universally optimal in the linear regression model (6.2).*

The following two results give optimal design in explicit form for polynomial regression models with two singular covariance kernels.

Theorem 6.16 (Dette et al. 2013) *Consider the linear regression model (6.1) with $f(x) = (1, x, x^2, \dots, x^{m-1})^T$ on the interval $\mathcal{X} = [-1, 1]$, and the covariance kernel $\rho(t) = -\ln(t^2)$. Then the arcsine design ξ_a with density*

$$p(x) = 1/(\pi\sqrt{1 - x^2}), \quad x \in (-1, 1)$$

is the universally optimal design.

Theorem 6.17 (Dette et al. 2013) *Consider the linear regression model (6.1) with $f(x) = (1, x, x^2, \dots, x^{m-1})^T$ on the interval $\mathcal{X} = [-1, 1]$, and covariance kernel $\rho(t) = 1/|t|^\alpha$. Then the design with generalized arcsine density*

$$p_\alpha(x) = \frac{2^{-\alpha}}{B(\frac{1+\alpha}{2}, \frac{1+\alpha}{2})} (1+x)^{\frac{\alpha-1}{2}} (1-x)^{\frac{\alpha-1}{2}},$$

is universally optimal.

Acknowledgments

The authors thank Martina Stein, who typed parts of this manuscript with considerable technical expertise. This work has been supported in part by the Collaborative Research Center *Statistical Modeling of Nonlinear Dynamic Processes* (SFB 823, Teilprojekt C2) of the German Research Foundation (DFG). The work of Pepelyshev was partly supported by Russian Foundation of Basic Research, project 12-01-00747.

Appendix 6.A Proofs

Proof of Lemma 6.1 Let $G^{\mathcal{N}} = (g_{i,j}^{\mathcal{N}})_{i,j=1}^{\mathcal{N}} = (\Sigma_{\mathcal{N}})^{-1}$ where $\Sigma_{\mathcal{N}} = (K(x_i, x_j))_{i,j \in \mathcal{N}}$. Also let $G^{\mathcal{A}} = (g_{i,j}^{\mathcal{A}})_{i,j \in \mathcal{A}} = (\Sigma_{\mathcal{A}})^{-1}$ where $\Sigma_{\mathcal{A}} = (K(x_i, x_j))_{i,j \in \mathcal{A}}$. Straightforward calculus gives

$$\begin{aligned}
 \mathbf{M}(\xi_N) &= \sum_{i,j \in \mathcal{N}} f(x_i) g_{ij}^N f^T(x_j) \\
 &= \sum_{k,l \in A} \sum_{i,j \in \mathcal{N}} a_k K(x_k, x_i) g_{ij}^N K(x_j, x_l) a_l^T \\
 &= \sum_{k,l \in A} a_k K(x_k, x_l) a_l^T \\
 &= \sum_{k,l \in A} f(x_k) g_{kl}^A f^T(x_l) = \mathbf{M}(\{x_i\}_{i \in A}),
 \end{aligned}$$

which completes the proof. □

Proof of Lemma 6.2 We have

$$\begin{aligned}
 D(\alpha \xi_2 + (1 - \alpha) \xi_1) &= \iint K(x, z) [\alpha \xi_2(dx) + (1 - \alpha) \xi_1(dx)] [\alpha \xi_2(dz) + (1 - \alpha) \xi_1(dz)] \\
 &= (1 - \alpha)^2 \iint K(x, z) \xi_1(dx) \xi_1(dz) + \alpha^2 \iint K(x, z) \xi_2(dx) \xi_2(dz) \\
 &\quad + 2\alpha(1 - \alpha) \iint K(x, z) \xi_1(dx) \xi_2(dz) \\
 &= \alpha^2 D(\xi_2) + (1 - \alpha)^2 D(\xi_1) + 2\alpha(1 - \alpha) \iint K(x, z) \xi_1(dx) \xi_2(dz) \\
 &= \alpha D(\xi_2) + (1 - \alpha) D(\xi_1) - \alpha(1 - \alpha) A,
 \end{aligned}$$

where

$$\begin{aligned}
 A &= \iint K(x, z) [\xi_2(dx) \xi_2(dz) + \xi_1(dx) \xi_1(dz) - 2\xi_2(dx) \xi_1(dz)] \\
 &= \iint K(x, z) \zeta(dx) \zeta(dz)
 \end{aligned}$$

and $\zeta(dx) = \xi_2(dx) - \xi_1(dx)$. Since the correlation function $K(x, z)$ is positive definite, it follows $A \geq 0$. If $K(x, z)$ is strictly positive definite, we have $A > 0$ whenever ζ is not trivial. Therefore, the functional $D(\cdot)$ is strictly convex. □

Proof of Theorem 6.5 Denote $K_{ij} = K(x_i, x_j)$, $f(x_i) = f_i$, $a_i = f_i w_i$, $i, j = 1, \dots, n$, $\mathbf{a} = (a_1, \dots, a_n)^T$. Then for any signed design $\xi = \{x_1, \dots, x_n; w_1, \dots, w_n\}$, we have

$$D(\xi) = \frac{\sum_i \sum_j K_{ij} f_i f_j w_i w_j}{(\sum_i f_i^2 w_i)^2} = \frac{\sum_i \sum_j K_{ij} a_i a_j}{(\sum_i f_i a_i)^2} = \frac{\mathbf{a}^T \Sigma \mathbf{a}}{(\mathbf{a}^T \mathbf{f})^2}.$$

Since Σ is symmetric and $\Sigma > \mathbf{0}$, there exists Σ^{-1} and a symmetric matrix $\Sigma^{1/2} > \mathbf{0}$ such that $\Sigma = \Sigma^{1/2} \Sigma^{1/2}$. Denote $\mathbf{b} = \Sigma^{1/2} \mathbf{a}$ and $\mathbf{d} = \Sigma^{-1/2} \mathbf{f}$. Then we can write $D(\xi)$ as $D(\xi) = \mathbf{b}^T \mathbf{b} / (\mathbf{b}^T \mathbf{d})^2$. The Cauchy–Schwarz inequality gives for any two vectors \mathbf{b} and \mathbf{d} : $(\mathbf{b}^T \mathbf{d})^2 \leq (\mathbf{b}^T \mathbf{b})(\mathbf{d}^T \mathbf{d})$, that is, $\mathbf{b}^T \mathbf{b} / (\mathbf{b}^T \mathbf{d})^2 \geq 1 / (\mathbf{d}^T \mathbf{d})$. This inequality with \mathbf{b} and \mathbf{d} as mentioned earlier is equivalent to

$$D(\xi) \geq \frac{1}{\mathbf{f}^T \Sigma^{-1} \mathbf{f}}$$

for all $\xi \in \Xi^{(S)}$. The equality in this inequality is attained if the vector \mathbf{b} is proportional to the vector \mathbf{d} ; that is, if $b_i = cd_i$ for all i and any $c \neq 0$. Then the equality $b_i = cd_i$ can be rewritten in the form $w_i = c(\Sigma^{-1}\mathbf{f})_i/f(x_i)$. \square

Appendix 6.B Common Correlation Functions

Definition 6.1 (Stationarity in the wide sense). *A random field is a stationary random field in the wide sense if $E[\varepsilon(\mathbf{x})] = \text{const}$ and $K(\mathbf{x}, \mathbf{z}) := \text{cov}(\varepsilon(\mathbf{x}), \varepsilon(\mathbf{z})) = \rho(\mathbf{x} - \mathbf{z})$.*

The covariance function $\rho(\mathbf{x})$ on \mathbb{R}^d is fully separable if $\rho(\mathbf{x}) = \rho_1(x_1) \cdots \rho_d(x_d)$. Note that the product and the sum of two covariance functions are also covariance functions; see Abrahamsen (1997), Sec 3.1.

Definition 6.2 *A stationary random field is an isotropic random field (in the wide sense) if the covariance function depends on distance alone, that is,*

$$K(\mathbf{x}, \mathbf{z}) = \rho(\|\mathbf{x} - \mathbf{z}\|),$$

and $\|\mathbf{x}\| = \sqrt{\mathbf{x}^T \mathbf{x}}$.

Definition 6.3 *A stationary random field is an anisotropic random field (in the wide sense) if the covariance function depends on a non-Euclidean norm of the difference of two points, that is,*

$$K(\mathbf{x}, \mathbf{z}) = \rho(\|\mathbf{x} - \mathbf{z}\|_{\mathbf{A}}),$$

where $\|\mathbf{x}\|_{\mathbf{A}} = \sqrt{\mathbf{x}^T \mathbf{A} \mathbf{x}}$ and \mathbf{A} is a positive semidefinite matrix. The function $\rho(\|\mathbf{x} - \mathbf{z}\|_{\mathbf{A}})$ is called an ellipsoidal correlation function.

In Table 6.1 we present correlation functions for the one-dimensional case. In the case of $d > 1$, most common isotropic correlation functions have the form $\rho(\|\mathbf{x}\|)$.

References

- Abrahamsen, P. (1997). A review of Gaussian random fields and correlation functions. Technical report 917, Norwegian Computing Center, Oslo Norway.
- Atkinson, A. C. (2008). Examples of the use of an equivalence theorem in constructing optimum experimental designs for random-effects nonlinear regression models. *Journal of Statistical Planning and Inference*, 138:2595–2606.

- Bates, R. A., Buck, R. J., Riccomagno, E., and Wynn, H. P. (1996). Experimental design and observation for large systems. *Journal of the Royal Statistical Society, Series B*, 58:77–94.
- Bettinger, R., Duchene, P., Pronzato, L., and Thierry, E. (2008). Design of experiments for response diversity. *Journal of Physics: Conference Series*, 135:12–17.
- Bickel, P. J. and Herzberg, A. M. (1979). Robustness of design against autocorrelation in time I: Asymptotic theory, optimality for location and linear regression. *Annals of Statistics*, 7:77–95.
- Bickel, P. J., Herzberg, A. M., and Schilling, M. F. (1981). Robustness of design against autocorrelation in time II: Optimality, theoretical and numerical results for the first-order autoregressive process. *Journal of the American Statistical Association*, 76:870–877.
- Bischoff, W. (1995a). Determinant formulas with applications to designing when the observations are correlated. *Annals of the Institute of Statistical Mathematics*, 47:385–399.
- Bischoff, W. (1995b). Lower bounds for the efficiency of designs with respect to the D -criterion when the observations are correlated. *Statistics*, 27:27–44.
- Boltze, L. and Näther, W. (1982). On effective observation methods in regression models with correlated errors. *Mathematische Operationsforschung und Statistik Series Statistics*, 13:507–519.
- Brimkulov, U. N., Krug, G. V., and Savanov, V. L. (1980). Numerical construction of exact experimental designs when the measurements are correlated (in Russian). *Zavodskaya Laboratoria*, 36: 435–442.
- Cambanis, S. (1985). Sampling designs for time series. In Hannan, E. J., Krishnaiah, P. R., and Rao, M. M., editors, *Handbook of Statistics*, p. 337–362. Elsevier, Amsterdam, the Netherlands.
- Cutler, D. R. (1993a). Efficient block designs for comparing test treatments to a control when the errors are correlated. *Journal of Statistical Planning and Inference*, 36:107–125.
- Cutler, D. R. (1993b). Efficient block designs for comparing test treatments to a control when the errors are correlated. *Journal of Statistical Planning and Inference*, 37:393–412.
- Dette, H. and Holland-Letz, T. (2009). A geometric characterization of c -optimal designs for heteroscedastic regression. *Annals of Statistics*, 37:4088–4103.
- Dette, H., Kunert, J., and Pepelyshev, A. (2008). Exact optimal designs for weighted least squares analysis with correlated errors. *Statistica Sinica*, 18:135–154.
- Dette, H., Leonenko, N. N., Pepelyshev, A., and Zhigljavsky, A. (2009). Asymptotic optimal designs under long-range dependence error structure. *Bernoulli*, 15:1036–1056.
- Dette, H. and Pepelyshev, A. (2010). Generalized Latin hypercube design for computer experiments. *Technometrics*, 52:421–429.
- Dette, H., Pepelyshev, A., and Holland-Letz, T. (2010). Optimal designs for random effect models with correlated errors with applications in population pharmacokinetics. *Annals of Applied Statistics*, 4:1430–1450.
- Dette, H., Pepelyshev, A., and Zhigljavsky, A. (2013). Optimal design for linear models with correlated observations. *Annals of Statistics*, 41:143–176.
- Doukhan, P., Oppenheim, G., and Taqqu, M. (2003). *Theory and Applications of Long-Range Dependence*. Birkhauser, Boston, MA.
- Fedorov, V. and Flanagan, D. (1997). Optimal monitoring network design based on Mercer’s expansion of covariance kernel. *Journal of Combinatorics, Information and System Sciences*, 23:237–250.
- Fedorov, V. and Müller, W. (2004). Optimum design for correlated processes via eigenfunction expansions. Research Report Series, Institut für Statistik und Mathematik, WU Vienna University of Economics and Business, Wien, Austria.
- Fedorov, V. and Müller, W. (2007). Optimum design for correlated fields via covariance kernel expansions. In López-Fidalgo, J., Rodríguez-Díaz, J. M., and Torsney, B., editors, *MODA 8—Advances in Model-Oriented Design*, p. 57–66. Physica-Verlag, Heidelberg, Germany.
- Fedorov, V., Wu, Y., and Zhang, R. (2012). Optimal dose-finding designs with correlated continuous and discrete responses. *Statistics in Medicine*, 31:217–234.
- Fedorov, V. V. (1996). Design of spatial experiments: Model fitting and prediction. In Gosh, S. and Rao, C. R., editors, *Handbook of Statistics*, p. 515–553. Elsevier, Amsterdam, the Netherlands.
- Glatzer, E. and Müller, W. G. (1999). A comparison of optimum design algorithms for regressions with correlated observations—A computational study. *Tatra Mountains Mathematical Publications*,

- 17:149–156. *Proceedings of the Third International Conference on Mathematical Statistics*, Smolenice, Slovakia, February 9–13, 1998 (Smolenice Castle).
- Hajek, J. (1956). Linear estimation of the mean value of a stationary random process with convex correlation function. *Czechoslovak Mathematical Journal*, 6:94–117.
- Harman, R. and Stulajter, F. (2010). Optimal prediction designs in finite discrete spectrum linear regression models. *Metrika*, 72:281–294.
- Harman, R. and Stulajter, F. (2011). Optimality of equidistant sampling designs for the Brownian motion with a quadratic drift. *Journal of Statistical Planning and Inference*, 141:2750–2758.
- Holland-Letz, T., Dette, H., and Pepelyshev, A. (2011). A geometric characterization of optimal designs for regression models with correlated observations. *Journal of the Royal Statistical Society, Series B*, 73:239–252.
- Holland-Letz, T., Dette, H., and Renard, D. (2012). Efficient algorithms for optimal designs with correlated observations in pharmacokinetics and dose finding studies. *Biometrics*, 68:138–145.
- Kanwal, R. (1997). *Linear Integral Equations*. Birkhauser, Boston, MA.
- Kiefer, J. and Wynn, H. (1981). Optimum balanced block and Latin square designs for correlated observations. *Annals of Statistics*, 9:737–757.
- Kiefer, J. and Wynn, H. (1984). Optimum and minimax exact treatment designs for one-dimensional autoregressive error processes. *Annals of Statistics*, 12:414–450.
- Kiselak, J. and Stehlík, M. (2008). Equidistant D -optimal designs for parameters of Ornstein-Uhlenbeck process. *Statistics and Probability Letters*, 78:1388–1396.
- Kunert, J., Martin, R. J., and Eccleston, J. (2010). A -optimal block designs for the comparison with a control for correlated errors and analysis with the weighted least squares estimate. *Journal of Statistical Planning and Inference*, 140:2719–2738.
- Müller, W. G. (2005). A comparison of spatial design methods for correlated observations. *Environmetrics*, 16:495–505.
- Müller, W. G. (2007). *Collecting Spatial Data: Optimum Design of Experiments for Random Fields*, revised edition. Contributions to statistics. Physica-Verlag, Heidelberg, Germany.
- Müller, W. G. and Pázman, A. (1999). An algorithm for the computation of optimum designs under a given covariance structure. *Computational Statistics*, 14:197–211.
- Müller, W. G. and Pázman, A. (2003). Measures for designs in experiments with correlated errors. *Biometrika*, 90:423–434.
- Müller, W. G. and Stehlík, M. (2004). An example of D -optimal designs in the case of correlated errors. In *COMPSTAT 2004—Proceedings in Computational Statistics*, p. 1543–1550. Physica-Verlag, Heidelberg, Germany.
- Müller, W. G. and Stehlík, M. (2010). Compound optimal spatial designs. *Environmetrics*, 21:354–364.
- Näther, W. (1985a). *Effective Observation of Random Fields*. Teubner Verlagsgesellschaft, Leipzig, Germany.
- Näther, W. (1985b). Exact design for regression models with correlated errors. *Statistics*, 16:479–484.
- Pázman, A. (2010). Information contained in design points of experiments with correlated observations. *Kybernetika*, 46:771–783.
- Pázman, A. and Müller, W. G. (1998). A new interpretation of design measures. In *MODA 5—Advances in Model-Oriented Data Analysis and Experimental Design (Marseilles, 1998)*, Contributions to Statistics, Atkinson, A., Pronzato, L., and Wynn, H. P. (Eds.), p. 239–246. Physica, Heidelberg, Germany.
- Pázman, A. and Müller, W. G. (2001). Optimal design of experiments subject to correlated errors. *Statistics and Probability Letters*, 52:29–34.
- Pázman, A. and Müller, W. G. (2010). A note on the relationship between two approaches to optimal design under correlation. In *MODA 9: Advances in Model-Oriented Design and Analysis*, Contributions to Statistics, Giovagnoli, A., Atkinson, A. C., Torsney, B., and May, C. (Eds.), p. 145–148. Physica-Verlag Heidelberg, Germany.
- Pisarenko, V. F. and Rozanov, J. A. (1963). On certain problems for stationary processes leading to integral equations related to Wiener-Hopf equations. *Problemy Peredaci Informacii*, 14: 113–135.

- Pronzato, L. and Müller, W. (2012). Design of computer experiments: Space filling and beyond. *Statistics and Computing*, 22:681–701.
- Pukelsheim, F. (1993). *Optimal Design of Experiments*. John Wiley & Sons, New York.
- Pukelsheim, F. (2006). *Optimal Design of Experiments*. SIAM, Philadelphia, PA.
- Puntanen, S., Styan, G., and Isotalo, J. (2011). *Matrix Tricks for Linear Statistical Models: Our Personal Top Twenty*. Heidelberg, Springer, Berlin, Germany.
- Sacks, J. and Ylvisaker, N. D. (1966). Designs for regression problems with correlated errors. *Annals of Mathematical Statistics*, 37:66–89.
- Sacks, J. and Ylvisaker, N. D. (1968). Designs for regression problems with correlated errors; many parameters. *Annals of Mathematical Statistics*, 39:49–69.
- Schmelter, T. (2007). Considerations on group-wise identical designs for linear mixed models. *Journal of Statistical Planning and Inference*, 137:4003–4010.
- Smit, J. C. (1961). Estimation of the mean of a stationary stochastic process by equidistant observations. *Trabajos de Estadística*, 12:34–45.
- Stehlík, M., Rodríguez-Díaz, J. M., Müller, W. G., and López-Fidalgo, J. (2008). Optimal allocation of bioassays in the case of parametrized covariance functions: An application to lung's retention of radioactive particles. *TEST*, 17:56–68.
- Ucinski, D. and Atkinson, A. (2004). Experimental design for time-dependent models with correlated observations. *Studies in Nonlinear Dynamics and Econometrics*, 8:13.
- Zagoraiou, M. and Baldi-Antognini, A. (2009). Optimal designs for parameter estimation of the Ornstein-Uhlenbeck process. *Applied Stochastic Models in Business and Industry*, 25:583–600.
- Zhigljavsky, A., Dette, H., and Pepelyshev, A. (2010). A new approach to optimal design for linear models with correlated observations. *Journal of the American Statistical Association*, 105:1093–1103.
- Zimmerman, D. (2006). Optimal network design for spatial prediction, covariance parameter estimation, and empirical prediction. *Environmetrics*, 17:635–652.

Section III

Designs Accommodating Multiple Factors

7

Regular Fractional Factorial Designs

Robert Mee and Angela Dean

CONTENTS

7.1	Introduction.....	280
7.1.1	Fractional Factorial Experiments.....	280
7.2	Regular Fractions and Models.....	281
7.2.1	Models for Factorial Experiments.....	281
7.2.2	Aliasing of Factorial Effects.....	283
7.2.3	Defining Contrast Subgroup and Construction of Designs.....	285
7.2.4	Criteria for Good Designs: Resolution and Aberration.....	287
7.3	Resolution III Designs.....	288
7.4	Resolution IV Designs.....	290
7.4.1	Constructing Even Resolution IV Designs.....	291
7.4.2	Constructing Resolution IV Designs with $k \leq 5n/16$	292
7.4.2.1	Key Results for Existence and Construction of Resolution IV SOS Designs.....	293
7.4.2.2	Summary of Resolution IV SOS Designs with $k > n/4$	294
7.4.2.3	Minimum Aberration Resolution IV Designs with $k \leq n/4$	294
7.4.3	Other Criteria for Choosing a Resolution IV Design.....	295
7.5	Resolution V Designs.....	296
7.6	Row Coincidence Matrix.....	297
7.7	Analysis of Factorial Experiments.....	300
7.8	Screening.....	304
7.9	Augmentation of Designs and Follow-Up Experiments.....	305
7.9.1	Augmenting Resolution III Designs by Foldover.....	305
7.9.2	Augmenting Resolution IV Designs by Foldover and Semifoldover.....	307
7.9.3	Other Design Augmentation Strategies.....	308
7.10	Regular Fractions for More Than Two Levels.....	309
7.11	Equivalence of Designs.....	312
7.12	Incorporation of Noise Factors.....	314
7.13	Summary and Discussion.....	315
	References.....	315

7.1 Introduction

7.1.1 Fractional Factorial Experiments

As explained in Chapter 1, the purpose of a factorial experiment is to investigate the average contribution that each treatment factor makes to the response (called the *factor main effect*) as well as the nonadditive (joint) effect of the factors on the response (called the *factor interactions*). In this chapter, we label the response variable as y and the k factors as F_1, F_2, F_3, \dots (except in Section 7.10 where we use A, B, C, \dots for ease of exposition), and we look at unblocked experiments.

If the i th factor, F_i , is observed at ℓ_i levels, then the total number of treatment combinations (combinations of factor levels) is $t = \prod_{i=1}^k \ell_i$. In experiments with numerous factors, t will be very large even when each factor has only two levels. For example, an experiment with 10 two-level factors has a total of $t = 2^{10} = 1024$ treatment combinations. If all of these were run sequentially in an experiment and each one required only 10 min to set up and obtain an observation, the entire experiment would take over 170 hours. Many experiments are run in industrial settings, where “time is money” and the allotted time frame for experimentation may be much less than this. The situation is worse if the time required for a single observation is considerably longer than 10 min! Consequently, it is frequently the case that only a fraction of the treatment combinations will be observed, resulting in a *fractional factorial experiment*. For similar reasons, fractions are usually run with factors at small numbers of levels (say, 2 or 3).

The use of fractional factorial experiments is widespread, especially for research and development in industry and engineering and in medical research and biotechnology (see, e.g., Schmidt and Launsby 1992; Dante et al. 2003; and Jaynes et al. 2012). Although a selection of a manageable number, n , of the t treatment combinations for the experiment could be made at random, a random choice is unlikely to lead to a “good” design; the main effects and important interaction estimates may turn out to be highly correlated, may have large variances, or may not even be estimable. Consequently, a huge amount of research has been devoted to the selection (construction) of good fractions (e.g., Draper and Mitchell 1967; Bailey 1977; Fries and Hunter 1980; Franklin 1984; Chen et al. 1993; Chen and Cheng 2004; and Xu 2005, among others).

The early work on construction of good fractions was mostly for the case of two-level factors and focused on half, quarter, eighth, and so on fractions, resulting in 2^{k-1} , 2^{k-2} , 2^{k-3} , \dots treatment combinations to be observed in the experiment. These are referred to as 2^{k-q} fractions, where $q = 1, 2, 3, \dots, k - 1$. In general, if every factor has s levels, then s^{k-q} fractions have been sought instead. If factors have nonprime numbers of levels, they can be replaced by *pseudofactors*; for example, the levels of a six-level factor F_i can be represented by a combination of levels of factors $F_i^{(1)}$ and $F_i^{(2)}$ with two and three levels, respectively. An experiment is called *symmetric* if all its factors have the same number of levels, and *asymmetric* otherwise.

In this chapter, we give special attention to recent work on the construction of the most useful 2^{k-q} fractions. We restrict our discussion to *regular fractions* that are governed by *defining relations* and that allow main effects and interactions either to be estimated independently of each other or to be completely *aliased* (indistinguishable). In Chapter 9, *nonregular fractions* will be discussed, where partial aliasing occurs and where the choice of the number of observations is more flexible.

In planning an experiment, it is often convenient to code the levels of each factor. If the factor is quantitative with two levels, then we can think of the two levels as being *low* and *high* where the low level can be defined as either the smaller level of the factor (e.g., 70°C as opposed to 80°C) or the level that is likely to give the smaller response. There are five common codings in the literature for the (low, high) levels; these are (1, 2), (0, 1), (−1, 1), (−, +), and (⊖, *a*). The first two are self-explanatory, depending on whether the levels are numbered starting from 1 or starting from 0. These codings also extend naturally to more than two levels. The use of the third coding will become apparent as we discuss main effect and interaction contrasts; the fourth is an abbreviation of the third. The last notation is now falling into disuse; it records which factors are at their high levels. If there are two factors at two levels each, we can record the four treatment combinations in one of the following ways.

Notation 1	11	12	21	22
Notation 2	00	01	10	11
Notation 3	(−1, −1)	(−1, 1)	(1, −1)	(1, 1)
Notation 4	−−	−+	+−	++
Notation 5	(1)	<i>b</i>	<i>a</i>	<i>ab</i>

We mostly use notations 2 and 3 in this chapter.

In Section 7.2, we present two ways to write the model for a factorial experiment and discuss criteria for good fractional factorial designs, such as minimum aberration that was introduced in Chapter 1. Methods for constructing and characterizing good regular 2^{k-q} fractions are given in Sections 7.3 through 7.6. Analysis issues are addressed in Section 7.7, including analysis of saturated models and count data. The topic of screening to find the most important factors is dealt with in Section 7.8.

Follow-up experiments are often run to untangle effects that are indistinguishable in the original fractional factorial experiment; suggestions for designing follow-up experiments are given in Section 7.9. Section 7.10 discusses some construction methods for designs when factors have more than two levels. Recent work on discovering which fractions are *equivalent* to each other is summarized in Section 7.11; these results are useful in compiling catalogs of designs and searching for the best fractions. Section 7.12 deals with incorporation of *noise factors* into the experiment in order to design a product or process that is *robust*. Finally, some other issues such as blocked designs, and minimizing the number of level changes are mentioned briefly in Section 7.13.

7.2 Regular Fractions and Models

7.2.1 Models for Factorial Experiments

Using terminology similar to that of Chapter 3, a general model for a *full* (not fractional) factorial experiment with one observation on each of the t treatment combinations is

$$y = \tau + \epsilon, \quad \epsilon \sim N(0, I\sigma^2), \tag{7.1}$$

where τ is a vector of parameters that represent the mean responses at the t treatment combinations, and y and ϵ are the corresponding vectors of response and error variables; matrix A_d of Chapter 3 becomes the identity matrix and there are no blocks. As discussed in Chapter 1, the functions $c'\tau$ of τ that are of interest are the main effect and interaction contrasts. If factor F_i has two levels, its main effect, γ_i , is a contrast (comparison) between the mean responses when the factor is at its high and low levels, while the interaction, γ_{ij} , of the two-level factors i and j is a contrast between the effect of the i th factor when the j th factor is at its two different levels (or vice versa). Higher-order interactions are defined in a similar way.

It was shown in Chapter 1 how to obtain the contrast vector c' for the main effect and interaction contrasts $c'\tau$. For example, in a 2^3 experiment with three two-level factors and treatment combinations coded as 000, 001, . . . , 111 (in lexicographical order), the main effect and interaction contrasts, written as functions of τ , are shown as the rows of (7.2):

$$y = \begin{bmatrix} \mu^* \\ \gamma_1 \\ \gamma_2 \\ \gamma_{12} \\ \gamma_3 \\ \gamma_{13} \\ \gamma_{23} \\ \gamma_{123} \end{bmatrix} = \frac{1}{4} \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ -1 & -1 & -1 & -1 & 1 & 1 & 1 & 1 \\ -1 & -1 & 1 & 1 & -1 & -1 & 1 & 1 \\ 1 & 1 & -1 & -1 & -1 & -1 & 1 & 1 \\ -1 & 1 & -1 & 1 & -1 & 1 & -1 & 1 \\ 1 & -1 & 1 & -1 & -1 & 1 & -1 & 1 \\ 1 & -1 & -1 & 1 & 1 & -1 & -1 & 1 \\ -1 & 1 & 1 & -1 & 1 & -1 & -1 & 1 \end{bmatrix} \begin{bmatrix} \tau_{000} \\ \tau_{001} \\ \tau_{010} \\ \tau_{011} \\ \tau_{100} \\ \tau_{101} \\ \tau_{110} \\ \tau_{111} \end{bmatrix} = \frac{1}{n/2} X'\tau. \tag{7.2}$$

It can be verified that each of the contrast vectors (rows of X') in (7.2) can be written as

$$c' = \frac{1}{n/2} c'_1 \otimes c'_2 \otimes c'_3,$$

where \otimes denotes Kronecker product, $c'_i = [-1 \ 1]$ when the main effect or interaction involves factor i , and $c'_i = [1 \ 1]$ when it does not. Since $X'X = nI$ where n is the number of observations and X is square and full rank, it follows that $X' = nX^{-1}$. Therefore, $\tau = (n/2)(X')^{-1}y = (1/2)Xy$. If we write $\beta = (1/2)y$, so that the parameters in vector β represent the factorial effects divided by 2, then model (7.1) can be expressed as a usual regression model,

$$y = X\beta + \epsilon, \quad \epsilon \sim N(0, I\sigma^2). \tag{7.3}$$

For the 2^3 experiment, this is,

$$Y_{ijk} = \mu + \beta_1x_1 + \beta_2x_2 + \beta_3x_3 + \beta_{12}x_1x_2 + \beta_{13}x_1x_3 + \beta_{23}x_2x_3 + \beta_{123}x_1x_2x_3 + \epsilon_{ijk},$$

where $x_i = -1$ when factor i is at its low level and $x_i = +1$ when factor i is at its high level ($i = 1, 2, 3$). We call X the *model matrix*.

If a factor has more than two levels, its main effect has more than one degree of freedom (df). In this case, it is usual to select a set of orthonormal contrasts (i.e., the contrasts are orthogonal and normalized to satisfy $c'c = 1.0$). For example, if factor 1 has three levels

and the main effect or interaction involves factor 1, then c'_1 in the Kronecker product may be taken to have two orthonormal rows, such as

$$c'_1 = \begin{bmatrix} -1/\sqrt{2} & 0 & 1/\sqrt{2} \\ 1/\sqrt{6} & -2/\sqrt{6} & 1/\sqrt{6} \end{bmatrix}. \tag{7.4}$$

Any two orthogonal contrasts could be chosen, but the two illustrated here tend to be used for quantitative factors with equally spaced levels since they model the linear and quadratic trends in the response as the level of the factor increases (see, e.g., Wu and Hamada 2009, Section 1.8). If factor i has ℓ_i levels, then c'_i would have $\ell_i - 1$ orthonormal rows. As stated earlier, the coefficients in c'_i could be used in the model matrix X of a regression model.

We end this section with some notation that will be needed later in this chapter. Throughout, we consider experiments with k factors and n runs, where a run of the experiment consists of observing one treatment combination. The least squares estimate of β is $\hat{\beta} = (X'X)^{-1}X'y$ and has variance-covariance matrix $(X'X)^{-1}\sigma^2$ (see Chapter 2). The inverse of this matrix, $(X'X)\sigma^{-2}$, is called the *information matrix* for estimating β , and when comparing designs, it is common to set $\sigma^2 = 1.0$ without loss of generality. For two-level designs, the main effect columns from the model matrix X form an $n \times k$ matrix D , which we refer to as the *design matrix*, since its entries x_{ij} define the design d ; that is, x_{ij} defines whether the j th factor is to be set at its high or low level in the i th run of the experiment, according to its value $+1$ or -1 . Because of the correspondence between D and the main effect columns of X for two-level factors, we talk about a column as representing either a factor or its main effect, interchangeably. For simplicity, sometimes the design matrix, D , is itself called the design. Lastly, if we partition the columns of the design matrix, D , into $D = [D_1 D_2]$, then the design D_1 is called a *projection* of D onto the factors included in D_1 .

7.2.2 Aliasing of Factorial Effects

Table 1.17 shows an example of a design with three two-level factors in which only treatment combinations 000, 011, 101, and 110 are to be observed in the experiment. This is a 2^{3-1} half fraction with $k = 3$ factors and $n = 2^2 = 4$ runs. The design matrix is

$$D = \begin{bmatrix} -1 & -1 & -1 \\ -1 & 1 & 1 \\ 1 & -1 & 1 \\ 1 & 1 & -1 \end{bmatrix}, \tag{7.5}$$

and the regression model (7.3) is

$$\begin{bmatrix} y_{000} \\ y_{011} \\ y_{101} \\ y_{110} \end{bmatrix} = \begin{bmatrix} 1 & -1 & -1 & -1 & 1 & 1 & 1 & -1 \\ 1 & -1 & 1 & 1 & -1 & -1 & 1 & -1 \\ 1 & 1 & -1 & 1 & -1 & 1 & -1 & -1 \\ 1 & 1 & 1 & -1 & 1 & -1 & -1 & -1 \end{bmatrix} \begin{bmatrix} \mu \\ \beta_1 \\ \beta_2 \\ \beta_3 \\ \beta_{12} \\ \beta_{13} \\ \beta_{23} \\ \beta_{123} \end{bmatrix} + \epsilon = [X_1 \ X_2] \begin{bmatrix} \beta_1 \\ \beta_2 \end{bmatrix} + \epsilon, \tag{7.6}$$

where $\beta_1' = [\mu \ \beta_1 \ \beta_2 \ \beta_3]'$ and $\beta_2 = [\beta_{12} \ \beta_{13} \ \beta_{23} \ \beta_{123}]'$ and correspondingly X_1 contains the first four columns of X and X_2 contains the last four columns. Since there are only four observations, the model matrix cannot have more than four columns and still have full column rank. For example, in this particular design, we can fit the *main effects model*

$$y = X_1\beta_1 + \epsilon$$

and estimate β_1 via its least squares estimator $\hat{\beta}_1 = (X_1'X_1)^{-1}X_1'y$. However, from (7.6), the expected value of y is $E[y] = [X_1\beta_1 + X_2\beta_2]$, and so we have

$$\begin{aligned} E[\hat{\beta}_1] &= \beta_1 + (X_1'X_1)^{-1}X_1'X_2\beta_2 \\ &= \beta_1 + \frac{1}{4} \begin{bmatrix} 0 & 0 & 0 & -4 \\ 0 & 0 & -4 & 0 \\ 0 & -4 & 0 & 0 \\ -4 & 0 & 0 & 0 \end{bmatrix} \beta_2 = \frac{1}{2} \begin{bmatrix} \mu - \beta_{123} \\ \beta_1 - \beta_{23} \\ \beta_2 - \beta_{13} \\ \beta_3 - \beta_{12} \end{bmatrix} = \frac{1}{2} \begin{bmatrix} \mu^* - \gamma_{123} \\ \gamma_1 - \gamma_{23} \\ \gamma_2 - \gamma_{13} \\ \gamma_3 - \gamma_{12} \end{bmatrix}. \end{aligned} \quad (7.7)$$

Strictly speaking, we cannot estimate any single element of β based just on the half fraction (7.5); we can estimate only certain linear combinations of the elements of β (see also Chapter 1, Section 1.7.1). In practice, it is common to take $\hat{\beta}_1$ as an estimate for β_1 , contingent on the assumption that $\beta_2 = 0$. Obviously, such an assumption may be false; if so, one's estimate of β_1 will be biased due to fitting a model that omits terms that are not in fact zero. The matrix

$$A = (X_1'X_1)^{-1}X_1'X_2 = (1/n)X_1'X_2$$

is called the *alias matrix* and is calculated for the current example in (7.7); it is useful for showing how omitted terms may bias estimates for β_1 .

The alias matrix is useful not just in the context of main effects models but also for *any* models of the form $y = X_1\beta_1 + \epsilon$ in which β_1 contains the mean and up to $n - 1$ estimable factorial effects and β_2 contains the remaining factorial effects. For regular 2^{k-q} fractions, the alias matrix consists of 1's, -1's, and 0's. A column of X_2 (corresponding to β_2) that is identical to a column in X_1 yields a 1 in $X_1'X_2/n$, and the two corresponding parameters are said to be *aliased*. For our main effects model example in (7.7), each row of the alias matrix contains one -1 and the rest zeros. For instance, the element of the first column and last row of the alias matrix is -1; this is because the column associated with β_{12} , the first column of X_2 , is the negative of the last column from X_1 , which is associated with β_3 . Thus, $-\beta_{12}$ is aliased with β_3 and, equivalently, $-\gamma_{12}$ is aliased with γ_3 . (We say that the interaction F_1F_2 is aliased with F_3 .) The other entries in the last row of the alias matrix are zero because the remaining columns of X_2 are orthogonal to the column for β_3 .

For our example, only two parameters appear in each linear combination in (7.7) because D is a half fraction. For regular 2^{k-q} designs, when all columns corresponding to the full factorial model are included in $[X_1 \ X_2]$, the 2^{-q} fraction leads to $n = 2^{k-q}$ alias sets of size 2^q . The first of these sets, the one associated with μ , is discussed first. The *defining*

relation for a regular 2^{k-q} design contains the $2^q - 1$ factorial effects whose columns of X are constant for the n treatment combinations in the fraction. In (7.6), only one column besides the first is constant, that being the last column corresponding to the three-factor interaction $F_1F_2F_3$. The defining relation for this fraction is written as $I = -123$. Here and later, I denotes a column of 1's corresponding to μ , any bold numeral denotes the column of X corresponding to the main effect of that factor, and a string of bold numbers indicates an interaction column from X . In the next two sections, we discuss examples where $q = 4$ (i.e., $2^{-4} = 1/16$ th fractions). There, the defining relation consists of 15 interactions in addition to the intercept.

From X in (7.6), each main effect contrast is the negative of the other two factors' interaction contrast. Thus, we write $1 = -23$, $2 = -13$, and $3 = -12$. These are referred to as alias chains or alias strings. They correspond to the linear combinations of parameters in (7.7). In general, a unique least squares estimate of β_1 can only be obtained if the effects in β_1 are from different alias chains.

Selection of a (regular) fraction of treatment combinations to be observed in an experiment is an art and forms the topic of Sections 7.3 through 7.6. Selection of good designs for situations where regular fractions do not exist or are not competitive with nonregular fractions is discussed in Chapter 9.

7.2.3 Defining Contrast Subgroup and Construction of Designs

The interactions that appear in the defining relation for a 2^{k-q} regular fraction are referred to as *words*. Together with the *identity word* I , this collection of interactions or words is labeled the *defining contrast subgroup*, since they form a subgroup of all possible words. This means that for any pair of interaction contrasts in the defining contrast subgroup, their elementwise product is also in the defining contrast subgroup.

For example, let us examine the defining relation for a 2^{9-4} design that we label as d_1 in this section and the next. Design d_1 is defined by the $q = 4$ generating words **1236**, **1247**, **1258**, and **13459**. The product of the first two generating words is

$$1236 \cdot 1247 = 1^22^23467 = 3467,$$

where “.” means the elementwise product of vectors (columns of X). The result follows from the fact that any column multiplied by itself is a vector of all +1's and so disappears from the product; **3467** is called the *generalized interaction* of **1236** and **1247**.

The entire defining contrast subgroup of d_1 of size $2^q = 16$ is

$$\begin{array}{cccccccc}
 I & 1236 & 1247 & 3467 & 1258 & 3568 & 4578 & 12345678 \\
 13459 & 24569 & 23579 & 15679 & 23489 & 14689 & 13789 & 26789.
 \end{array} \tag{7.8}$$

The second half of the first row is obtained by multiplying the generating word **1258** by the first half of the row. Similarly, the second row is obtained by multiplying the first row times the last generating word, **13459**.

The $q = 4$ generating words can also be expressed as alias chains as follows:

$$6 = 123, \quad 7 = 124, \quad 8 = 125, \quad 9 = 1345. \tag{7.9}$$

Note that the factors that appear on the left-hand sides in (7.9) do not appear in any of the interactions on the right-hand sides. When expressing the generators for a design in this way, the factors on the left are called *generated factors* and the others are called *basic factors*. We will also refer to the interactions **123**, ..., **1345** as *factor generators*. It is customary to take the first $k - q$ factors as basic factors, but there are other sets of $k - q$ generating words, basic factors, and factor generators that would result in the identical defining relation and defining contrast subgroup. For instance, in this example, the same defining relation could also be obtained by taking the last five factors as basic factors and using the alias chains **1 = 5679**, **2 = 6789**, **3 = 568**, and **4 = 578** to generate the first four factors. Whichever representation is chosen for a regular fraction, the defining relation and the alias scheme can be recovered.

The design matrix D for any regular 2^{k-q} design can easily be constructed by creating a full factorial design in $k - q$ basic factors and then adding the q additional columns to complete D based on interactions of the basic factors. We can also use the alias chain representation to create a series of fractions. For example, if we start with a full 2^5 factorial design for factors F_1, \dots, F_5 and add just the two generated factors F_6 and F_7 in (7.9) to the five basic factors, we have a 2^{7-2} design and its defining contrast subgroup consists of just the first four terms in (7.8). If we added the third generator, we would have a 1/8th fraction, and the defining contrast subgroup would consist of the first line in (7.8). Adding the fourth generated factor results in a 1/16th fraction, that is, a 2^{9-4} design, and so the defining contrast subgroup is then the entire (7.8). Each time a generator is added, a new word is multiplied by the previous defining contrast subgroup, which doubles its size. Note also that each factor that appears in the defining contrast subgroup appears in exactly half of the words.

Besides the defining contrast subgroup, there are $n - 1$ alias sets of size 2^q . Multiplying the defining contrast subgroup by any factorial effect not in that subgroup identifies a set of 2^q aliased effects. For example, the first four aliases of **12** are **12·1236 = 36**, **12·1247 = 47**, **12·3467 = 123467**, and **12·1258 = 58**. This set with **12** and its 15 aliases is a *coset* of (7.8). Often, the longer aliases (which represent the higher-order interactions) are just ignored, as we will do shortly in Table 7.1.

Rather than constructing a 2^{k-q} design by creating a full factorial in a set of basic factors augmented with q generated factors, the treatment combinations to be observed in the design d can be constructed using modulo 2 arithmetic. Define x_i to be the level (0 or 1) of the i th factor in a particular run, and let a_{ij} be 1.0 if the j th generator of the defining relation contains factor F_i and let a_{ij} be 0.0 otherwise. The design consists of all of those treatment

TABLE 7.1
Two-Factor Interaction Alias Chains for d_1 and d_2 , Assuming That Higher-Order Interactions Are Negligible

d_1	d_2
12 = 36 = 47 = 58	12 = 36 = 47
13 = 26 34 = 67	13 = 26 = 48
14 = 27 35 = 68	14 = 27 = 38
15 = 28 36 = 47	16 = 23 = 78
16 = 23 38 = 56	17 = 24 = 68
17 = 24 45 = 78	18 = 34 = 67
18 = 25 48 = 57	28 = 37 = 46

combinations $(x_1x_2 \dots x_k)$ that satisfy $\sum_{i=1}^k a_{ij}x_i = 0$ modulo 2 for $j = 1, \dots, q$. Fractions obtained in this way always contain the *zero* treatment combination $(000 \dots 0)$. Thus, for example, the 2^{3-1} fraction that has defining contrast subgroup $\{\mathbf{I}, -\mathbf{123}\}$ is obtained from the equation

$$x_1 + x_2 + x_3 = 0, \text{ modulo } 2.$$

and results in the fraction $d = (000, 011, 101, 110)$. To obtain an alternative fraction with the same defining contrast subgroup (up to sign changes), one can *add* any treatment combination not in d , where addition means addition of corresponding digits and reducing modulo 2. For example, addition of 100 to d results in the fraction $d^* = (100, 111, 001, 010)$. The primary advantage of this method is that it readily extends to the construction of s^{k-q} fractions using modulo s (see Section 7.10).

7.2.4 Criteria for Good Designs: Resolution and Aberration

As shown in Section 7.2.2, unique estimability of main effects and interaction effects in a regular fraction can only be achieved when at most one effect from an alias chain is included in the model (and other effects in the same chain are omitted from the model). Consequently, interpretation of the analysis requires either that certain effects are assumed negligible or that a follow-up experiment be performed to disentangle the aliased effects (Section 7.9). A widely held belief, supported in general by investigations of real experiments (see the study of Lin et al. 2006), is that it is much more likely for high-order effects (interactions between large numbers of factors) to be smaller than low-order effects such as main effects and two-factor interactions. Consequently, it is best if the defining contrast subgroup contains only high-order interactions. As defined in Chapter 1, the *resolution* of the design is the minimum number of factors involved in any of the interactions in the defining contrast subgroup, and a goal in designing the experiment is to choose a fraction with highest possible resolution.

There are often several designs of size n with maximum resolution, which may still differ in terms of usefulness for the experiment. For example, let us look at the following defining relations for two resolution IV 2^{9-4} designs (denoted by 2_{IV}^{9-4}) examined by Chen et al. (1993):

$$d_1 : \mathbf{I} = \mathbf{1236} = \mathbf{1247} = \mathbf{3467} = \mathbf{1258} = \mathbf{3568} = \mathbf{4578} = \dots \text{ and nine longer words,}$$

$$d_2 : \mathbf{I} = \mathbf{1236} = \mathbf{1247} = \mathbf{3467} = \mathbf{1348} = \mathbf{2468} = \mathbf{2378} = \mathbf{1678} = \dots \text{ and eight longer words.}$$

We have already considered d_1 in Section 7.2.3; its complete defining contrast subgroup is given in (7.8). Design d_2 has the generating words **1236**, **1247**, **1348**, and **23459**. Although both d_1 and d_2 have resolution IV, they differ in their alias properties as can be seen from the aliases among the two-factor interactions that are shown in Table 7.1.

A natural criterion for a good design is to minimize the numbers of short words (i.e., low-order interactions) in the defining relation among those for all fractions of the same size and maximum resolution (Fries and Hunter 1980). Let A_j denote the number of length- j words in a design's defining relation and define the word length pattern (wlp) for that design as

$$\text{wlp} = (A_3, A_4, \dots, A_k). \tag{7.10}$$

Referring to (7.8), in addition to the six length-4 words, there are eight length-5 words and one length-8 word, so design d_1 has $wlp_1 = (0,6,8,0,0,1,0)$. It may be verified that design d_2 has $wlp_2 = (0,7,7,0,0,0,1)$. The first nonzero entry of the wlp is A_r , where r is the resolution of the design. Design d_1 has smaller A_4 than design d_2 , and it is said to have *lower aberration*. In fact, d_1 is the *minimum aberration* 2^{9-4} design and d_2 is second best.

For some 2^{k-q} cases, two or more of the maximum resolution designs may be tied on the minimum value of A_r . The wlp's are then compared on subsequent entries until one design wins (i.e., has a smaller entry) and is declared to be the minimum aberration design; on rare occasions (e.g., $n = 128, k = 41-44$ and 50), there are two or more nonequivalent minimum aberration designs. All other designs tied on the first nonzero entry of the wlp receive the consolation prize of having *weak minimum aberration* (Chen and Hedayat 1996).

Interestingly, although the minimum aberration design d_1 mentioned above has fewer words than d_2 of length 4, from Table 7.1, it can be seen that it has only eight two-factor interactions that are not aliased with other two-factor interactions (all $F_iF_9, i = 1, \dots, 8$), whereas d_2 has 15 two-factor interactions (all F_iF_9 and all F_iF_5) not aliased with other two-factor interactions. Since the aliasing structure of designs of the same resolution can be very different and the minimum aberration design may or may not be the design with the properties most preferred by the experimenter, other criteria have also been found useful for comparing and selecting the best resolution IV designs. These are discussed in Section 7.4.3.

In Sections 7.3 through 7.6, we look at the construction, ranking, and equivalence of two-level fractions. Although the goal is usually to obtain designs with the highest possible resolution, we start by examining designs with resolution III, followed by resolution IV and then resolution V designs. Designs with factors having more than two levels will be discussed later, in Section 7.10.

7.3 Resolution III Designs

Resolution III fractions for two-level factors exist for number of factors $k \leq n - 1$. When $k = n - 1$, the design is said to be *saturated* and its design matrix is very easily constructed recursively. Define $S_1 = 1$ and

$$S_n = \begin{bmatrix} S_{n/2} & S_{n/2} \\ S_{n/2} & -S_{n/2} \end{bmatrix} \tag{7.11}$$

for $n = 2, 4, 8, \dots$ (i.e., for n any power of 2). For example,

$$S_8 = \begin{bmatrix} S_4 & S_4 \\ S_4 & -S_4 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 & 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 & 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 & 1 & -1 & -1 & 1 \\ 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 \\ 1 & -1 & 1 & -1 & -1 & 1 & -1 & 1 \\ 1 & 1 & -1 & -1 & -1 & -1 & 1 & 1 \\ 1 & -1 & -1 & 1 & -1 & 1 & 1 & -1 \end{bmatrix}.$$

Each S_n in this series is a square matrix of -1 's and $+1$'s with orthogonal columns, which implies $S_n' S_n = nI_n$. The matrix in (7.11) is the Sylvester-type Hadamard matrix; see Hedayat et al. (1999, Chapter 7) for additional details. After removing the first column of S_n (which we will call column 0, or c_0 for short), columns 1, 2, 4, 8, \dots , $n/2$ can be associated with basic factors since they form a full factorial. Labeling the basic factors as F_1, F_2, \dots, F_{n-q} , we thus associate the main effect of F_1 with c_1 (i.e., column 1), F_2 with c_2 , F_3 with c_4 , F_4 with c_8 , etc. The q generated factors will be defined by the other columns of S_n , which correspond to interactions of the basic factors; for example, if we associate F_5 with column c_3 , then $5 = 12$. Omitting c_0 from S_8 , one obtains a 2^{7-4} fraction. We note, in passing, that X' in (7.2) is S_8' with F_3, F_2 , and F_1 associated with columns c_1, c_2 , and c_4 , respectively, and low and high levels interchanged.

All regular, saturated two-level fractions of size n are *isomorphic*—two designs are isomorphic if one can be obtained from the other by interchanging rows (treatment combinations), interchanging columns (factor labels), and reversing the levels for one or more factors (see Section 7.11). Fractions with $k = n - 2$ factors are also unique up to isomorphism, as are those with $k = n - 3$ factors (Mukerjee and Wu (2006), p. 59). When $k \leq n - 4$, two 2^{k-q} designs of the same size may not be isomorphic. However, all regular 2^{k-q} designs with n rows are projections of a saturated n -run design. That is, one can form the design with fewer factors by taking a subset of the columns of S_n .

The primary criterion for ranking resolution III designs is the minimization of A_3 , the number of length-3 words in the defining relation. Each word of length-3 in the defining relation produces three aliases between main effects and two-factor interactions. (For instance, the product of c_1, c_2 , and c_3 in S_8 is equal to I , so the main effect of each of the three factors associated with these columns is aliased with the interaction of the other two.) Thus, $3A_3$ equals the number of two-factor interactions that are aliased with main effects; any design that minimizes A_3 necessarily minimizes the number of two-factor interactions that might bias the main effect estimates. For the saturated design where $k = n - 1$, all $k(k - 1)/2$ two-factor interactions are aliased with main effects and so $A_3 = (n - 1)(n - 2)/6$. If we delete any column from the saturated design S_n , then A_3 decreases to $(n - 2)(n - 4)/6$; drop another column and the resulting design with $k = n - 3$ will have $A_3 = (n - 4)(n - 5)/6$ (Angelopoulos et al. (2008)). Each time a factor is removed, A_3 declines, provided the factor being deleted appeared in one or more length-3 words in the defining relation. Thus, the smaller the number of factors for a given number of runs, n , the less risk of bias to main effects estimates. For $k \leq n/2$, it is possible to eliminate all aliasing between main effects and two-factor interactions; such designs are the subject of Section 7.4.

When $k \leq n - 4$, how does one obtain a 2^{k-q} design that minimizes A_3 ? One means is to utilize software, such as SAS Proc Factex, which uses a depth-first search for selecting generators. Alternatively, one may employ existing tables of designs. For instance, Mee (2009, Appendix G) lists all minimum aberration designs of sizes 8, 16, 32, 64, and 128. Here and in similar tables, the interactions to be used as generator columns are identified by their *Yates order* number (Yates 1937). Yates introduced factorial effects in the systematic order $\gamma_1, \gamma_2, \gamma_{12}, \gamma_3, \gamma_{13}, \gamma_{23}, \gamma_{123}, \gamma_4, \dots$, which puts main effects of individual factors in places 1, 2, 4, 8, and so on, in increasing powers of 2. Furthermore, each interaction column number can be decomposed as the sum of powers of 2. For instance, $7 = 1 + 2 + 4$, implying that the interaction contrast γ_{123} identified as Yates column 7 (c_7) is the elementwise product of c_1, c_2 , and c_4 , corresponding to the main effects of factors F_1, F_2 , and F_3 . Note that, after eliminating c_0 , the columns of S_n correspond to Yates order.

We illustrate the use of such tables to obtain the 16-run, minimum aberration 2^{9-5} design; Mee (2009, p. 488) lists columns 7 and 11–14 for generating this design. We need nine columns in total from S_{16} , so we take the four basic columns c1, c2, c4, and c8 for $F_1, F_2, F_3,$ and F_4 , plus the five generator columns, c7, c11, c12, c13, and c14, for $F_5, F_6, F_7, F_8,$ and F_9 , respectively. For instance, since $12 = 4 + 8$, with c4 and c8 corresponding to F_3 and F_4 , Yates column c12 corresponds to the γ_{34} interaction. Thus, $7 = 3 \cdot 4$ and so **347** is a word in the defining relation. Columns c7, c11, c13, and c14 correspond to three-factor interactions; these are used to create the other four generated factors $F_5, F_6, F_8,$ and F_9 so that the generating words **1235, 1246, 347, 1348,** and **2349** and their products form the defining contrast subgroup. The defining relation of this design includes just four length-3 words: **178, 279, 347,** and **567**. From these, it can be seen that while the main effect of factor F_7 is aliased with four two-factor interactions, each of the other main effects is aliased with just a single two-factor interaction, and this interaction involves F_7 . Thus, to bias as few main effects as possible, the factor deemed least likely to have active two-factor interactions should be the one labeled F_7 and assigned to column c12.

Tang and Wu (1996) defined the column complement of a design to be the contrast columns that are omitted from S_n . For instance, for $n = 16$ runs and $k = 9$ factors, the minimum aberration 2^{9-5} design {c1, c2, c4, c7, c8, c11, c12, c13, c14} has the complement consisting of the omitted contrast columns {c3, c5, c6, c9, c10, c15}. Let \bar{A}_j denote the number of length- j words in the defining relation for the complementary design with $n - k - 1$ factors, just as A_j denotes the number of length- j words for the 2^{k-q} design. Chen and Hedayat (1996) and Tang and Wu (1996) both showed that for a given n and k ,

$$A_3 = C_3(n, k) - \bar{A}_3,$$

$$A_4 = C_4(n, k) + \bar{A}_3 + \bar{A}_4,$$

where $C_3(n, k)$ and $C_4(n, k)$ are constants that depend only on n and k . Tang and Wu further showed that the minimum aberration 2^{k-q} design may be found by sequentially maximizing \bar{A}_3 , minimizing \bar{A}_4 , maximizing \bar{A}_5 , etc. For $n = 16$ and $k = 9$, $C_3(16, 9) = 8$, so $A_3 = 8 - \bar{A}_3$. The minimum aberration 2^{9-5} design is obtained if we omit six columns from S_{16} that form a replicated 2^{6-3} design; this complement has $\bar{A}_3 = 4$, so $A_3 = 4$ also. In general, the minimum aberration design is obtained by deleting $n - k - 1$ columns from S_n that together have the maximum aliasing of two-factor interactions with main effects.

In some instances, there exist several resolution III designs with the same A_3 as the minimum aberration design. These weak minimum aberration designs minimize the number of two-factor interactions aliased with main effects and so should be considered as desirable designs along with the minimum aberration design.

7.4 Resolution IV Designs

To appreciate the difference between the designs discussed here and those in Sections 7.3 and 7.5, consider the structure of the information matrix $X'X$. Let X_1 denote the $n \times (k + 1)$

model matrix for a main effects model, and let X_2 denote the $n \times k(k - 1)/2$ matrix of two-factor interaction contrasts. Then $X = [X_1, X_2]$ is the model matrix for the *two-factor interaction model*. For designs in the previous section, there is aliasing between main effects and two-factor interactions resulting from entries of $\pm n$ appearing in $X_1'X_2$ (Section 7.2.2). For designs in this and the next section, such aliasing is avoided by requiring the intercept and main effect columns to be orthogonal to two-factor interaction contrasts (i.e., $X_1'X_2 = \mathbf{0}$). Ignoring three-factor and higher-order interactions, the only aliasing is among two-factor interactions for resolution IV designs. In order to provide this *clear* estimation for main effects, for a given k , designs here will require more runs than those in Section 7.3. (The resolution V designs in Section 7.5 will be larger still.) If one expects only a few two-factor interactions, then a resolution IV design will be suitable. For example, suppose one has seven factors. If one thought all interactions were negligible, one might use either the eight-run 2^{7-4} design S_8 or seven columns from the 12-run Plackett–Burman design (see Chapter 9). Suppose instead that many two-factor interactions are expected to be important. To estimate the 7 main effects and all 21 two-factor interactions requires 64 runs for an orthogonal design (see Section 7.5). Such large run sizes are impractical for many applications. We now consider a compromise between the frugal resolution III designs and the large designs of the next section. Regular resolution IV designs for seven factors can be constructed with either 16 runs or 32 runs; although these designs do not permit estimation of the full two-factor interaction model, they do provide

$$\begin{bmatrix} X_1'X_1 & X_1'X_2 \\ X_2'X_1 & X_2'X_2 \end{bmatrix} = \begin{bmatrix} n\mathbf{I} & \mathbf{0} \\ \mathbf{0} & X_2'X_2 \end{bmatrix},$$

with $\text{rank}(X_2)$ of up to 7 and 18, respectively.

7.4.1 Constructing Even Resolution IV Designs

For n any power of 2 ($\geq 2^3$), the maximum number of factors for a resolution IV fraction is $k = n/2$. The unique (up to isomorphism) regular designs $2^{4-1}, 2^{8-4}, 2^{16-11}, 2^{32-26}$, etc. may be constructed by *foldover* as follows:

$$E_n = \begin{bmatrix} S_{n/2} \\ -S_{n/2} \end{bmatrix}, \tag{7.12}$$

where $S_{n/2}$ is the Sylvester-type Hadamard matrix defined in (7.11). These designs have only even-length words in their defining relations. We refer to fractions with only even words in the defining relation as *even designs*. The series of designs E_n , with $k = n/2$, are called *maximal even designs* because they maximize the possible number of factors k . Every even 2^{k-q} design with $k < n/2$ is a projection of E_n . Let $\ell = n/2 - k$. Butler (2003a) showed that a projection of the maximal design (7.12) with $k > 5n/16$ factors has minimum aberration if and only if the ℓ columns dropped from E_n form a design that has minimum aberration among all even designs with ℓ factors and n runs; see also Example 2 of Xu and Cheng (2008). If the omitted columns form a full factorial, the projection of E_n onto the remaining factors will obviously have minimum aberration.

The minimum aberration 2^{26-20} design having $n = 64$ runs can be constructed by creating a full 2^6 factorial and then adding the 20 interactions listed by Mee (2009, p. 490) as generators for the remaining factors. However, it is simpler to construct the minimum aberration 2^{26-20} design using the results of this section. First, construct E_{64} by folding over

S_{32} ; then, delete columns $c_1, c_2, c_4, c_8, c_{16},$ and c_0 . The remaining 26 columns form a minimum aberration 2^{26-20} design, since the six columns we deleted from E_{64} form a full 2^6 design. To obtain the minimum aberration 2^{25-19} , one must delete one more column such that the seven columns deleted from E_{64} form a resolution VI design; one choice is c_{31} , the five-factor interaction of $c_1, c_2, c_4, c_8,$ and c_{16} .

The maximal design E_n aliases the two-factor interactions in sets of size $k/2 = n/4$; every factor appears once in each set of aliased two-factor interactions. These designs have $A_4 = k(k - 1)(k - 2)/24$. Projections of these designs have shorter alias chains and smaller A_4 . If $k > n/4$, even designs have $n/2 - 1$ df for two-factor interactions and $n/2 - k$ df for error, assuming interactions involving three or more factors are negligible. For resolution IV designs with $k > 5n/16$, every two-factor interaction is aliased with other two-factor interactions (Chen and Hedayat 1998). The restriction $k > 5n/16$ also appeared in the minimum aberration result attributed to Butler (2003a) in the previous two paragraphs, since for $k = 5n/16$, resolution (\geq) IV designs with odd-length words exist. These designs are discussed in the next section. No even resolution IV design with $k \leq 5n/16$ has minimum aberration; for these cases, we must use different constructions or exhaustive searches for the best designs. (We will return to the idea of foldover in Section 7.9 when discussing follow-up designs.)

7.4.2 Constructing Resolution IV Designs with $k \leq 5n/16$

Only even resolution IV designs are projections of the maximal even designs introduced in the previous section. Here, we introduce resolution IV designs that are not even. In particular, we begin with the minimum aberration 2^{10-5} design, which may be constructed by augmenting a full 2^5 factorial with factor generators $6 = 1234, 7 = 125, 8 = 135, 9 = 145,$ and $10 = 2345$. Unlike the even designs described in the previous section, this design uses some four-factor interactions as generators, which results in odd-length words in the defining relation. This design is attractive for several reasons. First, this design has only 10 length-4 words in the defining relation versus $A_4 = 15$ for the best even 2^{10-5} design. These 10 length-4 words create the following aliases among two-factor interactions:

$$\begin{array}{cccc}
 1 \cdot 2 = 5 \cdot 7 & 1 \cdot 3 = 5 \cdot 8 & 1 \cdot 4 = 5 \cdot 9 & 1 \cdot 6 = 5 \cdot 10 \\
 1 \cdot 7 = 2 \cdot 5 & 1 \cdot 8 = 3 \cdot 5 & 1 \cdot 9 = 4 \cdot 5 & 1 \cdot 10 = 5 \cdot 6 \\
 2 \cdot 3 = 7 \cdot 8 & 2 \cdot 4 = 7 \cdot 9 & 2 \cdot 6 = 7 \cdot 10 & 2 \cdot 8 = 3 \cdot 7 \\
 2 \cdot 9 = 4 \cdot 7 & 2 \cdot 10 = 6 \cdot 7 & 3 \cdot 4 = 8 \cdot 9 & 3 \cdot 6 = 8 \cdot 10 \\
 3 \cdot 9 = 4 \cdot 8 & 3 \cdot 10 = 6 \cdot 8 & 4 \cdot 6 = 9 \cdot 10 & 4 \cdot 10 = 6 \cdot 9 \\
 1 \cdot 5 = 2 \cdot 7 = 3 \cdot 8 = 4 \cdot 9 = 6 \cdot 10.
 \end{array} \tag{7.13}$$

Each length-4 word in the defining relation creates aliasing between three pairs of two-factor interactions; there are 30 alias pairs in (7.13); these consist of the 20 separate alias pairs and an additional 10 pairs implied by the last chain of five aliases.

Second, besides having smaller A_4 than any even 2^{10-5} design, this design has more df for estimating two-factor interaction effects. The 45 two-factor interactions for the 10 factors are associated with 21 orthogonal contrasts, one for each alias chain in (7.13); that is, $\text{rank}(X_2) = 21$ for the minimum aberration design, whereas $\text{rank}(X_2) = 15$ for any

even 2^{10-5} design. This minimum aberration design is *second-order saturated* (SOS), since for its two-factor interaction model matrix $[X_1, X_2]$,

$$\text{rank}[X_1, X_2] = 1 + k + \text{rank}(X_2) = n;$$

that is, the two-factor interaction model utilizes all the df (Block and Mee 2003). Chen and Cheng (2004) refer to all resolution IV SOS designs as *maximal*, since no factor can be added to these designs without lowering their resolution; see their paper for connections with how *maximal* is used in projective geometry.

7.4.2.1 Key Results for Existence and Construction of Resolution IV SOS Designs

All resolution III designs with $k > n/2$ are SOS, while most resolution IV and resolution V designs are not SOS (Chen and Cheng 2004). However, every resolution IV design that is not SOS is the projection of a resolution IV SOS design. We now discuss both the existence and construction of resolution IV SOS designs by citing three key results.

Key result 1: As noted by Chen and Cheng (2006), there exist SOS designs with resolution IV (or higher) and of size $n (\geq 16)$ for $k = 1 + n/4$. When $n = 16$, this is the resolution V 2^{5-1} design; for $n = 32$, this is the 2^{9-4} design with $A_4 = 7$ identified as d_2 in Section 7.2.4; for $n = 64$, there are actually five nonisomorphic 2^{17-11} SOS designs. All regular SOS designs with $k = 1 + n/4$ may be constructed as follows. Let $E_{n/2} = [B \ C]$ denote a partition of the columns of $E_{n/2}$ into two nonempty sets. Then define

$$D_1 = \begin{bmatrix} \mathbf{1} & B & C \\ -\mathbf{1} & -B & C \end{bmatrix}. \tag{7.14}$$

The design D_1 constructed according to (7.14) has resolution IV (or higher). In addition, if B contains an odd number of columns, then D_1 must also be SOS (Cheng et al. 2008, Section 2.1). Consider two examples. First, starting with E_8 and taking any one or three columns for B , the resulting design is the resolution V 2^{5-1} fraction. Second, starting with E_{16} and taking a single column for B , one obtains the 2^{9-4} SOS design mentioned earlier.

Key result 2: From any SOS design D of resolution IV SOS (or higher), another resolution IV SOS design may be constructed with twice the number of runs and twice the number of factors via the construction

$$D_2 = \begin{bmatrix} D & D \\ D & -D \end{bmatrix}. \tag{7.15}$$

For instance, doubling the resolution V 2^{5-1} fraction in this way produces the SOS 2^{10-5} with the 21 alias sets shown earlier in (7.13). Repeated doubling produces a series of SOS designs with the same k/n ratio.

Key result 3: Every regular SOS design of resolution IV (or higher) with $k > n/4$ is produced either by (7.14) if k is odd or by (7.15) if k is even. This insight was first derived from a coding theory result in Bruen and Wehlau (1999).

7.4.2.2 Summary of Resolution IV SOS Designs with $k > n/4$

Combining these key results, we have the following (incomplete) summary regarding SOS designs with resolution IV or higher.

1. For $k = n/2$, we have the even SOS series $E_8, E_{16}, E_{32}, E_{64}, \dots$ given by (7.12).
2. For $k = 5n/16$, we have the SOS series with the maximum number of factors k for resolution IV or higher designs that are not even. The designs $2^{5-1}, 2^{10-5}, 2^{20-14}, \dots$ in this series all have minimum aberration.
3. For $k = 9n/32$, we have the SOS series $2^{9-4}, 2^{18-12}, \dots$. While these designs do not have minimum aberration, they have many degrees of freedom for short alias chains of two-factor interactions.
4. For $k = 17n/64$, the series begins with five nonisomorphic 2^{17-11} designs given by (7.14); these five \mathbf{D}_1 designs are constructed by taking different subsets of columns from E_{32} for \mathbf{B} . Doubling these leads to different doubled designs for $n = 128, 256$, etc.

Additional series begin with designs constructed according to (7.14) of sizes $n = 128$ (2^{33-26}), $n = 256$ (2^{65-57}), etc.

If there were no SOS designs with $k \leq n/4$, the search for the best designs would be simpler. However, Bruen and Wehlau's (1999) result does not address the existence of SOS designs with $k \leq n/4$. For $n \leq 32$, there are no SOS designs with $k \leq n/4$. However, Block (2003) identified a 64-run 2^{13-7} design that is SOS and 38 SOS designs of size 128 with k between 21 and 31 (including the minimum aberration designs for $k = 25$ and 29). At $n = 256$, Block's (naïve) search discovered over 34,000 nonisomorphic resolution IV SOS designs, one of which has just $k = 33$ factors. Thus, the search for the best resolution IV designs must entertain designs that are not projections of SOS designs constructed by (7.14) or (7.15).

Because SOS designs of resolution IV are known for $k > n/4$, minimum aberration designs are more simply characterized for this range. Complementary column theory facilitates the search for minimum aberration projections. For $k = n/2 - 1, n/2 - 2, \dots, 5n/16 + 1$, the minimum aberration designs are projections of E_n . Xu and Cheng (2008) developed a complementary column theory for projections of the $k = 5n/16$ SOS designs and proved that the best projections have minimum aberration among all designs for $17n/64 \leq k < 5n/16$. That is, the best projections from the $k/n = 5/16$ series have lower aberration than that of the SOS designs from the $9/32$ and $17/64$ series.

7.4.2.3 Minimum Aberration Resolution IV Designs with $k \leq n/4$

For up to $k = n/4$ factors, minimum aberration designs are found by computational searches. Start with all nonisomorphic designs of a given size; to each, add one generated factor from the set of all eligible generators, and keep only one version of designs that are isomorphic. This process is repeated again and again, so that all possible designs are obtained by adding factors sequentially. Chen et al. (1993) were the first to publish substantive results obtained in this manner. For the current state of the art in searching for larger two-level designs, see Xu (2009), who has enumerated all resolution IV designs of size 128 and all resolution IV designs of size 256 for $k \leq 17$; see also the lists in Tables 3, 5 and 6 of Lin and Sitter (2008).

We now turn our attention to criteria other than minimum aberration for defining “best” for resolution IV designs.

7.4.3 Other Criteria for Choosing a Resolution IV Design

For resolution III designs, the minimum aberration criterion was sufficient for defining the best designs. However, for resolution IV designs, the wlp is not sufficient for ranking designs because A_4 does not adequately characterize the two-factor interaction alias structure. Thus, besides minimum aberration, the following criteria involving the aliasing of two-factor interactions have been found useful for distinguishing the best designs:

- Maximize the number of two-factor interactions that are not aliased with any main effects or other two-factor interactions; such interactions are called *clear*.
- Maximize $\text{rank}(X_2)$, the df for two-factor interactions.
- Minimize the average length (or the maximum length) of the chains of aliased two-factor interactions.
- Maximize the design’s estimation capacity, which Cheng et al. (1999) defined as a vector of the proportions of models that can be estimated having 1, 2, 3, . . . two-factor interactions.

These criteria can all be expressed in terms of the alias length pattern, alp, for two-factor interactions. As in Block and Mee (2003), we define a_j to be the number of sets of aliased two-factor interactions of size j and

$$\text{alp} = (a_1, a_2, \dots, a_L),$$

where L is the longest chain of aliased two-factor interactions. For example, the minimum aberration 2^{10-5} design, whose aliasing was shown in (7.13), has $\text{alp} = (0, 20, 0, 0, 1)$. Linking the aforementioned criteria to the alp,

- a_1 is the number of clear two-factor interactions;
- $\text{Rank}(X_2) = a_1 + a_2 + \dots + a_L$;
- the maximum alias length is L ; the average length is $0.5k(k-1)/(a_1 + a_2 + \dots + a_L)$.

The estimation capacity vector of Cheng et al. (1999) can also be computed using the alp, although the formulae are not simple. Finally, $A_4 = \sum_{j=2}^L j(j-1)a_j/6$, so the weak minimum aberration criterion too is based on alp.

We conclude this section by considering the case of $k = 13$ factors to illustrate the need for criteria other than minimum aberration. The smallest resolution IV regular fraction is a 2^{13-8} design with 32 runs and $A_4 = 55$. However, this even design has only 15 df for two-factor interactions, with alias chains of length-5 and length-6. To learn much about interactions would require a larger design. The 64-run, minimum aberration 2^{13-7} design may be constructed using columns c21, c22, c31, c39, c41, c42, and c51 from S_{64} as generating columns (Mee 2009, p. 490). This design has $\text{wlp} = (0, 14, 28, 24, \dots)$ and $\text{alp} = (20, 18, 6, 1)$. Thus, by doubling the number of runs, we now have $\sum a_i = 20 + 18 + 6 + 1 = 45$ df for estimating the two-factor interactions, with 20 clear, and another 36 aliased in chains of length 2. A second resolution IV 2^{13-7} design worthy of consideration uses instead generating columns c7, c11, c13, c30, c46, c49, and c63 from S_{64} ; see design 13-7.2 from

Chen et al. (1993). Assigning these seven columns to the main effects of F_7, \dots, F_{13} , respectively, and the six basic columns to the main effects of F_1, \dots, F_6 , we obtain the following aliasing of two-factor interactions:

$$\begin{array}{ll}
 1 \cdot 2 = 3 \cdot 7 = 4 \cdot 8 & 1 \cdot 3 = 2 \cdot 7 = 4 \cdot 9 \\
 1 \cdot 4 = 2 \cdot 8 = 3 \cdot 9 & 1 \cdot 5 = 6 \cdot 12 = 11 \cdot 13 \\
 1 \cdot 6 = 5 \cdot 12 = 10 \cdot 13 & 1 \cdot 7 = 2 \cdot 3 = 8 \cdot 9 \\
 1 \cdot 8 = 2 \cdot 4 = 7 \cdot 9 & 1 \cdot 9 = 3 \cdot 4 = 7 \cdot 8 \\
 1 \cdot 10 = 6 \cdot 13 = 11 \cdot 12 & 1 \cdot 11 = 5 \cdot 13 = 10 \cdot 12 \\
 1 \cdot 12 = 5 \cdot 6 = 10 \cdot 11 & 1 \cdot 13 = 5 \cdot 11 = 6 \cdot 10 \\
 2 \cdot 9 = 3 \cdot 8 = 4 \cdot 7 & 5 \cdot 10 = 6 \cdot 11 = 12 \cdot 13.
 \end{array} \tag{7.16}$$

Notice that only $14 \times 3 = 42$ of the 78 two-factor interactions appear in (7.16). Thus, $78 - 42 = 36$ of the two-factor interactions are clear and so this design has $\text{alp} = (36, 0, 14)$. With 50 df for two-factor interactions (and 13 for main effects), this 64-run design is SOS. In fact, it is the only 64-run SOS design with $k < 17$ factors and so it has more df (50) for two-factor interactions than any other 2^{13-7} design. Although this design does not have minimum aberration, it seems to win every other prize for regular 2^{13-7} designs: with $\text{wlp} = (0, 14, 33, 16, \dots)$, it has weak minimum aberration. With $a_1 = 36$ it has 13 more clear two-factor interactions than any other design. With $L = 3$, every other design has some longer alias chains. The 36 clear two-factor interactions also have a special structure; as noted by Wu et al. (2012), if the six factors ($F_2, F_3, F_4, F_7, F_8, F_9$) are placed in group G_1 and the six factors ($F_5, F_6, F_{10}, F_{11}, F_{12}, F_{13}$) are placed in group G_2 , then every two-factor interaction involving one factor from each group is clear. This set of clear interactions is particularly useful for applications as discussed in Section 7.12, taking one group for control factors and the other for noise factors.

The criteria previously presented consider all two-factor interactions as of equal interest. If one has additional information about the likelihood or relevance of specific interactions, this information should be incorporated into the choice of the design. Articles that assist the user in estimating particular sets of two-factor interactions include Ke and Tang (2003); Ke et al. (2005); Wu and Chen (1992); Wu et al. (2012); and Grömping (2014).

7.5 Resolution V Designs

The resolution V 2^{5-1} (SOS) design is an especially efficient design in terms of run size, since its size $n (= 16)$ is equal to $1 + 0.5k(k + 1)$, which is the number of parameters in the two-factor interaction model with k factors. While all resolution V designs have orthogonal columns for terms in the two-factor interaction model and so are optimal in terms of variance efficiency, these designs generally require much larger n than the number of parameters to be estimated. While resolution III designs exist for $k = n - 1$ and resolution IV designs exist for $k = n/2$, resolution V designs generally allow only up to $k \approx \sqrt{n}$ factors. For instance, for $n = 32, 64, 128, 256$, and 512, the largest k 's for regular resolution V designs are 6, 8, 11, 17, and 23, respectively. In every case, $k \leq \sqrt{n} + 1$.

Generators for the smallest minimum aberration designs having resolution of V or more are listed in Table 7.2 for $k \leq 23$. Franklin (1984) and Xu (2009) found the larger designs

TABLE 7.2

Smallest Minimum Aberration Designs of Resolution V (or More) for 5-23 Factors

<i>k</i> Factors	<i>n</i> Runs	Design	Generator Columns ^a	<i>A</i> ₅	<i>A</i> ₆
5	16	5-1.1	15	1	
6	32	6-1.1	31	0	1
7	64	7-1.1	63	0	0
8	64	8-2.1	45, 51	2	1
9	128	9-2.1	31, 121	0	3
10	128	10-3.1	15, 51, 121	3	3
11	128	11-4.1	15, 51, 85, 120	6	6
12	256	12-4.1	31, 107, 205, 241	0	12
13	256	13-5.1	103, 121, 157, 179, 207	3	12
14	256	14-6.1	31, 39, 107, 169, 243, 254	9	18
15	256	15-7.1	78, 109, 135, 171, 181, 211, 246	15	30
16	256	16-8.1	23, 46, 92, 113, 139, 184, 197, 226	24	44
17	256	17-9.1	23, 46, 92, 113, 139, 184, 197, 226, 255	34	68
18	512	18-9.1	47, 93, 185, 227, 279, 369, 395, 453, 511	0	102
19	512	19-10.1	105, 127, 143, 181, 211, 285, 307, 327, 427, 473	12	84
20	512	20-11.1	Design 19-10.1, plus 485	16	120
21	512	21-12.1	Design 20-11.1, plus 510	21	168
22	512	22-13.1	105, 127, 155, 188, 206, 275, 298, 301, 350, 358, 369, 391, 507	63	189
23	512	23-14.1	23, 90, 99, 127, 155, 188, 206, 301, 340, 358, 391, 430, 435, 450	84	252

Source: With kind permission from Springer Science+Business Media: *A Comprehensive Guide to Factorial Two-Level Experimentation*, 2009, p. 285, Mee, R.W.

^a Column numbering according to Yates order.

in this table. The designs with $k \geq 7$ would not be recommended unless experimentation is inexpensive, since the run sizes can be two or three times the number of parameters to be estimated in the two-factor interaction model. Of course, if the error variance is large, so that more runs are needed to obtain adequate precision, then large designs may be justified. As alternatives to regular resolution V designs, one should consider nonregular orthogonal arrays or D-optimal designs that provide sufficient precision for estimates and sufficient protection from possible bias due to interactions omitted from the model; see Chapter 9.

7.6 Row Coincidence Matrix

We now discuss a tool that is useful for understanding the nature of two-level fractional factorial designs. Let **D** denote the $n \times k$ design matrix with ± 1 coding (cf. Section 7.2.2). The row coincidence matrix **T** is an $n \times n$ symmetric matrix defined as

$$\mathbf{T} = \mathbf{D}\mathbf{D}'.$$

The diagonal elements of **T** equal k , since each row of **D** coincides perfectly with itself. The off-diagonal elements range from $-k$ (if the design contains rows that are *mirror image pairs*) to k (if the design contains replicated rows). Mirror image pairs occur only

in even fractional factorial designs (those with no odd-length words in the defining relation) or full factorial designs; thus, resolution III designs will have a value no smaller than $-(k-2)$.

The row coincidence matrix reveals that saturated designs have a very simple structure. Since removing the column of $+1$'s from \mathbf{S}_n , defined in (7.11), produces a saturated design and $\mathbf{S}_n \mathbf{S}'_n = n\mathbf{I}_n$, the row coincidence matrix for any regular saturated fraction is $n\mathbf{I}_n - \mathbf{J}_n$, where \mathbf{J}_n is an $n \times n$ matrix of $+1$'s. (This result holds for all orthogonal saturated designs.) Thus, the treatment combinations for any saturated design have an invariance property; the levels for each pair of treatment combinations coincide on $n/2 - 1$ factors and differ on the other $n/2$ factors; see Xu and Wu's (2001) Lemma 2. Viewing each treatment combination as a point in k -dimensional space, the distance between every pair of points is the same.

The row coincidence matrices for maximal even designs also have a simple structure. Computing $\mathbf{T} = \mathbf{E}_n \mathbf{E}'_n$ defined by (7.12) reveals that every row of \mathbf{T} consists of the values k , $-k$, and the rest zeros. Projections of maximal even designs will have row coincidence values that are also symmetrical, with all odd moments zero, where the moments are defined as follows.

Let t_1, t_2, \dots, t_n denote the elements of any column of \mathbf{T} for a regular 2^{k-q} design; for regular 2^{k-q} designs, Block (2003, Theorem 5.1) proved that the distributions of values are identical for every column of \mathbf{T} . Thus, the r th moment of \mathbf{T} is defined to be

$$M_r = \frac{[t_1^r + t_2^r + \dots + t_n^r]}{n}. \quad (7.17)$$

For two-level designs, the *Hamming distance matrix* \mathbf{H} is related to the row coincidence matrix through

$$\mathbf{T} = k\mathbf{J}_n - 2\mathbf{H}.$$

Clark and Dean (2001) addressed the problem of equivalence of designs using the Hamming distance matrices for a sequence of projections (see Section 7.11). Xu and Wu (2001) derived explicit expressions for the generalized word length pattern (gwlp) of any design as functions of moments of the design's Hamming distance matrix. For regular fractions, the wlp and gwlp are identical. Thus, the wlp (7.10) for regular fractions, which is our concern here, can be written as moments of either \mathbf{H} or \mathbf{T} and so can be determined from a single column of either. It is easily shown that for designs of resolution III or more, $M_1 = 0$ (i.e., columns of \mathbf{T} sum to zero) and $M_2 = k$. Subsequent moments are used to compute A_3, \dots, A_6 as follows (Butler 2003b):

$$\begin{aligned} A_3 &= M_3/6, \\ A_4 &= [M_4 - (3k - 2)k]/24, \\ A_5 &= [M_5 - 10(k - 2)M_3]/120, \\ A_6 &= [M_6 - 5(3k - 8)M_4 + 2(15k^2 - 60k + 32)k]/720. \end{aligned}$$

Xu and Wu (2001) proved that sequentially minimizing A_3, A_4, \dots (i.e., minimum aberration) is equivalent to sequentially minimizing the M_r 's. Furthermore, since the frequency

TABLE 7.3
Row Coincidence Distributions for Five 2_{III}^{9-5} Designs

	Design 9-5.1	Design 9-5.2	Design 9-5.3	Design 9-5.4	Design 9-5.5
	-7	-3	-5	-3	-3
	-1	-3	-3	-3	-3
	-1	-3	-3	-3	-3
	-1	-3	-1	-3	-1
	-1	-3	-1	-1	-1
	-1	-3	-1	-1	-1
	-1	1	-1	-1	-1
	-1	1	-1	-1	-1
	-1	1	-1	-1	-1
	1	1	1	-1	-1
	1	1	1	1	-1
	1	1	1	1	1
	1	1	1	1	1
	1	1	1	3	1
	1	1	3	3	5
	9	9	9	9	9
M_3	24	36	36	42	48
M_4	561	441	465	441	465

Source: With kind permission from Springer Science+Business Media: *A Comprehensive Guide to Factorial Two-Level Experimentation*, 2009, p. 515, Mee, R.W.

distribution of values in **T** can be determined from the wlp, one can determine whether two regular fractions have the same wlp by comparing their row coincidence distributions, without even computing the moments. Finally, if the 2^{k-q} fraction contains the treatment combination $(+1, +1, \dots, +1)$, then its row coincidence distribution can be obtained simply by summing across the columns of the design.

We conclude this section by displaying the row coincidence distributions for the five non-isomorphic regular 2_{III}^{9-5} designs, as identified by Chen et al. (1993). Table 7.3 contains these distributions and their third and fourth moments. Comparing the coincidence distributions reveals the following insights. Since M_3 is proportional to A_3 , all resolution III designs have positively skewed coincidence distributions; the best design, 9-5.1, is the one with the least skewness, that is, $M_3 = 24$ and $A_3 = 4$. According to Butler’s (2003a) Theorem 4, all minimum aberration designs of resolution III project to the minimal even design, so design 9-5.1 must include the resolution IV 2^{8-4} fraction as a projection. The row coincidence distribution for this maximal even design is $-8, 0, \dots, 0, 8$. Since design 9-5.1 has coincidence values that just differ by ± 1 from the symmetric coincidence values of the resolution IV 2^{8-4} , this row coincidence distribution has the least skewness. In contrast, design 9-5.3 has a minimum row coincidence of -5 , so its least coincident treatment combination pairs still agree on two of the nine factors. Thus, dropping any one of design 9-5.3’s factors cannot produce the even 2_{IV}^{8-4} design with coincident values of -8 . (The other designs in Table 7.3, all with least coincidence values of -3 , can achieve a row coincidence value no smaller than -4 when one factor is dropped.) Note also how the best design maximizes the minimum distance between points (i.e., has the fewest coincidences among factors in different treatment combinations). For the best two designs, the maximum agreement between

two treatment combinations is on just five of the nine factors, resulting in a row coincidence value of 1. The poorer designs have pairs of runs that agree on six or seven factors, producing a t_i value of 3 or 5, respectively. This makes explicit the intuition that the best fractions more effectively spread out treatment combinations in the design space.

7.7 Analysis of Factorial Experiments

There is a simplicity to analyzing regular two-level fractional factorial designs due to the fact that (1) each main effect and interaction represents a single df and (2) the contrasts from different alias sets are orthogonal to one another. When the model in (7.3) is reasonable, the only complication is the estimation of σ^2 from data that lacks any replication. While the factorial effects can be estimated separately, an estimate of σ is needed to compute standard errors for the factorial effects. To estimate σ properly, some assumption must be made. One approach is to assume that all interactions above a certain order are negligible. For instance, for the 2^{6-1} fraction with resolution VI, if we assume no three-factor or higher-order interactions, then fitting the two-factor interaction model will leave 10 df for error. If this assumption is valid, then the resulting mean squared error (MSE) provides an unbiased estimator for σ^2 and one may construct valid t -tests for each of the lower-order effects. However, the assumption of no three-factor or higher-order interactions furnishes no df for error in cases where the fraction is SOS.

For the analysis of resolution III designs and most resolution IV designs, an assumption of *effect sparsity* is commonly invoked. Rather than specifying in advance of the analysis that particular interactions are negligible, effect sparsity supposes that the number of “active” (nonnegligible) main effects and interactions is no more than, say, 25% or 30% of $(n - 1)$. For instance, for a 16-run (32-run) fraction, this means that no more than about four (nine) effects are active. Numerous papers have been written proposing methods of analysis for two-level factorials based on effect sparsity. The most popular method is from Lenth (1989), due to its simplicity and overall good performance (also see Haaland and O’Connell 1995).

We now illustrate the use of both analysis approaches using the initial experiment of Barnett et al. (1997). These authors describe a wafer etching experiment with 18 runs, one wafer per run. The response is the standard deviation (SD) of nine thickness measurements on a single wafer. A 2^{6-2}_{IV} fractional factorial was used to define treatment combinations for 16 of the 18 runs; the generators for this fraction were $5 = 124$ and $6 = 123$. The remaining two runs were taken at the center; this is possible since the six factors are all quantitative. Table 7.4 shows the factors in coded levels, with runs sorted from best to worst SD. The two additional center runs furnish 1 df for pure error from replication and 1 df for identifying the presence of pure quadratic curvature, that is, effects of the form x_i^2 . We use the natural log, $\ln(\text{SD})$, to satisfy the constant error variance assumption.

We begin by assuming that no third-order terms (or higher order) are needed. The full second-order model cannot be estimated due to aliasing; but a model with 14 df besides the intercept can be fitted: 6 df for main effects, 7 df for two-factor interactions (one from each alias set), and 1 df for the aliased pure quadratic terms. The MSE is 0.0345, with 3 df. Least squares estimates, t -ratios, and p -values are shown in Table 7.5, sorted from the most significant to the least. (The pure quadratic estimate listed is the coefficient for $3x_i^2$; this rescaling is so that all estimates in Table 7.5 have the same standard error of $\sigma/\sqrt{16}$.) Four estimates are significant at the 0.01 level. What is perhaps surprising is that only one

TABLE 7.4

Barnett et al. (1997) Experiment

F_1	F_2	F_3	F_4	F_5	F_6	SD
1	1	1	-1	-1	1	1.60
-1	-1	1	-1	-1	1	1.71
-1	1	1	1	-1	-1	2.61
-1	1	1	-1	1	-1	3.10
1	1	1	1	1	1	3.10
-1	-1	1	1	1	1	3.43
-1	-1	-1	1	1	-1	4.05
0	0	0	0	0	0	4.83
1	1	-1	1	1	-1	6.18
0	0	0	0	0	0	6.37
-1	-1	-1	-1	-1	-1	6.49
1	-1	-1	1	-1	1	7.00
1	1	-1	-1	-1	-1	7.20
1	-1	-1	-1	1	1	7.25
-1	1	-1	1	-1	1	16.58
1	-1	1	-1	1	-1	19.11
1	-1	1	1	-1	-1	19.20
-1	1	-1	-1	1	1	38.50

Source: Reprinted with permission from Barnett, J. et al., Using fewer wafers to resolve confounding in screening experiments, in *Statistical Case Studies for Industrial Process Improvement*, Czitrom, V. and Spagon, P. (eds.), SIAM, Philadelphia, PA, 1997, 235–250. Copyright 1997 Society for Industrial and Applied Mathematics.

of the four is a main effect. Thus, this experiment indicates that high F_3 is beneficial for minimizing $\ln(\text{SD})$ and that at least three other factors are involved. The pragmatist might just declare a preference for factors 4 and 5 low and factors 3 and 6 high based on the first two rows of Table 7.4, requiring no more experimentation. The curious investigator would pursue a follow-up experiment. This example will be revisited for that purpose in Section 7.9.

These data may also be analyzed based on an assumption of effect sparsity, without specifying negligible high-order interactions. With a 16-run 2^{6-2} fraction (and ignoring the centerpoint runs), Lenth's method would be based on 15 contrasts (6 main effects, 7 for two-factor interactions, 2 for three-factor interactions). The addition of the two centerpoint runs adds two additional df. JMP's modeling screening analysis option implements Lenth's method using all 17 df that we describe here. Table 7.5 shows the 17 least squares estimates. In addition to two three-factor interaction contrasts, we see an effect labeled *null*, which is the pure error contrast, calculated as $\ln(6.37) - \ln(4.83) = 0.2768$ divided by $\sqrt{32}$ so that it has the same standard error as the factorial effects.

To compute Lenth's t statistics, one begins by computing the median magnitude of the estimates scaled as in Table 7.5 to have the same variance. In Table 7.5, this median is 0.0795. Lenth's method computes $s_0 = 1.5(0.0795) = 0.1192$ as a preliminary estimate of the standard error of the estimates. (This multiplier is based on the fact that for a standard normal random variable Z , $E(1.5|Z|)$ is very nearly 1.) If no effects were active, s_0 could serve as the estimated standard error. However, if a few effects are active, their corresponding estimates

TABLE 7.5
Sorted Estimates for ln(SD) from Two Models

Term	Estimate	<i>t</i> -Ratio ^a	<i>p</i> -Value ^a	Lenth's <i>t</i> ^b	<i>p</i> -Value ^b
$F_1 F_2 = F_3 F_6 = F_4 F_5$	0.4967	-10.69	0.0018	-5.38	0.0021
$F_1 F_6 = F_2 F_3$	-0.4400	-9.47	0.0025	-4.77	0.0029
F_3	-0.3802	- 8.18	0.0038	-4.12	0.0058
$F_1 F_3 = F_2 F_6$	0.3533	7.60	0.0047	3.83	0.0080
$F_3 F_4 = F_5 F_6$	0.1682	3.62	0.0362	1.82	0.0824
F_5	0.1111	2.39	0.0966	1.20	0.2120
F_1	0.1026	2.21	0.1143	1.11	0.2434
F_6	-0.0815	-1.75	0.1777	-0.88	0.3497
$F_3 F_5 = F_4 F_6$	0.0795	1.71	0.1857	0.86	0.3612
$F_1 F_4 = F_2 F_5$	0.0793	1.71	0.1866	0.86	0.3624
$F_1 F_5 F_6 = \dots$	-0.0615			-0.67	0.5315
F_2	-0.0577	-1.24	0.3023	-0.63	0.5693
Null	0.0489			0.53	0.6312
$F_1 F_5 = F_2 F_4$	-0.0438	-0.94	0.4159	-0.47	0.6686
Aliased quadratic terms	0.0293	0.63	0.5734	0.32	0.7723
F_4	-0.0195	-0.42	0.7032	-0.21	0.8472
$F_1 F_3 F_5 = \dots$	0.0173			0.19	0.8652

^a Based on standard error calculated as $[MSE/16]^{1/2} = 0.0465$, with 3 df.

^b Based on PSE = 0.0923.

should be much larger than the true standard error, which would tend to inflate s_0 . Consequently, Lenth's procedure is adaptive. Estimates larger than $2.5s_0$ are eliminated, and the median is recalculated. Here, four estimates in Table 7.5 exceed $2.5(0.1192) = 0.298$ in absolute value, and the median from the remaining 13 estimates is 0.0615. Thus, Lenth's pseudo standard error (PSE) is $1.5(0.0615) = 0.0923$. This estimate is double the standard error of $[MSE/16]^{1/2} = 0.0465$ used in the first analysis; this discrepancy is because the 3 df for error used there are among the smallest estimates in Table 7.5. Dong (1993) proposed a procedure similar to that of Lenth, but based on the mean square of the estimates after pruning rather than on their median. While the mean might be more efficient than the median, it is also less robust to the presence of active effects after pruning. Schoen and Kaul (2000) consider both methods, provide their own adjustments, and make recommendations that take into consideration the expected number of active effects.

Lenth's PSE is not the final goal of Lenth's procedure. Rather, it is just a step in the process of determining statistical significance. Here, we describe how the *p*-values in the last column of Table 7.5 were computed. Lenth (1989) proposed the simple procedure of computing *p*-values for Lenth's *t* statistics using a *t* distribution with $m/3$ df, where *m* is the total number of contrasts used. Here, $m = 17$, so Lenth (1989) would have us compute *p*-values using the *t* distribution with $17/3 = 5.67$ df. However, subsequent work has shown that using a *t* distribution with $m/3$ df to compute *p*-values is too conservative. While some software still uses $m/3$, it is preferable to abandon use of a *t* approximation and to compute *p*-values by simulating the null distribution of Lenth's *t*; see Edwards and Mee (2008). This was done to provide the last column of Table 7.5. Using Lenth's procedure, four estimates have *p*-values below 0.01. For references to the power of Lenth's method and other procedures based on effect sparsity, see Haaland and O'Connell (1995), Hamada and Balakrishnan (1998), and Mee (2009, Section 13.1).

TABLE 7.6

Irvine et al. (1996) 2^{13-9} Pulping Experiment

Run	F_1	F_2	F_3	F_4	F_5	F_6	F_7	F_8	F_9	F_{10}	F_{11}	F_{12}	F_{13}	Brightness
1	1	1	-1	1	-1	1	-1	-1	1	-1	-1	-1	1	20.5
2	-1	1	1	1	1	1	-1	1	1	-1	-1	1	-1	19.8
3	-1	-1	-1	1	1	1	-1	1	-1	1	1	-1	-1	17.3
4	-1	-1	1	1	1	-1	1	-1	1	1	-1	-1	1	29.0
5	1	-1	-1	1	-1	-1	1	1	1	1	-1	1	-1	19.5
6	-1	1	1	-1	-1	-1	-1	-1	1	1	1	1	-1	17.7
7	-1	-1	1	-1	-1	1	1	1	1	-1	1	-1	1	25.8
8	1	-1	1	1	-1	1	-1	-1	-1	1	1	1	1	23.8
9	1	-1	-1	-1	1	1	1	-1	1	-1	1	1	-1	23.4
10	1	1	1	-1	1	1	1	-1	-1	1	-1	-1	-1	22.1
11	1	1	-1	-1	1	-1	-1	1	1	1	1	-1	1	26.3
12	1	-1	1	-1	1	-1	-1	1	-1	-1	-1	1	1	23.4
13	-1	1	-1	-1	-1	1	1	1	-1	1	-1	1	1	26.4
14	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	-1	17.6
15	1	1	1	1	-1	-1	1	1	-1	-1	1	-1	-1	20.0
16	-1	1	-1	1	1	-1	1	-1	-1	-1	1	1	1	26.4

Consider now another experiment, as described by Irvine et al. (1996) that used a 2_{III}^{13-9} design to investigate the best method to remove lignin during the pulping stage of paper production. Table 7.6 presents the coded levels of the 13 factors used in the experiment, plus pulp brightness, which was one of 15 response variables reported for each treatment combination. The design is a minimum aberration design; for 13 factors in 16 runs, any 13 contrast columns from S_{16} provide a minimum aberration design.

Main effects account for 13 of the $n - 1 = 15$ df available. The MSE of the main effects model has just 2 df and provides an unbiased estimator for σ^2 only if the two contrast columns of S_{16} not assigned main effects have estimates with expected value of zero. For the design in Table 7.6, these two columns correspond to $F_{10}F_{13}$ and $F_{11}F_{13}$; each of these is also aliased with five more two-factor interactions. So, if any of these 12 two-factor interactions is not negligible, the MSE of the main effects model will tend to overestimate σ^2 . Due to this risk, it is preferable here to use Lenth's method.

Fitting a two-factor interaction model for this experiment results in $PSE = 0.281$ (see Table 7.7). Four of the 15 p -values estimated by simulation of Lenth's t null distribution are less than 0.02, while the others exceed 0.15. Thus, we have evidence for as many as three active main effects and at least one active two-factor interaction. Which particular interaction(s) is active cannot be determined from this resolution III design. In Table 7.7, the estimate of 0.85 has been identified with $F_{10}F_{13}$, from among the aliases, since factor 13 has the largest main effect estimate. We will revisit this experiment in Section 7.9 where we discuss augmenting fractional factorial designs.

If we had fit the main effects model here, then MSE would have been 6.185, and the estimates would have had standard error $(6.185/16)^{1/2} = 0.6217$, more than double the PSE. Although we do not know the true error variance here, the MSE is very likely much too large due to bias by an active interaction. Note that Lenth's method should be applied to a saturated model, not a reduced model, since it is the scatter of estimates near zero that furnishes information about σ .

TABLE 7.7

Irvine et al. (1996) 2^{13-9} Pulping Experiment

Term	Aliases	Estimate	Lenth t	p -Value
F_{13}	4·8 = 9·12 = 1·7 = 5·6 = 2·3	2.7625	9.82	0.0003
F_7	10·11 = 5·8 = 3·12 = 1·13 = 4·6 = 2·9	1.6375	5.82	0.0021
F_5	3·11 = 7·8 = 10·12 = 1·4 = 6·13	1.0250	3.64	0.0087
$F_{10}F_{13}$	5·9 = 3·4 = 1·11 = 6·12 = 2·8	0.8500	3.02	0.0175
F_4	9·11 = 8·13 = 1·5 = 6·7 = 2·10	-0.4000	-1.42	0.1521
F_{10}	7·11 = 3·8 = 5·12 = 6·9 = 2·4	0.3250	1.16	0.2281
F_9	4·11 = 12·13 = 1·3 = 6·10 = 2·7	0.3125	1.11	0.2450
F_3	5·11 = 8·10 = 7·12 = 1·9 = 2·13	0.2625	0.93	0.3189
$F_{11}F_{13}$	8·9 = 4·12 = 1·10 = 3·6 = 2·5	0.2250	0.80	0.3873
F_{11}	7·10 = 4·9 = 3·5 = 8·12 = 2·6	0.1500	0.53	0.6165
F_8	5·7 = 4·13 = 3·10 = 11·12 = 1·6	-0.1250	-0.44	0.6755
F_{12}	5·10 = 9·13 = 3·7 = 8·11 = 1·2	0.1125	0.40	0.7044
F_1	7·13 = 4·5 = 3·9 = 6·8 = 2·12	-0.0625	-0.22	0.8301
F_6	5·13 = 4·7 = 9·10 = 1·8 = 2·11	-0.0500	-0.18	0.8631
F_2	4·10 = 7·9 = 3·13 = 1·12 = 6·11	-0.0375	-0.13	0.8946

 $s_0 = 0.394$, Lenth's PSE = $1.5(0.225 + 0.150)/2 = 0.28125$.

7.8 Screening

Experiments for the purpose of screening for active factors are common, for example, in industrial research and development, in engineering, in genetic and medical screening, in drug discovery, and in simulation and computer experiments (see, e.g., articles in Dean and Lewis 2006). Screening involves the evaluation of a large number of potentially important factors to search for the few active factors (now defined as those that have substantially large main effects and/or are involved in substantially large interactions). Once detected, active factors are followed up in later studies for more detailed analysis. Screening experiments may involve more factors than can be accommodated by a regular fractional factorial design. Chapter 9 gives alternative designs, such as orthogonal arrays, search designs, and supersaturated designs that can be used for this purpose. Here, we mention an alternative strategy that allows regular fractions to be used in *two-stage group screening*. The technique, described by Dorfman (1943) for blood screening, was developed further by Watson (1961) for use in the factorial setting for screening main effects; it was later extended for screening interactions by Lewis and Dean (2001) and Vine et al. (2005). An overview of group screening is given by Morris (2006), and an example of its use in the automobile industry is given by Vine et al. (2008).

In two-stage group screening, factors are reordered and partitioned into groups so that, say, the g factors F_1, \dots, F_g belong to group G_1 , the h factors F_{g+1}, \dots, F_{g+h} belong to group G_2 , and so on. Then a fractional factorial experiment is designed for the *group factors*, G_1, G_2, \dots , and effects of individual factors within each of these groups are intentionally confounded. The experiment is analyzed as if each group factor is a single experimental factor and those group factors that are found to have active main effects or active interactions are passed to the second-stage experiment. At the second stage, the active groups are

dismantled and a new fractional factorial experiment is designed, run, and analyzed for the resulting individual factors.

For example, suppose there are 20 factors to be screened and each is to be observed at two levels. A 2_{V}^{20-11} design with 512 runs exists (see Table 7.2), allowing all main effects and two-factor interactions to be estimated independently. However, if only a few of the 20 factors are active (*factor sparsity*), these can likely be found in far fewer runs. For example, if, at stage 1, we partition the 20 factors into five groups G_1, \dots, G_5 , each containing four factors, then a 2_{V}^{5-1} fraction with 16 observations can be used to determine which of the five groups have active main effects and two-factor interactions. If two groups, say, are found to be active, then there are eight individual factors within these two groups that now need to be screened, and this can be done at stage 2 with a 2_{V}^{8-2} fraction having 64 observations (Table 7.2). So a total of 80 observations will have been used. And if three groups were found to be active at stage 1, then a 2_{VI}^{12-4} fraction with 256 observations could be used, for a total of 272 observations (as compared with 512 for the 2_{V}^{20-11} fraction). In fact, there is another possibility at stage 2 since an orthogonal array (Chapter 9) exists with 128 observations that allows independent estimation of main effects and interactions; so only 144 observations would be needed. There is a small possibility that all groups could be found to be active; then the experimenter may actually use more observations in the two-stage design than in the single-stage nongroup experiment, unless the choice is made to ignore factors in groups with smaller effects.

When factor G is at level +1 in the stage 1 design, this means that all the factors within the groups are held at level +1. Consequently, there is some unusual confounding at stage 1. In fact, this is a resolution II design for the individual factors since the main effect contrasts of all individual factors within the same group are identical; in addition, all two-factor interactions between individual factors within the same group are aliased with the mean. Also aliased are the two-factor interactions between any two factors, one from G_i and one from G_j (for all pairs of groups G_i and G_j). Therefore, group screening designs, as all fractional factorial designs, should be used only when factor sparsity is expected, and the hope is that any active effects that happen to be aliased will reinforce each other and not cancel. For main effects, reinforcement can be achieved whenever prior knowledge enables the assignment of factor levels so that +1 represents the level that gives the larger response.

Draguljić et al. (2014) investigate the performance of group screening with a regular design at stage 1 and a nonregular design at stage 2 and show that the technique works well with shrinkage regression analysis methods, such as the Dantzig selector.

7.9 Augmentation of Designs and Follow-Up Experiments

In this section, we consider a variety of follow-up designs that are commonly used to augment an initial two-level fractional factorial design in order to gain information about aliased interactions. We use the Irvine et al. (1996) and Barnett et al. (1997) data analyzed in Section 7.7 to motivate each augmentation choice.

7.9.1 Augmenting Resolution III Designs by Foldover

Lenth's analysis of the 2^{13-9} experiment of Irvine et al. (1996) identified three likely main effects and at least one interaction as indicated by the p -values in Table 7.7. Generators for

the design in Table 7.6 are

$$\begin{aligned} 5 &= -1 \cdot 4, & 6 &= -1 \cdot 2 \cdot 3 \cdot 4, & 7 &= 1 \cdot 2 \cdot 3, & 8 &= 2 \cdot 3 \cdot 4, & 9 &= -1 \cdot 3, \\ 10 &= -2 \cdot 4, & 11 &= 1 \cdot 3 \cdot 4, & 12 &= -1 \cdot 2, & 13 &= -2 \cdot 3. \end{aligned} \quad (7.18)$$

Of the 78 two-factor interactions, 66 are aliased with main effects. The fact that at least one interaction has been found among the 12 not aliased with main effects raises the question of whether other interactions exist that are biasing our main effect estimates. To answer this question, it would be helpful to run a second 2^{13-9} fraction obtained by simply reversing the signs of some of the original generators (a *foldover fraction*). Since $2^9 = 512$, Table 7.6 shows a 1/512th fraction, and the other 511 fractions from this same family of generators can be obtained by reversing the signs of one or more of the generated factors in (7.18). If we reverse the signs of just the six even-length aliases of the generated factors {5, 6, 9, 10, 12, 13}, the new fraction has generators

$$\begin{aligned} 5 &= 1 \cdot 4, & 6 &= 1 \cdot 2 \cdot 3 \cdot 4, & 7 &= 1 \cdot 2 \cdot 3, & 8 &= 2 \cdot 3 \cdot 4, & 9 &= 1 \cdot 3, \\ 10 &= 2 \cdot 4, & 11 &= 1 \cdot 3 \cdot 4, & 12 &= 1 \cdot 2, & 13 &= 2 \cdot 3. \end{aligned} \quad (7.19)$$

Note that even-length factor generators create odd-length generating words in the defining relation. For any 2^{13-9} fraction, there are 256 odd-length words and 255 even-length words in the defining relation, not counting the identity column. In the initial fraction here, the signs of the generators were chosen to make all the odd-length words in the defining relation to be negative and the even-length words to be positive. For the second fraction we have proposed, all words in the defining relation have a positive sign. This implies that if these two 2^{13-9} fractions (7.18) and (7.19) are put together, the resulting 2^{13-8} fractional factorial will have resolution IV. For instance, the defining relation for (7.18) includes **-145** and the contrast coefficients for the interaction $F_1F_4F_5$ are -1 for the 16 treatment combinations in Table 7.6, while the coefficients for the same interaction are $+1$ for all treatment combinations in the proposed follow-up design (7.19). Since the contrast coefficients for $F_1F_4F_5$ sum to zero over the combined 32 runs, this interaction does not appear in the defining relation for the design after augmentation. The same is true for all the odd-length words in the defining relation for (7.18).

One does not need to know the generators or the defining relation of a resolution III design to find a fraction that when added to the first design will increase the resolution. One simply has to reverse all the signs in the design columns. This simple rule ensures that the combined design will have resolution IV or more, since all odd-length words will change sign in the augmenting fraction and so sum to zero for the combined design.

Analysis of a regular fraction augmented by foldover is straightforward, as it is simply a regular $2^{k-(q-1)}$ design. If each portion is run as a completely randomized design, as we have assumed in our previous analysis, the combined design is a randomized block design run in two blocks. When the data are analyzed, one should include a main effect for blocks.

What is estimable after the *reverse all column signs* fraction is run? Combining two 2^{k-q} blocks breaks each alias chain (of length 2^q) in two, with odd-length effects in one set and even-length effects in the other. Thus, for the Irvine et al. design, if we assume no three-factor or higher-order interactions, the main effects in the 15 alias sets in Table 7.7 are separated from the two-factor interactions and can all be estimated independently.

All two-factor interactions aliased in the original design are still aliased. This is because no length-4 words were eliminated from the defining relation. When $k \geq 5n/8$, the *reverse*

all column signs fraction is the only one of the other 2^{k-q} fractions that will create a $2^{k-(q-1)}$ fraction with higher resolution; the resulting augmented design is always even, since it contains mirror image treatment combinations. However, for smaller k , there are sometimes additional choices that will both eliminate all length-3 words (increasing the resolution to IV) and eliminate some length-4 words; see Li and Mee (2002).

7.9.2 Augmenting Resolution IV Designs by Foldover and Semifoldover

The 2_{IV}^{6-2} fraction by Barnett et al. (1997) analyzed in Section 7.7 revealed that two-factor interactions are the most prominent effects (see Table 7.5). However, the resolution IV fraction aliases every two-factor interaction with one or two others as a result of the three length-4 words in the defining relation

$$\mathbf{I} = \mathbf{1245} = \mathbf{1236} = \mathbf{3456}. \quad (7.20)$$

This design is a quarter fraction so there are three other fractions that may be obtained by reversing the sign of one or both of the generators. If we reverse the sign of only the column for factor 6, the new defining relation is $\mathbf{I} = \mathbf{1245} = -\mathbf{1236} = -\mathbf{3456}$, and the combined 2^{6-1} design has defining relation $\mathbf{I} = \mathbf{1245}$. Augmenting the initial 2_{IV}^{6-2} design with any foldover fraction different from the initial design eliminates two of the three words in the defining relation (7.20). Here, the preferable foldover fraction is obtained by reversing the sign of either factor 3 or factor 6, since all two-factor interactions involving F_3 or F_6 will then be clear of aliasing with other two-factor interactions. This choice is motivated by the fact that factor F_3 has the largest main effect and so is often deemed the most likely factor to be involved in interactions. A foldover of the design in Table 7.4 obtained by reversing column 6 is shown in Table 7.8. If all 16 runs had been performed, five of the seven two-factor interaction alias sets (refer to Table 7.5) would be split. The only remaining aliasing would involve $\mathbf{12} = \mathbf{45}$, $\mathbf{14} = \mathbf{25}$, and $\mathbf{15} = \mathbf{24}$.

Barnett et al. (1997) ran only eight additional runs to augment their original experiment, not 16, as indicated by the responses in Table 7.8. Their follow-up experiment is called a semifold fraction, since it only includes 8 runs from a second 2^{6-2} design. Mee and Peralta (2000) introduced notation to describe semifold designs; Barnett et al.'s semifold would be designated $f_0 = \mathbf{6}$, $ss = \mathbf{6+}$, which means the foldover fraction is obtained by reversing the signs of the column for F_6 and then subsetting (ss) on the runs with positive values for F_6 . If we had been recommending a semifold design, we would have recommended the eight runs defined by $f_0 = \mathbf{6}$, $ss = \mathbf{3+}$, since the initial experiment indicated that the best results are obtained at the high level for factor 3.

For any even resolution IV design, there are at most $(n/2 - 1)$ df for two-factor interactions. For such even designs, if one semifolds by subsetting on a main effect, the semifold fraction permits estimation of as many two-factor interactions as could be estimated if the entire foldover were to be run (Mee and Xiao 2008). The gain in completing the entire foldover rather than the semifoldover is in the added precision. After foldover, the standard errors are $\sigma/\sqrt{2n}$, while after semifoldover, the standard errors for two-factor interactions can be as large as σ/\sqrt{n} . For more details, see Barnett et al. (1997) or Mee (2009, Section 9.5).

For the 2^{6-2} and the 2^{7-3} resolution IV designs, every possible foldover fraction can be obtained by reversing the signs of a single main effect column. For the 2^{8-4} and any design where $k < 2^q - 1$, some foldover fractions are obtained only by reversing the signs of two or

TABLE 7.8
Foldover Design by Barnett et al. (1997)

F_1	F_2	F_3	F_4	F_5	F_6	SD
1	1	1	-1	-1	-1	
-1	-1	1	-1	-1	-1	
-1	1	1	1	-1	1	6.63
-1	1	1	-1	1	1	6.60
1	1	1	1	1	-1	
-1	-1	1	1	1	-1	
-1	-1	-1	1	1	1	13.67
1	1	-1	1	1	1	9.23
-1	-1	-1	-1	-1	1	12.96
1	-1	-1	1	-1	-1	
1	1	-1	-1	-1	1	8.59
1	-1	-1	-1	1	-1	
-1	1	-1	1	-1	-1	
1	-1	1	-1	1	1	6.19
1	-1	1	1	-1	1	9.81
-1	1	-1	-1	1	-1	

Source: Design by Barnett, J. et al., Using fewer wafers to resolve confounding in screening experiments, in Statistical Case Studies for Industrial Process Improvement, Czitrom, V. and Spagon, P. (eds.), SIAM, Philadelphia, PA, 1997, 235–250.

more columns. If the signs of the column for a single factor are reversed, all two-factor interactions involving that factor are separated from their aliased two-factor interaction strings. However, for resolution IV designs with many sets of aliased two-factor interactions of size 4 or more, a larger number of length-4 words are eliminated by foldovers obtained by reversing column signs for more than one factor; for details, see Li and Lin (2003), Mee and Xiao (2008), and Ai et al. (2010). Generally, once the foldover fraction is determined, a judiciously chosen subset of $n/2$ runs will suffice (Edwards 2011); all n runs of the foldover fraction should be performed only if the standard errors after the initial 2^{k-q} experiment were deemed unsatisfactorily large.

7.9.3 Other Design Augmentation Strategies

The previous two sections highlighted augmentation of an experiment by adding all (or half) of a regular foldover fraction of size n to the original n -run design. Here, we give a brief survey of four other common choices for augmentation: confirmation runs, D-optimal augmentation, steepest ascent, and composite designs.

The analysis of data from a 2^{k-q} fraction inevitably requires some assumptions and/or educated guesses. At the very least, one ought to perform follow-up runs to confirm the tentative conclusions reached from the fractional factorial analysis. For instance, if Barnett et al. (1997) had not had the resources or time for running a semifold fraction, at least a couple of confirmation runs should have been performed with F_3 and F_6 at the high level and F_4 and F_5 at the low level, since this combination produced the two smallest SD. One confirmation run should be a repeat of row 1 or 2 in Table 7.4, while a second run should

take $(F_1, F_2) = (-1, 1)$ since the initial experiment hints that low F_1 is preferred. This pair of runs would help confirm that $SD < 2$ is repeatable.

A foldover or semifoldover of Barnett et al.'s 2^{6-2} design still has aliasing between three pairs of two-factor interactions. This is inevitable since three two-factor interactions are aliased together in the original fraction. Rather than restricting the follow-up runs to come from a regular fraction, one can simply specify both the model of interest and the number of follow-up runs and then use software to determine the additional treatment combinations that maximize the information about the parameters of this model. If the criterion is to maximize the determinant of $X'X$, this is D-optimal augmentation. This follow-up strategy is popular due to its flexibility, since one is free to specify any model of interest and any number of runs. If the run size is not sufficient to estimate the given model, one may use Bayesian D-optimal design; the simplest version of this maximizes the determinant of $X'X + \lambda I$ for a suitably small value of λ (Jones and duMouchel 1996).

If the objective is to optimize the response(s) and one can explore outside the region of experimentation for the initial fraction, then *steepest ascent* (or *descent*) is an approach that uses a fitted model to identify a search direction that optimizes the predicted response, subject to a constraint on the standard error of prediction or a constraint on distance from the center of the original design (see Box and Draper 2007, Chapter 6). Mee (2009, Section 9.3) presents a variety of situations, including optimizing two or more responses and optimization with multiple constraints.

Finally, one common model of interest for quantitative factors is the second-order model, which contains not only the k linear main effects and $k(k-1)/2$ two-factor interactions that represent the two-factor interaction model, but also the k pure quadratic terms, x_1^2, \dots, x_k^2 . Box and Wilson (1951) first proposed the *central composite design* for estimating the second-order model (see also Chapter 1, Section 1.8.3, and Chapter 5, Section 5.2.2). If the initial 2^{k-q} fractional factorial has resolution V or more, then the addition of a second block containing $2k$ axial points and at least one centerpoint run allows each factor to be observed at 3 or more levels and permits estimation of the second-order model. The pair of axial points for the i^{th} factor uses levels that are equally spaced above and below the midpoint and for every other factor uses the midpoint (between low and high levels). In coded units, the axial points for the first factor are thus $(-\alpha, 0, \dots, 0)$ and $(\alpha, 0, \dots, 0)$. Taking $\alpha = 1$ creates exactly three levels for each factor, but this provides poor precision for the pure quadratic coefficients unless the axial points are replicated. Alternatively, one may choose α near to \sqrt{k} , which (in coded units) places the axial points the same distance from the centerpoint as the factorial points; this improves the precision for the pure quadratic coefficients, but it requires the design to have five levels for each factor. For details regarding the choice for α and the number of centerpoint replicates, see Box and Draper (2007, Section 15.3). Box and Wilson (1951) also suggested a composite design where the augmenting runs are placed outside one corner of the factorial design. The notion of such asymmetric and noncentral composite designs for estimating the second-order model is further explored in Lucas (1974) and Mee (2001).

7.10 Regular Fractions for More Than Two Levels

When the experimental budget is tight so that the number of observations is limited, economical experiments are obtained by using fractions with all factors limited to two levels.

There are situations, however, in which more than two levels are needed as, for example, when a qualitative factor has a number of features that need to be compared, such as *batch of wafers* (with several distinct batches) or *susceptor rotation method* (clockwise, oscillating, fixed) in integrated-circuit fabrication (e.g., Shoemaker et al. 1991). In this section, we give a brief discussion of regular fractions with factors at more than two levels.

Suppose, first, that all factors have three levels. Then all main effects have 2 df, which can be represented by any pair of orthogonal contrasts between the three levels, as in (7.4), for example. The main effect of a factor is estimable if both contrasts are estimable, which implies that all linear combinations of these contrasts are also estimable. Any pair of estimable orthogonal contrasts can be included as columns of the model matrix \bar{X} in (7.3). Similarly, each two-factor interaction has $2 \times 2 = 4$ df, and an interaction is estimable if there exists a set of four orthogonal and estimable interaction contrasts (which can then be included in the model); three-factor interactions have $2 \times 2 \times 2 = 8$ df; four-factor interactions have 16 df, etc.

Regular s -level fractions, with s prime, have numbers of runs being a power of s , and similar to the two-level setting, we refer to s^{k-q} fractions. Also, similar to the 2-level case, defining relations and defining contrast subgroups of s^{k-q} fractions, with s prime, contain the identity I and $s^q - 1$ words that are generated by q generating words. The words in the defining contrast subgroup fall into $(s^q - 1)/(s - 1)$ sets of size $s - 1$, each set representing $s - 1$ df. For example, the defining relation of the resolution III 3^{4-2} design in Chapter 1, Section 1.7.1 is generated by the $q = 2$ words ABC^2 and ACD , and the full set of 3^2 words in the defining relation is

$$\begin{aligned} I &= ABC^2 = A^2B^2C \\ &= ACD = A^2BD = B^2C^2D \\ &= A^2C^2D^2 = BCD^2 = AB^2D^2. \end{aligned} \tag{7.21}$$

The words in (7.21) fall into the $(3^2 - 1)/(3 - 1)$ pairs (ABC^2, A^2B^2C) , $(ACD, A^2C^2D^2)$, (AB^2D^2, A^2BD) , and (BCD^2, B^2C^2D) , where the second word of each pair is the square of the first, with exponents reduced modulo 3. It is quite common to list only the first word in each pair, as was done in Chapter 1, with the understanding that it represents 2 df.

To understand how the generating words match the interaction contrasts, consider the pair $(A^{a_{1j}}B^{a_{2j}}C^{a_{3j}}D^{a_{4j}}, (A^{a_{1j}}B^{a_{2j}}C^{a_{3j}}D^{a_{4j}})^2)$. These represent any two orthogonal contrasts between three sets of treatment combinations, $(x_1x_2x_3x_4)$, satisfying $\sum_{i=1}^4 a_{ij}x_i = 0 \pmod 3$, $\sum_{i=1}^4 a_{ij}x_i = 1 \pmod 3$, and $\sum_{i=1}^4 a_{ij}x_i = 2 \pmod 3$, where the a_{ij} are the exponents of A, B, C, D in the j th generating word. Consider the first pair of words (ABC^2, A^2B^2C) in (7.21), for example, which represent any pair of orthogonal contrasts between three sets of treatment combinations, each set satisfying one of the three equations of (7.22):

$$\begin{aligned} x_1 + x_2 + 2x_3 + 0x_4 &= 0 \pmod 3, \\ x_1 + x_2 + 2x_3 + 0x_4 &= 1 \pmod 3, \\ x_1 + x_2 + 2x_3 + 0x_4 &= 2 \pmod 3, \end{aligned} \tag{7.22}$$

where we have taken the a_{ij} from the exponents of ABC^2 , but those from A^2B^2C would result in the same division of treatment combinations. Similarly, the second pair of generating

words $(ACD, A^2C^2D^2)$ represents a pair of orthogonal contrasts between three sets of treatment combinations, each set satisfying one of the three equations of (7.23):

$$\begin{aligned} x_1 + 0x_2 + x_3 + x_4 &= 0 \pmod 3, \\ x_1 + 0x_2 + x_3 + x_4 &= 1 \pmod 3, \\ x_1 + 0x_2 + x_3 + x_4 &= 2 \pmod 3. \end{aligned} \tag{7.23}$$

If we take the treatment combinations for the fraction satisfying, say, the first equation of both (7.22) and (7.23), we will have a 3^{4-2} fraction with defining relation as in (7.21) and the 9 treatment combinations

$$0000, 0112, 0221, 1011, 1120, 1202, 2022, 2101, 2210. \tag{7.24}$$

Any other pair of rows from (7.22) and (7.23) could be chosen instead, resulting in a different set of nine treatment combinations for the fraction but the same defining relation. These alternative fractions can also be obtained by adding (modulo 3) a treatment combination that is not in the fraction (7.24) to all of the treatment combinations that are in the fraction. For any of these fractions, all columns in the model matrix X for contrasts corresponding to the words in the defining relation will be a constant multiple of the column I of all 1's.

In general, we can use the same idea to obtain s^{k-q} fractions for any prime number of levels s by choosing q generating words. The defining contrast subgroup is of size s^q and is generated from the q generating words each raised to powers $u = 1, 2, \dots, s - 1$ together with their products, with exponents reduced modulo s . The treatment combinations for the fraction are those that satisfy equations of the form $\sum_{i=1}^k a_{ij}x_i = 0 \pmod s$, where a_{ij} is the exponent of the i th letter in the j th generating word ($i = 1, \dots, k, j = 1, \dots, q$). Alternative fractions with the same defining relation can be obtained as stated earlier.

Although a similar procedure can be used for s nonprime, care needs to be taken in the choice of generating words, so that all exponents are relatively prime to s . For example, suppose that a 4^{4-1} fraction is required and that the $q = 1$ generator $AB^2C^3D^2$ is chosen with two of the exponents not relatively prime to 4. The defining contrast subgroup contains $AB^2C^3D^2$ raised to every power $u = 1, 2, 3$ and exponents reduced modulo 4, so the defining relation is

$$I = AB^2C^3D^2 = A^2C^2 = A^3B^2CD^2,$$

and the four-factor interaction generating word results in a loss of 1 df for the two-factor interaction F_1F_3 . If generating word ABC^3D , for example, were chosen instead, then all exponents are relatively prime to 4 and the 3 df in the defining relation would belong to the four-factor interaction (i.e., $I = ABC^3D = A^2B^2C^2D^2 = A^3B^3CD^3$).

There are several alternative representations for obtaining s^{k-q} fractions for nonprime s ; one is to use the Galois fields when s is a power of a prime and another is to use pseudofactors as described in Section 7.1 that gives a larger class of designs if s has a prime powered divisor (see Bailey 1977, 1985; and Voss and Dean 1987, for more information).

For asymmetric designs (i.e., factors having different numbers of levels, s_1, s_2, \dots, s_k), the same ideas can be used, but the modulus arithmetic has to be done as though all factors have $m = LCM(s_1, s_2, \dots, s_k)$ levels, where LCM is the lowest common multiple

of s_1, s_2, \dots, s_k . Correspondingly, the a_{ij} are multiplied by m/s_i , $i = 1, \dots, k$, (cf. Dean and John 1975; and Bailey 1977). For example, for a $2^{3-1} \times 3^{2-1}$ fraction with three two-level factors, two three-level factors, and 12 treatment combinations, suppose that we choose the single generator $ABCD^2E$ for the defining contrast subgroup. Multiplying the exponents a_{i1} by m/s_i and reducing modulo $m = 6$, the defining relation will be

$$I = A^3B^3C^3D^4E^2 = D^2E^4 = A^3B^3C^3 = D^4E^2 = A^3B^3C^3D^4E^2,$$

and again, we see that a four-factor word has led to inclusion of lower-order interactions in the defining relation. Since $m = 6$, the $2^{2-1} \times 3^{2-1}$ treatment combinations satisfy

$$3x_1 + 3x_2 + 3x_3 + 4x_4 + 2x_5 = 0 \pmod{6},$$

that is,

$$00000, 00011, 00022, 11000, 11011, 11022, 10100, 10111, 10122, 01100, 01111, 01122.$$

There are few catalogs of designs for regular fractions with factors at more than two levels. Chen et al. (1993) gave a complete catalog of 27-run three-level fractions (which have at most 12 factors). Xu (2005) presented a catalog of all 81-run three-level fractions (with a maximum of 40 factors), 243-run three-level fractions of resolution IV and higher (with a maximum of 20 factors), and 729-run three-level fractions of resolution V and higher (with a maximum of 14 factors), as well as more information about the properties of the 27-run fractions. A subset of 3^{k-q} fractions is listed by Wu and Hamada (2009), as well as some fractions of $2^k \times 4$ and $2^k \times 4^2$ factorials.

7.11 Equivalence of Designs

In Section 7.3, we stated the fact that all saturated, regular, resolution III two-level fractions of size n are *isomorphic* in that the design matrix of one can be obtained from that of another by interchanging rows (treatment combinations), interchanging columns (factor names), and reversing the levels for one or more factors. In Sections 7.4 through 7.6, we described properties and classes of nonisomorphic designs of various resolution. In general, isomorphic designs have identical statistical properties under the model being fitted, and therefore, only nonisomorphic designs need be evaluated when searching for optimal designs or when preparing a design catalog; see, for example, the catalogs of nonisomorphic regular fractional factorial designs by Chen et al. (1993), Xu (2005, 2009), and Lin and Sitter (2008).

Unless there are theoretical results that can be used to cut down the search, checking for isomorphism is computationally intensive and a number of methods have been proposed for reducing the computational burden. A naïve check for isomorphism of two 2^{k-q} fractional factorial design matrices could require, in the worst case, an examination of up to $k!(2!)^k(2^{k-q}!)$ permutations of one design matrix to compare with the other, and Shrivastava and Ding (2010) calculated $\binom{2^{k-q}-1}{q} - \binom{k-q}{q}$ possible pairs of designs that need to be compared.

There are two ways to reduce the size of the isomorphism check. The first is to partition the set of all possible designs into small pools of possibly isomorphic designs, with designs in different pools known to be nonisomorphic, and the second way is to reduce the number of permutations that need to be examined. The first is accomplished either by using theoretical results or by examining the basic properties of the designs. For example, designs with different row coincidence distributions (Section 7.6), or different row coincidence moments (7.17), are clearly nonisomorphic, as are designs with different values under the criteria defined in Section 7.4.3. Then only designs that match on all of the basic properties remain in the same pool and need to be examined for isomorphism. A variety of other design checks for dividing the pools of possibly isomorphic designs were reviewed by Katsaounis and Dean (2008), Lin and Sitter (2008), Xu (2009), and Shrivastava and Ding (2010).

Algorithms for detecting isomorphic regular two-level fractions among those remaining in the candidate pools either operate directly on permutations of the design matrix or use the result of Chen (1992) that two-level regular fractional factorial designs are isomorphic if and only if the defining relation of one can be transformed to the other by a permutation of factor labels and reordering of words. Fast algorithms using the latter approach have been proposed by Xu (2009) and Shrivastava and Ding (2010).

When factors have more than two levels, there are two different types of isomorphism depending upon whether factors are qualitative or quantitative. For example, when modeling the effect of a three-level qualitative factor (say, *color*), it is irrelevant which of the levels (say, *red, blue, green*) are labeled 1, 2, and 3, since pairwise comparisons among their effects are not altered by the labeling. So there are $3!$ different permutations of levels of every three-level factor that result in an isomorphic design. Building on our earlier definition, designs with qualitative factors are said to be *combinatorially isomorphic* if one can be obtained from the other by (1) reordering the treatment combinations (rows of the design matrix), (2) relabeling the factors (reordering the columns) and (3) relabeling the levels of one or more factors.

On the other hand, if a three-level factor is quantitative, such as *temperature*, then the levels have a definite ordering (low, middle, high) and a design for fitting a second-order polynomial with the labeling 1, 2, 3, for three temperatures would not necessarily have the same properties as the design with the labeling 1, 3, 2. (To see this, consider how the graph of a straight line changes when the middle and high levels of the independent variable are interchanged.) A simple *reversing* of the labeling, however, would not change the nature of the fitted polynomial (other than switching signs of some of the polynomial coefficients). Using the terminology of Cheng and Ye (2004), two designs for quantitative factors are called *geometrically isomorphic* if one can be obtained from the other by (1) and (2) stated earlier and (3) *reversing* the levels of one or more factors.

When all factors have two levels, the two types of isomorphism coincide since the only possible relabeling of a factor's levels is a reversal. For factors with more than two levels, geometrically equivalent designs must also be combinatorially equivalent, since reversal of levels is one possible relabeling. Combinatorially equivalent designs, however, may or may not be geometrically equivalent.

The algorithm of Clark and Dean (2001), based on the Hamming distances of the treatment combinations, can detect isomorphic designs with qualitative or quantitative factors two or more levels, as well as identifying the permutations that transform one design into the other (see, e.g., Katsaounis et al. 2007; Katsaounis and Dean 2008; Katsaounis 2012). The algorithm of Cheng and Ye (2004) can detect geometrically isomorphic designs. Both of these algorithms are designed to handle fractions in general, not necessarily regular fractions or orthogonal arrays, and consequently, they can be slow. For the two-level regular

case, the algorithms mentioned earlier are faster. For factors at two or more levels, a fast method of generation of combinatorially nonequivalent orthogonal arrays, which include regular fractions, is given by Schoen et al. (2009) (see also references therein).

7.12 Incorporation of Noise Factors

A common requirement in industrial research and development is to determine suitable operating levels of *control* experimental factors (those that can be set by the manufacturer or operator such as automobile spark plug gap, factory conveyor belt speed, and concentration of chemical) that will result in a *robust* product or process. Here, *robust* means that a measured response is stable in the presence of fluctuations of nuisance environmental or uncontrollable variables (such as ambient temperature, humidity, driving speed, and incorrect user-controlled machine settings); see, also, Chapter 1, Section 1.11.5.

Even though nuisance variables cannot be controlled in normal operating conditions, they can sometimes be controlled in the laboratory that allows them to be incorporated into an experiment, where they are known as *noise* factors. Instead of, or in addition to, the average response, the effect of interest in such experiments is a function of the *dispersion*, or variance, of the response at each control factor combination as the noise factors change and the objective is to minimize the dispersion (cf. the first example in Section 7.7).

The focus on use of experimental design to minimize dispersion was strongly advocated by Genichi Taguchi, former director of the Japanese Academy of Quality, who first visited AT&T Bell Laboratories Quality Assurance Center in the United States in 1980. Taguchi called such experimentation *parameter design* to distinguish it from engineering design; see Taguchi (1990) for more information on his philosophy. During the 1980s, statisticians responded to the suggested change in focus by developing new methods of designing experiments as well as several approaches to their analysis. Reviews are given by Ankenman and Dean (2003) and Bursztyn and Steinberg (2006), and an interesting panel discussion on a variety of related issues can be found in Nair et al. (1992). One such method uses regular fractional factorial experiments to measure dispersion through interactions between control and noise factors. A large control \times noise interaction indicates that the settings of the control factors involved in the interaction have an effect on how much the response changes with noise factor fluctuations. These control factors then give scope for achieving a robust product through careful choice of their settings. Examples of experiments to determine robust control factor settings are given by, for example, Shoemaker et al. (1991); Engel (1992); Tuck et al. (1993); and Vine et al. (2008); see also Section 15.7 of Dean and Voss (1999) and Section 11.5 of Wu and Hamada (2009).

In such experiments, estimation of control \times noise interactions is of highest priority, and the criterion of minimum aberration used in Sections 7.3 through 7.5, which weights all interactions equally, no longer applies. Instead, a criterion that gives more prominence to the control \times noise interactions than to control \times control or noise \times noise interactions is required. Generators for regular fractions that involve no aliasing of control \times noise interactions with main effects or other two-factor interactions are presented by Borkowski and Lucas (1997) and Russell et al. (2004); see also Ke et al. (2005). These designs are sometimes known as *compromise plans* or *mixed resolution designs*. Nonregular designs based on orthogonal arrays (Chapter 9) for ensuring the estimation of control \times noise interactions are the *crossed arrays* or *direct product arrays* of Taguchi and more recent improvements

such as the *compound arrays* of Rosenbaum (1994, 1996), Hedayat and Stufken (1999), and Zhu et al. (2007).

7.13 Summary and Discussion

Regular fractional factorial designs are prevalent in research and development in industry as well as having widespread use in other areas such as engineering, drug discovery, consumer and behavior research, genetics, and health-care research. In this chapter, we have discussed in some detail the construction and properties of regular two-level fractions that achieve optimality under criteria such as minimum aberration, estimation capacity, and number of *clear* interactions (Sections 7.3 through 7.6). We also described, in Section 7.10, a construction method for regular fractions with more than two levels. Analysis of fractional factorial experiments, screening for active factors, follow-up experiments to resolve confounding, design equivalence, and incorporation of noise factors for measuring dispersion effects have also been discussed (Sections 7.7 through 7.9, and 7.11, and 7.12).

In general, in order to avoid biases, the runs of a fractional factorial design should be randomly ordered before they are observed. However, there are situations in which practical operation necessitates restrictions on randomization (see also Chapter 1, Section 1.3.5). For example, if it is costly to change the level of one or more factors, then an order of observation that reduces the number of factor level changes would be advantageous. Similarly, if the response is affected by uncontrollable variables that change over space or time, it is desirable to use an arrangement of factorial treatment combinations that ensures treatment estimates are independent of any such trends. Overview of work on trend-free arrangements of two-level fractions, including cost-optimal arrangements, published before 1990 has been given by Cheng (1990) and Jacroux (1990). Other relevant work includes that of Mount-Campbell and Neuhardt (1980); Cheng and Steinberg (1991); Bailey et al. (1992); Coster (1993a,b); Cui and John (1998); and Mee and Romanova (2010), among others. When the levels of one or more factors are held constant for a large portion of the experiment before being changed, the resulting design is a *split-plot design*. Chapter 8 gives constructions of fractional factorial experiments run as block designs and split-plot designs (see Chapter 1, Section 1.5).

Fractional factorial experiments in other contexts are presented in other chapters. For example, Chapter 9 discusses the construction of nonregular fractions, orthogonal arrays, and supersaturated designs, and there, the concept of minimum aberration is extended to the non-regular case. Chapter 21 discusses algorithms that can be used for searching for factorial designs that do not necessarily have a well-defined combinatorial structure such as those in this chapter. Finally, Chapter 22 includes the case of factorial experiments in the discrete choice setting.

References

- Ai, M., Xu, X., and Wu, C. (2010), Optimal blocking and foldover plans for regular two-level designs, *Statistica Sinica*, 20, 183–207.
- Angelopoulos, P., Koukouvinos, C., and Mee, R. (2008), Near-saturated projection designs from Hadamard matrices, *Journal of Statistical Theory and Applications*, 7, 141–150.

- Ankenman, B. E. and Dean, A. M. (2003), Quality improvement and robustness via design of experiments, in *Handbook of Statistics*, eds. Rao, C. R. and Khattree, R., Elsevier, Amsterdam, the Netherlands, vol. 22, pp. 263–317.
- Bailey, R. A. (1977), Patterns of confounding in factorial designs, *Biometrika*, 64, 597–604.
- Bailey, R. A. (1985), Factorial design and abelian groups, *Linear Algebra and Its Applications*, 70, 349–368.
- Bailey, R. A., Cheng, C.-S., and Kipnis, P. (1992), Construction of trend-resistant factorial designs, *Statistica Sinica*, 2, 393–411.
- Barnett, J., Czitrom, V., John, P. W. M., and Leon, R. V. (1997), Using fewer wafers to resolve confounding in screening experiments, in *Statistical Case Studies for Industrial Process Improvement*, eds. Czitrom, V. and Spagon, P., SIAM, Philadelphia, PA, pp. 235–250.
- Block, R. M. (2003), Theory and construction methods for large regular resolution IV designs, PhD thesis, University of Tennessee, Knoxville, TN. <http://etd/utk.edu/2003/BlockRobert.pdf>.
- Block, R. M. and Mee, R. W. (2003), Second order saturated resolution IV designs, *Journal of Statistical Theory and Applications*, 2, 96–112.
- Borkowski, J. J. and Lucas, J. M. (1997), Designs of mixed resolution for process robustness studies, *Technometrics*, 39, 63–70.
- Box, G. E. P. and Draper, N. R. (2007), *Response Surfaces, Mixtures, and Ridge Analysis*, 2nd edn., Wiley, Hoboken, NJ.
- Box, G. E. P. and Wilson, K. B. (1951), On the experimental attainment of optimum conditions, *Journal of the Royal Statistical Society, Series B*, 13, 1–45.
- Bruen, A. and Wehlau, D. (1999), Long binary linear codes and large caps in projective space, *Designs, Codes and Cryptography*, 17, 37–60.
- Bursztyń, D. and Steinberg, D. M. (2006), Screening experiments for dispersion effects, in *Screening: Methods for Experimentation in Industry, Drug Discovery and Genetics*, eds. Dean, A. M. and Lewis, S. M., Springer, New York, pp. 191–206.
- Butler, N. A. (2003a), Some theory for constructing minimum aberration fractional factorial designs, *Biometrika*, 90, 233–238.
- Butler, N. A. (2003b), Minimum aberration construction results for nonregular two-level fractional factorial designs, *Biometrika*, 90, 891–898.
- Chen, H. and Hedayat, A. S. (1996), 2^{n-1} designs with weak minimum aberration, *Annals of Statistics*, 24, 2536–2548.
- Chen, H. and Hedayat, A. S. (1998), 2^{n-m} designs with resolution III or IV containing clear two-factor interactions, *Journal of Statistical Planning and Inference*, 75, 147–158.
- Chen, H. H. and Cheng, C.-S. (2004), Aberration, estimation capacity and estimation index, *Statistica Sinica*, 14, 203–215.
- Chen, H. H. and Cheng, C.-S. (2006), Doubling and projection: A method of constructing two-level designs of resolution IV, *Annals of Statistics*, 34, 546–558.
- Chen, J. (1992), Some results on 2^{n-k} fractional factorial designs and search for minimum aberration designs, *Annals of Statistics*, 20, 2124–2141.
- Chen, J., Sun, D. X., and Wu, C. F. J. (1993), A catalogue of two-level and three-level fractional factorial designs with small runs, *International Statistical Review*, 61, 131–145.
- Cheng, C.-S. (1990), Construction of run orders of factorial designs, in *Statistical Design and Analysis of Industrial Experiments*, ed. Ghosh, S., Marcel Dekker Inc., New York, vol. 109, pp. 423–439.
- Cheng, C.-S., Mee, R. W., and Yee, O. (2008), Second order saturated orthogonal arrays of strength three, *Statistica Sinica*, 18, 105–119.
- Cheng, C.-S. and Steinberg, D. M. (1991), Trend robust two-level factorial designs, *Biometrika*, 78, 325–336.
- Cheng, C.-S., Steinberg, D. M., and Sun, D. X. (1999), Minimum aberration and model robustness for two-level fractional factorial designs, *Journal of the Royal Statistical Society, Series B*, 61, 85–93.
- Cheng, S. W. and Ye, K. (2004), Geometric isomorphism and minimum aberration for factorial designs with quantitative factors, *Annals of Statistics*, 32, 2168–2185.

- Clark, J. B. and Dean, A. M. (2001), Equivalence of fractional factorial designs, *Statistica Sinica*, 11, 537–547.
- Coster, D. C. (1993a), Tables of minimum cost, linear trend-free run sequences for two- and three-level fractional factorial designs, *Computational Statistics and Data Analysis*, 16, 325–336.
- Coster, D. C. (1993b), Trend-free run orders of mixed-level fractional factorial designs, *The Annals of Statistics*, 21, 2072–2086.
- Cui, X. and John, P. W. M. (1998), Time-trend free run orders with the minimum level changes, *Communications in Statistics: Theory and Methods*, 27, 55–68.
- Dante, R. C., Escamilla, J. L., Madrigal, V., Theuss, T., Calderon, J. D. D., Solorza, O., and Rivera, R. (2003), Fractional factorial design of experiments for PEM fuel cell performances improvement, *International Journal of Hydrogen Energy*, 28, 343–348.
- Dean, A. M. and Lewis, S. M. (eds.) (2006), *Screening: Methods for Experimentation in Industry, Drug Discovery and Genetics*, Springer, New York.
- Dean, A. and Voss, D. (1999), *Design and Analysis of Experiments*, Springer-Verlag Inc., New York.
- Dean, A. M. and John, J. A. (1975), Single replicate factorial experiments in generalized cyclic designs: II. asymmetrical arrangements, *Journal of the Royal Statistical Society, Series B*, 37, 72–76.
- Dong, F. (1993), On the identification of active contrasts in unreplicated fractional factorials, *Statistica Sinica*, 3, 209–217.
- Dorfman, R. (1943), The detection of defective members of large populations, *Annals of Mathematical Science*, 14, 436–440.
- Draguljić, D., Woods, D. C., Dean, A. M., Lewis, S. M., and Vine, A. J. (2014), Screening strategies in the presence of interactions, *Technometrics*, 56, 1–16.
- Draper, N. R. and Mitchell, T. J. (1967), The construction of saturated 2_R^{k-p} designs, *Annals of Mathematical Statistics*, 38, 1110–1126.
- Edwards, D. J. (2011), Optimal semifoldover plans for two-level orthogonal designs, *Technometrics*, 53, 274–284.
- Edwards, D. J. and Mee, R. W. (2008), Empirically determined p-values for Lenth t statistics, *Journal of Quality Technology*, 40, 368–380.
- Engel, J. (1992), Modelling variation in industrial experiments, *Applied Statistics*, 41, 579–593.
- Franklin, M. F. (1984), Constructing tables of minimum aberration p^{n-m} designs, *Technometrics*, 26, 225–232.
- Fries, A. and Hunter, W. G. (1980), Minimum aberration 2^{n-k} designs, *Technometrics*, 22, 601–608.
- Grömping, U. (1974), A note on dominating fractional factorial two-level designs with clear two-factor interactions, *Technometrics*, 56, 42–45.
- Haaland, P. and O’Connell, M. (1995), Inference for effect-saturated fractional factorials, *Technometrics*, 37, 82–93.
- Hamada, M. and Balakrishnan, N. (1998), Analyzing unreplicated factorial experiments: A review with some new proposals, *Statistica Sinica*, 8, 1–28.
- Hedayat, A. S., Sloane, N. J. A., and Stufken, J. (1999), *Orthogonal Arrays*, Springer-Verlag, New York.
- Hedayat, A. S. and Stufken, J. (1999), Compound orthogonal arrays, *Technometrics*, 41, 57–61.
- Irvine, G., Clark, N., and Recupero, C. (1996), Extended delignification of mature and plantation eucalypt wood. 2. The effects of chip impregnation factors, *Appita Journal*, 49, 347–352.
- Jacroux, M. (1990), Methods for constructing trend-resistant run orders of 2-level factorial experiments, in *Statistical Design and Analysis of Industrial Experiments*, ed. Ghosh, S., Marcel Dekker Inc., New York, vol. 109, pp. 441–457.
- Jaynes, J., Ding, X., Xu, H., Wong, W. K., and Ho, C.-M. (2012), Application of fractional factorial designs to study drug combinations, *Statistics in Medicine*, 32, 307–318.
- Jones, B. and DuMouchel, W. (1996), Follow-up designs to resolve confounding in multifactor experiments. Discussion, *Technometrics*, 38, 323–326.
- Katsaounis, T. I. (2012), Equivalence of factorial designs with qualitative and quantitative factors, *Journal of Statistical Planning and Inference*, 142, 79–85.

- Katsaounis, T. I. and Dean, A. M. (2008), A survey and evaluation of methods for determination of combinatorial equivalence of factorial designs, *Journal of Statistical Planning and Inference*, 138, 245–258.
- Katsaounis, T. I., Dingus, C. A., and Dean, A. M. (2007), Determination of geometric equivalence of symmetric factorial designs, *Journal of Statistical Theory and Practice*, 1, 101–115.
- Ke, W. and Tang, B. (2003), Selecting 2^{m-p} designs using a minimum aberration criterion when some two-factor interactions are important, *Technometrics*, 45, 352–360.
- Ke, W., Tang, B., and Wu, C. F. J. (2005), Compromise plans with clear two-factor interactions, *Statistica Sinica*, 15, 709–715.
- Lenth, R. V. (1989), Quick and easy analysis of unreplicated factorials, *Technometrics*, 31, 469–473.
- Lewis, S. M. and Dean, A. M. (2001), Detection of interactions in experiments on large numbers of factors (with discussion), *Journal of the Royal Statistical Society, B*, 63, 633–672.
- Li, H. and Mee, R. W. (2002), Better foldover fractions for resolution III designs, *Technometrics*, 44, 278–283.
- Li, W. and Lin, D. K. J. (2003), Optimal foldover plans for two-level fractional factorial designs, *Technometrics*, 45, 142–149.
- Lin, C. D. and Sitter, R. R. (2008), An isomorphism check for two-level fractional factorial designs, *Journal of Statistical Planning and Inference*, 138, 1085–1101.
- Lin, X., Sudarsanam, N., and Frey, D. (2006), Regularities in data from factorial experiments, *Complexity*, 11, 32–45.
- Lucas, J. M. (1974), Optimum composite designs, *Technometrics*, 16, 561–567.
- Mee, R. W. (2001), Noncentral composite designs, *Technometrics*, 43, 34–43.
- Mee, R. W. (2009), *A Comprehensive Guide to Factorial Two-Level Experimentation*, Springer, New York.
- Mee, R. W. and Peralta, M. (2000), Semifolding 2^{k-p} designs, *Technometrics*, 42, 122–134.
- Mee, R. W. and Romanova, A. V. (2010), Constructing and analyzing two-level trend-robust designs, *Quality Engineering*, 22, 306–316; erratum 23, 112.
- Mee, R. W. and Xiao, J. (2008), Optimal foldovers and semifolding for minimum aberration even fractional factorial designs, *Journal of Quality Technology*, 40, 448–460.
- Morris, M. D. (2006), An overview of group factor screening, in *Screening: Methods for Experimentation in Industry, Drug Discovery and Genetics*, eds. Dean, A. M. and Lewis, S. M., Springer, New York, pp. 191–206.
- Mount-Campbell, C. A. and Neuhardt, J. B. (1980), Selecting cost-optimal main-effect fractions of 3^n factorials, *AIIE Transactions*, 12, 80–86.
- Mukerjee, R. and Wu, C. F. J. (2006), *A Modern Theory of Factorial Designs*, Springer, New York.
- Nair, V. N., Abraham, B., MacKay, J., Box, G., Kacker, R., Lorenzen, T., Lucas, J., Myers, R., Vining, G., Nelder, J., Phadke, M., Sacks, J., Welch, W., Shoemaker, A., Tsui, K., Taguchi, S., Wu, C. F. J. (1992), Taguchi's parameter design: A panel discussion, *Technometrics*, 34, 127–161.
- Rosenbaum, P. R. (1994), Dispersion effects from fractional factorials in Taguchi's method of quality design, *Journal of the Royal Statistical Society, Series B*, 56, 641–652.
- Rosenbaum, P. R. (1996), Some useful compound dispersion experiments in quality design, *Technometrics*, 38, 354–364.
- Russell, K. G., Lewis, S. M., and Dean, A. M. (2004), Fractional factorial designs for the detection of interactions between design and noise factors, *Journal of Applied Statistics*, 31, 545–552.
- Schmidt, S. R. and Launsby, R. G. (eds.) (1992), *Understanding Industrial Designed Experiments*, Air Academy Press, Colorado Springs, CO.
- Schoen, E. D., Eendebak, P. T., and Nguyen, M. V. M. (2009), Complete enumeration of pure-level and mixed-level orthogonal arrays, *Journal of Combinatorial Designs*, 18, 123–140.
- Schoen, E. D. and Kaul, E. A. A. (2000), Three robust scale estimators to judge unreplicated experiments, *Journal of Quality Technology*, 32, 276–283.
- Shoemaker, A. C., Tsui, K., and Wu, C. F. J. (1991), Economical experimentation methods for robust design, *Technometrics*, 33, 415–427.
- Shrivastava, A. K. and Ding, Y. (2010), Graph based isomorph-free generation of two-level regular fractional factorial designs, *Journal of Statistical Planning and Inference*, 140, 169–179.

- Taguchi, G. (1990), Experimental design for product design, in *Statistical Design and Analysis of Industrial Experiments*, ed. Ghosh, S., Marcel Dekker Inc, New York, pp. 1–33.
- Tang, B. and Wu, C. F. J. (1996), Characterization of minimum aberration 2^{n-k} designs in terms of their complementary designs, *Annals of Statistics*, 24, 2549–2559.
- Tuck, M. G., Lewis, S. M., and Cottrell, J. I. L. (1993), Response surface methodology and Taguchi: A quality improvement study from the milling industry, *Applied Statistics*, 42, 671–681.
- Vine, A. E., Lewis, S. M., and Dean, A. M. (2005), Two-stage group screening in the presence of noise factors and unequal probabilities of active effects, *Statistica Sinica*, 15, 871–888.
- Vine, A. E., Lewis, S. M., Dean, A. M., and Brunson, D. (2008), A critical assessment of two-stage group screening through industrial experimentation, *Technometrics*, 50, 15–25.
- Voss, D. T. and Dean, A. M. (1987), A comparison of classes of single replicate factorial designs, *Annals of Statistics*, 15, 376–384.
- Watson, G. S. (1961), A study of the group screening method, *Technometrics*, 3, 371–388.
- Wu, C. F. J. and Chen, Y. (1996), A graph-aided method for planning two-level experiments when certain interactions are important, *Technometrics*, 34, 162–175.
- Wu, C. F. J. and Hamada, M. (2009), *Experiments: Planning, Analysis, and Parameter Design Optimization*, 2nd edn., Wiley, New York.
- Wu, H., Mee, R., and Tang, B. (2012), Fractional factorial designs with admissible sets of clear two-factor interactions, *Technometrics*, 54, 191–197.
- Xu, H. (2005), A catalogue of three-level regular fractional factorial designs, *Metrika*, 62, 259–281.
- Xu, H. (2009), Algorithmic construction of efficient fractional factorial designs with large run sizes, *Technometrics*, 51, 262–277.
- Xu, H. and Cheng, C.-S. (2008), A complementary design theory for doubling, *Annals of Statistics*, 36, 445–457.
- Xu, H. and Wu, C. F. J. (2001), Generalized minimum aberration for asymmetrical fractional factorial designs, *Annals of Statistics*, 29, 1066–1077.
- Yates, F. (1937), *The Design and Analysis of Factorial Experiments*, Imperial Bureau of Soil Science, Harpenden, England.
- Zhu, Y., Zeng, P., and Jennings, K. (2007), Optimal compound orthogonal arrays and single arrays for robust parameter design experiments, *Technometrics*, 49, 440–453.

8

Multistratum Fractional Factorial Designs

Derek Bingham

CONTENTS

8.1 Introduction.....	321
8.2 Motivating Examples.....	322
8.3 Notation and Background.....	325
8.4 Full Factorial Designs with Randomization Restrictions.....	327
8.5 Impact on Data Analysis.....	329
8.6 Fractional Factorial Designs with Randomization Restrictions.....	332
8.7 Choosing Designs.....	333
8.8 Concluding Remarks.....	336
References.....	337

8.1 Introduction

In the early stages of investigation, particularly in industrial applications, screening experiments are often used to identify the few important factors affecting a system's response. To this end, fractional factorial designs (Chapter 7) are common choices in settings that aim to investigate relatively many factors in relatively few trials. In many cases, such experiments are performed as completely randomized designs, and data analysis is conducted via regression and ANOVA approaches (e.g., Dean and Voss 1999; Montgomery 2008; Wu and Hamada 2011) or through the use of half-normal probability plots (Daniel 1959). However, complete randomization of the experiment trials may be impossible due to practical or cost constraints. Consequently, restrictions on the randomization are imposed, thereby complicating both the design and the analysis of the experiment.

Experimental plans such as block (see, e.g., Bisgaard 1994; Sitter et al. 1997; Cheng et al. 2004), split-plot (Addelman 1964; Huang et al. 1998; Bisgaard 2000; Bingham and Sitter 2001; Vine et al. 2008), strip-plot (Miller 1997), and split-plot (Mee and Bates 1998; Butler 2004) designs have all been adapted to the fractional factorial settings. In spite of the common treatment structure, a quick glance at this literature reveals very different approaches for design construction to accommodate the restricted randomization. It has been demonstrated that these designs can be viewed in a unified manner. In particular, Bingham et al. (2008) showed that each of the aforementioned designs can be constructed using an extension of the usual approach for blocking in fractional factorial designs. Furthermore, Cheng and Tsai (2011) demonstrated that the construction of regular fractional factorial designs with randomization restrictions fits within the theory of block structures proposed by Nelder (1965) and extended by Speed and Bailey (1982) as well as Tjur (1984).

In this chapter, a review of the design of fractional factorial experiments with randomization restrictions is presented. The framework used, to varying degrees, is most closely related to work by Bingham et al. (2008) and Cheng and Tsai (2011). The work presented here is related to other chapters in this volume—specifically, Chapters 7 and 10. This chapter is outlined as follows. A few motivating examples are presented in Section 8.2. Basic notation is introduced in Section 8.3, and Section 8.4 outlines the simple approach for imposing randomization restrictions on factorial designs. The impact of the randomization restrictions on the analysis of the experiment is considered in Section 8.5. Fractional factorial designs with randomization restrictions are presented in Section 8.6, and the ranking of designs is briefly discussed in Section 8.7. This chapter concludes with some final comments in Section 8.8.

8.2 Motivating Examples

Designs with randomization restrictions have a long history in agricultural experiments (see Federer and King 2007, for background). More recently, attention has turned to so-called multistratum designs in industrial settings (e.g., Miller 1997; Huang et al. 1998; Bisgaard 2000; Gilmour and Goos 2009; Jones and Goos 2009; Cheng and Tsai 2011). In this section, a few motivating examples are given to illustrate how randomization restrictions arise in factorial design settings.

Example 8.1 Blocking Full Factorial Designs

Montgomery (2008) outlines a setting where four factors (label them 1–4) are studied to determine their impact on a product filtration system. Each factor has 2 levels and the process must be run in 2 blocks of eight experiment trials. The design shown in Table 8.1

TABLE 8.1
2⁴ Factorial Design in 2 Blocks

1	2	3	4	$\delta = 1234$	Block
-1	-1	-1	1	-1	1
-1	-1	1	-1	-1	1
-1	1	-1	-1	-1	1
-1	1	1	1	-1	1
1	-1	-1	-1	-1	1
1	-1	1	1	-1	1
1	1	-1	1	-1	1
1	1	1	-1	-1	1
-1	-1	-1	-1	1	2
-1	-1	1	1	1	2
-1	1	-1	1	1	2
-1	1	1	-1	1	2
1	-1	-1	1	1	2
1	-1	1	-1	1	2
1	1	-1	-1	1	2
1	1	1	1	1	2

If the block effects are treated as fixed effects, then the significance of these effects cannot be separated from the impact of the blocks. If the block effects are treated as random effects, as will be done throughout this chapter, then it is sometimes possible to assess the significance of the confounded interactions. In either case, the factorial effects that are not confounded with block are not impacted by the blocks. This is known as *orthogonal blocking*.

Example 8.2 Fractional Factorial Split-Plot Design

Bingham and Sitter (1999) describe a fractional factorial split-plot experiment to identify factors that will help improve the efficiency of a ball mill. In their example, there are seven, 2-level factors to be investigated in 2^4 trials. So, a 2^{7-3} fractional factorial design was required. Complicating matters, however, were that factors 1–4 were expensive to adjust, and it was felt that only 8 unique settings of these factors would be considered. Furthermore, once one of these 8 settings was chosen, all treatments at that setting would be performed (in random order, of course) before moving on. The hard-to-change factors (1–4) are called *whole-plot* factors and, from a randomization perspective, serve the same role as blocking factors in the previous example. The easy-to-change factors (factors 5–7) are allowed to vary for fixed whole-plot settings and are called *subplot* factors. This design is an example of a 2^{7-3} fractional factorial split-plot design with 2^3 whole plots, each with two subplots.

The defining relation for the fraction in Bingham and Sitter’s illustration was $I = 123 = 456 = 157 = 123456 = 2357 = 1467 = 23467$, and the design is shown in Table 8.3. Notice that the 8 unique settings of factors 1–4 are grouped into blocks (or whole plots) of size 2. This is done to highlight how the randomization was performed; that is, a whole-plot setting is first randomly selected and then all trials with that setting are performed in random order before moving on to the next whole-plot setting. The analysis of such a design is slightly more complicated than that of a fractional factorial design because there are two types of error terms: whole plot and subplot. For example, the final two

TABLE 8.3
 2^{7-3} Fractional Factorial Split-Plot Design with 2^3 SubPlots of Size 2

1	2	3 = 12	4	5	6 = 45	7 = 15	Whole-Plot Error	SubPlot Error
-1	-1	1	-1	-1	1	1	$\epsilon_{1,1}$	$\epsilon_{2,1}$
-1	-1	1	-1	1	-1	-1	$\epsilon_{1,1}$	$\epsilon_{2,2}$
-1	-1	1	1	-1	-1	1	$\epsilon_{1,2}$	$\epsilon_{2,3}$
-1	-1	1	1	1	1	-1	$\epsilon_{1,2}$	$\epsilon_{2,4}$
-1	1	-1	-1	-1	1	1	$\epsilon_{1,3}$	$\epsilon_{2,5}$
-1	1	-1	-1	1	-1	-1	$\epsilon_{1,3}$	$\epsilon_{2,6}$
-1	1	-1	1	-1	-1	1	$\epsilon_{1,4}$	$\epsilon_{2,7}$
-1	1	-1	1	1	1	-1	$\epsilon_{1,4}$	$\epsilon_{2,8}$
1	-1	-1	-1	-1	1	-1	$\epsilon_{1,5}$	$\epsilon_{2,9}$
1	-1	-1	-1	1	-1	1	$\epsilon_{1,5}$	$\epsilon_{2,10}$
1	-1	-1	1	-1	-1	-1	$\epsilon_{1,6}$	$\epsilon_{2,11}$
1	-1	-1	1	1	1	1	$\epsilon_{1,6}$	$\epsilon_{2,12}$
1	1	1	-1	-1	1	-1	$\epsilon_{1,7}$	$\epsilon_{2,13}$
1	1	1	-1	1	-1	1	$\epsilon_{1,7}$	$\epsilon_{2,14}$
1	1	1	1	-1	-1	-1	$\epsilon_{1,8}$	$\epsilon_{2,15}$
1	1	1	1	1	1	1	$\epsilon_{1,8}$	$\epsilon_{2,16}$

columns of Table 8.3 display error terms that are associated with whole plots and subplot, respectively. Notice that each whole plot gets its own error term and so does each subplot. It turns out that the effects that change sign within whole-plot settings will have one error variance and the remaining effects will have another (see Section 8.5).

Still remaining is how the blocks of runs were chosen. The blocks were constructed using the following columns as *randomization restriction factors*: $\delta_1 = 1$, $\delta_1 = 2$, and $\delta_3 = 4$. The use of the randomization factors (the δ 's) is exactly the same as the blocking factors in the previous example. Letting Δ denote the identity column, the *randomization restriction defining relation (RRDR)* is $\Delta = 1\delta_1 = 2\delta_2 = 12\delta_1\delta_2 = 4\delta_3 = 14\delta_1\delta_3 = 24\delta_2\delta_3 = 124\delta_1\delta_2\delta_3$, and these effects will have a larger error variance than the remaining factorial effects (Section 8.5). The difference between a split-plot design and a blocked design is that in the former, the factors (or main effects) are used to specify blocks instead of interactions in the latter.

Example 8.3 Factorial Split-Lot Design

There is a wide variety of combinations of randomization restrictions that can arise in applications. Mee and Bates (1998), for example, considered a multistage split-lot experiment where some of the factors that are fixed at one stage of experimentation are allowed to vary at another stage. For this case, there were different sets of blocking factors to consider and different RRDRs for each stage. Consider the following simple illustration.

Suppose an industrial experiment that follows a 2^5 treatment structure is to be performed. Further suppose that at the first stage of processing, once the levels of factors 1 and 2 are set, 8 of the 32 experimental units are processed before moving on to the next level setting of these factors. At the next stage of processing, all of the 32 experimental units are available, but once the levels of factors 3 and 4 are fixed, all trials at this setting are run before choosing another setting for these factors. So, at the first stage of processing, there are randomization restrictions placed on factors 1 and 2; at the next stage, randomization restrictions are placed on factors 3 and 4; and no randomization restrictions are placed on factor 5. To run the experiment, one of the four setting combinations of factors 1 and 2 is randomly selected, and then 8 experimental units are processed in random order. Then a second setting combination of factors 1 and 2 is randomly chosen, and so on. After this first stage is completed, one randomly chooses one of the four setting combinations of factors 3 and 4, and each of the eight possible settings of factors 1, 2, and 5 appears within this block. The other three settings of factors 3 and 4 are processed similarly in random order. This is an example of a *factorial split-lot design*. To construct such a design, separate RRDRs are required for each stage of randomization. This example is revisited in Example 8.5 of Section 8.4 where the use of separate main effects and interactions for specifying randomization restrictions at different stages is formalized.

8.3 Notation and Background

A 2^k factorial design has k factors, each with 2 levels, and $n = 2^k$ trials that consist of all possible combinations of the treatment factors. Throughout this chapter, the level settings will be denoted ± 1 . Similarly, a 2^{k-q} regular fractional factorial design considers k 2-level factors, but in only $n = 2^{k-q}$ trials, and has an associated defining contrast subgroup (see Chapter 7). Nonregular fractional factorial designs (those without a defining

contrast subgroup will be discussed in Chapter 9) will not be considered in this chapter. The *treatment structure* of a design defines the set of treatments used in the experiment. The *randomization structure* of the design, on the other hand, speaks to the run order of the design—more on this shortly. However, it is worth noting that the restrictions in randomization that are discussed here form blocks of equal size and are based on the columns of the factorial model matrix (defined later).

A 2^k full factorial design consists of all possible treatment combinations of k factors at 2 levels each. For example, ignoring blocking for the moment, the rows of columns 1–4 in Table 8.1 give the 16 unique treatments for a 2^4 factorial design. These columns are used to measure the main effects of factors 1–4. Then p -factor interaction columns ($p = 2, \dots, k$) are constructed by taking the element-wise product of the p -factor columns. The $n \times n$ full factorial model matrix, X , is formed by including all main effect and interaction columns and augmenting them with a column of 1's for the grand mean. Without loss of generality, the columns of X can be denoted as $X = \{c_0, c_1, c_2, \dots, c_k, c_{k+1}, \dots, c_{n-1}\}$, with c_0 denoting the column of 1's for the grand mean, and main effect columns labelled c_1, \dots, c_k for factors 1, \dots , k . This matrix is referred to as the *model matrix*. From a design perspective, the main effect columns c_1, c_2, \dots, c_k of X form the $n \times k$ *design matrix*, D . That is, the rows of D identify the treatments used in the experiment.

The columns of the 2^{k-q} full factorial model matrix are used to construct the fractional factorial design. Specifically, the first $k - q$ factors are assigned to independent columns c_1, c_2, \dots, c_{k-q} of X and the remaining q factors are assigned to interaction columns. This specification of k factors to columns defines the treatments of the 2^{k-q} fractional factorial design. The first $k - q$ factors are called *basic factors* and the remaining q factors are referred to as *generated factors*. The interactions to which the generated factors are assigned are often called *factor generators* or just *generators*. Of course, not all assignments of the q factors to interaction columns are equally good due to the induced aliasing among the factorial effects. Consequently, designs are usually ranked using a criterion that emphasizes the estimation of lower-order effects (usually main effects and 2-factor interactions) such as the minimum aberration (Fries and Hunter 1980) or clear effects (Wu and Wu 2002) criteria (see Chapter 7).

The response model for a full factorial experiment is typically the multiple linear regression model:

$$y = X\beta + \epsilon, \quad (8.1)$$

where X is the $n \times 2^k$ model matrix, n is the number of trials, $\beta = (\beta_0, \beta_1, \dots, \beta_{2^k-1})$ denotes the regression parameters, and ϵ is the n -vector of independent and identically distributed (*iid*) normal random errors with mean zero and variance σ^2 .

In a fractional factorial experiment, there are more main effects and interactions than trials. In fact, there are only $2^{k-q} - 1$ factorial effects that can be estimated. Furthermore, all of these effects are aliased with other effects. To get around this issue, experimenters often appeal to a few assumptions that allow data analysis to proceed. Specifically, these are *effect sparsity* and *effect hierarchy*. The first assumption specifies that, relative to the run size of the experiment, the number of important effects impacting the response is small. The second assumption stipulates that lower-order effects are more likely to be important than higher-order effects and effects of the same order are equally likely to be active. In many cases, main effects and 2-factor interactions are the only effects that are considered.

An attractive feature of completely randomized 2-level factorial and fractional factorial designs with response model (8.1) is that the regression effect estimators are independent and normally distributed with the same variance. This fact, along with the effect sparsity assumption, allows for simple and straightforward data analyses. For example, half-normal probability plots (Daniel 1959) or more elaborate approaches (see, e.g., Hamada and Balakrishnan 1998, and Chapter 7) are often used to identify active lower-order effects.

8.4 Full Factorial Designs with Randomization Restrictions

In the context of full factorial and regular fractional factorial experiments, several different approaches have been proposed to study designs—indeed too many to list exhaustively here. Abelian groups and vector spaces have been used to elicit the properties of blocked factorial and fractional factorial experiments (e.g., Dean and John 1975; Bailey 1977; Collings 1989). Blocked designs in this context were also studied via finite geometries by Bose (1947) and later by, for example, Mukerjee et al. (1999). Regular split-plot designs were studied by Addelman (1964), Bingham and Sitter (1999), Bisgaard (2000), Huang et al. (1998), and Vine et al. (2008), to name only a few—all using apparently different approaches.

It has been shown that the construction of full factorial and regular fractional factorial designs with randomization restrictions falls under one unified framework (Bingham et al. 2008; Ranjan et al. 2009). Indeed, Cheng and Tsai (2011) showed that this approach fits squarely under the framework proposed by Nelder (1965) and extended by Speed and Bailey (1982) and Tjur (1984). This chapter adopts the notation used in Bingham et al. (2008) because it allows the main features of the design construction to be most practically illustrated. Full factorial designs are considered first and fractional factorial designs considered as a special case in Section 8.5.

Denote the number of stages of randomization as S . For a blocked factorial design, for example, $S = 2$, where the first stage of randomization corresponds to the order in which the blocks are run and the second stage corresponds to the order in which the treatments are run within each block. Similarly, a split-plot design has $S = 2$ stages of randomization—the first stage corresponds to a random selection of the whole-plot level settings and the second stage corresponds to the complete randomization of treatments on subplots within a whole plot. A blocked split-plot design (McLeod and Brewster 2004), on the other hand, has $S = 3$ stages, where the first stage of randomization is for blocks, the second corresponds to the selection of the whole plot level settings nested within blocks, and the final stage of randomization is at the subplot level of the design.

There are $S - 1$ stages corresponding to randomization restrictions, and the S th stage of randomization is a complete randomization within the final stage of nesting. As illustrated in Examples 8.1 through 8.3, the columns of the factorial model matrix can be used to define the randomization structure of an experiment in exactly the same way as blocking factorial designs.

Begin by considering the s th stage of randomization ($s \leq S - 1$) for a 2^k factorial design. To group the experiment trials into blocks of size $2^k/2^{r_s}$, r_s linearly independent columns of X are selected. Let the i th randomization restriction factor ($i = 1, 2, \dots, r_s$) at stage s be denoted as $\delta_i^{(s)}$. The 2^{r_s} unique randomization restriction factor level settings in the columns of X corresponding to $\delta_1^{(s)}, \dots, \delta_{r_s}^{(s)}$ are used to specify the sets of trials (e.g., blocks) where

restricted randomization takes places. Note that the δ 's are not factors in a practical sense, but the term is used in the same sense as one refers to a blocking factor (see, e.g., Sun et al. 1997; and Example 8.1).

Following the notation of Bingham et al. (2008), let $\mathcal{L}_s = \{l_{si}\}_{i=1}^{r_s}$ be the index set of columns of \mathbf{X} that are assigned to randomization restriction factors at stage s , so that $\delta_i^{(s)} = c_{l_{si}}$. Using the ± 1 coding of \mathbf{X} , the interaction between $\delta_i^{(s)}$ and $c_{l_{si}}$ is a column of ones. For the s th stage of randomization, denote this column of 1's as $\Delta^{(s)}$. Consequently, $\Delta^{(s)} = c_{l_{si}} \delta_i^{(s)}$ for $i = 1, \dots, r_s$. Each $c_{l_{si}}$ is called a *randomization defining contrast*. Analogous to the block defining contrast subgroup, taking all possible interactions between the defining contrasts creates the *randomization defining contrast subgroup (RDCSG)* for the s th stage of randomization. The RDCSG contains the list of columns of \mathbf{X} whose settings do not vary within each of the unique settings in the columns corresponding to the randomization defining words. For instance, in Example 8.2, this would be $\{\Delta^{(1)}, 1, 2, 12, 4, 14, 24, 124\}$. Similarly, we can equivalently write the RDCSG in terms of the RRDR as illustrated in that example.

This framework for specifying the randomization restrictions, one stage at a time, makes it easy to identify which main effects or interactions are impacted at the s th stage of randomization. For example, split-plot design randomization restriction factors are assigned directly to main effect columns. Furthermore, when the current stage of randomization, s , is nested within another stage(s), each of the randomization restrictions from the previous stage that describe the nesting must be included as randomization restrictions in the current stage, as shown in Example 8.4.

Example 8.4 Factorial Split-Split-Plot Design

Consider a 2^5 split-split-plot design, with whole-plot factors $\{1, 2\}$, subplot factors $\{3, 4\}$, and sub-subplot factor $\{5\}$. There are three stages of randomization to consider and, thus, randomization restriction factors to specify for stages $s = 1$ and 2. Since there are two whole-plot factors, the first-stage randomization restriction factors are assigned to factors 1 and 2. That is, $\delta_1^{(1)} = 1$ and $\delta_2^{(1)} = 2$. For the second stage of randomization, factors 1 and 2 are still fixed, along with factors 3 and 4. Therefore, $\delta_1^{(2)} = 1$, $\delta_2^{(2)} = 2$, $\delta_3^{(2)} = 3$, and $\delta_4^{(2)} = 4$. The corresponding RRDRs are

$$\begin{aligned} \Delta^{(1)} &= 1\delta_1^{(1)} = 2\delta_2^{(1)} = 12\delta_1^{(1)}\delta_2^{(1)}, \\ \Delta^{(2)} &= 1\delta_1^{(2)} = 2\delta_2^{(2)} = 3\delta_3^{(2)} = 4\delta_4^{(2)} = 12\delta_1^{(2)}\delta_2^{(2)} = 13\delta_1^{(2)}\delta_3^{(2)} = 14\delta_1^{(2)}\delta_4^{(2)}, \\ &= 23\delta_2^{(2)}\delta_3^{(2)} = 24\delta_2^{(2)}\delta_4^{(2)} = 34\delta_3^{(2)}\delta_4^{(2)} = 123\delta_1^{(2)}\delta_2^{(2)}\delta_3^{(2)} = 124\delta_1^{(2)}\delta_2^{(2)}\delta_4^{(2)}, \\ &= 134\delta_1^{(2)}\delta_3^{(2)}\delta_4^{(2)} = 234\delta_2^{(2)}\delta_3^{(2)}\delta_4^{(2)} = 1234\delta_1^{(2)}\delta_2^{(2)}\delta_3^{(2)}\delta_4^{(2)}. \end{aligned}$$

Example 8.5 Factorial Split-Lot Design (Example 8.3 Revisited)

Consider again the 2^5 split-lot design in Example 8.3. The randomization restriction factors for the first stage were $\delta_1^{(1)} = 1$ and $\delta_2^{(1)} = 2$, and the second-stage randomization restriction factors were $\delta_1^{(2)} = 3$ and $\delta_2^{(2)} = 4$. Unlike the split-split-plot design in Example 8.4, the first-stage randomization restriction factors are not randomization restriction factors at the second stage of randomization. This means that factors 3 and 4 are not nested within level settings of factors 1 and 2. Therefore, the level settings of

factors 1 and 2 vary within fixed settings of factors 3 and 4. The level settings of factor 5 vary at all stages of randomization under this specification. The RRDRs for this design are

$$\begin{aligned} \Delta^{(1)} &= 1\delta_1^{(1)} = 2\delta_2^{(1)} = 12\delta_1^{(1)}\delta_2^{(1)}, \\ \Delta^{(2)} &= 3\delta_1^{(2)} = 4\delta_2^{(2)} = 34\delta_1^{(2)}\delta_2^{(2)}. \end{aligned}$$

8.5 Impact on Data Analysis

Restricting the randomization of an experiment impacts the manner in which the data analysis is conducted. Typically, observations that arise from the same blocks in a block design receive a common random effect in the specified statistical model, and that is, the case for multistage designs as well. The multiple linear regression model is again the model of interest, and it is assumed that each stage of randomization introduces a new random error term with trials in the same block receiving the same random effect. For example, looking back at the split-plot design shown in Table 8.3, the final two columns list the error terms corresponding to the whole-plot (block) error and subplot error. Each observation within a plot is subject to the same whole-plot random error term. Consequently, observations that are randomized together within a block are correlated. For ease of notation, only unreplicated designs will be considered here and thus the number of trials in the experiment is 2^k .

The response model for designs with randomization restrictions in this chapter is specified as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}, \tag{8.2}$$

where $\mathbf{y} = (y_1, \dots, y_{2^k})'$ and $\boldsymbol{\beta} = (\beta_0, \beta_1, \dots, \beta_k)'$. The difference between models (8.1) and (8.2) is the error induced by the randomization structure of the experiment. When there are S stages of randomization, the error vector $\boldsymbol{\epsilon}$ can be written as the sum of S independent error vectors: $S - 1$ vectors for restricted randomization stages and 1 for complete randomization of experimental units. At the s th stage of randomization, there are $n_s = 2^{r_s}$ blocks and a single random effect, $\epsilon_{s,j}$ ($j = 1, \dots, n_s$), for each block. Let $\boldsymbol{\epsilon}_s$ ($s = 1, 2, \dots, S - 1$) be the $n_s \times 1$ vector of error terms associated with the s th stage of randomization – one random effect per block. The $\epsilon_{s,j}$'s are taken to be iid $N(0, \sigma_s^2)$ random variables. Let N_s be a $2^k \times n_s$ incidence matrix that associates an error term to observations within a block. Specifically, the (l, m) th element of N_s is 1 if the l th experimental unit belongs to the m th block at the s th stage of randomization and 0 otherwise. The error structure for (8.2) can then be written as

$$\boldsymbol{\epsilon} = N_1\boldsymbol{\epsilon}_1 + N_2\boldsymbol{\epsilon}_2 + \dots + N_{S-1}\boldsymbol{\epsilon}_{S-1} + \boldsymbol{\epsilon}_S, \tag{8.3}$$

where $\boldsymbol{\epsilon}_S$ is an $2^k \times 1$ vector of iid $N(0, \sigma_S^2)$ error terms associated with the 2^k experimental units and not associated with restricted randomization. It follows that the covariance matrix for the vector of responses is $\text{Var}(\mathbf{y}) = \sum_{i=1}^{S-1} N_i N_i' \sigma_i^2 + I_{2^k} \sigma_S^2$, where I_{2^k} is a $2^k \times 2^k$ identity matrix.

Looking back to Example 8.2, there are $n_1 = 8$ whole plots, and thus, the error vector for the first stage of randomization is $\epsilon_1 = (\epsilon_{1,1}, \epsilon_{1,2}, \dots, \epsilon_{1,8})'$. Furthermore, in (8.3), $S = 2$ and N_1 is

$$N_1 = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

Theorems 1 and 2 of Ranjan et al. (2009) are helpful in pointing to analysis strategies for these designs. Briefly, their Theorem 1 shows that the effect estimators arising from designs constructed in the manner specified here are independent and normally distributed. The theorem also demonstrates that the ordinary least squares and generalized least squares estimators for β are the same. Thus, $\hat{\beta} = (X'X)^{-1}Xy$. Theorem 2 of Ranjan et al. (2009) goes on to develop formulae for the variance of an estimator. Let $1_s(c_m)$ ($m = 0, 1, \dots, 2^k - 1$) denote the indicator function that is equal to 1 when column m of X does not change levels within blocks at the s th stage of randomization and 0 otherwise. Another way of looking at this is $1_s(c_m) = 1$ only if the effect associated with the m th column of X is in the s th stage RDCSG. The variance of the effect estimator associated with column m of X is $\text{Var}(\hat{\beta}_m) = \sigma_S^2/2^k + \sum_{s=1}^{S-1} 1_s(c_m)\sigma_s^2/2^s$.

So, under model (8.2) and the design construction described in this chapter, it turns out that the factorial effect estimators have an error variance that is a linear combination of the variances $\sigma_1^2, \sigma_2^2, \dots, \sigma_S^2$. However, only the effect estimators that appear in the s th stage RDCSG have an error variance that contains a multiple of σ_s^2 ($s = 1, 2, \dots, S - 1$) in addition to σ_S^2 in the linear combination of variance components. The variances of the remaining effects do not depend on σ_s^2 .

From a data analysis point of view, since the effect estimators are independent and normally distributed, one can use the usual data analysis techniques. However, separate analyses need to be performed on effects with the same variance. That is, one would require S half-normal plots to conduct the analysis, one for each stage of randomization.

Example 8.6 Factorial Split-Lot Design (Examples 8.3 and 8.5 Revisited)

Consider again the 2^5 unreplicated split-lot design in Example 8.5. There are $S = 3$ stages of randomization—two for restricted randomization. Looking at the first RRDR in Example 8.5, we see that the effect estimators corresponding to the main effects of factors 1 and 2, as well as the interaction between factors 1 and 2, will have a variance

TABLE 8.4
Variances of Regression Effects in Example 8.6

Effect	Variance
1, 2, 12	$\frac{\sigma_S^2}{2^5} + \frac{\sigma_1^2}{2^2}$
3, 4, 34	$\frac{\sigma_S^2}{2^5} + \frac{\sigma_2^2}{2^2}$
All other effects	$\frac{\sigma_\epsilon^2}{2^5}$

$\sigma_S^2/2^5 + \sigma_1^2/2^2$. Similarly, the RRDR for the second stage of randomization reveals that the effect estimators corresponding to the main effects of factors 3 and 4, as well as their interaction, will have a variance of $\sigma_S^2/2^5 + \sigma_2^2/2^2$. The remaining factorial effects are not impacted by the restrictions in randomization (i.e., they are orthogonal to both block structures) and have variance $\sigma_\epsilon^2/2^5$. The variances of the effect estimators are summarized in Table 8.4.

The approach outlined earlier is useful in determining which effects are estimated with which variance. This specific split-lot design deserves some further consideration. If a half-normal probability plot, for example, is used to assess the significance of the factorial effects, then 3 separate half-normal plots must be constructed—one for each set of effects that have the same variance. Therefore, the first two half-normal plots each show only three effects. Consequently, the importance of effects 1, 2, and 12, or 3, 4, and 34 cannot be assessed because there are too few degrees of freedom within each of the first two groups. It is unlikely that anyone would run an experiment where the significance of 4 out of 5 main effects could not be investigated.

To get around this issue, the structure of the design can be altered so that degrees of freedom are shifted from the final stage of randomization to other levels of the design. To do this, the number of blocks at the first 2 stages of randomization is changed from 4 blocks of 8 runs to 8 blocks of 4 trials. This is achieved by including two additional randomization restriction factors $\delta_3^{(1)}$ and $\delta_3^{(2)}$, one at each stage of restricted randomization. Operationally, this amounts to including additional randomization factors that have the specific goal of altering the block size at the first two stages of randomization. For example, let $\delta_3^{(1)} = 345$ and $\delta_3^{(2)} = 125$, with the resulting RRDRs

$$\begin{aligned} \Delta^{(1)} &= 1\delta_1^{(1)} = 2\delta_2^{(1)} = 12\delta_1^{(1)}\delta_2^{(1)} = 345\delta_3^{(1)} = 1345\delta_1^{(1)}\delta_3^{(1)} = 2345\delta_2^{(1)}\delta_3^{(1)} \\ &= 12345\delta_1^{(1)}\delta_2^{(1)}\delta_3^{(1)}, \\ \Delta^{(2)} &= 3\delta_1^{(2)} = 4\delta_2^{(2)} = 34\delta_1^{(2)}\delta_2^{(2)} = 125\delta_3^{(2)} = 1235\delta_1^{(2)}\delta_3^{(2)} = 1245\delta_2^{(2)}\delta_3^{(2)} \\ &= 12345\delta_1^{(2)}\delta_2^{(2)}\delta_3^{(2)}. \end{aligned}$$

The effects and their variances are summarized in Table 8.5. To analyze the design, separate half-normal plots are constructed for each set of effects in the rows of the table. Since the 12345 interaction appears in both RRDRs, then this effect inherits the variance from each stage of randomization and the importance of this effect cannot be assessed by the use of plots. This may be a sacrifice that an experimenter can live with since effect hierarchy suggests that this effect is unlikely to be important. The downside of

TABLE 8.5
Variances of Regression Effects in Second Design of Example 8.6

Effect	Variance
1, 2, 12, 345, 1345, 2345	$\frac{\sigma_5^2}{2^5} + \frac{\sigma_1^2}{2^3}$
3, 4, 34, 125, 1245, 1235	$\frac{\sigma_5^2}{2^5} + \frac{\sigma_2^2}{2^3}$
12345	$\frac{\sigma_5^2}{2^5} + \frac{\sigma_1^2}{2^3} + \frac{\sigma_2^2}{2^3}$
All other effects	$\frac{\sigma_5^2}{2^5}$

this design is that the number of blocks has increased at each stage of randomization, thereby potentially increasing the cost of the design.

Bingham et al. (2008) pointed out that another assignment ($\delta_3^{(1)} = 1234$ and $\delta_3^{(2)} = 1345$) could be made. In that example, the effect that appeared in both RRDRs is 34 instead of 12345. Thus, in the second design, the significance of the 34 interaction cannot be assessed. Again the effect hierarchy would suggest that the first design specification is likely to be preferable since sacrificing the 12345 interaction is more desirable than sacrificing the 34 interaction in most applications.

Finally, it is worth noting that not every assignment $\delta_3^{(1)}$ and $\delta_3^{(2)}$ is possible if the desired randomization structure is to be maintained. For example, if we assign $\delta_3^{(1)} = 24$, then $\delta_2^{(1)} \delta_3^{(1)} = 4$. As a consequence factor, 4 is fixed at both stages of restricted randomization instead of only at the second stage, thereby destroying the desired randomization structure. So inspection of the RDCSGs is an important task. Ranjan et al. (2009) outline an algorithmic approach to search for designs that preserves the specified randomization structure. If all nonisomorphic designs are desired, their algorithm can be used with the fast isomorphism check outlined in Spencer et al. (2013).

8.6 Fractional Factorial Designs with Randomization Restrictions

In this section, fractional factorial designs with randomization restrictions are considered. In this setting, the columns of X are used both for constructing the fractional factorial design and also for the RDCSGs specifying the randomization restrictions. The defining contrast subgroup of a fraction determines the treatment structure of the design, whereas the RDCSG defines the randomization structure of the design. One can think of these as specifying which effects (or sets of aliased effects) can be estimated and identifying their variances. The challenge for the experimenter is in determining whether the desired randomization structure is maintained (for split-plot designs, see, e.g., Bisgaard 2000) and whether or not the design allows assessment of the importance of the effects of interest.

For a given fractional factorial design with specified defining contrast subgroup (the treatment structure), it is fairly easy to see the impact of a chosen RDCSG (the randomization structure). This is done by “crossing” each RDCSG with the defining contrast subgroup of the factorial design or, equivalently, by “crossing” their defining relations. This is illustrated via the following example.

Example 8.7 Fractional Factorial Split-Lot Design

Consider a split-lot design with 6 factors to be investigated in 2^5 trials. Thus, the treatment structure of the design is a 2^{6-1} fractional factorial design and the randomization structure follows a 2-stage split-lot design. In this example, the RDCSG for the second design in Example 8.6 is used for illustration. Thus, factors 1 and 2 are fixed at the first stage of randomization, factors 3 and 4 are fixed at the second, and the levels of factors 5 and 6 are allowed to vary at each stage of randomization. Suppose that an experimenter decides to use the column of X that corresponds to the 12345 interaction to generate the levels of factor 6 (i.e., $6 = 12345$) or, equivalently, the defining relation for the fractional factorial design is $I = 123456$. This is the minimum aberration 2^{6-1} fractional factorial design (Fries and Hunter 1980).

Notice that the defining relation implies that several effects are aliased. For example, the 125 and 346 interactions cannot be separated. This is usually denoted as $125 = 346$ (see Chapter 7). Notice that, from $\Delta^{(2)}$ in Example 8.6, the 125 interaction appears in the second-stage RRDR. Therefore, this effect estimator ($125 = 346$) has an error variance as shown in the second row of Table 8.5. We also see a serious problem: the main effect of factor 6 is aliased with the 12345 interaction. This means that the main effect estimate for factor 6 has the same variance as the 12345 interaction and its significance cannot be assessed. In fact, since the 12345 interaction appears in both RRDRs, then the assignment $6 = 12345$ means that factor 6 does not change levels within blocks as required.

More generally, to see which effects are estimated with which variance, each of the RRDRs is multiplied by the words in the defining contrast subgroup of the fractional factorial design. The effects that appear in the products pick up an error term from that stage of randomization. While this may appear easy, care must be taken to maintain the desired randomization structure. For example, if the specified RRDR is used, factor 6 should not be assigned to any of the effects in rows 2, 3, or 4 of Table 8.5 to avoid destroying the desired randomization structure. A fractional factorial design that is better than the minimum aberration design for this setting is given by the defining relation $I = 12346$. To see the impact on factor 6, and interactions involving this factor, the RRDR is multiplied by 12346, giving

$$\begin{aligned}
 12346\Delta^{(1)} &= 2346\delta_1^{(1)} = 1346\delta_2^{(1)} = 346\delta_1^{(1)}\delta_2^{(1)} = 1256\delta_3^{(1)} = 256\delta_1^{(1)}\delta_3^{(1)} = 156\delta_2^{(1)}\delta_3^{(1)} \\
 &= 56\delta_1^{(1)}\delta_2^{(1)}\delta_3^{(1)}, \\
 12346\Delta^{(2)} &= 1246\delta_1^{(2)} = 1236\delta_2^{(2)} = 126\delta_1^{(2)}\delta_2^{(2)} = 3456\delta_3^{(2)} = 456\delta_1^{(2)}\delta_3^{(2)} = 356\delta_2^{(2)}\delta_3^{(2)} \\
 &= 56\delta_1^{(2)}\delta_2^{(2)}\delta_3^{(2)}.
 \end{aligned}$$

Notice that neither 6 nor its alias 1234 appears in either RDCSG. Therefore, the levels of factor 6 are not fixed at any stage of randomization—as desired. In addition, the product of the defining contrast subgroup and the RRDR identifies which interactions involving a given factor are impacted by the randomization restrictions. In this case, observe that, of the two-factor interactions, only 56 appears in both products.

8.7 Choosing Designs

For most experimenters, the task of design construction is just one of many decisions to make when exploring a system. Choosing the design that best fits the particular setting often plays a critical role in the success of the endeavor. In this section, some of the issues

facing the ranking of designs are illustrated. In the end, there is no one criterion that is best for all applications, and experimenters are typically encouraged to consider several “good” designs before making a final choice.

The main issues facing an experimenter who wishes to perform the type of design considered in this chapter are related to estimability and variability of the effects of interest. In unreplicated experiments, in particular, these issues present themselves rather quickly. For instance, consider the designs at the end of Example 8.5. The choice of RDCSG indicates the variance associated with each effect, and also that the significance of some effects cannot be assessed. In that example, there was a choice between a design where the 12345 interaction could not be assessed and another where the significance of the 34 interaction could not be investigated using, for example, half-normal plots.

For completely randomized experiments, fractional factorial design criteria usually emphasize the estimation of lower-order effects (e.g., main effects and two-factor interactions). To this end, criteria such as minimum aberration (Fries and Hunter 1980), clear effects (Wu and Wu 2002), and maximum estimation capacity (Cheng and Mukerjee 1998) all attempt to maximize the ability to estimate lower-order terms, or collections of main effects and some two-factor interactions, under the assumptions of effect sparsity and effect heredity (see also Chapter 7). Such criteria have been adapted to the ranking of fractional factorial designs with randomization restrictions. For example, blocked fractional factorial designs have been ordered using variations of the minimum aberration criterion by Sitter et al. (1997), Chen et al. (1999), Cheng and Wu (2002), and Xu and Lau (2006). Split-plot designs were ranked with adapted versions of the minimum aberration criterion by Huang et al. (1998), Bingham and Sitter (2001), and Yang et al. (2007), while Zhao and Chen (2012) used the clear effects criterion, and Ai and Zhang (2004) used the estimation capacity criterion to rank designs. Furthermore, similar to Cheng et al. (1999) and Tsai et al. (2000), Cheng and Tsai (2011), proposed a more generally applicable criterion that accounts for the different effect variances and also the estimability of two-factor interactions.

The key wrinkle for all of the proposed design selection criteria is that they account for the randomization restrictions and the corresponding impact on the effect variances beyond focusing only on the estimability of the effects. The following example helps to illustrate some of the main features of design selection.

Example 8.8 Fractional Factorial Split-Lot Design with $S = 4$

Suppose a 2^{6-1} fractional factorial split-lot design is to be performed with $S = 4$ stages of randomization and 2 factors defining each of the 3 levels of restricted randomization. Further, suppose that each stage is to have 8 blocks of 4 trials. As in Example 8.7, the minimum aberration treatment structure for the fractional factorial design assigns factor 6 to the 12345 interaction column of a 2^5 full factorial design (i.e., $I = 123456$). To form 8 blocks, there needs to be $r_s = 3$ randomization restriction factors for each of the 3 stages of restricted randomization. One possible assignment of randomization restriction factors is $\delta_1^{(1)} = 1$, $\delta_2^{(1)} = 2$, $\delta_3^{(1)} = 35$, $\delta_1^{(2)} = 3$, $\delta_2^{(2)} = 4$, $\delta_3^{(2)} = 15$, $\delta_1^{(3)} = 5$, $\delta_2^{(3)} = 6$, and $\delta_3^{(3)} = 13$, leading to the three RRDRs shown below. Keeping in mind that $1235 = 46$, $1345 = 26$, and $1356 = 24$ due to aliasing in the fraction, the variances of the effect estimators are summarized in Table 8.6.

$$\begin{aligned} \Delta^{(1)} &= 1\delta_1^{(1)} = 2\delta_2^{(1)} = 12\delta_1^{(1)}\delta_2^{(1)} = 35\delta_3^{(1)} = 135\delta_1^{(1)}\delta_3^{(1)} = 235\delta_2^{(1)}\delta_3^{(1)} \\ &= 1235\delta_1^{(1)}\delta_2^{(1)}\delta_3^{(1)}, \end{aligned}$$

TABLE 8.6

Variates of Regression Effects for the Fractional Factorial Split-Lot Design in Example 8.8 with $I = 123456$

Effect	Variance
1, 2, 12, 35, 46, and 1 higher-order interaction	$\frac{\sigma_5^2}{2^5} + \frac{\sigma_1^2}{2^3}$
3, 4, 34, 15, 26, and 1 higher-order interaction	$\frac{\sigma_5^2}{2^5} + \frac{\sigma_2^2}{2^3}$
5, 6, 56, 13, 24, and 1 higher-order interaction	$\frac{\sigma_5^2}{2^5} + \frac{\sigma_3^2}{2^3}$
1 higher-order interaction	$\frac{\sigma_5^2}{2^5} + \frac{\sigma_1^2}{2^3} + \frac{\sigma_2^2}{2^3} + \frac{\sigma_3^2}{2^3}$
14, 16, 23, 25, 36, 45, and 6 other effects	$\frac{\sigma_5^2}{2^5}$

$$\begin{aligned} \Delta^{(2)} &= 3\delta_1^{(2)} = 4\delta_2^{(2)} = 34\delta_1^{(2)}\delta_2^{(2)} = 15\delta_3^{(2)} = 135\delta_1^{(2)}\delta_3^{(2)} = 145\delta_2^{(2)}\delta_3^{(2)} \\ &= 1345\delta_1^{(2)}\delta_2^{(2)}\delta_3^{(2)}, \\ \Delta^{(3)} &= 5\delta_1^{(3)} = 6\delta_2^{(3)} = 56\delta_1^{(3)}\delta_2^{(3)} = 13\delta_3^{(3)} = 135\delta_1^{(3)}\delta_3^{(3)} = 136\delta_2^{(3)}\delta_3^{(3)} \\ &= 1356\delta_1^{(3)}\delta_2^{(3)}\delta_3^{(3)}. \end{aligned}$$

Notice that there are 6 effects in each of the first 3 rows of Table 8.6, and 3 separate half-normal plots (one for each row) are needed to assess the significance of these effects. Another half-normal plot is required to investigate the importance of the effects in the final row of the table, and the 135 = 246 interaction cannot be assessed since there is nothing with which to compare it.

Bingham et al. (2008) noted that it is desirable to have the proportion of lower-order terms on each plot be (roughly) equal for each half-normal plot. The reason is related to the effect sparsity and hierarchy assumptions. If the set of effects to be displayed on a half-normal plot are, for example, all main effects or two-factor interactions, then the effect sparsity assumption must also hold for the set of effects on this plot when viewed in isolation, in addition to the entire experiment. As a more extreme example, suppose there were to be another row in the table with only higher-order interactions. A half-normal plot of these effects would provide no information about the lower-order effects of interest. So there is some benefit to spreading the effects of interest across the different levels of the analysis.

In the aforementioned example, the treatment structure was a minimum aberration 2^{6-1} fractional factorial design. Suppose instead that factor 6 had been assigned to the 1234 interaction column ($I = 12346$). If interactions involving 3 or more factors are assumed to be negligible, this resolution *V* design would be just as good as the resolution *VI* design used earlier in terms of estimating lower-order effects. Now suppose that this design ($I = 12346$) is used with the same RRDRs as earlier. The effect variances are summarized in Table 8.7. The resulting design has 4 main effects and two-factor interactions in the first 2 rows of Table 8.7, 5 in the third row, and 8 in the final row.

Beforehand, it is tempting to view the two designs in Example 8.7 as essentially the same. The first design has defining relation $I = 123456$ and the second has defining relation $I = 12346$. So in neither case are the two-factor interactions aliased with other main effects or two-factor interactions. In addition, both designs have the same randomization restriction factors. However, the appropriate effect variances are determined by taking the product of the defining contrast subgroup for the fractional factorial design and the RRDRs. It turns out that the second design has fewer two-factor interactions

TABLE 8.7

Variances of Regression Effects for the Fractional Factorial Split-Lot Design in Example 8.8 with $I = 12356$

Effect	Variance
1, 2, 12, 35, and 2 higher-order interactions	$\frac{\sigma_{12}^2}{2^5} + \frac{\sigma_1^2}{2^3}$
3, 4, 34, 15, and 2 higher-order interactions	$\frac{\sigma_{34}^2}{2^5} + \frac{\sigma_2^2}{2^3}$
5, 6, 56, 13, 24, and 1 higher-order interaction	$\frac{\sigma_{56}^2}{2^5} + \frac{\sigma_3^2}{2^3}$
1 higher-order interaction	$\frac{\sigma_{12}^2}{2^5} + \frac{\sigma_1^2}{2^3} + \frac{\sigma_2^2}{2^5} + \frac{\sigma_3^2}{2^3}$
14, 16, 23, 25, 26, 36, 45, 46, and 4 other effects	$\frac{\sigma_{23}^2}{2^5}$

confounded with blocks than the first design. Therefore, comparing the design summarized in Table 8.7 with the design summarized in Table 8.6, there is some advantage to the design that is not minimum aberration. While both designs achieve the desired randomization structures and can estimate all main effects and two-factor interactions, the second design has more of these effects with the smallest variance, a clearly desirable property in terms of estimation efficiency and statistical power.

Bingham et al. (2008) proposed a criterion for selecting designs that attempts to spread the effects of interest uniformly among the levels of the analysis. Cheng and Tsai (2011) go even further, proposing an *information capacity* criterion for multistratum designs (see also Cheng et al. 1999; Tsai et al. 2000). The idea behind their criterion is to maximize an efficiency measure, averaged over the set of models containing all main effects and k two-factor interactions. Specifically, they maximize the average $V^{1/k}$ over their model set, where V is the determinant of the information matrix $X'\Sigma^{-1}X$ (where $\Sigma = \text{Var}(\mathbf{y})$) for a model with all main effects and k two-factor interactions. Indeed, one can consider a sequence of such averages for increasing k in resolution III and IV fractional factorial designs where 2-factor interactions are aliased with main effects and/or other 2-factor interactions. In addition to considering the variability of the effect estimators, their approach also tries to minimize the degree of aliasing among 2-factor interactions and other potentially important effects.

8.8 Concluding Remarks

In this chapter, randomization restrictions for factorial and fractional factorial designs have been discussed. The general approach presented here brings together the factorial design structure with many of the randomization structures developed for agriculture experiments. The methodology fits within the orthogonal block structures developed by Nelder (1965) – though it takes a fair amount of work to see how this can be so (Cheng and Tsai 2011). The key feature in design construction is to recognize that forming randomization restrictions is essentially another form of blocking (see Bingham et al. 2008) and to be clear about which effects are assigned to each level of the analysis. The examples developed here aim to illustrate some of the key issues impacting the choice of designs in practical settings.

References

- Addelman, S. (1964). Some two-level factorial plans with split-plot confounding. *Technometrics*, 6(3):253–258.
- Ai, M. and Zhang, R. (2004). Multistratum fractional factorial split-plot designs with minimum aberration and maximum estimation capacity. *Statistics and Probability Letters*, 69(2):161–170.
- Bailey, R. A. (1977). Patterns of confounding in factorial designs. *Biometrika*, 64(3):597–603.
- Bingham, D. and Sitter, R. (1999). Some theoretical results for fractional factorial split-plot designs. *Annals of Statistics*, 27:1240–1255.
- Bingham, D. and Sitter, R. (2001). Design issues in fractional factorial split-plot experiments. *Journal of Quality Technology*, 33(1):2–15.
- Bingham, D., Sitter, R., Kelly, E., Moore, L., and Olivas, J. D. (2008). Factorial designs with multiple levels of randomization. *Statistica Sinica*, 18(2):493.
- Bisgaard, S. (1994). A note on the definition of resolution for blocked 2^{k-p} designs. *Technometrics*, 36(3):308–311.
- Bisgaard, S. (2000). The design and analysis of $2^{k-p} \times 2^{q-r}$ split-plot experiments. *Journal of Quality Technology*, 32(1):39–56.
- Bose, R. C. (1947). Mathematical theory of the symmetrical factorial design. *Sankhyā: The Indian Journal of Statistics*, 3:107–166.
- Butler, N. A. (2004). Construction of two-level split-plot fractional factorial designs for multistage processes. *Technometrics*, 46(4):445–451.
- Chen, H., Cheng, C.-S. (1999). Theory of optimal blocking of 2^{n-m} designs. *Annals of Statistics*, 27(6):1948–1973.
- Cheng, C.-S., Steinberg, D. M., and Sun, D. X. (1999). Minimum aberration and model robustness for two-level fractional factorial designs. *Journal of the Royal Statistical Society: Series B*, 61(1):85–93.
- Cheng, C.-S. and Mukerjee, R. (1998). Regular fractional factorial designs with minimum aberration and maximum estimation capacity. *Annals of Statistics*, 26(6):2289–2300.
- Cheng, C.-S. and Tsai, P.-W. (2011). Multistratum fractional factorial designs. *Statistica Sinica*, 21(3):1001.
- Cheng, S.-W., Li, W., and Kenny, Q. Y. (2004). Blocked nonregular two-level factorial designs. *Technometrics*, 46(3).
- Cheng, S. W. and Wu, C. J. (2002). Choice of optimal blocking schemes in two-level and three-level designs. *Technometrics*, 44(3).
- Collings, B. J. (1989). Shorter communication: Quick confounding. *Technometrics*, 31(1):107–110.
- Daniel, C. (1959). Use of half-normal plots in interpreting factorial two-level experiments. *Technometrics*, 1(4):311–341.
- Dean, A. M. and John, J. A. (1975). Single replicate factorial experiments in generalized cyclic designs: II. Asymmetrical arrangements. *Journal of the Royal Statistical Society. Series B*, 37:72–76.
- Dean, A. M. and Voss, D. T. (1999). *Design and Analysis of Experiments*. Springer-Verlag, New York.
- Federer, W. T. and King, F. (2007). *Variations on Split Plot and Split Block Experiment Designs*. John Wiley & Sons, Hoboken, N.J.
- Fries, A. and Hunter, W. G. (1980). Minimum aberration 2^{k-p} designs. *Technometrics*, 22(4):601–608.
- Gilmour, S. G. and Goos, P. (2009). Analysis of data from non-orthogonal multistratum designs in industrial experiments. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 58(4):467–484.
- Hamada, M. and Balakrishnan, N. (1998). Analyzing unreplicated factorial experiments: A review with some new proposals. *Statistica Sinica*, 8(1):1–28.
- Huang, P., Chen, D., and Voelkel, J. O. (1998). Minimum-aberration two-level split-plot designs. *Technometrics*, 40(4):314–326.
- Jones, B. and Goos, P. (2009). D-optimal design of split-split-plot experiments. *Biometrika*, 96(1):67–82.

- McLeod, R. G. and Brewster, J. F. (2004). The design of blocked fractional factorial split-plot experiments. *Technometrics*, 46(2):135–146.
- Mee, R. W. and Bates, R. L. (1998). Split-plot designs: Experiments for multistage batch processes. *Technometrics*, 40(2):127–140.
- Miller, A. (1997). Strip-plot configurations of fractional factorials. *Technometrics*, 39(2):153–161.
- Montgomery, D. C. (2008). *Design and Analysis of Experiments*. John Wiley & Sons, Hoboken, N.J.
- Mukerjee, R., Wu, C. et al. (1999). Blocking in regular fractional factorials: A projective geometric approach. *The Annals of Statistics*, 27(4):1256–1271.
- Nelder, J. A. (1965). The analysis of randomized experiments with orthogonal block structure. I block structure and the null analysis of variance. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 283(1393):147–162.
- Ranjan, P., Bingham, D. R., and Dean, A. M. (2009). Existence and construction of randomization defining contrast subspaces for regular factorial designs. *The Annals of Statistics*, 37(6A): 3580–3599.
- Sitter, R. R., Chen, J., and Feder, M. (1997). Fractional resolution and minimum aberration in blocked 2^{n-k} designs. *Technometrics*, 39(4):382–390.
- Speed, T. and Bailey, R. A. (1982). On a class of association schemes derived from lattices of equivalence relations. In *Algebraic Structures and Applications* (eds. P. Schultz, C. E. Praeger and R. P. Sullivan), 55–74, Marcel Dekker, New York.
- Spencer, N., Ranjan, P., and Mendivil, F. (2013). Isomorphism check algorithm for two-level factorial designs with randomization restrictions. arXiv:1310.3574.
- Sun, D. X., Wu, C. J., and Chen, Y. (1997). Optimal blocking schemes for 2^n and 2^{n-p} designs. *Technometrics*, 39(3):298–307.
- Tjur, T. (1984). Analysis of variance models in orthogonal designs. *International Statistical Review*, 52:33–65.
- Tsai, P.-W., Gilmour, S. G., and Mead, R. (2000). Projective three-level main effects designs robust to model uncertainty. *Biometrika*, 87(2):467–475.
- Vine, A., Lewis, S., Dean, A., and Brunson, D. (2008). A critical assessment of two-stage group screening through industrial experimentation. *Technometrics*, 50(1):15–25.
- Wu, C. J. and Hamada, M. S. (2011). *Experiments: Planning, Analysis, and Optimization*. John Wiley & Sons, New York.
- Wu, H. and Wu, C. (2002). Clear two-factor interactions and minimum aberration. *The Annals of Statistics*, 30:1496–1511.
- Xu, H. and Lau, S. (2006). Minimum aberration blocking schemes for two- and three-level fractional factorial designs. *Journal of Statistical Planning and Inference*, 136(11):4088–4118.
- Yang, J., Zhang, R., and Liu, M. (2007). Construction of fractional factorial split-plot designs with weak minimum aberration. *Statistics and Probability Letters*, 77(15):1567–1573.
- Zhao, S. and Chen, X. (2012). Mixed-level fractional factorial split-plot designs containing clear effects. *Metrika*, 75(7):953–962.

9

Nonregular Factorial and Supersaturated Designs

Hongquan Xu

CONTENTS

9.1	Introduction.....	339
9.2	Examples of Nonregular Designs.....	340
9.3	Alias Structure.....	343
9.4	Optimality Criteria.....	345
9.4.1	Generalized Minimum Aberration.....	345
9.4.2	Minimum Moment Aberration.....	350
9.4.3	Generalized Resolution and Projection Properties for 2-Level Designs.....	351
9.4.4	Projection Aberration, Estimation Capacity and Design Efficiency.....	353
9.4.5	Uniformity and Space-Filling Property.....	354
9.5	Construction Methods.....	355
9.5.1	Algorithmic Methods.....	356
9.5.2	Quaternary Code Designs.....	356
9.6	Optimality Results.....	359
9.7	Supersaturated Designs.....	360
9.8	Analysis Strategies.....	363
9.9	Concluding Remarks.....	366
	Acknowledgment.....	366
	References.....	367

9.1 Introduction

Fractional factorial designs are classified into two broad types: *regular* designs and *nonregular* designs. Regular designs are constructed through defining relations among factors; they are introduced in Chapter 1, Section 1.7 and fully described in Chapter 7. These designs have a simple alias structure in that any two factorial contrasts are either orthogonal or fully aliased. The run sizes are always a power of two, three, or another prime, and thus the “gaps” between possible run sizes increase exponentially as the power increases.

Plackett and Burman (1946) first gave a large collection of two-level nonregular designs whose run sizes are not a power of two. These designs are often referred to as the Plackett–Burman designs in the literature and belong to a wide class of orthogonal arrays (OAs) (Rao 1947). Plackett–Burman designs and other nonregular designs are widely used in various screening experiments for their run size economy and flexibility (Wu and Hamada 2009). They fill the gaps between regular designs in terms of various run sizes and are flexible in accommodating various combinations of factors with different numbers of levels.

Unlike regular designs, nonregular designs may exhibit complex alias structures, that is, a large number of effects may be neither orthogonal nor fully aliased, which makes it difficult to identify and interpret significant effects. For this reason, nonregular designs have traditionally been used to estimate factor main effects only but not their interactions. However, in many practical situations it is often questionable whether the interaction effects are negligible.

Hamada and Wu (1992) went beyond the traditional approach and proposed an analysis strategy for nonregular designs in which some interactions could be entertained and estimated through their complex alias structure. They pointed out that ignoring interactions can lead to (1) important effects being missed, (2) unimportant effects being erroneously detected, and (3) estimated effects having reversed signs resulting in incorrectly recommended factor levels. Their pioneering work motivated the recent studies in design properties, optimality criteria, construction, and analysis of nonregular designs.

Supersaturated designs (SSDs) are factorial designs whose run sizes are too small to allow estimation of all factorial effects of interest. They have become increasingly popular in the last two decades because of their potential for reducing the number of runs. Broadly speaking, SSDs are a special class of nonregular designs; some of the optimality criteria and results developed for nonregular designs can easily be extended to the setting of SSDs. Since SSDs are typically used to estimate main effects only, the problems associated with SSDs are relatively simpler than those for other nonregular designs, so in this chapter, we emphasize nonregular fractions.

The remainder of the introduction gives some basic concepts and definitions. An OA of n runs, k factors, s levels, and strength t , denoted by $OA(n, s^k, t)$, is an $n \times k$ matrix in which each column has s symbols or levels, and for any t columns, all possible s^t combinations of symbols appear equally often as a row in the $n \times t$ subarray. A regular s^{k-q} design of resolution r is an $OA(n = s^{k-q}, s^k, t = r - 1)$, but not every OA with these parameters is a regular design. Further, let $OA(n, s_1^{k_1} \times \cdots \times s_m^{k_m}, t)$ denote a mixed-level OA of strength t with k_i columns of s_i levels for $i = 1, \dots, m$. Hedayat et al. (1999) gave a comprehensive account of theory and applications of OAs.

OAs of strength 2, such as Plackett–Burman designs, allow all the main effects to be estimated independently and they are universally optimal for the main effects model (Cheng 1980). A necessary condition for the existence of an $OA(n, s_1^{k_1} \times \cdots \times s_m^{k_m}, 2)$ is that $n - 1 \geq \sum_{i=1}^m k_i(s_i - 1)$. A design is called saturated if $n - 1 = \sum_{i=1}^m k_i(s_i - 1)$ and supersaturated if $n - 1 < \sum_{i=1}^m k_i(s_i - 1)$. In the literature, OAs of strength 2 are also called orthogonal designs or OAs without mentioning the strength explicitly. For convenience, an OA of strength 1 is also called a *balanced design*, where every level appears equally often for each factor.

9.2 Examples of Nonregular Designs

Example 9.1

Consider an experiment reported by Vander Heyden et al. (1999) who used the high-performance liquid chromatography (HPLC) method to study the assay of ridogrel and its related compounds in ridogrel oral film-coated tablet simulations. The researcher used a 12-run Plackett–Burman design to evaluate the importance of eight factors on

TABLE 9.1
Design and Data for the HPLC Experiment

Run	A	B	C	D	E	F	G	H	I	J	K	MC
1	+	+	+	-	+	+	-	+	-	-	-	101.6
2	+	+	-	+	-	-	-	+	+	+	-	101.7
3	+	-	+	+	-	+	-	-	-	+	+	101.6
4	+	-	-	-	+	+	+	-	+	+	-	101.9
5	+	-	+	-	-	-	+	+	+	-	+	101.8
6	-	+	+	+	-	+	+	-	+	-	-	101.1
7	-	+	-	-	-	+	+	+	-	+	+	101.1
8	-	-	-	+	+	+	-	+	+	-	+	101.6
9	-	-	+	+	+	-	+	+	-	+	-	98.4
10	-	+	+	-	+	-	-	-	+	+	+	99.7
11	+	+	-	+	+	-	+	-	-	-	+	99.7
12	-	-	-	-	-	-	-	-	-	-	-	102.3

Note: Columns C, G, and K were not used in the experiment.

several responses. One response was the percentage recovery of main compound, ridogrel. The eight factors were pH of the buffer (A), column manufacturer (B), column temperature (D), percent of organic solvent in the mobile phase at the start of the gradient (E), percent of organic solvent in the mobile phase at the end of the gradient (F), flow of the mobile phase (H), detection wavelength (I), and concentration of the buffer (J). Table 9.1 gives the design matrix and the observed data. Fitting a main effects model, we get

$$\hat{y} = 101.04 + 0.34A - 0.22B - 0.36D - 0.56E + 0.44F - 0.01H + 0.26I - 0.31J, \quad (9.1)$$

where each factor has two levels coded as +1 and -1 for + and -, respectively. This model has $R^2 = 0.78$ with $\hat{\sigma} = 1.045$ on 3 degrees of freedom. The most significant factors are E and F with *p*-values of 0.16 and 0.24, respectively. The researchers concluded there was no significant relationship between any of the factors and this response because none of the effects are significant at the 10% level.

For Plackett–Burman designs, main effects are partially aliased with two-factor interactions (2fi’s). Nonnegligible 2fi’s could bias the estimates of the main effects. Phoa et al. (2009) reanalyzed the data and found one very significant interaction. The interaction EF is more significant than the main effects E and F. They found the following model:

$$\hat{y} = 101.04 - 0.56E + 0.44F - 0.30H + 0.88EF, \quad (9.2)$$

where E and F are the percentages of organic solvent in the mobile phase at the start and the end of the gradient, respectively, and H is the flow of the mobile phase. This model has $R^2 = 0.96$, indicating a good fit. In the model (9.2), H is significant at the 5% level (*p*-value = 0.012) and E, F and EF are significant at the 1% level.

Example 9.2

Consider an experiment reported by Groten et al. (1996, 1997) who performed a 4-week oral/inhalatory study in which the toxicity of combinations of nine compounds was examined in male Wistar rats. The nine compounds were formaldehyde (A), dichloromethane (B), aspirin (C), cadmium chloride (D), stannous chloride (E), loperamide (F), spermine (G), butylated hydroxyanisole (H), and di ethylhexyl (J).

TABLE 9.2

Design and Data for the Chemical Toxicity Experiment

Run	A*	A	B	C	D	E	F	G	H	J	ASAT
1	1	1	-1	-1	-1	-1	-1	-1	-1	-1	70
2	1	-1	1	1	-1	1	1	-1	-1	-1	71
3	-1	-1	1	1	1	-1	-1	1	-1	-1	86
4	1	1	-1	-1	1	1	1	1	-1	-1	75
5	-1	-1	1	-1	-1	-1	1	1	1	-1	65
6	1	1	-1	1	-1	1	-1	1	1	-1	70
7	-1	1	-1	1	1	-1	1	-1	1	-1	96
8	-1	-1	1	-1	1	1	-1	-1	1	-1	65
9	-1	-1	-1	1	-1	-1	1	1	-1	1	77
10	-1	1	1	-1	-1	1	-1	1	-1	1	71
11	1	1	1	-1	1	-1	1	-1	-1	1	88
12	-1	-1	-1	1	1	1	-1	-1	-1	1	80
13	1	1	1	1	-1	-1	-1	-1	1	1	68
14	-1	-1	-1	-1	-1	1	1	-1	1	1	69
15	1	-1	-1	-1	1	-1	-1	1	1	1	72
16	1	1	1	1	1	1	1	1	1	1	82

Note: Columns A–J form a regular 2^{9-5} design; columns A* and B–J form a nonregular 2^{9-5} design.

Their experiment used a regular 2^{9-5} design with design generators $E = ABCD$, $F = AD$, $G = AE$, $H = AC$, $J = AB$. For each factor, the low level corresponds to no compound. One of the responses measured was aspartate aminotransferase (ASAT) activity. The design and data are given in Table 9.2. Ignoring three-factor or higher-order interactions, the alias relations among main effects and 2fi's are $A = BJ = CH = DF = EG$, $B = AJ$, $C = AH$, $D = AF$, $E = AG$, $F = AD$, $G = AE$, $H = AC$, $J = AB$, $BC = DG = EF = HJ$, $BD = CG = EH = FJ$, $BE = CF = DH = GJ$, $BF = CE = DJ = GH$, $BG = CD = EJ = FH$, $BH = CJ = DE = FG$. The researchers believed that formaldehyde (A) did not interact with other compounds, so the main effects of B–J can be estimated with confidence. The estimate of A is fully aliased with four 2fi's.

Groten et al. (1996, 1997) first analyzed the main effects and then analyzed the significant main effects together with their 2fi's in a subsequent analysis. They concluded that C, D, E, F, DE, and DF were significant effects and obtained the following model:

$$\hat{y} = 75.31 + 3.44C + 5.19D - 2.44E + 2.56F - 2.56DE + 2.19DF.$$

Recall that DF is fully aliased with A. When two effects are fully aliased in a regular design, it is impossible to distinguish between them based on the data only in the analysis. When a 2fi is aliased with a main effect, we often assume the 2fi is negligible. However, Groten et al. (1996, 1997) ignored the main effect of formaldehyde (A) and concluded that the DF interaction was significant in their model, based on their expert opinion. This contradicts the conventional statistical practice. This problem could be avoided if a nonregular design had been used.

Phoa et al. (2009) demonstrated that we could estimate A and DF simultaneously if a nonregular design had been used. Consider a new design in which formaldehyde has level settings as column A* and all other factors have the same level settings as before. Column A* differs from column A in the runs 2, 7, 10, and 15, where the high and low levels are switched. For the new design, A* is not fully aliased with any 2fi, but is partially

aliased with 16 2fi's with correlation ± 0.5 . Here, and throughout this chapter, correlation is defined as $1/n$ times the inner product of the columns corresponding to the factorial effects. For example, the correlation between A^* and CD is -0.5 , while the correlation between A^* and DF is 0.5 . Since A^* is not fully aliased with any 2fi's, we can separate A^* and DF . One problem is that we do not have data from the new design. If formaldehyde (A) were negligible as Groten et al. (1996, 1997) suggested, the changes we made in the levels for formaldehyde (A) would not affect the responses; so it is reasonable to use the same data for the new design. Then we can estimate A^* and DF simultaneously, yielding the following fitted model:

$$\hat{y} = 75.31 + 3.44C + 5.19D - 2.44E + 2.56F - 2.56DE + 3.46DF - 2.54A^* - 1.94H,$$

where all effects are significant at the 5% level. The model has $R^2 = 0.96$. The estimates of A^* and DF have opposite signs and the estimate of DF becomes larger (3.46 vs. 2.19) in the modified design. One possible interpretation is that formaldehyde (A) was important and the effect of DF could be underestimated in the original design when the main effect of A was ignored. For more discussions on this experiment, see Phoa et al. (2009).

The advantage of the nonregular design is that it is possible to estimate partially aliased effects without adding extra runs. The disadvantage is that the analysis becomes more complicated.

Example 9.3

Ding et al. (2013) reported an experiment studying a system with Herpes simplex virus type 1 (HSV-1) and five antiviral drugs, namely: interferon-alpha (A), interferon-beta (B), interferon-gamma (C), ribavirin (D), and acyclovir (E). Their original experiment used a composite design that consists of a 16-run factorial design with 2 levels and an 18-run OA with 3 levels. Two researchers conducted the experiment independently with different random orders, yielding two replicates. Here we look at the 18-run OA only, which corresponds to columns 2–6 of the commonly used $OA(18, 3^7, 2)$; see Table 9.3. For each drug, the 3 levels 0, 1, and 2 correspond to no drug, intermediate drug dosage, and high drug dosage, respectively. The observed data, readout, were the percentage of infected cells after the combination drug treatment.

Following Ding et al. (2013), we use the square root of readout as response. Since this is an $OA(18, 3^5, 2)$, we can fit both linear and quadratic main effects. A main effects model identifies D and E as the most significant drugs; both their linear and quadratic effects are significant. With only 18 runs, we do not have sufficient degrees of freedom to estimate all the interactions among five drugs. Nevertheless, we can perform stepwise variable selection. We find two significant bilinear effects DE and AC . The results are similar to those obtained by Ding et al. (2013) using the entire 34-run composite design.

9.3 Alias Structure

For regular designs, we can find the alias relationships among factorial effects from defining relations. For nonregular designs, the alias structure is complicated and we need to use the general regression method. Suppose we fit a model

$$\mathbf{y} = \mathbf{X}_1\boldsymbol{\theta}_1 + \boldsymbol{\epsilon}, \tag{9.3}$$

TABLE 9.3

Design and Data of the Antiviral Drug Experiment

Run	A B C D E							Readout	
	1	2	3	4	5	6	7	Replicate1	Replicate2
1	0	0	0	0	0	0	0	78.6	81.9
2	0	1	1	1	1	1	1	13.3	16.7
3	0	2	2	2	2	2	2	3.4	3.8
4	1	0	0	1	1	2	2	21.4	25.2
5	1	1	1	2	2	0	0	8.6	4.4
6	1	2	2	0	0	1	1	18.0	27.3
7	2	0	1	0	2	1	2	7.3	2.4
8	2	1	2	1	0	2	0	17.9	23.7
9	2	2	0	2	1	0	1	52.9	54.3
10	0	0	2	2	1	1	0	13.2	8.8
11	0	1	0	0	2	2	1	2.1	4.5
12	0	2	1	1	0	0	2	73.4	73.9
13	1	0	1	2	0	2	1	19.6	14.6
14	1	1	2	0	1	0	2	59.1	41.7
15	1	2	0	1	2	1	0	1.4	2.6
16	2	0	2	1	2	0	1	7.3	4.8
17	2	1	0	2	0	1	2	22.3	24.0
18	2	2	1	0	1	2	0	14.1	18.3

where \mathbf{y} is an $n \times 1$ vector of the responses, \mathbf{X}_1 is an $n \times p_1$ matrix corresponding to the fitted model, $\boldsymbol{\theta}_1$ is a $p_1 \times 1$ vector of the model parameters, and $\boldsymbol{\epsilon}$ is an $n \times 1$ vector of normal errors. The least squares estimator of $\boldsymbol{\theta}_1$ is

$$\hat{\boldsymbol{\theta}}_1 = (\mathbf{X}'_1 \mathbf{X}_1)^{-1} \mathbf{X}'_1 \mathbf{y},$$

which is unbiased under the model (9.3). However, if the **true** model is

$$\mathbf{y} = \mathbf{X}_1 \boldsymbol{\theta}_1 + \mathbf{X}_2 \boldsymbol{\theta}_2 + \boldsymbol{\epsilon}, \quad (9.4)$$

where \mathbf{X}_2 is an $n \times p_2$ matrix corresponding to the additional variables that are not in the fitted model, and $\boldsymbol{\theta}_2$ is a $p_2 \times 1$ vector of the additional model parameters. It is easy to show that

$$E(\hat{\boldsymbol{\theta}}_1) = \boldsymbol{\theta}_1 + \mathbf{C} \boldsymbol{\theta}_2,$$

where $\mathbf{C} = (\mathbf{X}'_1 \mathbf{X}_1)^{-1} \mathbf{X}'_1 \mathbf{X}_2$ is the **alias matrix**; see Chapter 7, Section 7.2.2, Box and Draper (1987), and Wu and Hamada (2009).

Example 9.4

Consider the HPLC experiment in Example 9.1. Suppose the model (9.3) contains the intercept and eight main effects and the model (9.4) contains all $\binom{8}{2} = 28$ 2fi's besides the main effects. The matrix \mathbf{X}_1 is a 12×9 matrix and \mathbf{X}_2 is a 12×28 matrix. The alias

matrix $C = 12^{-1}X_1'X_2$ is a 9×28 matrix whose elements are $-1/3, 0$, or $1/3$. Except for the first row, each row represents a factor and has exactly seven 0's, corresponding to the seven 2fi's involving the factor. Each main effect is partially aliased with $\binom{7}{2} = 21$ 2fi's (which do not involve the factor itself) with correlation $1/3$ or $-1/3$. For example, we have

$$E(\hat{H}) = H + \frac{1}{3}(AB - AD - AE - AF + AI - AJ - BD - BE + BF - BI + BJ + DE - DF + DI + DJ + EF - EI - EJ - FI - FJ - IJ). \tag{9.5}$$

If EF is the only significant interaction, this simplifies to $E(\hat{H}) = H + \frac{1}{3}EF$. Under the model (9.2), $H + \frac{1}{3}EF$ is estimated as $-0.30 + \frac{1}{3}(0.88) \approx -0.01$, which agrees well with the estimate of H in the main effects model (9.1). The main effect of H is not significant in the main effects model because it is partially canceled by the significant EF interaction.

Example 9.5

Consider the chemical toxicity experiment in Example 9.2. We again look at the alias relations between the main effects and 2fi's. In this case, X_1 is a 16×10 matrix and X_2 is a 16×36 matrix. It is easy to see that $X_1'X_1 = 16I$ so the alias matrix $C = 16^{-1}X_1'X_2$ is a 10×36 matrix. For the regular design with column A , we have $E(\hat{A}) = A + BJ + CH + DF + EG$, $E(\hat{B}) = B + AJ$, etc. These are indeed the same as the alias relations among the main effects and 2fi's. For the nonregular design with column A^* , we have $E(\hat{A}^*) = A^* + 0.5(BC + DG + EF + HJ + BF + CE + DJ + GH - BG - CD - EJ - FH + BJ + CH + DF + EG)$, $E(\hat{B}) = B + 0.5(A^*C + A^*F - A^*G + A^*J)$, etc. For the regular design, A is fully aliased with four 2fi's with correlation 1, and the other eight main effects are each fully aliased with one 2fi (involving factor A) with correlation 1. For the nonregular design, A^* is partially aliased with 16 2fi's with correlation ± 0.5 , and the other eight main effects are each partially aliased with four 2fi's (involving factor A^*) with correlation ± 0.5 .

9.4 Optimality Criteria

The main objective of the construction of optimal nonregular designs is to minimize the aliasing of higher-order interactions on the main effects. The minimum aberration (MA) criterion, defined in Chapter 1, Section 1.7.2, and Chapter 7, Section 7.2.4, is the standard criterion for comparing regular designs. This criterion can be extended to nonregular designs. This section introduces some of these extensions.

9.4.1 Generalized Minimum Aberration

For a factorial design D with n runs and k factors, the full ANOVA model is

$$y = X_0\theta_0 + X_1\theta_1 + \dots + X_k\theta_k + \epsilon, \tag{9.6}$$

where y is the vector of n observations, θ_0 is the general mean, θ_j is the vector of j th-order factorial effects for $j = 1, \dots, k$, X_0 is the vector of 1's, X_j is the matrix of contrast coefficients

for θ_j , and ϵ is the vector of independent random errors. Note that j th-order factorial effects represent main effects when $j = 1$ and interactions when $j \geq 2$. Here we consider only the cases where the contrast coefficient of a j -factor interaction is the product of its corresponding contrast coefficients of j main effects. As in Xu and Wu (2001), the main effect contrasts are normalized so that they have the same length \sqrt{n} for a balanced design. In particular, for a two-level factor, the contrast vector of its main effect is coded as $(-1, 1)$; for a three-level factor, the contrast vectors of the linear and quadratic main effects are coded as $(-\sqrt{3}/2, 0, \sqrt{3}/2)$ and $(1/\sqrt{2}, -\sqrt{2}, 1/\sqrt{2})$, respectively.

For $j = 1, \dots, k$, Xu and Wu (2001) defined A_j , a function of X_j , to measure the overall aliasing between all j th-order factorial effects and the general mean. Specifically, let $X_j = [x_{il}^{(j)}]$ and define

$$A_j(\mathbf{D}) = n^{-2} \mathbf{1}' X_j X_j' \mathbf{1} = n^{-2} \sum_{l=1}^{n_j} \left(\sum_{i=1}^n x_{il}^{(j)} \right)^2, \tag{9.7}$$

where $\mathbf{1}$ is the $n \times 1$ vector of ones, and n_j is the number of all j th-order factorial effects. The value of A_j is independent of the choice of the orthonormal contrasts used. The vector (A_1, \dots, A_k) is called the *generalized wordlength pattern*, because for a 2-level regular design, A_j is the number of words of length j . The generalized minimum aberration (GMA) criterion (Xu and Wu 2001) is to sequentially minimize A_1, A_2, A_3, \dots . A design that does this is said to have GMA.

Example 9.6

Consider two 2-level designs with 4 runs and 3 factors in Table 9.4. The first design D_1 is a regular 2^{3-1} design with defining relation $I = ABC$ and the second design D_2 is called a one-factor-at-a-time design. Table 9.4 also shows elements of the corresponding X_1, X_2, X_3 matrices. For D_1 , we have $A_1 = (0^2 + 0^2 + 0^2)/4^2 = 0$, $A_2 = (0^2 + 0^2 + 0^2)/4^2 = 0$, and $A_3 = 4^2/4^2 = 1$. For D_2 , we have $A_1 = [(-2)^2 + 0^2 + 2^2]/4^2 = 0.5$, $A_2 = [2^2 + 0^2 + 2^2]/4^2 = 0.5$, and $A_3 = 0^2/4^2 = 0$. Since $A_1(D_1) < A_1(D_2)$, D_1 has less aberration than D_2 ; so the regular design D_1 is preferred to the one-factor-at-a-time design D_2 with respect to the GMA criterion. This agrees with the well-known result that factorial designs are more efficient than one-factor-at-a-time designs.

TABLE 9.4
Two 2-Level Designs with 4 Runs and 3 Factors

(a) D_1 : A Regular 2^{3-1} Design								(b) D_2 : One-Factor-at-a-Time Design							
Run	X_1			X_2			X_3	Run	X_1			X_2			X_3
	A	B	C	AB	AC	BC	ABC		A	B	C	AB	AC	BC	ABC
1	+	+	+	+	+	+	+	1	+	+	+	+	+	+	+
2	+	-	-	-	-	+	+	2	-	+	+	-	-	+	-
3	-	+	-	-	+	-	+	3	-	-	+	+	-	-	+
4	-	-	+	+	-	-	+	4	-	-	-	+	+	+	-
Sum	0	0	0	0	0	0	4	Sum	-2	0	2	2	0	2	0

Example 9.7

Consider the 3-level design given in Table 9.5. This design is a regular 3^{3-1} design with defining relation $I = ABC^2$ and has one word of length 3, that is, ABC^2 ; for details of 3-level regular designs, see Chapter 1, Section 1.7.1, and Chapter 7, Section 7.10. Table 9.5 also shows the orthogonal polynomial contrasts. From the definition (9.7), we have $A_1(\mathbf{D}) = A_2(\mathbf{D}) = 0$ and

$$A_3(\mathbf{D}) = [(-3a^3)^2 + (-3a^2b)^2 + (3a^2b)^2 + (-9ab^2)^2 + (3a^2b)^2 + (-9ab^2)^2 + (9ab^2)^2 + (9b^3)^2]/9^2,$$

with $a = \sqrt{3/2}$ and $b = 1/\sqrt{2}$. This simplifies to $A_3(\mathbf{D}) = 2$.

For a regular 2-level design, as in Example 9.6, the generalized word length pattern is the same as the word length pattern defined in Chapter 7, Section 7.2.4. For a regular s -level design, $A_j(\mathbf{D})$ defined in (9.7) is the total degrees of freedom associated with words of length j in the generating relation, that is, $A_j(\mathbf{D})$ is $s - 1$ times the number of words of length j . Example 9.7 illustrates this for $s = 3$. Hence, GMA reduces to MA for regular designs. The *minimum G_2 -aberration* criterion, proposed by Tang and Deng (1999), for 2-level designs is a special case of the GMA criterion.

Suppose \mathbf{D} is a balanced design, that is, an OA of strength 1. For each factor, each symbol appears the same number of times. It is easy to see that $\mathbf{X}'_1 \mathbf{1}$ is a vector of 0s so $A_1(\mathbf{D}) = 0$. Let $\mathbf{X}'_1 \mathbf{X}_1 = (a_{ij})$. Then a_{ij}/n is the correlation between the i th and j th columns of \mathbf{X}_1 . Since each contrast coefficient of a 2-factor interaction is the product of its corresponding two main effect contrast coefficients, $\mathbf{X}'_2 \mathbf{1}$ is a column vector whose elements are a_{ij} , so $A_2(\mathbf{D}) = n^{-2} \mathbf{1}' \mathbf{X}_2 \mathbf{X}'_2 \mathbf{1} = n^{-2} \sum_{i < j} a_{ij}^2$. That is, $A_2(\mathbf{D})$ measures the overall aliasing among all possible main effects. For an OA of strength 2, $A_2(\mathbf{D}) = 0$; the reverse is also true. Xu and Wu (2001) showed the following important property regarding the generalized word length pattern.

Theorem 9.1 *A design \mathbf{D} is an OA of strength t if and only if $A_j(\mathbf{D}) = 0$ for $1 \leq j \leq t$.*

Therefore, following the GMA criterion, among all possible designs, we shall choose balanced designs and among them choose orthogonal designs with maximum strength.

Example 9.8

Consider choosing five columns from the commonly used $OA(18, 3^7, 2)$ given in Table 9.3. There are 21 possible choices. For illustration, consider three choices. Let \mathbf{D}_1 be the design formed by columns 2 to 6, \mathbf{D}_2 be the design formed by columns 1 and 4-7, and \mathbf{D}_3 be the design formed by columns 1-5, respectively. The generalized word length patterns for the three designs are $(0, 0, 5, 7.5, 0)$, $(0, 0, 6.5, 4.5, 1.5)$, and $(0, 0, 7, 3.5, 2)$, respectively. Hence, \mathbf{D}_1 is the best according to the GMA criterion. It can be easily verified that \mathbf{D}_1 has GMA among all 21 5-factor designs.

The GMA criterion has a sound statistical justification. Suppose we fit the model

$$\mathbf{y} = X_0 \theta_0 + \mathbf{X}_1 \boldsymbol{\theta}_1 + \boldsymbol{\epsilon}, \tag{9.8}$$

TABLE 9.5
3³-1 Design

			X ₁			X ₂						X ₃						
A	B	C	A	B	C	A × B	A × C	B × C	A × B × C	A × B	A × C	B × C	A × B × C	A × B	A × C	B × C	A × B × C	
0	0	0	-1	1	-1	1	-1	-1	1	1	-1	-1	1	-1	1	-1	1	-1
0	1	1	-1	0	-2	0	2	0	-2	0	0	0	4	0	0	-4	0	0
0	2	2	-1	1	1	-1	-1	1	1	1	1	1	1	-1	-1	-1	1	1
1	0	1	0	-2	-1	0	0	2	-2	0	0	4	0	0	0	0	0	-4
1	1	2	0	-2	0	0	0	0	-2	0	0	-2	-2	0	0	0	0	0
1	2	0	0	-2	1	0	0	-2	-2	0	0	2	-2	0	0	0	2	-2
2	0	2	1	1	-1	-1	1	1	1	-1	-1	1	1	-1	1	1	-1	1
2	1	0	1	1	0	-2	-1	1	-2	0	0	2	-2	0	0	2	0	0
2	2	1	1	1	1	1	0	-2	0	-2	0	-2	0	-2	0	-2	0	-2
Sum	0	0	0	0	0	0	0	0	0	0	0	0	0	-3	-3	3	-9	3
scale	a	b	a	b	a	b	a ²	ab	ab	b ²	a ²	ab	ab	b ²	a ²	ab	ab	b ²
	a	b	a	b	a	b	a ²	ab	ab	b ²	a ²	ab	ab	b ²	a ²	ab	ab	b ²
	a	b	a	b	a	b	a ²	ab	ab	b ²	a ²	ab	ab	b ²	a ²	ab	ab	b ²
	a	b	a	b	a	b	a ²	ab	ab	b ²	a ²	ab	ab	b ²	a ²	ab	ab	b ²

a = √3/2 and b = 1/√2.

which contains the general mean and main effects. For a balanced design, the least squares estimator of main effects θ_1 , $\hat{\theta}_1 = (X_1'X_1)^{-1}X_1'y$, is unbiased if (9.8) is correct. However, under the full model (9.6),

$$E(\hat{\theta}_1) = \theta_1 + C_2\theta_2 + \dots + C_k\theta_k,$$

where $C_j = (X_1'X_1)^{-1}X_1'X_j$ is the alias matrix between the main effects and j -factor interactions. The estimation of the main effects is biased or contaminated by (nonnegligible) interactions. A good design should have a small contamination. For a matrix $C = (c_{ij})$, let $\|C\|^2 = \sum_{i,j} |c_{ij}|^2$ be its squared norm. The value $\|C_j\|^2$ is an overall measure of the contamination of j -factor interactions on the estimation of main effects. It can be shown that $\|C_j\|^2$ is independent of the choice of orthonormal contrasts. In the spirit of the *hierarchical ordering principle* discussed by Wu and Hamada (2009, Section 4.6), a good design should sequentially minimize $\|C_j\|^2$ for $j = 2, \dots, k$. Xu and Wu (2001) showed that if all k factors have s levels,

$$\|C_j\|^2 = (j + 1)A_{j+1} + j(s - 2)A_j + (k - j + 1)(s - 1)A_{j-1} \quad \text{for } j = 2, \dots, k.$$

It is easy to see that sequentially minimizing A_3, A_4, \dots is equivalent to sequentially minimizing $\|C_2\|^2, \|C_3\|^2, \dots$

Theorem 9.2 *The GMA criterion sequentially minimizes the contamination of j -factor interactions on the estimation of main effects for $j = 2, 3, \dots$*

This result was first proved by Tang and Deng (1999) for 2-level designs and extended by Xu and Wu (2001) for general designs.

An issue of the GMA criterion is computation. It is cumbersome to compute the generalized word length pattern according to the definition of $A_j(D)$ in (9.7). Fortunately, Xu and Wu (2001) provided an efficient method for computing the generalized word length patterns via the coding theory. An alternative approach is to use the minimum moment aberration to be discussed in Section 9.4.2.

The *Hamming distance* between two vectors, (x_1, \dots, x_k) and (y_1, \dots, y_k) , is the number of places where they differ, that is, the number of l 's such that $x_l \neq y_l$. For an $n \times k$ design D , let $d_{ij}(D)$ be the Hamming distance between rows i and j and $B_l(D) = n^{-1}|\{(i, j) : d_{ij}(D) = l, i, j = 1, \dots, n\}|$ for $l = 0, 1, \dots, k$. In coding theory, the vector $(B_0(D), B_1(D), \dots, B_k(D))$ is the *distance distribution*. It is obvious that $\sum_{l=0}^k B_l(D) = n$. Xu and Wu (2001) showed that the generalized word length patterns are linear combinations of the distance distributions and vice versa.

Theorem 9.3 *For an $n \times k$ design D with s levels and $j = 0, 1, \dots, k$,*

$$A_j(D) = n^{-1} \sum_{i=0}^k P_j(i; k, s) B_i(D), \tag{9.9}$$

$$B_j(\mathbf{D}) = ns^{-k} \sum_{i=0}^k P_j(i; k, s) A_i(\mathbf{D}), \quad (9.10)$$

where $P_j(x; k, s) = \sum_{i=0}^j (-1)^i (s-1)^{j-i} \binom{x}{i} \binom{k-x}{j-i}$ are the Krawtchouk polynomials.

The equations (9.9) and (9.10) are known as the generalized *MacWilliams identities*, which play a pivotal role in the theoretical development of nonregular designs.

Two designs are called (combinatorially) *isomorphic* if one design can be obtained from the other by permutations of rows, columns, and levels in the columns. Isomorphic designs have the same generalized word length pattern.

Cheng and Ye (2004) proposed another extension of the MA criterion. The Cheng and Ye extension is intended for designs with quantitative factors and depends on the contrasts used in the model. On the other hand, the Xu and Wu extension described in this section is intended for designs with qualitative factors and does not depend on the contrasts used in the model.

9.4.2 Minimum Moment Aberration

Based on coding theory, Xu (2003) proposed the minimum moment aberration criterion for assessing nonregular designs. For an $n \times k$ design \mathbf{D} with s levels and a positive integer m , define the m th power moment to be

$$K_m(\mathbf{D}) = [n(n-1)/2]^{-1} \sum_{1 \leq i < j \leq n} [\delta_{ij}(\mathbf{D})]^m, \quad (9.11)$$

where $\delta_{ij}(\mathbf{D}) = k - d_{ij}(\mathbf{D})$ is the number of coincidences between rows i and j . The *minimum moment aberration* criterion proposed by Xu (2003) is to sequentially minimize the power moments K_1, K_2, \dots

The power moments measure the similarity among runs (i.e., rows). The first and second power moments measure the average and variance of the similarity among runs. Minimizing the power moments makes runs to be as dissimilar as possible. As we will see, the power moments also measure the non orthogonality among columns, and orthogonal designs have small power moments.

Note that the computation of the power moments involves the numbers of coincidences between rows. By applying generalized MacWilliams identities and Pless power moment identities, two fundamental results in the coding theory (MacWilliams and Sloane 1977, Chapter 5; Xu 2003) showed that the power moments K_m defined in (9.11) are linear combinations of the generalized word length patterns A_1, \dots, A_m defined in (9.7).

Theorem 9.4 For an $n \times k$ design \mathbf{D} with s levels and $m = 1, 2, \dots, k$,

$$K_m(\mathbf{D}) = c_m A_m(\mathbf{D}) + c_{m-1} A_{m-1}(\mathbf{D}) + \dots + c_1 A_1(\mathbf{D}) + c_0, \quad (9.12)$$

where $c_m = m!s^{-m}n(n-1)^{-1}$ and c_i are constants depending only on i, n, k, s for $i < m$.

It is not difficult to see now that sequentially minimizing K_1, K_2, \dots is equivalent to sequentially minimizing A_1, A_2, \dots . Therefore, the minimum moment aberration is equivalent to the GMA.

The minimum moment aberration provides a useful tool for efficient computation and theoretical development. For an $n \times k$ design with s levels, the complexity of computing the generalized word length pattern according to the definition (9.7) is $O(ns^k)$, whereas the complexity of computing k power moments is $O(n^2k^2)$. The saving in computation is tremendous when the number of factors k is large. This observation led to successful algorithmic constructions of mixed-level OAs (Xu 2002), a catalog of 3-level regular designs (Xu 2005b), and blocked regular designs with MA (Xu and Lau 2006). As a theoretical tool, Xu (2003) developed a unified theory for nonregular and SSDs. Xu and Lau (2006) and Xu (2006) further used the concept of minimum moment aberration to develop a theory for blocked regular designs and constructed MA blocked regular designs with 32, 64, and 81 runs.

For mixed-level designs, Xu (2003) suggested to use weights to reflect the importance of the factors. For a design $D = (x_{il})$, assign weight $w_l > 0$ to the l th column and let

$$\delta_{ij}(D) = \sum_{l=1}^k w_l \delta(x_{il}, x_{jl}), \tag{9.13}$$

be the weighted coincidence number between the i th and j th rows, where $\delta(x, y) = 1$ if $x = y$ and 0 otherwise. Then define $K_m(D)$ as in (9.11). A special choice for a mixed-level design is to choose weight proportional to its number of levels, say, $w_l = s_l$. For this choice, Xu (2003) showed that if D is an OA of strength t , the identity in (9.12) holds for $m = 1, \dots, t + 1$. Therefore, the minimum moment aberration is weakly equivalent to the GMA for mixed-level designs.

9.4.3 Generalized Resolution and Projection Properties for 2-Level Designs

The resolution, defined in Chapter 1, Section 1.7.2, and Chapter 7, Section 7.2.4, is an important concept for regular designs. To define the generalized resolution, we need to consider projections of a design. For an $n \times k$ design D , a j -factor projection design is the $n \times j$ submatrix representing the j factors. For clarity, we use d to denote a projection.

Suppose that d is an $n \times j$ matrix (d_{ij}) with two levels denoted by -1 and $+1$, and let

$$\rho(d) = \frac{1}{n} \sum_{i=1}^n d_{i1} \times \dots \times d_{ij}. \tag{9.14}$$

The quantity $\rho(d)$ is called a *design moment* in the response surface design literature. If d is a balanced design, $\rho(d)$ is the correlation between the main effect of a column and the interaction involving the other $j-1$ columns. For illustration, consider the 12-run Plackett–Burman design in Table 9.1. For $d = \{A, B\}$, $\rho(d) = 0$ since A and B are orthogonal. For $d = \{A, B, C\}$, $\rho(d) = -1/3$ is the correlation between the main effect A (or B or C) and the 2fi BC (or AC or AB). For $d = \{A, B, C, D\}$, $\rho(d) = -1/3$ is the correlation between the main effect A (or B or C or D) and 3fi BCD (or ACD or ABD or ABC), as well as the correlation between 2fi's AB and CD (or AC and BD , or AD and BC).

The quantity $|\rho(\mathbf{d})|$ is called the normalized J -characteristics by Tang and Deng (1999) or aliasing index by Cheng et al. (2004) and Phoa and Xu (2009) because $0 \leq |\rho(\mathbf{d})| \leq 1$. When $|\rho(\mathbf{d})| > 0$, the k columns in \mathbf{d} form a *word* of length k with aliasing index $|\rho(\mathbf{d})|$. A word is called *complete* if $|\rho(\mathbf{d})| = 1$ or *partial* if $|\rho(\mathbf{d})| < 1$. When $\rho(\mathbf{d}) = 0$, the k columns in \mathbf{d} do not form a word.

It is not difficult to see that if \mathbf{D} is a two-level regular design then $|\rho(\mathbf{d})| = 0$ or 1 where \mathbf{d} is a projection of \mathbf{D} . Ye (2004) showed that the reverse is also true. Therefore, for a non-regular design, there always exists some projection \mathbf{d} such that $0 < |\rho(\mathbf{d})| < 1$. It can be shown that $A_j(\mathbf{D})$, as defined in (9.7), can be expressed as

$$A_j(\mathbf{D}) = \sum_{|\mathbf{d}|=j} \rho^2(\mathbf{d}), \quad (9.15)$$

where the summation is over all j -factor projections \mathbf{d} .

Suppose that r is the smallest integer such that $A_r(\mathbf{D}) > 0$. Then the *generalized resolution* (Deng and Tang 1999) is defined to be

$$\tilde{R} = \tilde{R}(\mathbf{D}) = r + \delta, \quad \text{where } \delta = 1 - \max_{|\mathbf{d}|=r} |\rho(\mathbf{d})|, \quad (9.16)$$

where the maximization is over all projections \mathbf{d} with r factors. Grömping and Xu (2014) recently extended the generalized resolution to OAs with mixed levels.

Example 9.9

Consider the 12-run Plackett–Burman design in Table 9.1. It is an OA of strength 2, so $A_1 = A_2 = 0$ but $A_3 > 0$. It is straightforward to verify that $|\rho(\mathbf{d})| = 1/3$ for any 12×3 subdesign \mathbf{d} . So the generalized resolution is $\tilde{R} = 3 + (1 - 1/3) = 11/3$.

Like the GMA, the generalized resolution has a good statistical justification. For a 2-level OA with strength 2, maximizing the generalized resolution is equivalent to minimizing the maximum bias of any individual 2fi on the estimation of the main effects (Deng and Tang 2002). In contrast, the GMA criterion minimizes the overall bias of all 2fi's on the estimation of the main effects.

The generalized resolution has a nice geometric property. It is easy to see that for an $OA(n, 2^k, t)$, $\rho(\mathbf{d}) = 0$ for any projection \mathbf{d} with t factors or fewer and therefore $r \leq \tilde{R} < r + 1$, where $r = t + 1$. If $\delta > 0$, any projection with $r = t + 1$ factors contains at least $n\delta/2^r$ copies of a full 2^r factorial (Deng and Tang 1999). This result was first proved by Cheng (1995). Cheng (1995, 1998) and Bulutoglu and Cheng (2003) further studied some hidden projection properties of nonregular designs; see Xu et al. (2009) for a review.

Box and Tyssedal (1996) defined a design to be of *projectivity* p if the projection onto every subset of p factors contains a full factorial design, possibly with some points replicated. It follows from these definitions that an OA of strength t is of projectivity at least t . A regular design of resolution r has projectivity $r - 1$, while a nonregular design of generalized resolution $r + \delta$ has projectivity at least r if $\delta > 0$. A result of Cheng (1995) implies that, as long as the run size n is not a multiple of 2^{t+1} , an $OA(n, 2^k, t)$ with $k \geq t + 2$ has projectivity $t + 1$, even though the strength is only t .

Example 9.10

Consider the two 2^{9-5} designs in the chemical toxicity experiment given in Table 9.2. Both designs have the same generalized word length pattern $(0, 0, 4, 14, 8, 0, 4, 1, 0)$. For the regular design, there are four words of length 3 with aliasing index 1; for the nonregular design, there are 16 words of length 3 with aliasing index 0.5. The regular design has generalized resolution 3.0 and projectivity 2, while the nonregular design has generalized resolution 3.5 and projectivity 3.

Deng and Tang (1999) went beyond the generalized resolution and defined another version of the GMA, which we refer to as the *minimum G-aberration* criterion for clarity. Roughly speaking, the minimum G-aberration criterion always chooses a design with the smallest confounding frequency among designs with maximum generalized resolution. Formally, the minimum G-aberration criterion is to sequentially minimize the components in the confounding frequency vector

$$\text{CFV} = [(f_{11}, \dots, f_{1n}); (f_{21}, \dots, f_{2n}); \dots; (f_{k1}, \dots, f_{kn})],$$

where f_{ji} denotes the frequency of j -factor projections \mathbf{d} with $|\rho(\mathbf{d})| = 1 + (1 - i)/n$.

The minimum G_2 -aberration criterion, proposed by Tang and Deng (1999) and mentioned in Section 9.4.1, is a relaxed version of the minimum G-aberration criterion. For 2-level regular designs, both criteria reduce to the traditional MA criterion. However, these two criteria can result in selecting different nonregular designs. Minimum G-aberration nonregular designs always have maximum generalized resolution, whereas minimum G_2 -aberration nonregular designs may not. This is in contrast to the regular case where MA regular designs always have maximum resolution among all regular designs.

9.4.4 Projection Aberration, Estimation Capacity and Design Efficiency

The GMA criterion cannot distinguish designs when they have the same generalized word length pattern. It is useful to examine projections by using a criterion like the minimum G-aberration criterion. There are $\binom{k}{p}$ projected designs with p factors. Each of these designs has an A_p value, which is referred to as the *projected A_p value*, to distinguish the *overall A_p value* calculated from the whole k -factor design. The frequency of the projected A_p values is called the *p -dimensional projection frequency*. For an $OA(n, s^k, t)$, when projecting onto any t factors, we always get a full factorial design. So it is sufficient to consider projection frequencies with $p = t + 1, t + 2$, and so on. The larger the projected A_p values are, the more severe the aliasing is. One approach is to sequentially minimize the $(t + 1)$ -dimensional projection frequency starting from the largest projected A_{t+1} value. If there is a tie, one can further compare the $(t + 2)$ -dimensional projection frequency and so on. This criterion is referred to as the *projection aberration criterion* by Xu et al. (2004) and can serve as an extension of the minimum G-aberration for general designs.

Example 9.11

Consider choosing six columns from the commonly used $OA(18, 3^7, 2)$ given in Table 9.3. There are seven possible choices. For illustration, consider three choices. Let D_1, D_2 , and D_3 be the resulting designs from omitting the first, second, and third columns, respectively. The generalized word length patterns for the three designs are $(0, 0, 10, 22.5, 0, 7)$, $(0, 0, 13, 13.5, 9, 4)$, and $(0, 0, 13, 13.5, 9, 4)$, respectively. Hence, D_1 is the best according to

TABLE 9.6
Overall and Projected A_3 Values

Design	Overall A_3	Frequency of Projected A_3			
		1/2	2/3	1	2
D_1	10	20	0	0	0
D_2	13	16	0	3	1
D_3	13	14	0	6	0

the GMA criterion. Note that D_2 and D_3 have the same generalized word length pattern but they have different projection patterns. The frequencies of projected A_3 values are listed in Table 9.6. Among the three designs, D_2 is the worst under the projection aberration criterion since one of its 3-factor projections has projected $A_3 = 2$; D_1 is again the best because all its 3-factor projections have projected $A_3 = 0.5$.

Xu and Deng (2005) proposed another projection aberration criterion. When considering the projection frequency, they replaced $A_j(\mathbf{d})$, as defined in (9.7), with $K_j(\mathbf{d})$, as defined in (9.11), for a j -factor projection \mathbf{d} . They referred to the resulting criterion as the *moment aberration projection* criterion. The moment aberration projection criterion works the same as the projection aberration criterion based on the A_j values. Both criteria will often select the same best design. However, the moment aberration projection criterion can distinguish more designs than the projection aberration criterion based on the A_j values, even for the 2-level case. The concept of moment aberration projection turns out to be very useful in the algorithmic construction of regular designs; see Xu (2005b, 2009).

For regular designs, Cheng et al. (1999) justified the MA criterion by showing that it is a good surrogate for some model-robustness criteria. Following their approach, Cheng et al. (2002) considered the situation where (1) the main effects are of primary interest and their estimates are required and (2) the experimenter would like to have as much information about $2fi$'s as possible, under the assumption that higher-order interactions are negligible. Without knowing which $2fi$'s are significant, they considered the set of models containing all of the main effects and f $2fi$'s for $f = 1, 2, 3, \dots$. Let E_f be the number of estimable models and D_f be the average of D -efficiencies of all models that contain main effects plus f $2fi$'s. Cheng et al. (2002) showed that two-level GMA designs tend to have large E_f and D_f values, especially for small f ; therefore, the GMA criterion provides a good surrogate for the traditional model-dependent efficiency criteria. Ai et al. (2005) and Mandal and Mukerjee (2005) extended their approach to mixed-level designs.

9.4.5 Uniformity and Space-Filling Property

Uniformity or space filling is a desirable design property for physical and computer experiments (Fang et al. 2006). Various discrepancies have been used to assess the space-filling property or uniformity (Fang and Wang 1994; Fang et al. 2000). These discrepancies all have geometrical meanings and can be interpreted as the difference between the empirical distribution of points in the design and the uniform distribution. Among them, the centered L_2 -discrepancy, proposed by Hickernell (1998), is the most frequently used. For an $n \times k$ design $D = (x_{ij})$ over the unit cube $[0, 1]^k$, the squared centered L_2 -discrepancy (CD) has

an analytic expression as follows:

$$\begin{aligned}
 [CD(\mathbf{D})]^2 &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \prod_{l=1}^k \left(1 + \frac{1}{2} \left| x_{il} - \frac{1}{2} \right| + \frac{1}{2} \left| x_{jl} - \frac{1}{2} \right| - \frac{1}{2} |x_{il} - x_{jl}| \right) \\
 &\quad - \frac{2}{n} \sum_{i=1}^n \prod_{l=1}^k \left(1 + \frac{1}{2} \left| x_{il} - \frac{1}{2} \right| - \frac{1}{2} \left| x_{il} - \frac{1}{2} \right|^2 \right) + \left(\frac{13}{12} \right)^k. \tag{9.17}
 \end{aligned}$$

The centered L_2 -discrepancy is defined over the unit cube $[0, 1]^k$, but the s levels in a factorial design are normally denoted as $0, 1, \dots, s - 1$. Thus, as it is often done in the literature, whenever the centered L_2 -discrepancy is calculated, level i ($i = 0, 1, \dots, s - 1$) should first be transformed to $(2i + 1)/(2s)$. Note that this is a useful relationship only when the s levels of each factor actually represent equally spaced values of an underlying continuous variable.

Fang and Mukerjee (2000) found a connection between aberration and uniformity for 2-level regular designs. This connection was extended by Ma and Fang (2001) for general two-level designs. The basic result states that for a two-level $n \times k$ design \mathbf{D} , regular or nonregular, the centered L_2 -discrepancy can be expressed in terms of its generalized word length pattern $A_i(\mathbf{D})$ as follows:

$$[CD(\mathbf{D})]^2 = \left(\frac{13}{12} \right)^k - 2 \left(\frac{35}{32} \right)^k + \left(\frac{9}{8} \right)^k \left(1 + \sum_{i=1}^k \frac{A_i(\mathbf{D})}{9^i} \right).$$

Since the coefficient of $A_i(\mathbf{D})$ decreases exponentially with i , one can anticipate that designs with small $A_i(\mathbf{D})$ for small values of i should have small $CD(\mathbf{D})$; in other words, GMA designs tend to be uniform over the design region.

However, the result cannot be generalized to multi level designs directly, as level permutations of one or more factors can alter the centered L_2 -discrepancy, but keep the generalized word length pattern unchanged. By considering level permutations of three-level designs, Tang et al. (2012) established a relationship between average centered L_2 -discrepancy and generalized word length pattern. Tang and Xu (2013) generalized the relationship to designs with arbitrary levels. Zhou and Xu (2014) further generalized their results for any discrepancy defined by a reproducing kernel and showed that GMA designs have good space-filling properties on average in terms of distance as well.

Hickernell and Liu (2002) showed that the GMA criterion could be defined and generalized using discrepancy. Tang (2001) and Ai and Zhang (2004a) showed that GMA designs have good low-dimensional projection properties.

9.5 Construction Methods

The construction of optimal nonregular designs is challenging for two simple reasons: (1) nonregular designs do not have a unified mathematical description and (2) the class of nonregular designs is much larger than the class of regular designs.

9.5.1 Algorithmic Methods

A simple strategy for constructing GMA or other optimal designs is to search over all possible projection designs from existing OAs, such as Plackett–Burman designs or Hadamard matrices. A Hadamard matrix of order n is an $n \times n$ matrix of entries $+1$ and -1 whose columns (and rows) are orthogonal to each other; see Chapter 7 of Hedayat et al. (1999). A Hadamard matrix is said to be normalized if all of the entries of the first row and column are $+1$. From a normalized Hadamard matrix of order n , one obtains a saturated $OA(n, 2^{n-1}, 2)$ by deleting the first column. Deng and Tang (2002) presented a catalog of GMA designs by searching over Hadamard matrices of order 16, 20, and 24. Xu and Deng (2005) searched optimal designs over all Hadamard matrices of order 16 and 20 and 3-level designs from 68 saturated $OA(27, 3^{13}, 2)$'s. A limitation of this strategy is that we could miss the optimal design in some cases because the optimal design cannot be expressed as such a projection.

Much effort has been devoted to the complete enumeration of all non isomorphic designs with a small number of runs. Sun et al. (2002) proposed an algorithm for sequentially constructing non isomorphic orthogonal designs. They successfully enumerated all 2-level OAs with 12, 16, and 20 runs. An important result is that all 16-run OAs are projections of one of the five 16-run Hadamard matrices. However, such a result does not hold for 20-run designs. Ye et al. (2008) presented a complete set of combinatorially nonisomorphic OAs of types $OA(12, 2^k 3^1, 2)$, $OA(18, 3^k, 2)$, $OA(18, 2^1 3^k, 2)$, and $OA(20, 2^k 5^1, 2)$. Schoen (2009) also presented all OAs with 18 runs.

Schoen et al. (2010) proposed a general algorithm that can also handle mixed-level designs. They successfully enumerated most nontrivial mixed-level OAs up to 28 runs with strength 2, 64 runs with strength 3, and 168 runs with strength 4. They completely enumerated all 24-run OAs with strength 2 and 28-run OAs up to 7 columns. The number of non isomorphic designs $OA(28, 2^k, 2)$ is 4, 7, 127, 17,826, and 5,882,186, respectively, for $k = 3, 4, 5, 6, 7$.

Algorithmic constructions are typically limited to small run sizes (≤ 32) or small number of factors due to the existence of a large number of designs and the difficulty of determining whether two designs are isomorphic or not. Algebraic or combinatorial methods are necessary to construct larger designs. A good construction method is the quaternary code (QC) method introduced by Xu and Wong (2007).

9.5.2 Quaternary Code Designs

A QC is a linear subspace over $Z_4 = \{0, 1, 2, 3\} \pmod{4}$, the ring of integers modulus 4. A key device is the so-called Gray map:

$$\phi : 0 \rightarrow (0, 0), 1 \rightarrow (0, 1), 2 \rightarrow (1, 1), 3 \rightarrow (1, 0), \quad (9.18)$$

which maps each symbol in Z_4 to a pair of symbols in Z_2 . Let G be an $a \times b$ matrix with elements from Z_4 and let C consist of all possible linear combinations of the row vectors of G over Z_4 . Applying the Gray map to C , one obtains a $4^a \times 2b$ matrix or a two-level design with 4^a runs and $2b$ factors, denoted by D . Although C is linear over Z_4 , D may or may not be linear over Z_2 . From D , we can construct a two-level design with 2^{2a+1} runs and $4b$ factors via the *doubling method* as follows:

$$D^* = \begin{pmatrix} D & D \\ D & D + 1 \end{pmatrix} \pmod{2}. \quad (9.19)$$

TABLE 9.7

Example of Quaternary Code and Nonregular Design

(a) Quaternary Code C						(b) Nonregular Design D													
Run	1	2	3	4	5	6	Run	1	2	3	4	5	6	7	8	9	10	11	12
1	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0
2	0	1	1	2	1	3	2	0	0	0	1	0	1	1	1	0	1	1	0
3	0	2	2	0	2	2	3	0	0	1	1	1	1	0	0	1	1	1	1
4	0	3	3	2	3	1	4	0	0	1	0	1	0	1	1	1	0	0	1
5	1	0	2	1	1	1	5	0	1	0	0	1	1	0	1	0	1	0	1
6	1	1	3	3	2	0	6	0	1	0	1	1	0	1	0	1	1	0	0
7	1	2	0	1	3	3	7	0	1	1	1	0	0	0	1	1	0	1	0
8	1	3	1	3	0	2	8	0	1	1	0	0	1	1	0	0	0	1	1
9	2	0	0	2	2	2	9	1	1	0	0	0	0	1	1	1	1	1	1
10	2	1	1	0	3	1	10	1	1	0	1	0	1	0	0	1	0	0	1
11	2	2	2	2	0	0	11	1	1	1	1	1	1	1	1	0	0	0	0
12	2	3	3	0	1	3	12	1	1	1	0	1	0	0	0	0	1	1	0
13	3	0	2	3	3	3	13	1	0	0	0	1	1	1	0	1	0	1	0
14	3	1	3	1	0	2	14	1	0	0	1	1	0	0	1	0	0	1	1
15	3	2	0	3	1	1	15	1	0	1	1	0	0	1	0	0	1	0	1
16	3	3	1	1	2	0	16	1	0	1	0	0	1	0	1	1	1	0	0

Source: Xu, H. and Wong, A., *Stat. Sinica*, 17, 1191, 2007.

Example 9.12

Consider a 2×6 matrix

$$G = \begin{bmatrix} 1 & 0 & 2 & 1 & 1 & 1 \\ 0 & 1 & 1 & 2 & 1 & 3 \end{bmatrix}.$$

All linear combinations of the two rows of G form a 16×6 linear code C over Z_4 . Applying the Gray map, we obtain a 16×12 matrix $D = \phi(C)$, which is a 2^{12-8} design. See Table 9.7 for the matrices C and D . It is straightforward to verify that D has generalized resolution 3.5; therefore, it is a nonregular design. Moreover, the 32×24 matrix D^* obtained via doubling as in (9.19) is a 2^{24-19} design and has generalized resolution 3.5 too. For comparison, in both cases the best regular design of the same size has resolution 3.

Example 9.13

Consider a 4×8 matrix

$$G = \begin{bmatrix} 1 & 0 & 0 & 0 & 2 & 1 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 & 3 & 1 & 2 \\ 0 & 0 & 1 & 0 & 1 & 2 & 3 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 & 2 & 3 \end{bmatrix}.$$

All linear combinations of the rows of G over Z_4 form a 256×8 quaternary linear code C . Applying the Gray map, we obtain a 256×16 matrix $D = \phi(C)$, which is isomorphic to the (extended) Nordstrom–Robinson code (Xu 2005a). The resulting design D is an $OA(256, 2^{16}, 5)$ with many remarkable properties: it has generalized resolution 6.5 and projectivity 7. For comparison, for a regular design to achieve the same resolution and

projectivity, it would require at least 512 runs. For more statistical properties and results from the Nordstrom–Robinson code, see Xu (2005a).

Lemma 9.1 *Let G be an $a \times b$ matrix over Z_4 , C be the quaternary linear code generated by G , and $D = \phi(C)$ be the binary image. Then D is an OA of strength 2 if and only if G satisfies the following conditions:*

- (i) *It does not have any column containing entries 0 and 2 only.*
- (ii) *None of the columns is a multiple of another column over Z_4 .*

Xu and Wong (2007) further showed that if the two conditions in Lemma 9.1 are satisfied, then D has generalized resolution at least 3.5 and G has a maximum of $(4^a - 2^a)/2$ columns. Such a matrix can be constructed as follows:

1. Write down all possible columns of a elements over Z_4 .
2. Delete columns that do not contain any 1s.
3. Delete columns whose first non zero and non-two entries are 3s.

Theorem 9.5 *For an integer $a > 1$, let G be the generator matrix obtained from the aforementioned procedure. Then the binary image D generated by G has 4^a rows, $4^a - 2^a$ columns, and generalized resolution 3.5. The double D^* of D has 2^{2a+1} rows, $2^{2a+1} - 2^{a+1}$ columns, and generalized resolution 3.5.*

Note that the nonregular designs constructed in Theorem 9.5 have generalized resolution 3.5. It is known that a regular 2-level design with n runs and k factors has resolution at most 3 when $k > n/2$; see Chapter 7, Section 7.4.1. Therefore, nonregular designs constructed from QCs have higher resolution than corresponding regular designs when resolution 4 designs do not exist.

Since QC designs are linear over Z_4 , we can enumerate QC designs sequentially in a similar manner as enumerating regular designs. Xu and Wong (2007) developed a sequential algorithm, similar to those by Chen et al. (1993) and Xu (2005b). They also presented a collection of nonregular designs with 32, 64, 128, and 256 runs and up to 64 factors, many of which are better than regular designs of the same size in terms of resolution, aberration, and projectivity.

The linear structure of a QC also facilitates the derivation and analytical study of properties of QC designs. Phoa and Xu (2009) studied quarter-fraction QC designs, which are defined by a generator matrix that consists of an identity matrix plus an extra column. They showed that the generalized resolution, generalized word length pattern, and projectivity can be calculated in terms of the frequencies of the numbers 1, 2, and 3 that appear in the extra column.

Specifically, consider an $a \times (a + 1)$ generator matrix $G = (v, I_a)$, where v is an $a \times 1$ column vector over Z_4 and I_a is the $a \times a$ identity matrix. The binary image D generated by G is a $2^{(2a+2)-2}$ design. It is easy to verify that the identity matrix I_a generates a full 2^a design; therefore, the properties of D depend on the column v only. For $i = 0, 1, 2, 3$, let f_i be the number of times that number i appears in column v . Phoa and Xu (2009) showed

that the number of words of D , their lengths and aliasing indexes can be expressed in terms of the frequency f_i .

Theorem 9.6 *The $2^{(2a+2)-2}$ design D generated by $G = (v, I_a)$ has 1 complete word of length $2f_1 + 2f_3 + 2$ and $2/\rho^2$ words of length $f_1 + 2f_2 + f_3 + 1$ with aliasing index $\rho = 2^{-\lfloor (f_1+f_3)/2 \rfloor}$, where $\lfloor x \rfloor$ is the integer value of x .*

Since D has a complete word of length $2(f_1 + f_3) + 2$, its projectivity is at most $2(f_1 + f_3) + 1$. The following theorem from Phoa and Xu (2009) shows that the projectivity of D is not affected by the partial words.

Theorem 9.7 *Suppose that D is the $2^{(2a+2)-2}$ design generated by $G = (v, I_a)$:*

- (a) *If $f_2 > 0$, the projectivity of D is $2(f_1 + f_3) + 1$.*
- (b) *If $f_2 = 0$ and $f_1 + f_3 > 0$, the projectivity of D is $2(f_1 + f_3) - 1$.*

Based on these theoretical results, Phoa and Xu (2009) constructed optimal quarter-fraction QC designs under the maximum resolution, GMA, and maximum projectivity criteria. These optimal QC designs are often better than regular designs of the same size in terms of the design criterion. The GMA QC designs have the same aberration as the MA regular designs and frequently with larger resolution and projectivity. A maximum projectivity QC design is often different from a MA or maximum resolution design but can have much larger projectivity than a MA regular design. They further showed that some of these QC designs have GMA and maximum projectivity among all possible designs. Zhang et al. (2011) and Phoa et al. (2012) investigated 1/8 and 1/16 fraction QC designs, which are defined by a generator matrix that consists of an identity matrix plus two additional columns. Phoa (2012) further studied 1/64 fraction QC designs.

9.6 Optimality Results

It is infeasible to search over all possible designs in many situations. Theoretical results are extremely useful to determine whether a design is optimal under the GMA or other criteria. Xu (2003) gave several sufficient conditions for a design to have GMA among all possible designs using the concept of minimum moment aberration.

One sufficient condition is that the numbers of coincidences between distinct rows are constant or differ by at most one.

Theorem 9.8 *Design D has GMA among all possible designs if the differences among all $\delta_{ij}(D)$, $i < j$, do not exceed one.*

In other words, a design has GMA if its design points are equally or nearly equally spaced over the design region. As an example, the 12-run Plackett–Burman design in

Table 9.1 has GMA because the numbers of coincidences between any two distinct rows are 5. It is easy to see that deleting any column yields an $OA(12, 2^{10}, 2)$, which has GMA among all possible designs too.

An important class of designs that satisfy the conditions in Theorem 9.8 are saturated OAs of strength 2. Mukerjee and Wu (1995) showed the following result.

Lemma 9.2 *The numbers of coincidences between any distinct pair of rows of a saturated $OA(n, s^k, 2)$ are constant; specifically, $\delta_{ij}(\mathbf{D}) = (n - s)/(s^2 - s)$ for any $i \neq j$.*

Another sufficient condition relates to projections of a design.

Theorem 9.9 *Design \mathbf{D} has GMA among all possible designs if \mathbf{D} is an $OA(n, s^k, t)$ and there are no repeated runs in any $(t + 1)$ -factor projection.*

For example, consider the $OA(18, 3^6, 2)$ given by columns 2–7 in Table 9.3. It is easy to verify that its projection onto any three columns does not have repeated runs. Thus, this design (and any of its projections) has GMA among all possible designs.

Another general technique for constructing optimal designs is linear programming, which employs the generalized MacWilliams identities (9.9) and (9.10). The linear programming technique has been used to establish bounds on the maximum size of a code for given length and distance in coding theory (MacWilliams and Sloane 1977, Section 17.4) and bounds on the minimum size of an OA for given number of factors and strength (Hedayat et al. 1999, Section 4.5). Xu (2005a) used linear programming to show that several nonregular designs derived from the Nordstrom–Robinson code have GMA among all possible designs. The following result is from Xu (2005a).

Theorem 9.10 *Any regular 2^{k-2} MA design has GMA among all possible designs.*

Butler (2003, 2004) presented a number of construction results that allow 2-level GMA designs to be found for many of the cases with $n = 16, 24, 32, 48, 64,$ and 96 runs. Butler (2005) further developed theoretical results and presented methods that allow GMA designs to be constructed for more than two levels. A key tool of Butler (2003, 2004, 2005) is the use of some identities that link the generalized word length patterns with moments of the inner products or Hamming distances between the rows; see also Chapter 7, Section 7.6. These identities can be derived easily from the generalized Pless power moment identities developed by Xu (2003).

9.7 Supersaturated Designs

The study of SSDs dates back to Satterthwaite (1959) and Booth and Cox (1962). The former suggested the use of random balanced designs and the latter proposed an algorithm to

construct systematic SSDs. Many methods have been proposed for constructing two-level SSDs in the last two decades after Lin (1993) and Wu (1993). The early construction methods use Hadamard matrices or balanced incomplete block designs; see, among others, Lin (1993), Wu (1993), Nguyen (1996), Cheng (1997), and Tang and Wu (1997). Early algorithmic construction includes Lin (1995), Nguyen (1996), and Li and Wu (1997). Chapter 1, Section 1.7.4 gives some other references. Georgiou (2014) gave a review of construction methods and provided many additional references.

Lin (1993) used half fractions of Hadamard matrices to construct two-level SSDs. First obtain a saturated $OA(n, 2^{n-1}, 2)$ from a normalized Hadamard matrix of order n (by deleting the first column that is a column of ones). For example, we obtain the 12×11 Plackett–Burman design in Table 9.1 from a normalized Hadamard matrix of order 12. Now for each column, half of the entries are 1 and the other half are -1 . Use any column as the branching column and take those rows whose entries are 1 in the branching column. Deleting the branching column yields an SSD with $n/2$ runs and $n - 2$ columns (provided that the resulting design has no repeated columns).

Wu (1993) proposed another construction method by utilizing partial aliasing of $2fi$'s among Plackett–Burman designs or Hadamard matrices. Consider the 12×11 design matrix in Table 9.1 again for illustration. There are $\binom{11}{2} = 45$ $2fi$'s. None of the 45 $2fi$'s are fully aliased with the original 11 columns or other $2fi$'s. By combining the 45 columns with the original 11 columns, we obtain an SSD with 12 runs and 66 columns.

A popular criterion in the SSD literature is the $E(s^2)$ criterion (Booth and Cox 1962; Lin 1993). For a balanced $n \times k$ design D with two levels denoted by 1 and -1 ,

$$E(s^2) = \sum_{1 \leq i < j \leq k} s_{ij}^2 / [k(k - 1)/2],$$

with $s_{ij} = c_i'c_j$, where c_i and c_j are the i th and j th columns of D . Nguyen (1996) and Tang and Wu (1997) independently derived the following lower bound for two-level SSDs with n runs and k factors:

$$E(s^2) \geq n^2(k - n + 1) / [(k - 1)(n - 1)]. \tag{9.20}$$

The GMA criterion can be used to assess general SSDs, including mixed-level designs. Following the discussion of $A_2(D)$ in Section 9.4.1, it is easy to see that $E(s^2) = 2n^2A_2(D) / [k(k - 1)]^2$ for two-level SSDs. Therefore, the GMA criterion can be viewed as a refinement of the $E(s^2)$ criterion and the general optimality results on the GMA criterion can be applied to the SSDs directly. For example, using the connection between the minimum moment aberration and GMA, Xu and Wu (2005) obtained the following result regarding multi-level SSDs, which include many previous results as special cases.

Theorem 9.11 For an SSD D with n runs and k factors at s levels,

$$A_2(D) \geq \frac{k(s - 1)(ks - k - n + 1)}{2(n - 1)} + \frac{(n - 1)s^2\eta(1 - \eta)}{2n}, \tag{9.21}$$

where $\eta = k(n-s)/(ns-s) - \lfloor k(n-s)/(ns-s) \rfloor$ is the fractional part of $k(n-s)/(ns-s)$. The lower bound is achieved if and only if the numbers of coincidences, $\delta_{ij}(\mathbf{D})$, differ by at most one for all $i < j$. Furthermore, an SSD achieving the lower bound is optimal under the GMA criterion.

When applied to two-level SSDs, Theorem 9.11 improves the lower bound (9.20) whenever $\eta > 0$. Butler et al. (2001) and Bulutoglu and Cheng (2004) gave further improvements on the lower bound (9.20) for two-level SSDs.

Many optimal SSDs that achieve the lower bound in Theorem 9.11 can be derived from saturated OAs. A key property of saturated OAs, stated in Lemma 9.2, is that the numbers of coincidences, $\delta_{ij}(\mathbf{D})$, are constant for any pair of rows $i \neq j$.

Tang and Wu (1997) proposed construction of optimal two-level SSDs by juxtaposing saturated OAs derived from Hadamard matrices. This method can be easily extended to construct optimal multilevel SSDs. Suppose $\mathbf{D}_1, \dots, \mathbf{D}_f$ are f saturated $OA(n, s^k, 2)$ with $k = (n-1)/(s-1)$. Let $\mathbf{D} = \mathbf{D}_1 \cup \dots \cup \mathbf{D}_f$ be the $n \times fk$ array obtained by column juxtaposition, which may have duplicated or fully aliased columns. It is evident that $\delta_{ij}(\mathbf{D})$ are constant for any $i < j$. Then by Theorem 9.11, \mathbf{D} is an optimal SSD under the GMA criterion. The conclusion may no longer be valid if repeated columns are removed.

When $n = s^2$, Lemma 9.2 implies that the numbers of coincidences between any two rows are equal to 1 for a saturated $OA(n, s^k, 2)$. Then removing any number of orthogonal columns from \mathbf{D} also results in an optimal SSD under GMA, because the resulting design has the property that the numbers of coincidences between any two rows differ by at most one. In particular, for any k , the lower bound in Theorem 9.11 is tight when $n = s^2$.

The half fraction method of Lin (1993) can be easily extended to construct multilevel SSDs as follows. Taking any column of a saturated $OA(n, s^k, 2)$ as the branching column, we obtain s fractions according to the levels of the branching column. After removing the branching column, the fractions have the properties that all columns are balanced and the numbers of coincidences between any two rows are constant. The row juxtaposition of any f fractions form an SSD with fn/s rows and $k-1$ columns, of which the numbers of coincidences between any two rows differ by at most one. By Theorem 9.11, such a design is optimal under GMA. For $n = s^2$, any subdesign is also optimal, because the numbers of coincidences between any two rows are either 0 or 1.

For any prime power s and integer $p > 0$, a saturated $OA(n, s^k, 2)$ exists where $n = s^p$ and $k = (n-1)/(s-1)$. The following result is from Xu and Wu (2005).

Theorem 9.12 *Let s be a prime power. There exists an optimal s -level $n \times k$ SSD under the GMA criterion that achieves the lower bound in Theorem 9.11 when any of the following conditions hold:*

- (i) $n = s^p$ and $k = f(s^p - 1)/(s - 1)$ where $p > 0$ and $f > 0$ are integers.
- (ii) $n = fs^{p-1}$ and $k = (s^p - 1)/(s - 1) - 1$ where $p > 0$ and $0 < f < s$ are integers.
- (iii) $n = s^2$ and any integer $k > 0$.
- (iv) $n = fs$ and any integer $0 < k \leq s$ where $0 < f < s$ is an integer.

The aforementioned optimal SSDs may contain fully aliased columns, which are not useful in practice. To further distinguish designs with the same (overall) A_2 values, we consider their 2-factor projections and apply the generalized resolution for 2-level designs or the projection aberration idea in general; see Sections 9.4.3 and 9.4.4. To maximize the generalized

resolution is equivalent to minimize the maximum absolute correlation between any two columns, or the $\max(s^2) = \max_{i < j} s_{ij}^2$ criterion in the literature. While there are abundant results on the $E(s^2)$ criterion or its extensions, there are relatively few results on the $\max(s^2)$ criterion. Cheng and Tang (2001) studied the maximum number of factors that an SSD can have under the constraint on $\max(s^2)$.

Xu and Wu (2005) presented explicit construction methods that produce optimal SSDs without fully aliased columns using linear and quadratic functions over finite fields. The construction method was closely related to the Addelman and Kempthorne (1961) construction method of $OA(2s^p, s^k, 2)$ with $k = 2(s^p - 1)/(s - 1)$. Numerical comparisons for small 3-, 4-, and 5-level SSDs indicate that their algebraic method produces good SSDs.

9.8 Analysis Strategies

We begin with a discussion of analysis of nonregular factorial experiments and toward the end give references on analysis for SSDs.

The analysis strategy proposed by Hamada and Wu (1992) consists of three steps:

- Step 1: Entertain all the main effects and interactions that are orthogonal to the main effects. Use standard analysis methods such as ANOVA and half-normal plots to select significant effects.
- Step 2: Entertain the significant effects identified in the previous step and 2fi's that include at least one factor that has a significant main effect. Identify significant effects using a forward selection regression procedure.
- Step 3: Entertain the significant effects identified in the previous step and all the main effects. Identify significant effects using a forward selection regression procedure.

Iterate between steps 2 and 3 until the selected model stops changing. Note that the traditional analysis of Plackett–Burman or other nonregular designs ends at step 1.

Hamada and Wu (1992) based their analysis strategy on two empirical principles, *effect sparsity* and *effect heredity* (Wu and Hamada 2009, Section 4.6). Effect sparsity implies that only few main effects and even fewer 2fi's are relatively important in a factorial experiment. Effect heredity means that in order for an interaction to be included, at least one of the main effects associated with its parent factors should be included. In other words, effect heredity excludes models that contain an interaction but none of its parent main effects. Hamada and Wu (1992) wrote that the strategy works well when both principles hold and the correlations between partially aliased effects are small to moderate. The effect sparsity assumption suggests that only a few iterations will be required.

Using this procedure, Hamada and Wu (1992) reanalyzed data from three real experiments: a cast fatigue experiment using a 12-run Plackett–Burman design with seven 2-level factors, a blood glucose experiment using an 18-run mixed-level OA with one 2-level and seven 3-level factors, and a heat exchange experiment using a 12-run Plackett–Burman design with ten 2-level factors. They demonstrated that the traditional main effects analysis was limited and the results were misleading. Phoa et al. (2009) gave three more real examples to show the importance of considering interactions for screening experiments.

Hamada and Wu (1992) discussed limitations of their analysis strategy and provided solutions. Wu and Hamada (2009, Chapter 8) suggested some further extensions such as the use of all subset variable selection if possible.

Example 9.14

Consider the HPLC experiment in Example 9.1. The traditional main effects analysis shows that the two most important factors are E and F . The model that consists of only the main effects of E and F has $R^2 = 0.41$. Using the Hamada–Wu analysis strategy, we find a significant EF interaction in step 2. Adding EF to E and F increases R^2 from 0.41 to 0.89. In step 3, we further identify factor H , which is missed in the traditional approach, as significant at the 5% level. We repeat steps 2 and 3 iteratively until no more new significant effects are identified and the model does not change anymore. When this happens, we stop the procedure and report the final model, which is given in (9.2).

Box and Meyer (1993) proposed a Bayesian approach by considering all the possible explanations (models including interactions) of the data from a screening experiment and identifying those that fit the data well. The assumptions for prior distributions in their approach are as follows:

1. Each factor has independent prior probability π being active.
2. All effects from a model are assigned independent prior normal distributions with mean 0 and variance $\gamma^2\sigma^2$.
3. A noninformative prior distribution is employed for experimental noise σ .

The prior probability of a model with f active factors is $\pi^f(1-\pi)^{k-f}$ for a k -factor design. The model with f active factors includes main effects for each active factor and all of their interactions (up to any desired order). The parameter γ captures the magnitude of the effects relative to experimental noise σ . Box et al. (2005) suggested to choose $\pi = 0.25$ and γ between 2 and 3, based on a survey of a number of published analyses of factorial experiments. The results are not very sensitive to moderate changes in π and γ when active factors are present.

A Bayesian framework is used to assign posterior probabilities to all the models considered; see Box and Meyer (1993) or Box et al. (2005). These posterior probabilities are then accumulated to marginal posterior probabilities for each factor. A factor that has a relatively high posterior probability implies that either its main effect or an interaction involving it or both are important.

Example 9.15

We analyze the HPLC experiment in Example 9.1 via the Bayesian approach. The posterior probability plot in Figure 9.1a shows the marginal posterior probabilities for each factor with $\pi = 0.25$ and $\gamma = 2$. The posterior probability is high for factors E and F , moderate for factor H , and small for other factors. This suggests that factors E , F , and H are active. However, the marginal posterior probabilities do not tell which 2fi's are significant. Since the frequentist approach identifies the EF interaction as significant, we perform a second Bayesian analysis by treating the EF interaction as a new factor. The resulting posterior probability plot in Figure 9.1b shows that EF is as significant as factors E and F . Factor H is also significant, but not as significant as E , F , and EF . The finding is consistent with the frequentist approach.




FIGURE 9.1 The posterior probability plot for the HPLC experiment (a) with the original factors and (b) with the *EF* interaction.

Chipman et al. (1997) proposed a more sophisticated Bayesian approach for analyzing data with complex aliasing. They employed a Gibbs sampler to perform an efficient stochastic search of the model space, whereas Box and Meyer (1993) evaluated all possible models, which could require intensive computation for large data sets. In addition, Chipman et al. (1997) carefully implemented the effect sparsity and effect heredity principles with hierarchical models. They further introduced two types of effect heredity: weak and strong heredity. Under weak heredity, a 2-factor interaction is important only if at least one of its component factors is significant, while under strong heredity, both of its component factors have to be significant.

Yuan et al. (2007) proposed an efficient variable selection approach based on the least angle regression (LARS) algorithm of Efron et al. (2004). They modified the LARS algorithm so that heredity principles can be taken into account in the variable selection process.

When all factors are quantitative, it is natural to consider a polynomial model to explore the relationship between the response and factors. For k quantitative factors, denoted by x_1, \dots, x_k , the second-order model is

$$y = \beta_0 + \sum_{i=1}^k \beta_i x_i + \sum_{i=1}^k \beta_{ii} x_i^2 + \sum_{1 \leq i < j}^k \beta_{ij} x_i x_j + \epsilon, \tag{9.22}$$

where $\beta_0, \beta_i, \beta_{ii}, \beta_{ij}$ are unknown parameters, and ϵ is the error term. For the pure quadratic terms β_{ii} to be estimated, all the factors must have more than two levels. The second-order model (9.22) has $p = (k + 1)(k + 2)/2$ parameters. When the run size n is less than p , we cannot estimate all the parameters in (9.22). A traditional approach is to assume that the bilinear (or interaction) terms β_{ij} are negligible and fit a model with the linear and pure quadratic terms only, which is the main effects model for three-level designs. However, nonnegligible interaction terms will bias the estimate of linear and pure quadratic terms. A better approach is to use the Hamada–Wu strategy and perform variable selection guided by the effect heredity principle.

Cheng and Wu (2001) proposed an alternative analysis strategy in order to achieve the dual purpose of factor screening and response surface exploration using a single design. Their analysis strategy has two stages:

Stage 1: Perform factor screening and identify important factors.

Stage 2: Fit a second-order model for the factors identified in stage 1.

Various screening analyses can be utilized in stage 1, such as the conventional ANOVA or half-normal plots of the main effects, which include both linear and pure quadratic terms for 3-level factors. Their analysis strategy again assumes that effect sparsity and effect heredity principles hold. They reanalyzed a PVC insulation experiment reported by Taguchi (1987) that used a regular 27-run design with nine 3-level factors. They identified a significant linear-by-linear interaction effect that was missed by Taguchi. Xu et al. (2004) gave another example, which uses an 18-run OA with one 2-level factor and seven 3-level factors.

Finally, for SSDs we typically consider the main effects only with the assumption that all 2fi's are negligible. In principle, any variable selection procedures can be used for analyzing SSDs. Many analysis strategies have been used to analyze SSDs. The list includes forward stepwise regressions (Lin 1993; Westfall et al. 1998), all subset regressions (Abraham et al. 1999), Bayesian variable selections (Chipman et al. 1997), penalized least squares (Li and Lin 2003), partial least squares (Zhang et al. 2007), and the Dantzig selector (Phoa et al. 2009).

9.9 Concluding Remarks

We give an overview of recent developments in nonregular fractional designs and SSDs in this chapter. In summary, nonregular designs are more flexible, require smaller numbers of runs, and have better statistical properties than regular designs. Yet the analysis of nonregular designs is more complicated due to the partial aliasing among the effects. Xu et al. (2009) highlighted some future directions of research, from applications and analysis of nonregular designs to construction of good nonregular designs with large run sizes and optimality results with respect to the generalized resolution.

One underlying assumption for the GMA criterion in Section 9.4.1 is that factor levels are regarded as nominal symbols. This is appropriate for experiments with qualitative factors where there is no ordering among the levels. However, for experiments with quantitative factors, polynomial models such as (9.22) or other models are often used to describe the relationship between the response and the factors. In these circumstances, permuting levels for one or more factors can lead to designs with different geometrical structures and statistical properties. An important question is how the levels should be permuted for quantitative factors. There are some recent developments on this topic; see Tang et al. (2012), Tang and Xu (2013, 2014), and Zhou and Xu (2014). However, more work needs to be done.

Acknowledgment

The research was supported by National Science Foundation grant DMS-1106854.

References

- Abraham, B., Chipman, H., and Vijayan, K. (1999). Some risks in the construction and analysis of supersaturated designs. *Technometrics*, **41**, 135–141.
- Addelman, S. and Kempthorne, O. (1961). Some main-effect plans and orthogonal arrays of strength two. *The Annals of Mathematical Statistics*, **32**, 1167–1176.
- Ai, M. Y., Li, P. F., and Zhang, R. C. (2005). Optimal criteria and equivalence for nonregular fractional factorial designs. *Metrika*, **62**, 73–83.
- Ai, M. Y. and Zhang, R. C. (2004). Projection justification of generalized minimum aberration for asymmetrical fractional factorial designs. *Metrika*, **60**, 279–285.
- Booth, K. H. V. and Cox, D. R. (1962). Some systematic supersaturated designs. *Technometrics*, **4**, 489–495.
- Box, G. E. P. and Draper, N. R. (1987). *Empirical Model-Building and Response Surfaces*. New York: Wiley.
- Box, G. E. P. and Meyer, R. D. (1993). Finding the active factors in fractionated screening experiments. *Journal of Quality Technology*, **25**, 94–105.
- Box, G. E. P. and Tyssedal, J. (1996). Projective properties of certain orthogonal arrays. *Biometrika*, **83**, 950–955.
- Box, G. E. P., Hunter, W. G., and Hunter, J. S. (2005). *Statistics for Experimenters: Design, Innovation, and Discovery*, 2nd edn. New York: Wiley.
- Bulutoglu, D. A. and Cheng, C. S. (2003). Hidden projection properties of some nonregular fractional factorial designs and their applications. *Annals of Statistics*, **31**, 1012–1026.
- Bulutoglu, D. A. and Cheng, C. S. (2004). Construction of $E(s^2)$ -optimal supersaturated designs. *Annals of Statistics* **32**, 1662–1678.
- Butler, N. A. (2003). Minimum aberration construction results for nonregular two-level fractional factorial designs. *Biometrika*, **90**, 891–898.
- Butler, N. A. (2004). Minimum G_2 -aberration properties of two-level foldover designs. *Statistics and Probability Letters*, **67**, 121–132.
- Butler, N. A. (2005). Generalised minimum aberration construction results for symmetrical orthogonal arrays. *Biometrika*, **92**, 485–491.
- Butler, N. A., Mead, R., Eskridge, K. M., and Gilmour, S. G. (2001). A general method of constructing $E(s^2)$ -optimal supersaturated designs. *Journal of Royal Statistical Society: Series B*, **63**, 621–632.
- Chen, J., Sun, D. X., and Wu, C. F. J. (1993). A catalogue of two-level and three-level fractional factorial designs with small runs. *International Statistical Review*, **61**, 131–145.
- Cheng, C. S. (1980). Orthogonal arrays with variable numbers of symbols. *Annals of Statistics*, **8**, 447–453.
- Cheng, C. S. (1995). Some projection properties of orthogonal arrays. *Annals of Statistics*, **23**, 1223–1233.
- Cheng, C. S. (1997). $E(s^2)$ -optimal supersaturated designs. *Statistica Sinica*, **7**, 929–939.
- Cheng, C. S. (1998). Some hidden projection properties of orthogonal arrays with strength three. *Biometrika*, **85**, 491–495.
- Cheng, C. S., Deng, L. Y., and Tang, B. (2002). Generalized minimum aberration and design efficiency for nonregular fractional factorial designs. *Statistica Sinica*, **12**, 991–1000.
- Cheng, C. S., Steinberg, D. M., and Sun, D. X. (1999). Minimum aberration and model robustness for two-level fractional factorial designs. *Journal of the Royal Statistical Society: Series B*, **61**, 85–93.
- Cheng, S. W. and Wu, C. F. J. (2001). Factor screening and response surface exploration (with discussion). *Statistica Sinica*, **11**, 553–604.
- Cheng, S. W. and Ye, K. Q. (2004). Geometric isomorphism and minimum aberration for factorial designs with quantitative factors. *Annals of Statistics*, **32**, 2168–2185.
- Cheng, S. W., Li, W., and Ye, K. Q. (2004). Blocked nonregular two-level factorial designs. *Technometrics*, **46**, 269–279.

- Chipman, H., Hamada, M. and Wu, C.F.J. (1997). A Bayesian variable-selection approach for analyzing designed experiments with complex aliasing. *Technometrics*, **39**, 372–381.
- Deng, L. Y. and Tang, B. (1999). Generalized resolution and minimum aberration criteria for Plackett-Burman and other nonregular factorial designs. *Statistica Sinica*, **9**, 1071–1082.
- Deng, L. Y. and Tang, B. (2002). Design selection and classification for Hadamard matrices using generalized minimum aberration criteria. *Technometrics*, **44**, 173–184.
- Ding, X., Xu, H., Hopper, C., Yang, J., and Ho, C.-M. (2013). Use of fractional factorial designs in antiviral drug studies. *Quality and Reliability Engineering International*, **29**, 299–304.
- Efron, B., Hastie, T., Johnstone, I., and Tibshirani, R. (2004). Least angle regression. *Annals of Statistics*, **32**, 407–499.
- Fang, K. T. and Mukerjee, R. (2000). A connection between uniformity and aberration in regular fractions of two-level factorials. *Biometrika*, **87**, 193–198.
- Fang, K. T. and Wang, Y. (1994). *Number-Theoretic Methods in Statistics*. London, U.K.: Chapman & Hall.
- Fang, K. T., Li, R., and Sudjianto, A. (2006). *Design and Modeling for Computer Experiments*. London, U.K.: Chapman & Hall/CRC.
- Fang, K. T., Lin, D. K. J., Winker, P., and Zhang, Y. (2000). Uniform design: Theory and application. *Technometrics*, **42**, 237–248.
- Georgiou, S. D. (2014). Supersaturated designs: A review of their construction and analysis. *Journal of Statistical Planning and Inference*, **144**, 92–109.
- Grömping, U. and Xu, H. (2014). Generalized resolution for orthogonal arrays. *Annals of Statistics*, **42**, 918–939.
- Groten, J.P., Schoen, E.D., and Feron, V.J. (1996). Use of factorial designs in combination toxicity studies. *Food and Chemical Toxicology*, **34**, 1083–1089.
- Groten, J.P., Schoen, E.D., Kuper, C.F., van Bladeren, P.J., Van Zorge, J.A., and Feron, V.J. (1997). Subacute toxicity of a mixture of nine chemicals in rats: Detecting interactive effects with a fractionated two-level factorial design. *Fundamental and Applied Toxicology*, **36**, 13–29.
- Hamada, M. and Wu, C. F. J. (1992). Analysis of designed experiments with complex aliasing. *Journal of Quality Technology*, **24**, 130–137.
- Hedayat, A. S., Sloane, N. J. A., and Stufken, J. (1999). *Orthogonal Arrays: Theory and Applications*. New York: Springer.
- Hickernell, F. J. (1998). A generalized discrepancy and quadrature error bound. *Mathematics of Computation*, **67**, 299–322.
- Hickernell, F. J. and Liu, M. Q. (2002). Uniform designs limit aliasing. *Biometrika*, **89**, 893–904.
- Li, R. and Lin, D. K. J. (2003). Analysis methods for supersaturated design: Some comparisons. *Journal of Data Science*, **1**, 249–260.
- Li, W. W. and Wu, C. F. J. (1997). Columnwise-pairwise algorithms with applications to the construction of supersaturated designs. *Technometrics*, **39**, 171–179.
- Lin, D. K. J. (1993). A new class of supersaturated designs. *Technometrics*, **35**, 28–31.
- Lin, D. K. J. (1995). Generating systematic supersaturated designs. *Technometrics*, **37**, 213–225.
- Ma, C. X. and Fang, K. T. (2001). A note on generalized aberration in factorial designs. *Metrika*, **53**, 85–93.
- MacWilliams, F. J. and Sloane, N. J. A. (1977). *The Theory of Error-Correcting Codes*. Amsterdam, the Netherlands: North-Holland.
- Mandal, A. and Mukerjee, R. (2005). Design efficiency under model uncertainty for nonregular fractions of general factorials. *Statistica Sinica*, **15**, 697–707.
- Mukerjee, R. and Wu, C. F. J. (1995). On the existence of saturated and nearly saturated asymmetrical orthogonal arrays. *Annals of Statistics*, **23**, 2102–2115.
- Nguyen, N.-K. (1996). An algorithmic approach to constructing supersaturated designs. *Technometrics*, **38**, 69–73.
- Phoa, F. K. H. (2012). A code arithmetic approach for quaternary code designs and its application to $(1/64)$ th-fractions. *Annals of Statistics*, **40**, 3161–3175.

- Phoa, F. K. H. and Xu, H. (2009). Quarter-fraction factorial designs constructed via quaternary codes. *Annals of Statistics*, **37** (5A), 2561–2581.
- Phoa, F. K. H., Mukerjee, R., and Xu, H. (2012). One-eighth- and one-sixteenth-fraction quaternary code designs with high resolution. *Journal of Statistical Planning and Inference*, **142**, 1073–1080.
- Phoa, F. K. H., Pan, Y.-H., and Xu, H. (2009). Analysis of supersaturated designs via the Dantzig selector. *Journal of Statistical Planning and Inference*, **139**, 2362–2372.
- Phoa, F. K. H., Wong, W. K., and Xu, H. (2009). The need of considering the interactions in the analysis of screening designs. *Journal of Chemometrics*, **23**, 545–553.
- Phoa, F. K. H., Xu, H., and Wong, W. K. (2009). The use of nonregular fractional factorial designs in combination toxicity studies. *Food and Chemical Toxicology*, **47**, 2183–2188.
- Plackett, R. L. and Burman, J. P. (1946). The design of optimum multifactorial experiments. *Biometrika*, **33**, 305–325.
- Rao, C. R. (1947). Factorial experiments derivable from combinatorial arrangements of arrays. *Journal of Royal Statistical Society: Series B*, **9**, 128–139.
- Satterthwaite, F. E. (1959). Random balance experimentation. *Technometrics*, **1**, 111–137.
- Schoen, E. D. (2009). All orthogonal arrays with 18 runs. *Quality and Reliability Engineering International*, **25**, 467–480.
- Schoen, E. D., Eendebak, P. T., and Nguyen, M. V. M. (2010). Complete enumeration of pure-level and mixed-level orthogonal arrays. *The Journal of Combinatorial Designs*, **18**, 123–140.
- Sun, D. X., Li, W., and Ye, K. Q. (2002). An algorithm for sequentially constructing nonisomorphic orthogonal designs and its applications. Technical report SUNYSB-AMS-02-13, Department of Applied Mathematics and Statistics, SUNY at Stony Brook, NY.
- Tang, B. (2001). Theory of J -characteristics for fractional factorial designs and projection justification of minimum G_2 -aberration. *Biometrika*, **88**, 401–407.
- Tang, B. and Deng, L. Y. (1999). Minimum G_2 -aberration for non-regular fractional factorial designs. *Annals of Statistics*, **27**, 1914–1926.
- Tang, B. and Wu, C. F. J. (1997). A method for constructing supersaturated designs and its $E(s^2)$ optimality. *Canadian Journal of Statistics*, **25**, 191–201.
- Tang, Y. and Xu, H. (2013). An effective construction method for multi-level uniform designs. *Journal of Statistical Planning and Inference*, **143**, 1583–1589.
- Tang, Y. and Xu, H. (2014). Permuting regular fractional factorial designs for screening quantitative factors. *Biometrika*, **101**, 333–350.
- Tang, Y., Xu, H., and Lin, D. K. J. (2012). Uniform fractional factorial designs. *Annals of Statistics*, **40**, 891–907.
- Vander Heyden, Y., Jimidar, M., Hund, E., Niemeijer, N., Peeters, R., Smeyers-Verbeke, J., Massart, D.L., and Hoogmartens, J. (1999). Determination of system suitability limits with a robustness test. *Journal of Chromatography A*, **845**, 145–154.
- Westfall, P. H., Young, S. S., and Lin, D. K. J. (1998). Forward selection error control in the analysis of supersaturated designs. *Statistica Sinica*, **8**, 101–117.
- Wu, C. F. J. (1993). Construction of supersaturated designs through partially aliased interactions. *Biometrika*, **80**, 661–669.
- Wu, C. F. J. and Hamada, M. (2009). *Experiments: Planning, Analysis and Optimization*, 2nd edn. New York: Wiley.
- Xu, H. (2002). An algorithm for constructing orthogonal and nearly-orthogonal arrays with mixed levels and small runs. *Technometrics*, **44**, 356–368.
- Xu, H. (2003). Minimum moment aberration for nonregular designs and supersaturated designs. *Statistica Sinica*, **13**, 691–708.
- Xu, H. (2005a). Some nonregular designs from the Nordstrom and Robinson code and their statistical properties. *Biometrika*, **92**, 385–397.
- Xu, H. (2005b). A catalogue of three-level regular fractional factorial designs. *Metrika*, **62**, 259–281.
- Xu, H. (2006). Blocked regular fractional factorial designs with minimum aberration. *Annals of Statistics*, **34**, 2534–2553.

- Xu, H. (2009). Algorithmic construction of efficient fractional factorial designs with large run sizes. *Technometrics*, **51**, 262–277.
- Xu, H. and Deng, L. Y. (2005). Moment aberration projection for nonregular fractional factorial designs. *Technometrics*, **47**, 121–131.
- Xu, H. and Lau, S. (2006). Minimum aberration blocking schemes for two- and three-level fractional factorial designs. *Journal of Statistical Planning and Inference*, **136**, 4088–4118.
- Xu, H. and Wong, A. (2007). Two-level nonregular designs from quaternary linear codes. *Statistica Sinica*, **17**, 1191–1213.
- Xu, H. and Wu, C. F. J. (2001). Generalized minimum aberration for asymmetrical fractional factorial designs. *Annals of Statistics*, **29**, 1066–1077.
- Xu, H. and Wu, C. F. J. (2005). Construction of optimal multi-level supersaturated designs. *Annals of Statistics*, **33**, 2811–2836.
- Xu, H., Cheng, S. W., and Wu, C. F. J. (2004). Optimal projective three-level designs for factor screening and interaction detection. *Technometrics*, **46**, 280–292.
- Xu, H., Phoa, F. K. H., and Wong, W. K. (2009). Recent developments in nonregular fractional factorial designs. *Statistics Surveys*, **3**, 18–46.
- Ye, K. Q. (2004). A note on regular fractional factorial designs. *Statistica Sinica*, **14**, 1069–1074.
- Ye, K. Q., Park, D., Li, W., and Dean, A. M. (2008). Construction and classification of orthogonal arrays with small numbers of runs. *Statistics and Applications*, **6** (New Series), 5–15.
- Yuan, M., Joseph, V. R., and Lin, Y. (2007). An efficient variable selection approach for analyzing designed experiments. *Technometrics*, **49**, 430–439.
- Zhang, Q.Z., Zhang, R.C., and Liu, M.Q. (2007). A method for screening active effects in supersaturated designs. *Journal of Statistical Planning and Inference*, **137**(9), 235–248.
- Zhang, R.-C., Phoa, F. K. H., Mukerjee, R., and Xu, H. (2011). A trigonometric approach to quaternary code designs with application to one-eighth and one-sixteenth fractions. *Annals of Statistics*, **39**(2), 931–955.
- Zhou, Y.-D. and Xu, H. (2014). Space-filling fractional factorial designs. *Journal of the American Statistical Association*, **109**, 1134–1144.

10

Structures Defined by Factors

R.A. Bailey

CONTENTS

10.1	Factors.....	372
10.1.1	What Is a Factor?.....	372
10.1.2	Some Definitions Associated with Factors.....	374
10.1.3	Vectors and Matrices.....	375
10.2	Poset Block Structures.....	376
10.2.1	Panel Diagrams.....	376
10.2.2	Hasse Diagrams.....	378
10.2.3	Crossing and Nesting Operators.....	380
10.2.4	Subspaces and Degrees of Freedom.....	381
10.2.5	Randomization.....	382
10.2.6	Null Analysis of Variance.....	383
10.3	Orthogonal Block Structures.....	384
10.3.1	Working Directly with Genuine Factors.....	385
10.3.2	Randomization and the Null Analysis of Variance.....	386
10.4	Orthogonal Treatment Structures.....	387
10.4.1	Treatment Factors.....	387
10.4.2	Expectation Subspaces.....	390
10.5	Orthogonal Designs.....	393
10.5.1	Definition.....	394
10.5.2	Skeleton Analysis of Variance.....	395
10.5.3	Combined Hasse Diagram.....	396
10.5.4	Pseudofactors.....	399
10.6	Design Key.....	401
10.7	Other Designations of Factors as Fixed and Random.....	404
10.8	MultiStage Experiments.....	405
10.9	MultiPhase Experiments.....	408
10.10	Orthogonality.....	411
10.11	Notes on Terminology and Conventions.....	412
	References.....	413

This chapter is about factors. Section 10.1 discusses what a factor is and introduces relevant concepts and notation. Sections 10.2 and 10.3 deal with structure on the observational units and show how to derive the null analysis of variance (ANOVA). Hasse diagrams are introduced as a way to visualize some relationships between factors. Section 10.4 deals with structure on the set of treatments and relates this to the set of candidate models for

the expectation of the response variable. Section 10.5 puts these together by introducing the design, which shows which treatment is allocated to which observational unit; the skeleton ANOVA is derived.

The remaining sections build on Sections 10.1 through 10.5 but are fairly independent of each other. Section 10.6 gives an algorithm for constructing the design and its skeleton ANOVA. Section 10.7 discusses fixed and random factors. Sections 10.8 and 10.9 deal with multistage and multiphase experiments, respectively.

Throughout, structures and designs are restricted to be orthogonal. This restriction is discussed in Section 10.10. Finally, Section 10.11 gives some comments on the conventions and wording used in this chapter.

10.1 Factors

This section deals with the vexed question of what a factor is: is it a variable, a function or a partition? In particular, the definition of factor used in many statistical algorithms conflicts with the definition that gives unambiguous interpretation of data sheets: I prefer the second. The concept of one factor being finer than another is introduced in Section 10.1.2, as are the infimum and supremum of two factors. In Section 10.1.3, each factor is associated with a vector space and with two square matrices.

10.1.1 What Is a Factor?

A factor may be thought of as a variable whose levels can be adjusted for the purposes of an experiment. On the other hand, some factors, such as day of the week or position on the bench, are inherent in the experimental material: their levels cannot be changed. Both sorts of factor may be regarded as functions from the set of observational units to the set of its levels: thus, if F is a factor, then $F(\omega)$ denotes the level of F on observational unit ω .

Every factor F gives rise to a partition. I shall write $F[\omega]$ for the part of F containing observational unit ω . The distinction between these two concepts—function and partition—is rather clearer for treatment factors, whose levels can be deliberately applied by the experimenter, than it is for inherent factors. Suppose that V is a factor denoting variety of wheat. Then $V(\omega)$ means the variety of wheat sown on plot ω , while $V[\omega]$ means the set of all plots which receive the same variety of wheat as plot ω .

Note: Within this chapter, the notations $()$, $[\]$ and $\{ \}$ are never used interchangeably.

The set of observational units for an experiment is denoted Ω throughout. Individual observational units are denoted by lower-case Greek letters. The size of Ω is denoted N .

In the first three sections, we are concerned primarily with structure inherent in the observational units *before* treatments are applied. Treatments are introduced in Section 10.4.

Structure on the set of observational units is often defined by factors, but authors do not agree on what a factor is, even when they regard factors as functions. Two examples will clarify this.

Example 10.1 (Blocks and plots)

The observational units are 32 plots in a field, which have been grouped into four blocks of eight alike plots. There is one factor B , for blocks, and it has four levels. We may regard

B as a function from Ω to $\{1, 2, 3, 4\}$, so that $B(\omega)$ is the label of the block containing plot ω . Alternatively, we may regard B as a partition of Ω into four parts, called blocks, which are subsets of Ω . The block containing plot ω is denoted $B[\omega]$.

The real disagreement comes in deciding what the factor for plots should be. If the plots are labelled 1–32, then we can use the function E from Ω to $\{1, \dots, 32\}$, or the partition E which has $E[\omega] = \{\omega\}$. Alternatively, the plots may be numbered 1–8 within each block. This specifies a function P from Ω to $\{1, \dots, 8\}$. However, the four plots with $P = 1$ have nothing in common. To identify the plots, we need $B \wedge P$, whose levels are all combinations of levels of B and P (this notation is defined more precisely in Section 10.1.2).

One approach to Example 10.1 uses the second representation for the plots, calling B and P *factors* and $B \wedge P$ a *generalized factor*. The second approach is equivalent to the first representation for the plots: it calls B and P *pre-factors*, because they are used to construct the *genuine factors*, which are B and $B \wedge P$. Both approaches need the notion that P is nested in B , which is formalized in Section 10.2.1. The former approach saves a little computer space, because fewer levels are needed: in Example 10.1, P has eight levels but $B \wedge P$ has 32 levels. Statisticians were among the earliest users of computers, and they used the former approach to build powerful general algorithms such as those developed by Nelder (1965a) and Wilkinson and Rogers (1973). It covers everything in Section 10.2 and feels natural to statisticians who are familiar with these algorithms.

However, the first approach cannot cope with Example 10.2 or the more general structures in Sections 10.3 and 10.4. Moreover, it has a serious disadvantage in general use. Many scientists keep their data in spreadsheets. In Example 10.1 there is one column for blocks and one for plots (there will be further columns for treatments and for data). If the entries in the column for plots take the values 1–8 within each block, then there is a danger that these entries will be interpreted as meaning that all plots labelled 1 in that column have something in common. This mistake is made quite often in practice, especially in long-term experiments where those who analyse the data have not been involved from the start of the experiment.

I shall use the second approach, partly because it is more widely applicable and partly because it avoids this danger of mis-interpretation. That is, a genuine factor is either a function whose levels are meaningful in isolation or it is the corresponding partition.

Example 10.2 (A superimposed experiment)

One year, an experiment was conducted on 64 individual apple trees, arranged in an 8×8 square. There were eight different pruning methods, and they were applied in a Latin square (see Chapter 3). The following year, the same trees will be used for a different experiment, but there may still be some residual effect of the different pruning methods, so we need to consider the eight-level factors R , for rows; C , for columns; and P , for pruning methods. Only 64 of the 8^3 possible combinations of their levels occur.

There is an algorithmic framework which covers Examples 10.1 and 10.2, as well as other structures that we shall meet in Section 10.4. It is more general than the one that will be described in Section 10.2. For this, we need to take the view that a factor on Ω is a partition of Ω into meaningful parts, such as the blocks in Example 10.1 but not the set of plots labelled 1 in all blocks. If F is such a factor, then $F[\omega]$ denotes the part of F containing ω ; that is, $F[\omega]$ is a subset of Ω . It may be the case that the factor is defined by a function:

in Example 10.2, $P(\omega)$ denotes the previous year's pruning method on tree ω , while $P[\omega]$ denotes the set of all trees which received the same pruning method as tree ω .

10.1.2 Some Definitions Associated with Factors

Here we establish some definitions and notation that are crucial to the rest of the chapter, both conceptually and in algorithms such as the one for calculating degrees of freedom.

If F is a factor on Ω , the number of parts of F will be denoted n_F . This is the same as the number of levels of F if they all occur on Ω . Factor F is defined to be *uniform* if all its parts have the same size (some authors call this *balanced*, but the usage is different from *balanced* in *balanced incomplete-block design*: see Chapter 3). If F is uniform, denote the common size of its parts by k_F : then $n_F k_F = N = |\Omega|$.

If F and G are two factors on Ω , we say that F is *aliased* with G , written $F \equiv G$, if F and G have the same parts; that is, they are the same apart possibly from the labelling of their levels. Aliasing occurs in fractional factorial designs: see Chapter 7. We say that F is *finer* than G , or that G is *coarser* than F , written $F < G$ or $G > F$, if every part of F is contained in a part of G but at least one part of G contains two or more parts of F . In Example 10.1, $E < B$. The notation $F \preceq G$ means that either $F < G$ or $F \equiv G$. If $F \preceq G$ and $F[\alpha] = F[\beta]$, then $G[\alpha] = G[\beta]$.

We will always use E to denote the *equality* factor: it splits Ω into N parts of size one. At the other extreme is the *universal* factor U , which has a single part of size N . Thus, $E \preceq F \preceq U$ for all factors F on Ω .

If F and G are factors on Ω , their *infimum* $F \wedge G$ is defined to be the factor whose levels are all combinations (that occur on Ω) of levels of F with levels of G . That is, $F \wedge G(\omega) = (F(\omega), G(\omega))$ for ω in Ω . The parts of $F \wedge G$ are all the non-empty intersections of an F -part with a G -part. Thus, $F \wedge G[\omega] = F[\omega] \cap G[\omega]$ for ω in Ω . In Example 10.2, $R \wedge C \equiv R \wedge P \equiv C \wedge P \equiv E$. In particular, $F \wedge F = F$, $F \wedge G \preceq F$ and $F \wedge G \preceq G$ for all F and G . If $F \preceq G$, then $F \wedge G \equiv F$.

Example 10.3 (A consumer experiment)

A consumer organization plans an experiment to compare new vacuum cleaners during two notional *months* of 4 weeks each. Four volunteer homemakers participate during the first month, while another four take part during the second. Each week, each participating volunteer tests two vacuum cleaners and rates them. Thus, Ω consists of 64 tests, and two of the relevant factors are W (week) and H (homemaker), both with eight levels. Then $W \wedge H$ is a factor with 32 parts, one for each week-homemaker combination that occurs: see Figure 10.1.

The *supremum* of two factors is a dual notion to that of infimum, but it is less well known. The supremum $F \vee G$ of F and G satisfies $F \preceq F \vee G$ and $G \preceq F \vee G$ and has as many parts as possible subject to this. Thus, the parts of $F \vee G$ are the smallest subsets of Ω which are simultaneously unions of F -parts and unions of G -parts. In particular, $F \vee F = F$, and if $F \preceq G$, then $F \vee G = G$. In Example 10.2, $R \vee C = U$.

To find the supremum $F \vee G$ by hand, write the levels of F and G as the rows and columns, respectively, of a rectangle. In row i and column j , show all units ω with $F(\omega) = i$ and $G(\omega) = j$. Rearrange the rows, if necessary, so that the non-empty cells lie in rectangles along the main diagonal and so that no further such subdivision of the rectangles is possible. Then the rectangles on the diagonal are the parts of $F \vee G$. Figure 10.1 shows this diagram for Example 10.3: $W \vee H$ is the two-level factor M whose parts are the months.

		Homemaker							
		1	2	3	4	5	6	7	8
Week	1	●●	●●	●●	●●				
	2	●●	●●	●●	●●				
	3	●●	●●	●●	●●				
	4	●●	●●	●●	●●				
	5					●●	●●	●●	●●
	6					●●	●●	●●	●●
	7					●●	●●	●●	●●
	8					●●	●●	●●	●●

FIGURE 10.1
Calculating the supremum in Example 10.3: ● denotes one test.

Most statistical software provides the facility to specify and calculate the infimum of two factors, but it is rare to find explicit facilities for the supremum, even though Tjur (1984) pointed out its importance; a recent exception is Großmann (2014). Why is the supremum important? To take an extreme case, suppose that in Example 10.3 all tests in the first month are given a score of 100, while all those in the second month score 0. If you ignore months, and fit H to the data before fitting W , then you will conclude that there is a large difference between homemakers but no difference between weeks; if you fit W and H in the other order, then you will draw the opposite conclusion. If you fit M before either W or H , then you will reach the correct conclusion that there are differences between months, but no further differences between weeks within a month or between homemakers within a month.

10.1.3 Vectors and Matrices

The data from the experiment are a vector \mathbf{y} in the N -dimensional vector space \mathbb{R}^Ω of real vectors whose coordinates are indexed by Ω . We denote by V_F the vector subspace of \mathbb{R}^Ω consisting of those vectors which take a constant value on each part of F . Thus, $\dim(V_F) = n_F$. In particular, $V_E = \mathbb{R}^\Omega$, while V_U is the one-dimensional subspace of constant vectors. Moreover, if $F \preceq G$, then $V_G \leq V_F$, which means that V_G is a subspace of V_F . More generally, a vector is constant on each part of F and on each part of G if and only if it is constant on each part of $F \vee G$, and so $V_F \cap V_G = V_{F \vee G}$.

Each factor F defines some $N \times N$ real matrices with rows and columns indexed by Ω . The *relation* matrix \mathbf{R}_F has (α, β) -entry equal to 1 if α and β are in the same part of F ; otherwise, it is 0. The *averaging* matrix \mathbf{P}_F has (α, β) -entry equal to $1/|F[\alpha]|$ if α and β are in the same part of F ; otherwise, it is 0. In particular, if F is uniform, then $\mathbf{R}_F = k_F \mathbf{P}_F$. If \mathbf{z} is any vector in \mathbb{R}^Ω , then the α -coordinate of $\mathbf{P}_F \mathbf{z}$ is the average of the β -coordinates of \mathbf{z} for all β in the same F -part as α . It follows that \mathbf{P}_F is the matrix of orthogonal projection onto V_F , and so $\text{rank}(\mathbf{P}_F) = n_F$.

There are two special cases: \mathbf{R}_E is the $N \times N$ identity matrix \mathbf{I}_N and \mathbf{R}_U is the $N \times N$ matrix \mathbf{J}_N whose entries are all 1. Moreover, $\mathbf{P}_E = \mathbf{I}_N$ and $\mathbf{P}_U = N^{-1} \mathbf{J}_N$.

If $F \preceq G$, then $\mathbf{P}_F \mathbf{P}_G = \mathbf{P}_G \mathbf{P}_F = \mathbf{P}_G$. In particular, \mathbf{P}_F is idempotent (which means that $\mathbf{P}_F^2 = \mathbf{P}_F$), and $\mathbf{P}_F \mathbf{P}_U = \mathbf{P}_U \mathbf{P}_F = \mathbf{P}_U$, for all factors F .

10.2 Poset Block Structures

This section introduces panel diagrams as a way of describing many common structures. The list in each panel forms a partially ordered set, which is often abbreviated to *poset*, and so these structures are called poset block structures.

Hasse diagrams are introduced in Section 10.2.2, partly to give a way of visualizing the factors in each poset block structure and partly to give algorithms. Section 10.2.4 explains how each poset block structure gives a decomposition of \mathbb{R}^Ω into mutually orthogonal subspaces, one for each genuine factor.

Randomization justifies the simple assumption (10.4) for the variance–covariance matrix $\text{Cov}(Y)$ of a response Y on Ω . This in turn leads to the null ANOVA described in Section 10.2.6.

10.2.1 Panel Diagrams

Very many structures which occur in practice can be defined by a list of pre-factors, as in Example 10.1, their numbers of levels, and their nesting relationships. These can succinctly be shown in a panel diagram: see Brien and Bailey (2006, 2009, 2010) and Brien et al. (2011).

Figure 10.2 gives the panel diagram for Example 10.1. There are two pre-factors: Blocks, with four levels, and Plots, with eight levels. It is helpful to use pre-factor names with different initial letters, where possible, so that each may be abbreviated to its initial letter after its first mention. The line 4 *Blocks* tells us that there is a pre-factor called Blocks with four levels. The line 8 *Plots in B* tells us that there is another pre-factor called Plots, that it has eight levels and that it is nested in *B*. There is one observational unit for each combination of levels of *B* and *P*. To say that *P* is nested in *B* means that the relation $P(\alpha) = P(\beta)$ does not signify anything meaningful unless $B(\alpha) = B(\beta)$.

More generally, if there are several pre-factors, then nesting means that the relation $P(\alpha) = P(\beta)$ does not signify anything unless $Q(\alpha) = Q(\beta)$ for all pre-factors Q which nest P . Thus, nesting is a partial order: it satisfies the following:

- (N1) Every pre-factor is deemed to be nested in itself;
- (N2) If P is nested in Q and Q is nested in R , then P is nested in R ;
- (N3) If P is nested in Q and Q is nested in P , then $P = Q$.

In general, a *poset block structure* is specified by a finite set \mathcal{M} of pre-factors, partially ordered by nesting, together with the number n_i of levels for each pre-factor P_i in \mathcal{M} . Put $|\mathcal{M}| = m$: then $n_i \geq 2$ for $i = 1, \dots, m$. Let Ω_i be a set of size n_i , for $i = 1, \dots, m$. The product set $\Omega_1 \times \Omega_2 \times \dots \times \Omega_m$ gives a convenient way of indexing the observational units. Thus, Ω is



FIGURE 10.2

Panel diagram for Example 10.1.

identified with $\Omega_1 \times \Omega_2 \times \dots \times \Omega_m$ and each element ω of Ω is regarded as an m -tuple $(\omega_1, \dots, \omega_m)$, where ω_i is the level of P_i on ω .

Let \mathcal{L} be any subset of \mathcal{M} . Then \mathcal{L} is defined to be *ancestral* if

$$\text{whenever } P_i \in \mathcal{L} \text{ and } P_i \text{ is nested in } P_j, \text{ then } P_j \in \mathcal{L}. \tag{10.1}$$

In Example 10.1, $\mathcal{M} = \{B, P\}$ and there are three ancestral subsets: \emptyset , $\{B\}$ and $\{B, P\}$.

Each ancestral subset \mathcal{L} defines a genuine factor $F(\mathcal{L})$ on Ω . Given units $\alpha = (\alpha_1, \dots, \alpha_m)$ and $\beta = (\beta_1, \dots, \beta_m)$, then α and β are in the same part of $F(\mathcal{L})$ if and only if $\alpha_i = \beta_i$ for all P_i in \mathcal{L} . Thus, $F(\emptyset)$ is just the universal factor U , while $F(\mathcal{M})$ is the equality factor E . The number $n_{F(\mathcal{L})}$ of levels of $F(\mathcal{L})$ is the product of those n_i for which $P_i \in \mathcal{L}$; that is,

$$n_{F(\mathcal{L})} = \prod_{i:P_i \in \mathcal{L}} n_i. \tag{10.2}$$

Moreover, $F(\mathcal{L})$ is a uniform factor, because each part of $F(\mathcal{L})$ has size $k_{F(\mathcal{L})}$, which is the product of those n_i for which $P_i \notin \mathcal{L}$; that is,

$$k_{F(\mathcal{L})} = \prod_{i:P_i \in \mathcal{M} \setminus \mathcal{L}} n_i. \tag{10.3}$$

These numbers are illustrated for Example 10.3 in Table 10.1.

For each ancestral subset \mathcal{L} , the relation matrix $R_{F(\mathcal{L})}$ is a Kronecker product of m matrices: if $P_i \in \mathcal{L}$, then the i th matrix is I_{n_i} ; otherwise, it is J_{n_i} . Hence,

$$P_{F(\mathcal{L})} = \left(\bigotimes_{i:P_i \in \mathcal{L}} I_{n_i} \right) \otimes \left(\bigotimes_{i:P_i \in \mathcal{M} \setminus \mathcal{L}} n_i^{-1} J_{n_i} \right).$$

In the poset block structure defined by \mathcal{M} , the set \mathcal{F} of genuine factors is given by

$$\mathcal{F} = \{F(\mathcal{L}) : \mathcal{L} \text{ is an ancestral subset of } \mathcal{M}\}.$$

The notation $F(\mathcal{L})$ is useful for general statements, but in specific examples, it may be clearer to write the factor $F(\mathcal{L})$ as the *concatenation of the initial letters* of the pre-factors in \mathcal{L} , with

TABLE 10.1

Genuine Factors Defined by the Panel in Figure 10.3

Ancestral Subset \mathcal{L}	Factor $F(\mathcal{L})$	Number of Parts $n_{F(\mathcal{L})}$	Size of Each Part $k_{F(\mathcal{L})}$
\emptyset	U	1	64
$\{M\}$	$M = W \vee H$	2	32
$\{M, V\}$	$MV = H$	8	8
$\{M, P\}$	$MP = W$	8	8
$\{M, V, P\}$	$MVP = W \wedge H$	32	2
$\{M, V, P, T\}$	$MVPT = E$	64	1



FIGURE 10.3
Panel diagram for Example 10.3.

$F(\emptyset)$ shown as U . Thus, in Example 10.1 the genuine factors are U, B and BP ; moreover, $n_U = 1, n_B = 4, n_{BP} = 32, k_U = 32, k_B = 8$ and $k_{BP} = 1$.

The structure in Example 10.3 can be expressed as a poset block structure, but to do so, we need to change notation slightly, so that a four-level pre-factor V labels the volunteer homemakers within each month, another four-level pre-factor P labels the 4 weekly periods within each month, and a two-level pre-factor T labels the two tests within each combination of levels of M, V and P . The panel diagram is in Figure 10.3. Table 10.1 lists the six ancestral subsets, the corresponding genuine factors, their numbers of levels and the sizes of their parts.

The following results are proved in Bailey (2004, Chapter 9).

Proposition 10.1 *Let \mathcal{L} and \mathcal{K} be ancestral subsets of \mathcal{M} .*

- (i) *If $\mathcal{L} \subset \mathcal{K}$, then $F(\mathcal{K}) \prec F(\mathcal{L})$.*
- (ii) *$\mathcal{L} \cup \mathcal{K}$ is an ancestral subset, and $F(\mathcal{L}) \wedge F(\mathcal{K}) = F(\mathcal{L} \cup \mathcal{K})$.*
- (iii) *$\mathcal{L} \cap \mathcal{K}$ is an ancestral subset, and $F(\mathcal{L}) \vee F(\mathcal{K}) = F(\mathcal{L} \cap \mathcal{K})$.*
- (iv) *$P_{F(\mathcal{L})}P_{F(\mathcal{K})} = P_{F(\mathcal{K})}P_{F(\mathcal{L})} = P_{F(\mathcal{L} \cap \mathcal{K})}$.*

10.2.2 Hasse Diagrams

The *is finer than or aliased with* relation \preceq introduced in Section 10.1 is also a partial order, because

- (F1)** If F is a factor, then $F \preceq F$;
- (F2)** If $F \preceq G$ and $G \preceq H$, then $F \preceq H$;
- (F3)** If $F \preceq G$ and $G \preceq F$, then $F \equiv G$.

Any partial order can be visualized using a *Hasse diagram*. In particular, any collection \mathcal{F} of factors on Ω can be shown on a Hasse diagram, as in Bailey (2004, 2008). In Section 10.2.4, this leads to an algorithm for calculating degrees of freedom and sums of squares.

In the Hasse diagram, there is one dot for each factor. Beside the dot for factor F is shown the integer n_F , which is the number of parts of F . If $F \prec G$, then the dot for G is further up the diagram than the dot for F . If $F \prec G$ but there is no H in \mathcal{F} with $F \prec H \prec G$, then a line joins the dots for F and G . Thus, $F \prec G$ if and only if there is a path from the dot for F to the dot for G traversing lines in an upwards direction only. Figures 10.4 and 10.5 show the Hasse diagrams for the poset block structures defined by the panels in Figures 10.2 and 10.3, respectively.

The following algorithm can be used to derive the Hasse diagram.




FIGURE 10.4
Hasse diagram for the poset block structure defined by Figure 10.2.

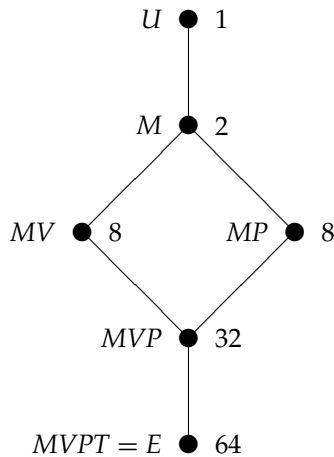


FIGURE 10.5
Hasse diagram for the poset block structure defined by Figure 10.3.

Algorithm 10.1 (From Panel Diagram to Hasse Diagram)

1. List the ancestral subsets of \mathcal{M} in increasing order of size. Thus, the list starts with \emptyset , and if $\mathcal{L} \subset \mathcal{K}$, then \mathcal{L} is listed before \mathcal{K} .
 2. For each \mathcal{L} in the list
 - (a) Name the corresponding factor by concatenating the initial letters of the pre-factors in \mathcal{L} .
 - (b) Calculate the number $n_{F(\mathcal{L})}$ of parts from Equation 10.2.
 - (c) Calculate the size $k_{F(\mathcal{L})}$ of each part from Equation 10.3.
 3. Start to draw the Hasse diagram by putting a dot for \emptyset at the top level, all dots for ancestral subsets of size one on the same level below this, those of size two on the next level, and so on.
 4. For each \mathcal{L} in the list, write the name of the factor $F(\mathcal{L})$ and the number $n_{F(\mathcal{L})}$ beside the dot for \mathcal{L} .
 5. If \mathcal{L} and \mathcal{K} are in the list, and $\mathcal{L} \subset \mathcal{K}$, and $|\mathcal{K}| = |\mathcal{L}| + 1$, then draw a line between the dots for \mathcal{L} and \mathcal{K} .
-




FIGURE 10.6
Panel diagram for a crossover trial.




FIGURE 10.7
Hasse diagram for a crossover trial.

Starting with the panel in Figure 10.3, Steps 1 and 2 of Algorithm 10.1 give Table 10.1; then Steps 3–5 give the Hasse diagram in Figure 10.5. The diamond shape in the centre of this shows how easy it is to read suprema and infima off the Hasse diagram: here, $MV \vee MP = M$ and $MV \wedge MP = MVP$.

Recall that the number of pre-factors is m . When $m = 1$, the only factors on Ω are U and E . When $m = 2$, there are two possibilities. Either one pre-factor is nested in the other, as in Example 10.1, or there is no nesting. An example of the latter occurs in crossover trials. Example 4.4 has 12 subjects for three periods, with each subject able to have a different treatment in each period. Figures 10.6 and 10.7 give the corresponding panel diagram and Hasse diagram, respectively.

10.2.3 Crossing and Nesting Operators

Many statistical algorithms use the symbolic crossing and nesting of pre-factors described by Nelder (1965a). Various different notations are in use: here we use $/$ for nesting and \times for crossing, which means lack of nesting. Thus, the panels in Figures 10.2 and 10.6 correspond to the formulas

$$(4 \text{ Blocks}) / (8 \text{ Plots})$$

and

$$(12 \text{ Subjects}) \times (3 \text{ Periods}),$$

respectively. The operators can be iterated. For three pre-factors, there are essentially five possibilities (up to change of names and numbers of levels):

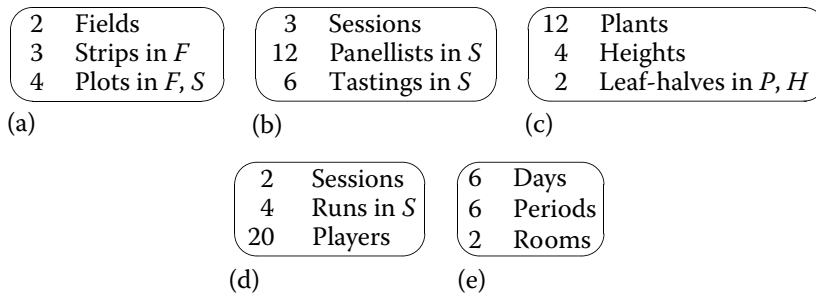


FIGURE 10.8
Examples of the five poset block structures with three pre-factors.

$$\begin{aligned}
 & (2 \text{ Fields}) / (3 \text{ Strips}) / (4 \text{ Plots}); \\
 & (3 \text{ Sessions}) / ((12 \text{ Panellists}) \times (6 \text{ Tastings})); \\
 & ((12 \text{ Plants}) \times (4 \text{ Heights})) / (2 \text{ Leaf-halves}); \\
 & ((2 \text{ Sessions}) / (4 \text{ Runs})) \times (20 \text{ Players}); \\
 & (6 \text{ Days}) \times (2 \text{ Periods}) \times (2 \text{ Rooms}).
 \end{aligned}$$

These correspond to the panels in Figure 10.8a through e, respectively. The structures in (a) and (d) are described by Bailey (2008), those in (b) and (e) by Brien and Bailey (2006) and that in (c) by Bailey and Monod (2001).

However, the formulas obtained from iterated crossing and nesting do not cover all possibilities when $m \geq 4$. Kempthorne et al. (1961) listed all poset block structures for $m \leq 4$ and included one like that in Figure 10.9, also discussed by Bailey (1991), that cannot be obtained from iterated crossing and nesting. When $m = 5$, only 48 of the 63 poset block structures can be obtained from iterated crossing and nesting. As m increases, so does the proportion of poset block structures that cannot be described by crossing and nesting alone. However, the methods in this section apply to all poset block structures.

10.2.4 Subspaces and Degrees of Freedom

Each genuine factor F in a poset block structure has an associated vector space V_F . These subspaces overlap, as shown in Section 10.1. In particular, if F and G are two factors, then $V_F \cap V_G = V_{FVG}$. However, it can be shown that part (iv) of Proposition 10.1 implies that if F and G are two factors in a poset block structure, then $V_F \cap V_{FVG}^\perp$ is orthogonal to $V_G \cap V_{FVG}^\perp$.




FIGURE 10.9
Panel diagram for a poset block structure that cannot be defined by iterated crossing and nesting.

Thus, if $v \in V_F$, then there are orthogonal vectors w and x such that $v = w + x$, $w \in V_{F \vee G}$, $x \in V_F$, and x is orthogonal to V_G . Put another way, the three subspaces $V_{F \vee G}$, $V_F \cap V_{F \vee G}^\perp$ and $V_G \cap V_{F \vee G}^\perp$ are mutually orthogonal and their vector space sum is $V_F + V_G$. Continuing in this way, this means that we can define further vector subspaces W_F , one for each genuine factor F in the poset block structure, in such a way that

- (S1) If $F \neq G$, then W_F is orthogonal to W_G ;
- (S2) For each F in the poset block structure, $V_F = \bigoplus_{G \succ F} W_G$.

Let $d_F = \dim(W_F)$, the so-called *degrees of freedom* for F . Let Q_F be the matrix of orthogonal projection onto W_F . The following algorithm shows how to calculate all of these recursively, starting at the top of the Hasse diagram and working downwards.

Algorithm 10.2 (From Hasse Diagram to Subspaces and Degrees of Freedom)

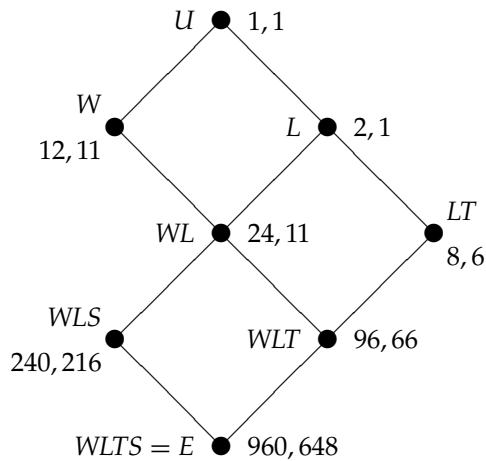
1. Start at the top of the Hasse diagram.
 2. Put $W_U = V_U$, $d_U = n_U = 1$ and $Q_U = P_U$.
 3. If W_G, d_G and Q_G have been calculated for all G in \mathcal{F} with $G \succ F$, then
 - (a) Put $W_F = V_F \cap \left(\bigoplus_{G \succ F} W_G \right)^\perp$.
 - (b) Put $d_F = n_F - \sum_{G \succ F} d_G$.
 - (c) Put $Q_F = P_F - \sum_{G \succ F} Q_G$.
-

Notice that although V_F and n_F are completely defined by the factor F , knowledge of the other factors present is needed to calculate d_F and Q_F or to interpret the notation W_F . Although it would be simpler to retain the factor names as given in the Hasse diagram, various other notations are commonly used for the W -subspaces to try to incorporate relevant information about other factors. For example, W_{BP} in Example 10.1 may be denoted $P[B]$ because P is nested in B ; and W_{SP} in the example shown in Figure 10.6 may be written as $S.P$ or $S\#P$, because there is no nesting between S and P .

It is useful to show each d_F on the Hasse diagram next to the corresponding n_F : there should be no ambiguity, because $n_F \geq d_F$ always. This is done in Figure 10.10, which shows the Hasse diagram corresponding to the panel in Figure 10.9. For example, $n_{LT} = n_L n_T = 8$ and $d_{LT} = n_{LT} - d_U - d_L = 6$. Some authors also give a brief notation for W_F next to the dot for F : see Brien and Bailey (2009, 2010).

10.2.5 Randomization

Suppose that a systematic design for the experiment has been decided. This might be as simple as having all the treatments in the same order in each block, or it might be constructed using an algorithm like the design key in Section 10.6. If the structure on Ω is a poset block structure, then the appropriate randomization can be deduced from the panel

**FIGURE 10.10**

Hasse diagram for the poset block structure defined by Figure 10.9.

diagram. For each P_i in \mathcal{M} independently, randomize the levels of P_i independently within each combination of levels of all pre-factors (apart from P_i itself) which nest P_i .

For example, in the structure shown in Figure 10.8c, randomize the (labelling of the) plants, randomize the (labelling of the) heights and independently randomize the (labels of the) two leaf-halves in each plant–height combination.

Applying this randomization to the systematic design gives the randomized layout in space or time for the experiment.

10.2.6 Null Analysis of Variance

Denote by Y the random vector for the responses, under the hypothetical assumption that the same treatment is applied to all observational units. The more realistic situation is developed in Section 10.4.2.

It was shown by Bailey et al. (1983) and Bailey (1991) that the method of randomization in Section 10.2.5 justifies the assumption that all the entries in $E(Y)$ are the same and that $\text{Cov}(Y)$ is a linear combination of the relation matrices R_F for the genuine factors F . That is, not only does this randomization justify the assumption that there are no fixed effects other than the overall mean; it also justifies the assumption that the pattern of entries in the variance–covariance matrix $\text{Cov}(Y)$ is the same as it would be if every factor in \mathcal{F} had random effects.

As shown on p. 377, the genuine factors are all uniform, and so $\text{Cov}(Y)$ is also a linear combination of the averaging matrices P_F . Hence, it is also a linear combination of the orthogonal idempotent matrices Q_F defined in Section 10.2.4. Thus, there are non-negative constants ξ_F such that

$$\text{Cov}(Y) = \sum_F \xi_F Q_F. \quad (10.4)$$

TABLE 10.2
Null Analysis of Variance for the
Structure in Figure 10.8b

Stratum	df
U (overall mean)	1
Sessions	2
Panellists[S]	33
Tastings[S]	15
$P\#T$ [S]	165
Total	216

If $v \in W_G$, then $Q_G v = v$ but $Q_F v = \mathbf{0}$ if $F \neq G$. Thus, it follows from Equation 10.4 that, for each factor F in \mathcal{F} , the subspace W_F is an eigenspace of $\text{Cov}(Y)$ with eigenvalue ξ_F . Therefore, $\text{Cov}(Q_F Y) = \xi_F Q_F$, which acts like a scalar matrix on vectors in W_F . These eigenspaces were called *strata* by Nelder (1965a), and the constants ξ_F , for F in \mathcal{F} , are called *spectral components of variance*. Under this randomization model, there is no linear dependence among the spectral components of variance. Data analysis within stratum W_F proceeds as if the variance–covariance matrix were $\xi_F I$.

The table listing the strata and their degrees of freedom, with W_G shown above W_F if $G \succ F$, is called the *null analysis-of-variance table*. Table 10.2 shows this for the structure in Figure 10.8b.

When data y are available, the preliminary sum of squares for genuine factor F is $y' Q_F y$.

Algorithm 10.3 (Calculation of Preliminary Sums of Squares)

1. For each factor F , calculate $y' P_F y$ as the sum, over all parts of F , of

$$\frac{(\text{total of } y \text{ on the part})^2}{\text{size of the part}}.$$

2. Put $y' Q_U y = y' P_U y = (\text{total of } y)^2 / N$.
3. If $y' Q_G y$ has been calculated for all G in \mathcal{F} with $G \succ F$, then put

$$y' Q_F y = y' P_F y - \sum_{G \succ F} y' Q_G y.$$

These preliminary sums of squares will be used in Section 10.5.2 to complete the ANOVA table.

10.3 Orthogonal Block Structures

This section introduces some slightly more general block structures where almost all of the results in Section 10.2 hold.

10.3.1 Working Directly with Genuine Factors

Sometimes we start not with a list of pre-factors, but with a given set Ω and some genuine factors on Ω , in the sense of Section 10.1.

Let F and G be factors on Ω . Then F is defined to be *orthogonal* to G if $P_F P_G = P_G P_F$. Thus, Proposition 10.1 shows that all factors in a poset block structure are orthogonal to each other. Part (iv) of Proposition 10.1 generalizes as follows.

Proposition 10.2 *If factors F and G are orthogonal to each other, then $P_F P_G = P_{F \vee G}$.*

There are some surprising special cases. Every factor F is orthogonal to itself, because P_F obviously commutes with itself. Likewise, F is orthogonal to every factor with which it is aliased. If $F \preceq G$, then F is orthogonal to G , because $P_F P_G = P_G P_F = P_G$ (see Section 10.1.3).

It can be shown that an equivalent definition of orthogonality is that

$$|F[\omega]| \times |G[\omega]| = |F \wedge G[\omega]| \times |F \vee G[\omega]| \quad \text{for all } \omega \text{ in } \Omega. \quad (10.5)$$

This means that, within each part of $F \vee G$ separately, the size of the part of $F \wedge G$ which contains ω is proportional to the product of the sizes of the parts of F and G which contain ω . If F , G , $F \wedge G$ and $F \vee G$ are all uniform, then this is equivalent to the fact that, within each part of $F \vee G$ separately, every part of F has non-empty intersection with every part of G .

A collection \mathcal{F} of factors on Ω is called an *orthogonal block structure* if it satisfies the following six conditions:

- (O1) $U \in \mathcal{F}$;
- (O2) $E \in \mathcal{F}$;
- (O3) If $F \in \mathcal{F}$, then F is uniform;
- (O4) If F and G are in \mathcal{F} , then F is orthogonal to G ;
- (O5) If F and G are in \mathcal{F} , then $F \vee G \in \mathcal{F}$;
- (O6) If F and G are in \mathcal{F} , then $F \wedge G \in \mathcal{F}$.

In Section 10.2.1, we noted that if \mathcal{F} is a poset block structure, then \mathcal{F} satisfies conditions (O1) and (O2): $U = F(\emptyset)$ and $E = F(\mathcal{M})$. Moreover, every factor in a poset block structure is uniform. Proposition 10.1 shows that every poset block structure satisfies conditions (O4)–(O6). Hence, every poset block structure is an orthogonal block structure. In other cases, conditions (O4)–(O6) can be tricky and time-consuming to check. However, Example 10.2 is an orthogonal block structure with $\mathcal{F} = \{U, R, C, P, E\}$ but it is not a poset block structure.

To see that Example 10.2 is an orthogonal block structure, notice that, for every pair of R , C and P , all combinations of their levels occur, so their supremum is U : thus, condition (O5) is satisfied. We saw in Section 10.1.2 that the pairwise infima of R , C and P are aliased with E , and so condition (O6) is satisfied. Since $|R[\omega]| = |C[\omega]| = |P[\omega]| = 8$ for all ω , it is now clear that Equation 10.5 is true for all pairs of R , C and P : hence, conditions (O3) and (O4) are also satisfied.

To see that Example 10.2 is not a poset block structure, suppose that \mathcal{J} , \mathcal{K} and \mathcal{L} are ancestral subsets of some poset \mathcal{M} . Then $\mathcal{J} \cap (\mathcal{K} \cup \mathcal{L}) = (\mathcal{J} \cap \mathcal{K}) \cup (\mathcal{J} \cap \mathcal{L})$. Hence, Proposition 10.1 shows that

$$\begin{aligned}
 F(\mathcal{J}) \vee (F(\mathcal{K}) \wedge F(\mathcal{L})) &= F(\mathcal{J}) \vee F(\mathcal{K} \cup \mathcal{L}) \\
 &= F(\mathcal{J} \cap (\mathcal{K} \cup \mathcal{L})) \\
 &= F((\mathcal{J} \cap \mathcal{K}) \cup (\mathcal{J} \cap \mathcal{L})) \\
 &= F(\mathcal{J} \cap \mathcal{K}) \wedge F(\mathcal{J} \cap \mathcal{L}) \\
 &= (F(\mathcal{J}) \vee F(\mathcal{K})) \wedge (F(\mathcal{J}) \vee F(\mathcal{L})).
 \end{aligned}$$

However, in Example 10.2, $R \vee (C \wedge P) = R \vee E = R$, while $(R \vee C) \wedge (R \vee P) = U \vee U = U$.

If \mathcal{F} is an orthogonal block structure, then the factors in \mathcal{F} can be shown on a Hasse diagram, as described in the first part of Section 10.2.2. Then conditions (O2), (O4) and (O5) show that subspaces W_F , for F in \mathcal{F} , can be defined in such a way that conditions (S1) and (S2) in Section 10.2.4 are satisfied and that $\mathbb{R}^\Omega = \bigoplus_{F \in \mathcal{F}} W_F$. Algorithm 10.2 can be used, just as in Section 10.2.4, to calculate W_F , d_F and Q_F . However, if \mathcal{F} is not a poset block structure, then there is no longer any way of labelling W_F so that it can be interpreted without knowing the whole of \mathcal{F} .

Figure 10.11 shows the Hasse diagram for Example 10.2. Here $W_R = V_R \cap V_U^\perp = V_R \cap W_U^\perp$, which has dimension $8 - 1 = 7$: it consists of all vectors which are constant on each row but whose entries sum to zero overall. Similarly, $W_E = V_E \cap W_U^\perp \cap W_R^\perp \cap W_P^\perp \cap W_C^\perp$: it consists of all vectors whose entries sum to zero on each row, each column and each pruning method. Its dimension is $64 - 1 - 7 - 7 - 7 = 42$.

10.3.2 Randomization and the Null Analysis of Variance

If \mathcal{F} is an orthogonal block structure which is not a poset block structure, then there is no longer a simple algorithm for randomization. One approach is to choose a random permutation g of Ω from among all those that satisfy $F[\alpha] = F[\beta]$ if and only if $F[g(\alpha)] = F[g(\beta)]$ for all F in \mathcal{F} . It is not straightforward to find out what the group of such permutations is: in Example 10.2, it depends on the Latin square used to allocate the pruning treatments in the previous year’s experiment. Further, there is no guarantee that the group is transitive on Ω , which is needed for the assumption that all entries in $E(\mathbf{Y})$ are the same. Even if the group is transitive on Ω , it still may not be the case that there is a permutation g with $g(\alpha) = \gamma$ and $g(\beta) = \delta$ if and only if $\{F \in \mathcal{F} : F[\alpha] = F[\beta]\} = \{F \in \mathcal{F} : F[\gamma] = F[\delta]\}$, which is needed for assumption (10.4). In Example 10.2, there is one 8×8 Latin square for which these conditions hold (Bailey 1982), but most Latin squares do not satisfy this.




FIGURE 10.11
Hasse diagram for Example 10.2.

One way forward is to simply assume that $E(\mathbf{Y})$ has constant entries and that $\text{Cov}(\mathbf{Y})$ is given by Equation 10.4. This corresponds to considering that each factor in \mathcal{F} gives random effects. This usually implies that $\xi_G > \xi_F$ if $G > F$, but Nelder (1954) argued that, for example, correlations within a block can sometimes be negative, so that the only constraint on the spectral components of variance ξ_F is that $\xi_F \geq 0$ for all F in \mathcal{F} . Then the null ANOVA follows exactly as in Section 10.2.6. In particular, Algorithm 10.3 gives preliminary sums of squares.

The second approach to randomization is specific to the type of combinatorial design used for assigning treatments. A design is chosen at random from among all those satisfying certain combinatorial conditions. This does not lead to any simple expressions for $E(\mathbf{Y})$ or $\text{Cov}(\mathbf{Y})$, but can sometimes lead to the result that the expected mean squares for treatments and residual, within a given stratum, are equal if the treatment effects are zero. For the general version of Example 10.2, Preece, Bailey and Patterson (1978) showed that this happens if the design is chosen at random from a set of Latin squares which, when augmented by the original Latin square, comprise a complete set of mutually orthogonal Latin squares.

From now on, Equation 10.4 is simply assumed.

10.4 Orthogonal Treatment Structures

Here we consider the set of treatments *before* they are applied to the observational units. We ignore the observational units until Section 10.5. Weaker conditions than those assumed in Sections 10.2 and 10.3 still allow us to decompose the appropriate vector space into a sum of mutually orthogonal subspaces, one for each factor. These allow us to build candidate models for $E(\mathbf{Y})$ in Section 10.4.2.

10.4.1 Treatment Factors

Let Θ be the set of treatments to be applied in an experiment. Here we consider factors on Θ and subspaces of \mathbb{R}^Θ .

Many treatment structures are poset block structures, in spite of the name. For example, a full factorial treatment structure with m treatment factors is a poset block structure with m pre-factors and no nesting. The next two examples show some other possibilities.

Example 10.4 (Full factorial plus control)

An experiment on controlling scab disease in potatoes, reported by Cochran and Cox (1957), compared the seven treatments shown in Table 10.3. Six treatments comprised all combinations of three non-zero quantities of sulphur (S) and two times of spraying (T); the seventh treatment was a *do nothing* control, for which the level of S was zero and the level of T was *not applicable* because nothing was ever sprayed. Thus, the treatments comprise a full factorial plus control.

Figure 10.12 displays the seven treatments in a rectangle whose rows are labelled by levels of T and whose columns are labelled by levels of S . This shows that there is a two-level factor C which differentiates the control treatment from the rest, that $C = S \vee T$ and that none of the three factors C , S and T is uniform. However, Figure 10.12 also shows that S and T do satisfy the *proportionality* condition (10.5), and so S is orthogonal to T .

TABLE 10.3
Seven Treatments in Example 10.4

Factor		Treatment						
		1	2	3	4	5	6	7
C	Control	1	2	2	2	2	2	2
S	Amount of sulphur	0	300	600	1200	300	600	1200
T	Timing	N/A	autumn	autumn	autumn	spring	spring	spring

		Sulphur			
		0	300	600	1200
Timing	N/A	1			
	autumn		2	3	4
	spring		5	6	7

FIGURE 10.12
Factors Sulphur and Timing in Example 10.4.

Example 10.5 (Methods with incomparable levels)

An experiment on reducing feed for chickens, reported by Bailey (2008), had the ten treatments shown in Table 10.4. Method 1 was the *no change* control treatment. Method 2 consisted of reducing the protein content, by three different amounts. Method 3 consisted of changing to a cheaper diet, at three different time points. Method 4 consisted of replacing 5% of the protein by an equal quantity of roughage, of three different types.

Different instances of different methods have nothing in common, and so the factors P , T and R are needed to label different treatments within each of the three non-control methods. None of the five factors shown is uniform, and all pairs are orthogonal. Moreover, $P \vee T = P \vee R = T \vee R = M$. Figure 10.13 shows this for P and T .

A collection \mathcal{G} of factors on Θ is called an *orthogonal treatment structure* if it satisfies the following three conditions:

- (T1) $U \in \mathcal{G}$;
- (T2) If $F \in \mathcal{G}$ and $G \in \mathcal{G}$, then F is orthogonal to G ;
- (T3) If $F \in \mathcal{G}$ and $G \in \mathcal{G}$, then $F \vee G \in \mathcal{G}$.

TABLE 10.4
Ten Treatments in Example 10.5

Factor		Treatment									
		1	2	3	4	5	6	7	8	9	10
C	Control	1	2	2	2	2	2	2	2	2	2
M	Method	1	2	2	2	3	3	3	4	4	4
P	Protein	1	2	3	4	5	5	5	6	6	6
T	Time	1	2	2	2	3	4	5	6	6	6
R	Roughage	1	2	2	2	3	3	3	4	5	6

		Time					
		1	2	3	4	5	6
Protein	1	1					
	2		2				
	3		3				
	4		4				
	5			5	6	7	
	6						8, 9, 10

FIGURE 10.13
Factors Protein and Time in Example 10.5.




FIGURE 10.14
Hasse diagram for treatment factors in Example 10.4.

These conditions are the counterparts of conditions **(O1)**, **(O4)** and **(O5)** in Section 10.3. Uniformity is no longer required: previously, it was needed to justify expressing $\text{Cov}(\mathbf{Y})$ as in Equation 10.4 as well as a linear combination of relation matrices. Nor is the equality factor mandatory: when it is included, we shall write it as \tilde{E} , to avoid ambiguity in the next section. Condition **(O6)** was needed in Sections 10.2 and 10.3 partly for the randomization argument: if a permutation of Ω preserves partitions F and G , then it also preserves both $F \vee G$ and $F \wedge G$. The counterpart of **(O6)** is not needed here, as all we seek is a unique decomposition of \mathbb{R}^Θ (or a subspace of \mathbb{R}^Θ if $\tilde{E} \notin \mathcal{G}$) into mutually orthogonal subspaces. See also the discussion in Section 10.7.

Each factor F in \mathcal{G} defines a subspace V_F of \mathbb{R}^Θ , as in Section 10.1. Since $V_F \cap V_G = V_{F \vee G}$, conditions **(T2)** and **(T3)** imply that further subspaces W_F of \mathbb{R}^Θ , for F in \mathcal{G} , can be defined to satisfy conditions **(S1)** and **(S2)** in Section 10.2.4. Then d_F and Q_F are defined as before.

The collection of factors can again be shown on a Hasse diagram. For clarity, we use open circles for treatment factors, to distinguish them from factors inherent on the observational units. Algorithm 10.2 can be used to calculate d_F , W_F and Q_F . Then Algorithm 10.3 can be used to calculate treatment sums of squares $\mathbf{y}'Q_G\mathbf{y}$ for G in \mathcal{G} .

The Hasse diagrams for Examples 10.4 and 10.5 are in Figures 10.14 and 10.15. Because we no longer insist on infima, these cannot necessarily be read off the Hasse diagram: in Example 10.5, $P \wedge T$ is an eight-level factor which is not included in the Hasse diagram.




FIGURE 10.15
Hasse diagram for treatment factors in Example 10.5.




FIGURE 10.16
Two possible Hasse diagrams for factors in a main-effects-only plan. (a) Assuming no interactions. (b) Allowing the possibility of interactions.

Figure 10.15 also shows that there are zero degrees of freedom for \tilde{E} in Example 10.5. Zero degrees of freedom are always possible if the structure is not a poset block structure. In this case, all treatment contrasts are covered by the factors C, M, P, T and R .

It is not obligatory to include \tilde{E} in an orthogonal treatment structure. For example, consider a main-effects-only plan for three four-level treatment factors A, B and C in a quarter replicate. Then $|\Theta| = 16$. If we are certain that there are no interactions among A, B and C , then we obtain the Hasse diagram in Figure 10.16a. If we think that there might be interactions but we are content not to estimate them, then we have the Hasse diagram in Figure 10.16b; the six degrees of freedom for $W_{\tilde{E}}$ will be excluded from the residual subspace in Section 10.5.2, so that the interaction effects do not inflate the estimation of spectral components of variance.

10.4.2 Expectation Subspaces

For the rest of this section, it is supposed, temporarily, that there is one observation on each treatment, so that $Y \in \mathbb{R}^\Theta$. Put $E(Y) = \tau$. The usual linear-model assumption is that τ is

an unknown vector in a known subspace of \mathbb{R}^Θ . More generally, we consider a family of known subspaces as candidate models for the expectation, or *expectation models* in brief.

For example, if the treatments consist of the six non-control treatments in Table 10.3, then one of the candidate models is the subspace V_U consisting of constant vectors: if $\tau \in V_U$, then there is no difference between treatments. Another candidate is V_S : if $\tau \in V_S$, then only the level of S affects the response and so the level of T is irrelevant. The space W_S , and the corresponding vector $Q_S\tau$, show how the expected responses differ with different levels of S : this is called the *main effect* of S . The spaces V_T and W_T have a similar interpretation. The space $V_S + V_T$, which consists of all vectors of the form $x + z$ with x in V_S and z in V_T , is another candidate model: if $\tau \in V_S + V_T$, then the effects of S and T are additive. If $\tau \notin V_S + V_T$, then we need to use the largest candidate model, which is $V_{\bar{E}}$, which can here be written as $V_{S \wedge T}$. The space $W_{S \wedge T}$ consists of all those vectors orthogonal to $V_S + V_T$. The projection $Q_{S \wedge T}\tau$ is the part of τ which is not explained by the additive effects of S and T : it is called the *S-by-T interaction*.

If V_1 and V_2 are any two subspaces of the same vector space, then $\dim(V_1 + V_2) = \dim(V_1) + \dim(V_2) - \dim(V_1 \cap V_2)$. In this example, $\dim(V_S) = 3$, $\dim(V_T) = 2$ and $\dim(V_S \cap V_T) = \dim(V_U) = 1$, and so $\dim(V_S + V_T) = 3 + 2 - 1 = 4$.

Now assume that \mathcal{G} is an orthogonal treatment structure. For each F in \mathcal{G} , put $\mathcal{A}(F) = \{G \in \mathcal{G} : F \preceq G\}$. If $F \in \mathcal{G}$, then V_F should be a candidate model. Because \mathcal{G} is an orthogonal treatment structure, $V_F = \bigoplus_{G \in \mathcal{A}(F)} W_G$. The relation \preceq is a partial order, and so (10.1) shows that a subset \mathcal{H} of \mathcal{G} is ancestral if $G \in \mathcal{H}$ whenever $F \preceq G$ and $F \in \mathcal{H}$. In particular, $\mathcal{A}(F)$ is ancestral. For each subset \mathcal{H} of \mathcal{G} , put $\mathcal{A}(\mathcal{H}) = \bigcup_{F \in \mathcal{H}} \mathcal{A}(F)$: then $\mathcal{A}(\mathcal{H})$ is also ancestral.

If two subspaces are candidate models, then it is usual to include their sum as a further candidate model (except in some experiments for screening a large number of factors). The W -subspaces are mutually orthogonal, and so

$$\sum_{F \in \mathcal{H}} V_F = \bigoplus_{G \in \mathcal{A}(\mathcal{H})} W_G$$

for every subset \mathcal{H} of \mathcal{G} . If \mathcal{H} is itself ancestral, then $\sum_{F \in \mathcal{H}} V_F = \bigoplus_{G \in \mathcal{H}} W_G$, so we obtain all models of the form $\bigoplus_{G \in \mathcal{H}} W_G$ for ancestral subsets \mathcal{H} of \mathcal{G} . Write the latter as $V_{\mathcal{H}}$. If \mathcal{H}_1 and \mathcal{H}_2 are ancestral subsets, then so is $\mathcal{H}_1 \cup \mathcal{H}_2$, and $V_{\mathcal{H}_1} + V_{\mathcal{H}_2} = V_{\mathcal{H}_1 \cup \mathcal{H}_2}$. In the 3×2 factorial experiment just discussed, $\mathcal{G} = \{U, S, T, S \wedge T\}$ and the ancestral subsets are $\emptyset, \{U\}, \{U, S\}, \{U, T\}, \{U, S, T\}$ and $\{U, S, T, S \wedge T\}$. These give the model subspaces $\{0\}, V_U, V_S, V_T, V_S + V_T$ and $V_{S \wedge T}$, respectively. The zero subspace $\{0\}$ is often ignored, unless data are measured on a ratio scale.

To enable data to distinguish between overlapping candidate models, it is also advisable for the set of model subspaces to include the intersection of any two of them. However, if \mathcal{H}_1 and \mathcal{H}_2 are ancestral subsets of \mathcal{G} , then so is $\mathcal{H}_1 \cap \mathcal{H}_2$, and $V_{\mathcal{H}_1} \cap V_{\mathcal{H}_2} = V_{\mathcal{H}_1 \cap \mathcal{H}_2}$. Therefore, the set $\{V_{\mathcal{H}} : \mathcal{H} \text{ is an ancestral subset of } \mathcal{G}\}$ of candidate models is closed under both sum and intersection, and this remains true if the zero model is omitted.

The notation $V_1 \leq V_2$ indicates that V_1 is a subspace of V_2 . Now, \leq is another partial order, and so the model subspaces can also be shown on a Hasse diagram, with V_1 joined by upwards lines to V_2 if $V_1 < V_2$. Here we use small black squares to show the subspaces, to avoid confusion with the previous types of Hasse diagram. We usually omit the zero model. It is helpful to show the dimension of V_1 next to the square for V_1 . Figures 10.17 through 10.20 show such Hasse diagrams for the 3×2 factorial experiment, Examples 10.4 and 10.5, and the fractional factorial design with the factors in Figure 10.16a.




FIGURE 10.17
Family of expectation models for a 3 × 2 factorial.




FIGURE 10.18
Family of expectation models for Example 10.4.

These Hasse diagrams are the opposite way up to the previous ones and usually include models such as $V_S + V_T$ which do not correspond to a single factor.

Under orthogonal treatment structure, each edge in the Hasse diagram of models corresponds to a factor in \mathcal{G} . If $V_1 < V_2$ and there is an edge between V_1 and V_2 , then there is a unique factor F in \mathcal{G} such that $V_2 = V_1 \oplus W_F$. Moreover, $d_F = \dim(V_2) - \dim(V_1)$. Given data \mathbf{y} , $Q_F \mathbf{y}$ is the difference between the vectors of fitted values in the models V_2 and V_1 .

In every case, $Q_F \boldsymbol{\tau}$ is the difference between $P_F \boldsymbol{\tau}$ and the best that $\boldsymbol{\tau}$ can be explained by models smaller than V_F . The inclusion of suprema and the presence of orthogonality guarantee that fitting factors in different orders does not produce different results, so long as G is fitted before F if $G > F$. Thus, in the 3 × 2 factorial, $Q_S \boldsymbol{\tau}$ is equal to both $P_S \boldsymbol{\tau} - P_U \boldsymbol{\tau}$ and $P \boldsymbol{\tau} - P_T \boldsymbol{\tau}$, where $P \boldsymbol{\tau}$ is the projection of $\boldsymbol{\tau}$ onto $V_S + V_T$.

In Example 10.4, W_S corresponds to any differences between non-zero quantities of sulphur. Similarly, in Example 10.5, W_P corresponds to differences among the three feeding regimes whose basic method is to reduce the amount of protein. Some authors write this




FIGURE 10.19
Family of expectation models for Example 10.5.




FIGURE 10.20
Family of expectation models for the treatment structure in Figure 10.16a.

as $P[M]$ (for P within M); others as $P \mid M$ (for P given M), but there is no generally agreed terminology or notation. In Example 10.4, $W_{S \wedge T}$ might still be referred to as the interaction between S and T . In Example 10.5, there is no need for such terminology, because $V_P + V_T = V_{P \wedge T}$, and similarly for P and R and for T and R .

10.5 Orthogonal Designs

Sections 10.2 and 10.3 gave suitable conditions on the set \mathcal{F} of factors on the set Ω of observational units which make assumption (10.4) about $\text{Cov}(\mathbf{Y})$ reasonable. This leads to the

decomposition of \mathbb{R}^Ω shown in the null ANOVA. Section 10.4 gave conditions on the set \mathcal{G} of factors on the set Θ of treatments which ensure a unique orthogonal decomposition of \mathbb{R}^Θ . To put this all together, we now need the actual *design*, which tells us which treatment is applied to which observational unit. Under two further conditions, this leads to the result in Proposition 10.3 and hence to the expansion of the null ANOVA to the skeleton ANOVA in Section 10.5.2. Hasse diagrams continue to be useful.

10.5.1 Definition

A design is a function $\phi: \Omega \rightarrow \Theta$ allocating treatment $\phi(\omega)$ to observational unit ω . We shall say that the design is *orthogonal* if there is a collection \mathcal{F} of factors on Ω which form an orthogonal block structure, there is a collection \mathcal{G} of factors on Θ which form an orthogonal treatment structure, and, in addition, ϕ satisfies the following further conditions:

- (D1) If G_1 and G_2 are in \mathcal{G} , then G_1 and G_2 remain orthogonal to each other when considered as factors on Ω ;
- (D2) If $G \in \mathcal{G}$ and $F \in \mathcal{F}$ then G is orthogonal to F .

Strictly speaking, when considered as a function, factor G in \mathcal{G} gives the factor $G(\phi)$ on Ω , but we usually retain the name G . Condition (D1) is always satisfied if ϕ is equireplicate, because, if the replication is r , the numbers in Equation 10.5 are all multiplied by r . However, many equireplicate designs are not orthogonal, because they do not satisfy condition (D2).

We need to consider suprema of the form $F \vee G$, where $F \in \mathcal{F}$ and $G \in \mathcal{G}$. Since $V_{F \vee G} \leq V_G$, which is an expectation subspace, we need to consider $F \vee G$ to be a treatment factor or pseudofactor (see Section 10.5.4). Put $\mathcal{G}^* = \{F \vee G : F \in \mathcal{F}, G \in \mathcal{G}\}$. Then $\mathcal{G} \subseteq \mathcal{G}^*$, because $E \in \mathcal{F}$. Condition (D2) shows that if $F \in \mathcal{F}$ and $G \in \mathcal{G}$, then $P_{F \vee G} = P_F P_G$ and so $F \vee G$ is orthogonal to all factors in $\mathcal{F} \cup \mathcal{G}^*$. Moreover, \mathcal{G}^* is closed under taking suprema. Therefore, we may replace \mathcal{G} by \mathcal{G}^* without violating any of conditions (T1), (T2), (T3) in Section 10.4 or conditions (D1) and (D2) above. In the extreme case that \mathcal{G} has not been given but ϕ is orthogonal to all factors in \mathcal{F} , it suffices to put $\mathcal{G}^* = \{\phi \vee F : F \in \mathcal{F}\}$.

As we are about to show, these conditions ensure that each treatment W -subspace is contained in a single stratum. Moreover, the combined Hasse diagram gives us an algorithm for finding the appropriate stratum. If $\mathcal{G} = \mathcal{G}^*$, then this process is straightforward. If \mathcal{G} has to be augmented by any supremum $F \vee G$ which is not already in \mathcal{G} , then there are subtle complications: see Section 10.5.4.

Given an orthogonal design, the subspaces W_F , for F in \mathcal{F} , are calculated with no reference to \mathcal{G}^* , and the subspaces W_G , for G in \mathcal{G}^* , are calculated without reference to \mathcal{F} . There then comes the problem of matching them up. The following is proved in Bailey (2008, Chapter 10).

Proposition 10.3 *Suppose that $\phi: \Omega \rightarrow \Theta$ is an orthogonal design, with orthogonal block structure \mathcal{F} on Ω and orthogonal treatment structure \mathcal{G} on Θ . Then there is a function $h: \mathcal{G}^* \rightarrow \mathcal{F}$ such that, for all G in \mathcal{G}^* ,*

- (i) $h(G) \preceq G$;
- (ii) If $F \in \mathcal{F}$ and $F \preceq G$, then $F \preceq h(G)$;
- (iii) $W_G \leq W_{h(G)}$.

For example, in a randomized complete-block design with treatment factor T , we have $\mathcal{G} = \{U, T\}$ and $\mathcal{F} = \{U, B, BP\}$. Since all treatment contrasts are orthogonal to blocks, $W_T \perp W_{BP}$. Thus, $h(U) = U$ and $h(T) = BP$.

If $F \in \mathcal{F}$, $G \in \mathcal{G}$ and $h(G) = F$, then G is said to be *confounded* with F .

For G in \mathcal{G}^* , write $\mathbf{Q}_G \boldsymbol{\tau}$ as $\boldsymbol{\tau}_G$, so that $E(\mathbf{Q}_G \mathbf{Y}) = \boldsymbol{\tau}_G$. The sum of squares for G is $\mathbf{y}' \mathbf{Q}_G \mathbf{y}$, and the corresponding expected mean square is $E(\|\mathbf{Q}_G \mathbf{Y}\|^2)$. Recall from Section 10.2.4 that $d_G = \text{rank}(\mathbf{Q}_G) = \dim(W_G)$ and from Section 10.2.6 that $\xi_{h(G)}$ is the spectral component of variance for stratum $W_{h(G)}$. It follows from Proposition 10.3 that

$$E(\|\mathbf{Q}_G \mathbf{Y}\|^2) = \|\boldsymbol{\tau}_G\|^2 + d_G \xi_{h(G)}. \quad (10.6)$$

10.5.2 Skeleton Analysis of Variance

The function h given in Proposition 10.3 shows how to expand the null ANOVA to the skeleton ANOVA, which consists of as much of the ANOVA table as is possible to show in advance of obtaining the data. Given F in \mathcal{F} , the row for F is expanded to have a row for every G in \mathcal{G}^* for which $h(G) = F$, and another for residual, which corresponds to the space

$$W_F \cap \left(\bigoplus_{G \in \mathcal{G}} W_G \right)^\perp,$$

which has dimension

$$d_F - \sum_{G: h(G)=F} d_G.$$

If this last dimension is zero, then no residual is shown and there is no estimator for ξ_F ; this always happens when $F = U$. If F has the property that there is no G in \mathcal{G}^* with $h(G) = F$, then the row for F does not need to be subdivided: see the row for field Locations in Example 10.6.

A column for expected mean squares may be added to show the values which follow from the right-hand side of Equation 10.6.

If preliminary sums of squares have been calculated from data as in Section 10.2.6, then the skeleton ANOVA table can be expanded to include sums of squares, mean squares and ratios of mean squares: this gives the complete ANOVA table. For G in \mathcal{G}^* , the relevant sum of squares is $\mathbf{y}' \mathbf{Q}_G \mathbf{y}$. For F in \mathcal{F} , the relevant sum of squares for residual is

$$\mathbf{y}' \mathbf{Q}_F \mathbf{y} - \sum_{G: h(G)=F} \mathbf{y}' \mathbf{Q}_G \mathbf{y}.$$

Skeleton ANOVA tables give a useful way of comparing proposed designs with the same \mathcal{F} and \mathcal{G} . They show how many residual degrees of freedom are available for estimating each spectral component of variance ξ_F . If $h(G) = F$, then the variance of a normalized contrast in W_G is just ξ_F . A preliminary assessment of the likely relative magnitudes of the spectral components of variance can be used to ensure that variances of important contrasts are as small as possible.

10.5.3 Combined Hasse Diagram

Of course, conditions (D1)–(D2) need to be checked before we carry out the procedures in Section 10.5.2 or recommend the design. Then the function h must be identified. One way to do this is to show the factors in $\mathcal{F} \cup \mathcal{G}^*$ on a combined Hasse diagram, using one symbol for factors in \mathcal{F} and another symbol for factors in \mathcal{G}^* . Here we use a filled circle for each factor in \mathcal{F} and an open circle for each factor in \mathcal{G}^* . Some factors, such as U , belong to both \mathcal{F} and \mathcal{G}^* , and so a combined symbol must be used for them. For simplicity, the numbers like n_F and d_F are omitted from the combined Hasse diagram.

Now parts (i) and (ii) of Proposition 10.3 show how to read $h(G)$ off the combined Hasse diagram, for G in \mathcal{G}^* : find the filled circle which is as far up the diagram as possible subject to being below (or equal to) the dot for G . If the dot for G has the combined symbol, then G is aliased with the factor $h(G)$ in \mathcal{F} , which is usually a sign of false replication.

Example 10.6 (A split-plot experiment)

Twelve treatments were applied to the 24 plots with the structure in Figure 10.8a. For clarity, the fields have been renamed as locations. The treatments were all combinations of Cultivar (with three levels), Nitrogen fertilizer (with two levels) and Potash (K, with two levels). Thus, both structures are poset block structures, with the panel diagrams shown in Figure 10.21.

Each level of C was applied to one Strip per field Location; each combination of levels of N and K was applied to one Plot per Strip. The design is orthogonal, and $\mathcal{G} = \mathcal{G}^*$. If $G \in \mathcal{G}$, then all levels of G occur on each field location, so $G \vee L = U$. Apart from U , no other treatment factor is coarser than L . If G does not involve C , then all levels of G occur on each strip, so $G \vee LS = U$; otherwise, $G \vee LS = C$. The only treatment factors which are coarser than LS are C and U . These considerations give the combined Hasse diagram in Figure 10.22. This shows that $h(U) = U$, $h(C) = LS$ and that $h(G) = LSP$ for all other factors G in \mathcal{G} . Hence, we obtain the skeleton ANOVA in Table 10.5.

Example 10.7 (A strip-plot experiment)

Sometimes it is convenient to apply more than one treatment factor to units larger than the observational units. An agricultural experiment described by Clarke and Kempson (1997) had four rectangular Blocks, each with four Rows and three Columns. The treatments were all combinations of four Times of cultivation with three Pesticides. In each block, one row was cultivated at each time, and one pesticide was applied to each column.

Again, both structures are poset block structures; the panel diagrams are in Figure 10.23. The design is orthogonal, and $\mathcal{G} = \mathcal{G}^*$. Figure 10.24 gives the combined Hasse diagram, which shows that $h(T) = BR$, $h(P) = BC$ and $h(TP) = BRC$. This leads to the skeleton ANOVA in Table 10.6.




FIGURE 10.21 Panel diagrams for Example 10.6.

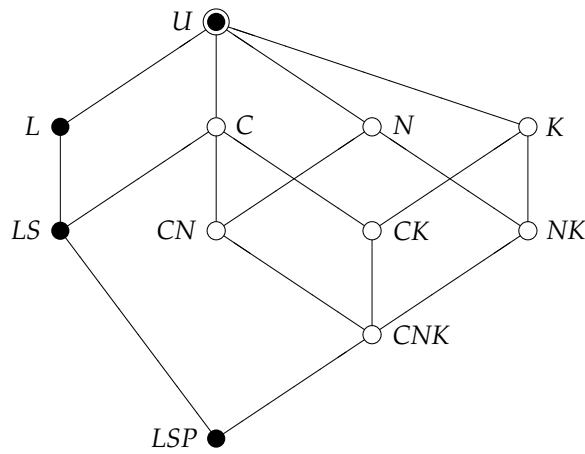


FIGURE 10.22
Combined Hasse diagram for Example 10.6.

TABLE 10.5

Skeleton Analysis of Variance for Example 10.6

Stratum	df	Source	df	EMS
U	1	mean	1	$\ \tau_U\ ^2 + \xi_U$
Field Locations	1		1	ξ_L
Strips[L]	4	Cultivar	2	$\ \tau_C\ ^2/2 + \xi_{LS}$
		residual	2	ξ_{LS}
Plots[$S \wedge L$]	18	Nitrogen	1	$\ \tau_N\ ^2 + \xi_{LSP}$
		K (potash)	1	$\ \tau_K\ ^2 + \xi_{LSP}$
		$C\#N$	2	$\ \tau_{CN}\ ^2/2 + \xi_{LSP}$
		$C\#K$	2	$\ \tau_{CK}\ ^2/2 + \xi_{LSP}$
		$N\#K$	1	$\ \tau_{NK}\ ^2 + \xi_{LSP}$
		$C\#N\#K$	2	$\ \tau_{CNK}\ ^2/2 + \xi_{LSP}$
		residual	9	ξ_{LSP}
Total			24	




FIGURE 10.23
Panel diagrams for Example 10.7.




FIGURE 10.24
Combined Hasse diagram for Example 10.7.

TABLE 10.6
Skeleton Analysis of Variance for Example 10.7

Stratum	df	Source	df	EMS
U	1	mean	1	$\ \tau_U\ ^2 + \xi_U$
Blocks	3		3	ξ_B
Rows[B]	12	Times	3	$\ \tau_T\ ^2/3 + \xi_{BR}$
		residual	9	ξ_{BR}
Columns[B]	8	Pesticides	2	$\ \tau_P\ ^2/2 + \xi_{BC}$
		residual	6	ξ_{BC}
$R\#C[B]$	24	$T\#P$	6	$\ \tau_{TP}\ ^2/6 + \xi_{BRC}$
		residual	18	ξ_{BRC}
Total			48	

Example 10.8 (Repeated splitting of the experimental units)

Sometimes apparent replication is achieved by repeated splitting, or subsampling, of the units to which the treatments are applied. An example in environmental toxicology is described by Gardner and Grue (1996). There were three treatments, one of which was a *do nothing* control. Thus, $\mathcal{G} = \{U, T, C\}$ where T has three levels and C has two levels, distinguishing the control treatment from the other two.

The experiment was performed on two Lakes. Each lake was divided into four Quadrants. The control treatment was applied to two quadrants per lake, the other two treatments each to one quadrant per lake. Four experimental Sites were established in each quadrant, and three Jars of daphnia were put into the water at each site. After a prescribed time, the amount of daphnia was measured in each jar. Figure 10.25 shows the panel diagram for the observational units.

The combined Hasse diagram in Figure 10.26 shows clearly that $h(C) = h(T) = LQ$. Thus, the repeated subdivision did not increase the numbers of degrees of freedom for testing these effects, even though it decreased the size of the standard errors: see Table 10.7.




FIGURE 10.25
Panel diagram for the observational units in Example 10.8.




FIGURE 10.26
Combined Hasse diagram for Example 10.8.

TABLE 10.7
Skeleton Analysis of Variance for Example 10.8

Stratum	df	Source	df	EMS
U	1	mean	1	$\ \tau_U\ ^2 + \xi_U$
Lakes	1		1	ξ_L
Quadrants[L]	6	Control	1	$\ \tau_C\ ^2 + \xi_{LQ}$
		Treatments[C]	1	$\ \tau_T\ ^2 + \xi_{LQ}$
		residual	4	ξ_{LQ}
Sites[$L \wedge Q$]	24		24	ξ_{LQS}
Jars[$L \wedge Q \wedge S$]	64		64	ξ_{LQSJ}
Total			96	

10.5.4 Pseudofactors

Although the preceding three examples may look complicated, in none of them did the original treatment structure \mathcal{G} need to be augmented. This complication arises most often in factorial experiments where treatment factors have numbers of levels which are not co-prime. This is now illustrated on an artificially small example.

Example 10.9 (Confounded factorial)

There are two treatment factors, A and B , both with two levels, which are coded by the integers modulo 2. As Table 10.8 shows, these define the four-level genuine factor $A \wedge B$, and also the two-level pseudofactor $A + B$, whose levels are obtained by adding levels of A and B modulo 2. It is a pseudofactor because its levels are not inherently meaningful.

Suppose that four People each conduct two Tests. Two of the people both test the two treatments where $A + B$ has level 0, and the other two people both test the other two treatments. The structure on the tests is a poset block structure, with $\mathcal{F} = \{U, P, PT\}$. The initial treatment structure \mathcal{G} is the factorial one $\{U, A, B, A \wedge B\}$. Figure 10.27 gives the panel diagrams. Conditions (D1) and (D2) are satisfied. However, $P \vee (A \wedge B) = A + B$, and so the pseudofactor $A + B$ needs to be included in \mathcal{G}^* . Although \mathcal{G} is a poset block structure, \mathcal{G}^* is not.

Figure 10.28 gives the Hasse diagram for the augmented treatment structure \mathcal{G}^* . Now W_{A+B} is orthogonal to W_U, W_A and W_B , so it corresponds to the A -by- B interaction. It has $d_{A+B} = 1$. Now that the pseudofactor $A + B$ is included, the space $W_{A \wedge B}$ is zero and has $d_{A \wedge B} = 0$. As always, interpretation of the W -subspaces depends on knowledge of all factors in \mathcal{G}^* .

The combined Hasse diagram is in Figure 10.29. This shows that $h(A) = h(B) = PT$, while $h(A + B) = P$. Thus, $A + B$ is confounded with P . Since $W_{A \wedge B} = \{0\}$, $h(A \wedge B)$

TABLE 10.8
Factors and a Pseudofactor in Example 10.9

		Treatment			
		1	2	3	4
Factor	A	0	0	1	1
Factor	B	0	1	0	1
Factor	$A \wedge B$	(0,0)	(0,1)	(1,0)	(1,1)
Pseudofactor	$A + B$	0	1	1	0




FIGURE 10.27
Panel diagrams for Example 10.9.




FIGURE 10.28
Hasse diagram for augmented treatment structure in Example 10.9.




FIGURE 10.29
Combined Hasse diagram for Example 10.9.

is not relevant. (However, in the analogous case with all powers of 2 replaced by powers of 3, $W_{A \wedge B}$ would consist of two degrees of freedom for the A -by- B interaction, and $h(A \wedge B) = PT$.)

The pseudofactor $A + B$ is needed to construct the design. It is needed for the analysis, because the standard error of the difference between levels of $A + B$ is not the same as the standard error of the difference between levels of A . However, this does not imply that V_{A+B} should be included in the list of expectation models, which should remain similar to that in Figure 10.17.

In general, whenever pseudofactors have to be included in \mathcal{G}^* , they have a role to play in the construction of the design and in the ANOVA. However, the family of candidate models for expectation should be based on the original treatment structure \mathcal{G} , excluding pseudofactors.

10.6 Design Key

Conditions **(T2)**, **(T3)** and **(D2)** can be tedious to check, and the combined Hasse diagram becomes more difficult to work with as the number of factors increases. When the structures on the observational units and the treatments are both poset block structures, an alternative way of constructing and understanding orthogonal designs is to use the design key, which was introduced by Patterson (1965): see also Patterson and Bailey (1978); Kobilinsky et al. (2014); and Cheng (2014).

In the simplest case, there is a prime p such that the number of levels of every pre-factor on Ω , and every pre-factor on Θ , is a power of p . If pre-factor F has p^r levels and $r > 1$, replace F by pseudofactors F_1, \dots, F_r , each with p levels, so that each level of F is a combination of levels of F_1, \dots, F_r . If $r = 1$, regard F itself as a pseudofactor. Then treat all levels as integers modulo p , so that they may be added modulo p , as in Example 10.9, to give further pseudofactors.

A design key Φ equates each treatment pseudofactor G with $\Phi(G)$, which is a linear combination of the pseudofactors on Ω . In Example 10.10, $\Phi(T) = R_2 + C_2$, which means that the level of T on observational unit ω is given by $T(\omega) = R_2(\omega) + C_2(\omega)$ for all ω in Ω .




FIGURE 10.30
Panel diagrams for Example 10.10.

Because all addition is performed modulo p , if F and G are two treatment pseudofactors, then $\Phi(F + G) = \Phi(F) + \Phi(G)$.

Example 10.10 (An experiment in plant houses)

An experiment on plants in pots is to be conducted in two plant Houses. Each has the pots arranged in four Rows by four Columns, with two Pots in each row–column intersection. The treatments are all combinations of levels of four two-level factors: Irrigation, Soil-type, Time of sowing and Variety of plant. Panel diagrams for treatments and for observational units are in Figure 10.30.

Water is supplied from overhead pipes which run along the rows, so each level of irrigation must be applied to whole rows. To avoid disturbance to plants already there, the two pots in each row–column intersection must be sown at the same time.

The pseudofactors for the observational units are H (for Houses), R_1 and R_2 (for Rows in H), C_1 and C_2 (for Columns in H) and P (for Pots in H, R, C), all with two levels. Those for treatments are I, S, T and V , all with two levels.

The houses are labelled 0 and 1. Within each house, the rows are labelled by the combinations of levels of R_1 and R_2 and the columns are labelled by the combinations of levels of C_1 and C_2 . The two pots within each row–column combination are distinguished by the levels of P .

The constraints on the design key are that $\Phi(I)$ must not involve any of C_1, C_2 or P and that $\Phi(T)$ must not involve P . One possibility is that

$$\Phi(I) = R_1, \quad \Phi(S) = P, \quad \Phi(T) = R_2 + C_2, \quad \Phi(V) = R_1 + R_2 + C_1 + P. \quad (10.7)$$

Now the design key Φ in (10.7) gives the algorithm for constructing the design. For example, on each pot ω , the variety is given by $V(\omega) = R_1(\omega) + R_2(\omega) + C_1(\omega) + P(\omega)$, all arithmetic being done modulo 2.

The construction using a design key gives an orthogonal design. There is a simple algorithm for identifying the confounding. Each pseudofactor L defines a W -subspace with $p - 1$ degrees of freedom. To find out which genuine effect $e(L)$ this belongs to, write down all the letters in the pseudofactor and then find the smallest ancestral subset which contains them all. In Example 10.10, the pseudofactor $R_1 + R_2 + C_2$ gives the subset $\{R, C\}$, and the smallest ancestral subset containing this is $\{H, R, C\}$, so $e(R_1 + R_2 + C_2) = HRC$. Similarly, $e(I + T)$ is part of the I -by- T interaction (in fact, the whole of the interaction, since $p = 2$). Using the design key Φ guarantees that $h(G) = e(\Phi(G))$ for each treatment pseudofactor G : in other words, G is confounded with $e(\Phi(G))$. This confounding can be calculated systematically for each treatment pseudofactor. In Example 10.10, $\Phi(I + T) = R_1 + R_2 + C_2$, and so $h(I + T) = HRC$. Working in this way gives the skeleton ANOVA in Table 10.9 (expected mean squares are not shown).

TABLE 10.9
Skeleton Analysis of Variance for Example 10.10

Stratum	df	Pseudofactor			df
		Units	Treatments	Source	
<i>U</i>	1	0	0	mean	1
Houses	1				1
Rows[<i>H</i>]	6	R_1	I	I	1
				residual	5
Columns[<i>H</i>]	6	$C_1 + C_2$	$I + S + T + V$	$I\#S\#T\#V$	1
				residual	5
<i>R\#C[H]</i>	18	$R_2 + C_2$	T	T	1
		$R_1 + R_2 + C_2$	$I + T$	$I\#T$	1
		$R_1 + R_2 + C_1$	$S + V$	$S\#V$	1
		$R_2 + C_1$	$I + S + V$	$I\#S\#V$	1
		$R_1 + C_1 + C_2$	$S + T + V$	$S\#T\#V$	1
				residual	13
Pots[$H \wedge R \wedge C$]	32	P	S	S	1
		$R_1 + R_2 + C_1 + P$	V	V	1
		$R_1 + P$	$I + S$	$I\#S$	1
		$R_2 + C_1 + P$	$I + V$	$I\#V$	1
		$R_2 + C_2 + P$	$S + T$	$S\#T$	1
		$R_1 + C_1 + C_2 + P$	$T + V$	$T\#V$	1
		$R_1 + R_2 + C_2 + P$	$I + S + T$	$I\#S\#T$	1
		$C_1 + C_2 + P$	$I + T + V$	$I\#T\#V$	1
		residual	24		
Total					64

If the numbers of levels of the factors are not all powers of the same prime, then use all the primes which divide the number of levels of any factor. Use pseudofactors to separate the primes: for example, a factor with six levels gives two pseudofactors, one with two levels and one with three levels. Then use a separate design key for each prime. Construct a separate design for each prime, and then take all combinations of the results.

If p_1 and p_2 are two different primes, then effects are combined as follows. For $i = 1, 2$, let T_i be a linear combination of treatment pseudofactors for prime p_i . Suppose that $e(T_i) = F(\mathcal{K}_i)$, where \mathcal{K}_i is an ancestral subset of the set of all treatment pre-factors, so that the effect of T_i accounts for $p_i - 1$ degrees of freedom in $W_{F(\mathcal{K}_i)}$, and that key Φ_i gives $h(T_i) = F(\mathcal{L}_i)$, where \mathcal{L}_i is an ancestral subset of the set of all pre-factors on the observational units. Then $T_1 \wedge T_2$ has $(p_1 - 1)(p_2 - 1)$ degrees of freedom for the treatment subspace $W_{F(\mathcal{K}_1 \cup \mathcal{K}_2)}$, and these are confounded with stratum $W_{F(\mathcal{L}_1 \cup \mathcal{L}_2)}$. This extends to three or more primes in the obvious way.

Example 10.11 (An experiment on turf grass)

Blouin et al. (2009) describe an experiment to investigate the effects of factorial treatments on weed infestation in turf. The observational units were defined by the panel diagram on the right of Figure 10.31. The treatments were all combinations of Grass type (with two levels), Mowing height (with three levels) and Fertilizer quantity (with three levels).




FIGURE 10.31
Panel diagrams for Example 10.11.

TABLE 10.10
Skeleton Analysis of Variance for Example 10.11

Stratum	df	Source	df
<i>U</i>	1	mean	1
Rows	1		1
Columns	3		3
<i>R</i> # <i>C</i>	3	Grass type	1
		residual	2
Strips[<i>R</i>]	4	Mowing height	2
		residual	2
Lines[<i>C</i>]	8	Fertilizer quantity	2
		residual	6
<i>C</i> # <i>S</i> [<i>R</i>]	12	<i>G</i> # <i>M</i>	2
		residual	10
<i>R</i> # <i>L</i> [<i>C</i>]	8	<i>G</i> # <i>F</i>	2
		residual	6
<i>S</i> # <i>L</i> [<i>R</i> \wedge <i>C</i>]	32	<i>M</i> # <i>F</i>	4
		<i>G</i> # <i>M</i> # <i>F</i>	4
		residual	24
Total			72

For the prime 2, the pseudofactors are *R* (for Rows), *C*₁ and *C*₂ (for Columns) and *G* (for Grass types). For the prime 3, they are *S* (for Strips in *R*), *L* (for Lines in *C*), *M* (for Mowing height) and *F* (for Fertilizer quantity). The design described can be obtained from the design keys Φ_2 and Φ_3 for the primes 2 and 3, respectively, where

$$\Phi_2(G) = R + C_1, \quad \Phi_3(M) = S, \quad \Phi_3(F) = L.$$

The three-level pseudofactors for the *M*-by-*F* interaction are *M* + *F* and *M* + 2*F*, both of which are confounded with *RCSL*. It is clear that $h(G) = RC$ and $h(M) = RS$. Since different primes are involved, it follows that $h(GM) = RCS$. Similar arguments give the skeleton ANOVA in Table 10.10.

10.7 Other Designations of Factors as Fixed and Random

So far, we have assumed that structure on the observational units gives rise to a variance-covariance structure justified by randomization. For poset block structures, this is

equivalent to the assumption that each genuine factor in \mathcal{F} gives random effects, possibly with negative components of variance. We have also assumed that the treatment structure defines the expectation part of the linear model—in other words, that its factors have fixed effects.

In practice, this is not always the correct dichotomy. For example, if sex is an inherent factor on the observational units, then it may be deemed to have fixed effects. On the other hand, in plant or animal breeding experiments, the treatments are different breeding lines, and these are often regarded as random.

To stay within the framework of this chapter, let \mathcal{F} denote the set of random factors and \mathcal{G} the set of fixed factors. We assume that $\text{Cov}(\mathbf{Y})$ is given by Equation 10.4. In order for the matrices Q_F to be defined, it is necessary that \mathcal{F} satisfy conditions **(O2)**, **(O3)**, **(O4)** and **(O5)**. In practical situations, it is usually wise to consider explicitly whether **(O6)** should be satisfied: for example, should the correlation between responses on the two tests in a homemaker-week combination in Example 10.3 be allowed to differ from the sum of the within-homemaker and within-week correlations? However, **(O6)** is not needed to derive the subspaces W_F for F in \mathcal{F} .

For the fixed factors, in order to obtain a family of expectation models with no ambiguity caused by the order of fitting effects, \mathcal{G} must satisfy **(T2)** and **(T3)**. Although **(T1)** is not always needed, it is usual to include U in \mathcal{G} so that the overall mean can be excluded from treatment differences. Again, if \mathcal{G} contains G_1 and G_2 but not $G_1 \wedge G_2$, then it is usually wise to consider whether this is reasonable in the given context; in other words, do not simply assume that the interaction is zero without explicitly considering it.

For orthogonality, the fixed and random factors together need to satisfy **(D1)** and **(D2)**. As in Section 10.5.4, it may be necessary to introduce pseudofactors in \mathcal{G}^* for some fixed effects, and then care is needed in interpreting the data.

10.8 MultiStage Experiments

In a multistage experiment, the same experimental material is used throughout but different sets of treatments are applied in different stages. For example, samples of blood may be collected under various conditions in Stage 1 and then stored under various conditions in Stage 2. For practical reasons, it may be necessary to apply treatments in Stage i to whole batches, rather than to individual observational units. See also Chapter 8.

Let B_i be the factor for Batches in Stage i . If the batch factors correspond to inherent differences in the observational units, then no new concepts are involved. However, there are other situations where the factors B_i are chosen simply so that the treatments can be applied economically, easily or quickly.

Example 10.12 (Batch processing of silicon wafers)

Mee and Bates (1998) describe an experiment where three two-level treatment factors T_1 , T_2 and T_3 are applied to silicon wafers, factor T_i being applied to batches in Stage i , for $i = 1, 2, 3$. There is freedom to choose the partition into batches at each stage. Two designs for 16 silicon wafers are compared.

In the first design, the wafers are grouped into two Lots of eight. Within each lot, a completely randomized design is used for the eight treatments, and the i th batch factor B_i is aliased with $T_i \wedge L$. This gives the partition into batches shown in Table 10.11a.

TABLE 10.11

Two Possible Partitions of 16 Wafers into Four Batches of Size Four in Each of the Three Stages in Example 10.12

	Lot	1	1	1	1	1	1	1	1	2	2	2	2	2	2	2	2
(a)	Stage 1 batch	1	1	1	1	2	2	2	2	3	3	3	3	4	4	4	4
	Stage 2 batch	1	1	2	2	1	1	2	2	3	3	4	4	3	3	4	4
	Stage 3 batch	1	2	1	2	1	2	1	2	3	4	3	4	3	4	3	4
	Stage 1 batch	1	1	1	1	2	2	2	2	3	3	3	3	4	4	4	4
(b)	Stage 2 batch	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4
	Stage 3 batch	1	3	4	2	4	2	1	3	2	4	3	1	3	1	2	4




FIGURE 10.32

Hasse diagrams for two choices for partitions into batches in the three-stage design in Example 10.12. (a) Partition (a) in Table 10.11. (b) Partition (b) in Table 10.11.

The Hasse diagram for the non-treatment factors is in Figure 10.32a. This is not a poset block structure, so it cannot be shown on a panel diagram. Infima such as $B_1 \wedge B_2$ are not included, but conditions (O1), (O2), (O3), (O4) and (O5) in Section 10.3 are satisfied.

In the second design, the wafers are imagined to be in a 4×4 array. The first-stage batches are the rows; the second-stage batches are the columns; the third-stage batches are given by the letters of a 4×4 Latin square: see Table 10.11b for the layout given by Mee and Bates (1998). Figure 10.32b shows the Hasse diagram. Now $\{U, B_1, B_2, B_3, E\}$ is an orthogonal block structure, similar to the one in Example 10.2.

The structure in Figure 10.32(b) can be generalized. Suppose that there are n^2 observational units, which need to be grouped into batches of size n in each of m stages, where $2 \leq m \leq n + 1$. The observational units are considered to form an abstract $n \times n$ array: factors B_1 and B_2 correspond to the partitions into rows and columns, while B_3, \dots, B_m are defined by the letters of $m - 2$ mutually orthogonal Latin $n \times n$ squares. See Mee and Bates (1998) and Butler (2004).

With conditions (O1)–(O5) satisfied, it is reasonable to assume that $\text{Cov}(Y)$ has the form given in (10.4). Multistage experiments are often full factorials with no replication, or fractional factorials. If there are too few degrees of freedom in the various strata for a meaningful ANOVA, then the relative importance of different treatment effects may be assessed

from half-normal plots, one for each stratum. Then the purpose of the skeleton ANOVA table is to show which treatment effect is in which stratum.

In some experiments, practical constraints imply that not every combination of levels of batches from different stages can occur. Then we need to consider suprema such as $B_i \vee B_j$.

Example 10.13 (An experiment on production of battery cells)

In an experiment reported by Vivacqua and Bisgaard (2009), electric battery cells were assembled in 16 batches in Stage 1. Four two-level treatment factors were applied at this stage, one treatment combination per batch. In the second stage, the battery cells were cured in batches of size eight, because of space limitations in the curing room: two further two-level treatment factors were applied to these batches.

Figure 10.33 shows two possible configurations for the batches. In (a), there are eight batches in Stage 2, and it is possible to have the full factorial set of treatments. In (b), there are only four batches in Stage 2, and so a half-fraction is required overall, even though each stage individually can have full replication. Vivacqua and Bisgaard (2009) call this *post-fractionation*. In both cases, we need to take account of the two-level factor $B_1 \vee B_2$, so the design must specify what treatment effect is confounded with $B_1 \vee B_2$.

When constraints on batch-size force B_1 and B_2 to have a non-trivial supremum, this must be included in the set of factors on Ω even though it is not, strictly speaking, inherent. For both possibilities in Example 10.13, put $H = B_1 \vee B_2$ and $\mathcal{F} = \{U, H, B_1, B_2, E\}$. Then \mathcal{F} is a poset block structure. For $i = 1, 2$, let S_i be the pre-factor for Stage i batches within H . Then the panel diagrams for the two possibilities are shown in Figure 10.34, and the Hasse diagrams in Figure 10.35.

Since both possibilities in Example 10.13 give poset block structures, the randomization argument in Section 10.2.5 gives $\text{Cov}(Y)$ of the form (10.4). On the other hand, the

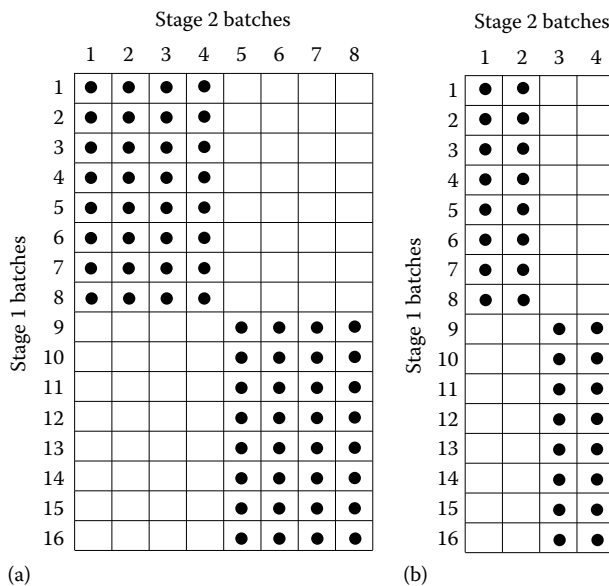


FIGURE 10.33 Two possible configurations for batches in Example 10.13: ● denotes one battery.




FIGURE 10.34
Panel diagrams for two possible configurations in Example 10.13.




FIGURE 10.35
Hasse diagrams for two possible configurations in Example 10.13.

assumption that E , B_1 and B_2 (but not H) have random effects gives (10.4) with the constraint that ξ_H is a particular linear combination of ξ_E , ξ_{B_1} and ξ_{B_2} . If the spectral components of variance (see Section 10.2.6) are to be estimated, or to be used in hypothesis testing, then this linear dependence can be problematical. However, half-normal plots are not meaningful for strata with very few degrees of freedom, so the stratum W_H would be omitted from such an analysis, and then the question of linear dependence among the spectral components of variance does not arise.

10.9 MultiPhase Experiments

In the first phase of a two-phase experiment, the set Θ of treatments is applied to a set Ω_1 of units using a design $\phi_1: \Omega_1 \rightarrow \Theta$. In the second phase, produce from the units of Ω_1 (e.g., food made from the crops, or machine parts manufactured during the runs) is applied to a second set Ω_2 of units, where responses are measured. A design $\phi_2: \Omega_2 \rightarrow \Omega_1$ is used for the second phase. Here we assume that there is an orthogonal treatment structure \mathcal{G} on Θ , and, for $i = 1, 2$, an orthogonal block structure \mathcal{F}_i on Ω_i . Further, both designs are orthogonal.

Example 10.14 (An experiment on food processing)

Figure 10.36 shows the panel diagrams for the structures on Θ , Ω_1 and Ω_2 in a simplified version of an experiment discussed by Brien and Bailey (2009). The six treatments consisted of three quantities of Irradiation in combination with two levels of Rosemary




FIGURE 10.36
Panel diagrams for Example 10.14.

(the herb was present or not). In the first phase, these treatments were applied in the baking of meat loaves: in each of three Batches, one Loaf received each treatment. The second phase had one Session per batch: in each session, pieces of the six meat-loaves were tasted, using a Latin square in which the rows corresponded to Order in the tasting sequence and the columns corresponded to Tasters.

As in Section 10.5, \mathcal{G} is augmented to \mathcal{G}^* and \mathcal{F}_1 is augmented to \mathcal{F}_1^* . Then Proposition 10.3 shows that there is a function $h_1: \mathcal{G}^* \rightarrow \mathcal{F}_1$ and a function $h_2: \mathcal{F}_1^* \rightarrow \mathcal{F}_2$ such that $W_G \leq W_{h_1(G)}$ for G in \mathcal{G}^* and $W_F \leq W_{h_2(F)}$ for F in \mathcal{F}_1^* .

When $\mathcal{F}_1 = \mathcal{F}_1^*$, it follows that $W_G \leq W_{h_1(G)} \leq W_{h_2(h_1(G))}$ for all G in \mathcal{G}^* . For H in \mathcal{F}_2 , the subspace W_H is the direct sum of the subspaces W_F for those factors F in \mathcal{F}_1 for which $h_2(F) = H$ and (usually) a residual subspace. In turn, W_F is the direct sum of those subspaces W_G for which $G \in \mathcal{G}^*$ and $h_1(G) = F$ and (usually) a residual, as in Section 10.5.2.

In Example 10.14, $h_1(I) = h_1(R) = h_1(IR) = BL$, $h_2(B) = S$ and $h_2(BL) = STO$. Thus, $W_I < W_{BL} < W_{STO}$ and similarly for W_R and W_{IR} . The residual in W_{BL} has $15 - 5 = 10$ degrees of freedom. The further residual in W_{STO} has $75 - 15 = 60$ degrees of freedom.

Assume that $E(\mathbf{Y})$ is determined by \mathcal{G} , as in Section 10.4.2, and that $\text{Cov}(\mathbf{Y})$ is a sum of two matrices like that on the right-hand side of (10.4), one for \mathcal{F}_1 and one for \mathcal{F}_2 . This is a natural assumption if both designs are randomized as in Section 10.2.5. For G in \mathcal{G}^* , three quantities now contribute to the expected mean square for W_G : the expectation part, τ_G ; the spectral component of variance, $\xi_{h_1(G)}$; and the spectral component of variance, $\xi_{h_2(h_1(G))}$. Thus, in Example 10.14 the standard error for the estimator of each treatment contrast depends on ξ_{BL} as well as ξ_{STO} .

Matters are not quite so straightforward when $\mathcal{F}_1^* \neq \mathcal{F}_1$. For overall orthogonality, we now need to insist that the design ϕ_1 is orthogonal when taking account of \mathcal{G} and \mathcal{F}_1^* , not just \mathcal{G} and \mathcal{F} . On the other hand, the extra W -subspaces for \mathcal{F}_1^* do not contribute any extra spectral components of variance.

Example 10.15 (Laboratory measurement)

In an experiment on laboratory animals, the structures on Θ , Ω_1 and Ω_2 had the panel diagrams shown in Figure 10.37. The 32 Animals were housed in Cages of four animals each. In the first phase, the animals were fed one of two Diets: these had to be allocated




FIGURE 10.37
Panel diagrams for Example 10.15.

to whole cages. They were also Injected, or not, with a special vaccine: two animals per cage were vaccinated.

In the second phase, some of a certain bodily fluid was extracted from each animal: samples of this were subject to analysis. Each run of the machine could accommodate eight samples, and these were distinguished by *labelling* each sample with a drop of coloured liquid, using eight colours per run.

An 8×8 Latin square was used to allocate the cages to combinations of Runs and Labels. With 64 observations, 32 animals and 8 runs, it is impossible to have every animal represented in each run. Instead, the runs were partitioned into four pairs by a four-level pseudofactor P_R , and each animal was allocated to one such pair. Similarly, the labels were partitioned into four pairs by a pseudofactor P_L . The experimental layout consisted of a 4×4 grid of 2×2 subsquares of the following form.

Cage <i>a</i>	Cage <i>b</i>
Animal 1	Animal 2
Cage <i>b</i>	Cage <i>a</i>
Animal 2	Animal 1

Adjacent subsquares had different levels of I . Since P_R and P_L had no inherent meaning, being merely devices to construct the design and being ignored in the randomization, this confounding of I with the 2×2 subsquares was not problematic.

Now $\mathcal{G} = \mathcal{G}^* = \{U, D, I, DI\}$, $\mathcal{F}_1 = \{U, C, CA\}$ and $\mathcal{F}_2 = \{U, R, L, RL\}$. Also, $CA \vee R = P_R$ and $CA \vee L = P_L$, and so $\mathcal{F}_1^* = \{U, C, CA, P_R, P_L\}$. Figure 10.38 gives the combined Hasse diagram, using white squares for factors in \mathcal{F}_2 . This gives the skeleton ANOVA in Table 10.12.

For those who prefer formulae to pictures, this design may also be constructed using the following pair of design keys (see Section 10.6).

$$\Phi_1(D) = C_3 \quad \Phi_1(I) = C_2 + A_1$$

$$\Phi_2(C_1) = R_1 + L_1 \quad \Phi_2(C_2) = R_2 + L_2 \quad \Phi_2(C_3) = R_3 + L_3 \quad \Phi_2(A_1) = R_1 \quad \Phi_2(A_2) = R_2$$

The confounding shown in Table 10.12 follows directly from this. For example, $\Phi_2(C_1 + A_1) = L_1$, $\Phi_2(C_2 + A_2) = L_2$ and $\Phi_2(C_1 + C_2 + A_1 + A_2) = L_1 + L_2$,




FIGURE 10.38
Combined Hasse diagram for Example 10.15.

TABLE 10.12
Skeleton Analysis of Variance for Example 10.15

Second-Phase Units		First-Phase Units		Treatments	
U	1	U	1	Mean	1
Runs	7	P_R (part of $A[C]$)	3		3
		residual	4		4
Labels	7	P_L (part of $A[C]$)	3		3
		residual	4		4
$R\#L$	49	Cages	7	Diets	1
				residual	6
		rest of animals[C]	18	Irradiation	1
				$D\#I$	1
				residual	16
		residual	24		24
Total					64

so these three degrees of freedom for Animals[C] are confounded with Labels. Likewise, $\Phi_1(D + I) = C_2 + C_3 + A_1$ and $\Phi_2(C_2 + C_3 + A_1) = R_1 + R_2 + R_3 + L_2 + L_3$, and so the D -by- I interaction is confounded with the part of Animals[C] that is confounded with $R\#L$.

This approach extends to three or more phases in the obvious way.

10.10 Orthogonality

This chapter has concentrated on orthogonality. Why is orthogonality desirable?

If the inherent factors on Ω (or, more generally, the factors with random effects) are not uniform and mutually orthogonal, then the eigenspaces of $\text{Cov}(\mathbf{Y})$ are not known unless all of the variance components are already known. This means that there is no uniformly best estimator of treatment effects.

If the structure on Θ (or, more generally, on the factors with fixed effects) is not orthogonal, then, in the Hasse diagram of expectation models, different edges which correspond to the same factor G can have different sums of squares associated with them, in each case being the difference between the sums of squares of fitted values in the larger and smaller models at the ends of the edge. This has the consequence that different routes down from the top of the Hasse diagram can lead to different decisions about which model to fit to the data.

Even if the structures on Ω and Θ are both orthogonal, the design ϕ may not be. Let \mathbf{X}_ϕ be the $N \times |\Theta|$ matrix whose (ω, i) -entry is equal to 1 if $\phi(\omega) = 1$ and is equal to 0 otherwise. If treatment effects are estimated using only $\mathbf{Q}_F \mathbf{y}$, which is the data projected onto stratum W_F , then the variance-covariance matrix of the estimators is $\xi_F(\mathbf{X}'_\phi \mathbf{Q}_F \mathbf{X}_\phi)^-$. The matrix $\mathbf{X}'_\phi \mathbf{Q}_F \mathbf{X}_\phi$ is called the *information matrix* for stratum W_F . As discussed in Chapter 3, knowledge of the eigenvectors and eigenvalues of the information matrix is important for understanding the design and assessing its properties.

If there are more than two strata apart from W_U , the eigenvectors of the information matrices for different strata may be different. There are some simplifications when all of the information matrices have common eigenvectors. In this case, the design is said to have *general balance*: see Nelder (1965b) and Houtman and Speed (1983). If, moreover, those common eigenspaces are the subspaces W_G for G in \mathcal{G}^* , the design is said to have *structure balance*: see Brien and Bailey (2009).

Of course, many experiments in practice lack orthogonality in some aspect. For example, many designs in Chapter 3 are not orthogonal, and some of the supersaturated designs in Chapter 9 are not orthogonal. Each pair of factors in an orthogonal array (also Chapter 9) is orthogonal with supremum U , but often infima are not included. However, it is helpful to understand the orthogonal case fully before considering non-orthogonal ones.

10.11 Notes on Terminology and Conventions

The theory of poset block structures was fully developed by Kempthorne et al. (1961) and Zyskind (1962) under the name *complete balanced response structures*. However, their diagrams show E and U as well as the pre-factors, and so are somewhere in between the panel diagrams and Hasse diagrams shown here. Nelder (1965a) developed a much more algorithmic approach with iterated crossing and nesting, calling his structures *simple orthogonal block structures*. Speed and Bailey (1982) generalized both of these to *lattices of commuting uniform equivalence relations*, which are here called orthogonal block structures. These approaches, and more, were discussed and named by Bailey (1996).

Houtman and Speed (1983) generalized the family of variance-covariance matrices given by (10.4). They defined such a family to be an *orthogonal block structure* if it consists of all positive semi-definite linear combinations of a known set of matrices like $\{Q_F : F \in \mathcal{F}\}$ which are mutually orthogonal symmetric idempotents summing to the identity. In general, their definition does not require the idempotents to be defined by factors, unlike the definition in Section 10.3.

Whenever there is a partial order, there are two opposite conventions for what is meant by “bigger.” For example, the partial order \subseteq on the set of ancestral subsets in Example 10.3 gives a Hasse diagram the opposite way up to that in Figure 10.5. The partial order \preceq used in this chapter is consistent with the usual notion of refinement of partitions, but some statistical authors prefer to think of a factor as “bigger” if it has more levels. Either approach works perfectly well if used consistently, but the words *infimum* and *supremum* have to be interchanged if the partial order is given in the other direction.

Some authors call ancestral subsets *upsets*; some call them *filters*.

This chapter introduces two partial orders (in addition to *is a subset of* and *is a subspace of*): nesting of pre-factors and refinement of genuine factors. Some statistical authors mix these up and call them both *nesting*.

As pointed out in Sections 10.2 through 10.4, each factor F defines two subspaces— V_F and W_F —and it can be hard to provide a succinct but clear notation for W_F in ANOVA tables. It may be because of this that some authors do not distinguish between (1) $F \wedge G$, which is a factor, (2) the interaction between F and G , which consists of the values in the vector $Q_{F \wedge G} \tau$ when the levels of $F \wedge G$ have fixed effects, and (3) the effects of the levels of $F \wedge G$ when these are random. When pseudofactors are used, it is common to write FG

rather than $F + G$, and this may be confused with all of (1)–(3) as well as the vector of fitted values in $Q_{F+G}\mathbf{y}$, or just one of those values if all factors have two levels.

References

- Bailey, R. A. (1982), Latin squares with highly transitive automorphism groups, *Journal of the Australian Mathematical Society, Series A*, **33**, 18–22.
- Bailey, R. A. (1991), Strata for randomized experiments (with discussion), *Journal of the Royal Statistical Society, Series B*, **53**, 27–78.
- Bailey, R. A. (1996), Orthogonal partitions in designed experiments, *Designs, Codes and Cryptography*, **8**, 45–77.
- Bailey, R. A. (2004), *Association Schemes: Designed Experiments, Algebra and Combinatorics*, Cambridge Studies in Advanced Mathematics, 84, Cambridge, U.K.: Cambridge University Press.
- Bailey, R. A. (2008), *Design of Comparative Experiments*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge, U.K.: Cambridge University Press.
- Bailey, R. A. and Monod, H. (2001), Efficient semi-Latin rectangles: Designs for plant disease experiments, *Scandinavian Journal of Statistics*, **28**, 257–270.
- Bailey, R. A., Praeger, C. E., Rowley, C. A., and Speed, T. P. (1983), Generalized wreath products of permutation groups, *Proceedings of the London Mathematical Society*, **47**, 69–82.
- Blouin, D. C., Taverner, J. D., and Beasley, J. S. (2009), A composite Latin rectangle and nonstandard strip block design, *Journal of Agricultural, Biological and Environmental Statistics*, **14**, 484–494.
- Brien, C. J. and Bailey, R. A. (2006), Multiple randomizations (with discussion), *Journal of the Royal Statistical Society, Series B*, **68**, 571–609.
- Brien, C. J. and Bailey, R. A. (2009), Decomposition tables for experiments. I. A chain of randomizations, *Annals of Statistics*, **37**, 4184–4213.
- Brien, C. J. and Bailey, R. A. (2010), Decomposition tables for experiments. II. Two–one randomizations, *Annals of Statistics*, **38**, 3164–3190.
- Brien, C. J., Harch, B. D., Correll, R. L., and Bailey, R. A. (2011), Multiphase experiments with at least one later laboratory phase. I. Orthogonal designs, *Journal of Agricultural, Biological and Environmental Statistics*, **16**, 422–450.
- Butler, N. A. (2004), Construction of two-level split-lot fractional factorial designs for multistage processes, *Technometrics*, **46**, 445–451.
- Cheng, C.-S. (2014), *Theory of Factorial Design: Single- and Multi-Stratum Experiments*, Monographs on Statistics and Applied Probability, 131, Boca Raton, FL: CRC Press.
- Clarke, G. M. and Kempson, R. E. (1997), *Introduction to the Design and Analysis of Experiments*, London, U.K.: Arnold.
- Cochran, W. G. and Cox, G. M. (1957), *Experimental Designs*, 2nd edn., New York: Wiley.
- Gardner, S. C. and Grue, C. E. (1996), Effects of Rodeo[®] and Garlon[®] 3A on nontarget wetland species in Central Washington, *Environmental Toxicology and Chemistry*, **15**, 441–451.
- Großmann, H. (2014), Automating the analysis of variance of orthogonal designs, *Computational Statistics and Data Analysis*, **70**, 1–18.
- Houtman, A. M. and Speed, T. P. (1983), Balance in designed experiments with orthogonal block structure, *Annals of Statistics*, **11**, 1069–1085.
- Kemphorne, O., Zyskind, G., Addelman, S., Throckmorton, T. N., and White, R. F. (1961), Analysis of Variance Procedures, Aeronautical Research Laboratory, Wright-Patterson AFB, OH, Report No. 149.
- Kobilinsky, A., Monod, H., and Bailey, R. A., (2014), Automatic generation of generalised regular factorial designs, submitted for publication.

- Mee, R. W. and Bates, R. L., (1998) Split-lot designs: Experiments for multistage batch processes, *Technometrics*, **40**, 127–140.
- Nelder, J. A. (1954), The interpretation of negative components of variance, *Biometrika*, **41**, 544–548.
- Nelder, J. A. (1965a), The analysis of randomized experiments with orthogonal block structure. I. Block structure and the null analysis of variance, *Proceedings of the Royal Society (London), Series A*, **273**, 147–162.
- Nelder, J. A. (1965b), The analysis of randomized experiments with orthogonal block structure. II. Treatment structure and the general analysis of variance, *Proceedings of the Royal Society (London), Series A*, **273**, 163–178.
- Patterson, H. D. (1965), The factorial combination of treatments in rotation experiments, *Journal of Agricultural Science*, **65**, 171–182.
- Patterson, H. D. and Bailey, R. A. (1978), Design keys for factorial experiments, *Journal of the Royal Statistical Society, Series C (Applied Statistics)*, **27**, 335–343.
- Preece, D. A., Bailey, R. A., and Patterson, H. D. (1978), A randomization problem in forming designs with superimposed treatments, *Australian Journal of Statistics*, **20**, 111–125.
- Speed, T. P. and Bailey, R. A. (1982), On a class of association schemes derived from lattices of equivalence relations, in *Algebraic Structures and Applications* (eds. P. Schultz, C. E. Praeger, and R. P. Sullivan), Marcel Dekker, New York, pp. 55–74.
- Tjur, T. (1984), Analysis of variance models in orthogonal designs (with discussion), *International Statistical Review*, **52**, 33–81.
- Vivacqua, C. A. and Bisgaard, S. (2009), Post-fractionated strip-block designs, *Technometrics*, **51**, 47–55.
- Wilkinson, G. and Rogers, C. E. (1973), Symbolic description of factorial models for analysis of variance, *Journal of the Royal Statistical Society, Series C (Applied Statistics)*, **22**, 392–399.
- Zyskind, G. (1962), On structure, relation, Σ , and expectation of mean squares, *Sankhyā, Series A*, **24**, 115–148.

11

Algebraic Method in Experimental Design

Hugo Maruri-Aguilar and Henry P. Wynn

CONTENTS

11.1	Introduction.....	415
11.2	Ideals and Varieties: Introducing Algebra.....	417
11.3	Gröbner Bases.....	419
11.3.1	Term Orderings.....	420
11.3.2	Matrix- and Vector-Based Term Orderings.....	421
11.3.3	Software Computations: CoCoA.....	422
11.3.4	Other Software Tools.....	422
11.3.5	Monomial Ideals and Hilbert Series.....	423
11.3.6	Gröbner Bases.....	425
11.4	Experimental Design.....	428
11.5	Examples.....	431
11.5.1	Response Surface Designs.....	431
11.5.2	Two-Level Designs.....	435
11.5.3	Other Designs.....	438
11.6	Understanding Aliasing.....	440
11.7	Indicator Functions and Orthogonality.....	443
11.8	Fans, State Polytopes and Linear Aberration.....	446
11.8.1	Algebraic Fan of a Design.....	446
11.8.2	State Polytope and Linear Aberration.....	449
11.9	Other Topics and References.....	451
	References.....	452

11.1 Introduction

Most readers will be familiar with the use of polynomials throughout statistics. Much of this handbook is concerned with design for polynomial regression (see, e.g., Chapters 2 and 5). Thus, we have polynomial terms in variables x_1, x_2, \dots , such as

$$x_1, x_2^2, x_1x_2, \dots$$

The first, but very important, algebraic point is that polynomials are made up of linear combinations of monomials. The next example will help to outline the scope of this chapter.

Example 11.1

Consider a central composite design of nine points in two factors (see also Chapter 5) composed of factorial points $(\pm 1, \pm 1)$ and axial points $(0, \pm\sqrt{2}), (\pm\sqrt{2}, 0)$ together with the origin $(0, 0)$. The second-order polynomial response surface in two factors is

$$f(x_1, x_2) = \theta_{00} + \theta_{10}x_1 + \theta_{01}x_2 + \theta_{20}x_1^2 + \theta_{11}x_1x_2 + \theta_{02}x_2^2, \tag{11.1}$$

where the coefficients $\theta_{00}, \dots, \theta_{02}$ are fixed but unknown constants. A direct computation of the 9×6 design-model matrix X using the aforementioned design and the second-order model (11.1) shows that the model is identifiable by the design, leaving three degrees of freedom available for estimation of the error. We now add monomial terms to the model to exhaust these three degrees of freedom while respecting the hierarchy principle of polynomial terms, that is, we can only add a term if its divisors are already in the model. After a few attempts by trial and error, we note that the design has five levels in each factor, so we could augment terms x_1^3, x_1^4 and, finally, the term $x_1^2x_2$ (see Table 11.1). A different possibility, suggested by the design symmetry, would be to instead add x_2^3, x_2^4 and $x_1x_2^2$. No other hierarchical models are found, and thus for this design, we have two possible saturated models. The first model has terms $1, x_1, x_2, x_1^2, x_1x_2, x_2^2, x_1^3, x_1^4, x_1^2x_2$, while the second model is $1, x_1, x_2, x_1^2, x_1x_2, x_2^2, x_2^3, x_2^4, x_1x_2^2$.

The search earlier is known as a direct problem, where the design is assumed known in advance and an identifiable model is sought. The algebraic method of this chapter provides a structured framework to solve the direct problem of identification of saturated hierarchical polynomial models for a given design. In the example earlier, we built the design-model matrix while balancing the hierarchy principle. This same example will be expanded using algebraic techniques in Example 11.13. However, the algebraic method goes beyond this initial task and provides a comprehensive way of understanding aliasing of model terms. The scope of designs that can be analysed with algebraic techniques is vast (see Section 11.4 for a sample of applications to design). For analysing designs, we only consider the minimal support set as replications are not considered.

The central object required for this analysis is the *design ideal*. This is a polynomial ideal, familiar from commutative algebra. The design ideal is generated by a finite set of polynomials called a *Gröbner basis*. This special basis is used to identify models and also to study confounding, as it encodes information about confounding induced by the design. Computation of Gröbner bases requires the specification of a *term order*.

TABLE 11.1

Design-Model Matrix for Central Composite Design of Example 11.1

Points	1	x_1	x_2	x_1^2	x_1x_2	x_2^2	x_1^3	x_1^4	$x_1^2x_2$
$(-1, -1)$	1	-1	-1	1	1	1	-1	1	-1
$(-1, 1)$	1	-1	1	1	-1	1	-1	1	1
$(1, -1)$	1	1	-1	1	-1	1	1	1	-1
$(1, 1)$	1	1	1	1	1	1	1	1	1
$(0, 0)$	1	0	0	0	0	0	0	0	0
$(0, -\sqrt{2})$	1	0	$-\sqrt{2}$	0	0	2	0	0	0
$(0, \sqrt{2})$	1	0	$\sqrt{2}$	0	0	2	0	0	0
$(-\sqrt{2}, 0)$	1	$-\sqrt{2}$	0	2	0	0	$-2\sqrt{2}$	4	0
$(\sqrt{2}, 0)$	1	$\sqrt{2}$	0	2	0	0	$2\sqrt{2}$	4	0

In the motivating example earlier, we already see several elements that are studied later in this chapter. First, there is the general problem of identifying a model for a given design. Second, we may obtain different models. Note that the main interest is in linking a set of monomial terms with the design in such a way that the corresponding design-model matrix will be square and full rank. A statistician may not necessarily want to use a saturated model in practice, and therefore, this may be the starting point to apply model selection techniques. For example, the first interpolator model stemming from Example 11.1 has the form

$$f(x_1, x_2) = \theta_{00} + \theta_{10}x_1 + \theta_{01}x_2 + \theta_{20}x_1^2 + \theta_{11}x_1x_2 + \theta_{02}x_2^2 + \theta_{30}x_1^3 + \theta_{40}x_1^4 + \theta_{21}x_1^2x_2, \quad (11.2)$$

from which terms could be dropped in a later model selection stage. Throughout this chapter, one main interest is for a given design, to produce a list of estimable polynomial terms which form the support of a saturated model.

This chapter starts by outlining the basic theory behind ideals in Section 11.2. The core background material of this chapter is that of the Gröbner basis which is covered in Section 11.3. Throughout the chapter we show some examples using the software CoCoA, which is introduced in Section 11.3.3. We collect the results as an algorithm in Section 11.4 and present a series of examples of the algebraic method using different designs in Section 11.5. Aliasing of monomial terms is discussed in Section 11.6, whereas the role of indicator functions and orthogonality is presented in Section 11.7. In Section 11.8, we discuss the impact of different term orderings in the collection of models. In Section 11.9, we give a brief survey of other topics including recent work on inverse problems on design.

11.2 Ideals and Varieties: Introducing Algebra

The basic reference for material covering algebra, definitions of objects used in this chapter (such as fields, polynomials, rings, ideals and Gröbner bases), is Cox et al. (2007), Chapters 1, 2, 6 and 9. We first define monomials and polynomials.

Definition 11.1 Consider a set of k variables x_1, \dots, x_k and a row vector of non-negative integers $\alpha = (\alpha_1, \dots, \alpha_k)$; a monomial is

$$x^\alpha = x_1^{\alpha_1} x_2^{\alpha_2} \dots x_k^{\alpha_k}.$$

The degree of the monomial x^α is the sum of its exponents, which we denote as $|\alpha| := \sum_{i=1}^k \alpha_i$. A polynomial is a finite linear combination of m monomials

$$f(x) = \sum_{i=1}^m \theta_{\alpha^{(i)}} x^{\alpha^{(i)}},$$

where the coefficients $\theta_{\alpha^{(1)}}, \dots, \theta_{\alpha^{(m)}}$ take values over a base field K and $\alpha^{(i)} = (\alpha_{(i)1}, \dots, \alpha_{(i)k})$.

When we use the term polynomial we typically mean a polynomial, in one or more variables, and unless it is required, we refer to a polynomial as f rather than $f(x)$. Note that contrary to most of the material in the book, we will not use bold font for the vector of exponents α , and we will distinguish between a vector and a component in that the component will have a subindex. Variables such as x_1, a, a^2c are used throughout this chapter, and no bold font has been used to refer to them, unless we specifically refer to a point such as x which is to be interpreted as a row vector.

A monomial x^α can be represented by its exponent vector α , and we can list the monomials in a model either directly or by listing a set of exponents. We shall often use the notation $\{x^\alpha, \alpha \in M\}$, for some set of exponents, M . For example, the set M for the polynomial model in (11.1) is $\{(0, 0), (1, 0), (0, 1), (2, 0), (1, 1), (0, 2)\}$.

This chapter is largely concerned with the link between a design and the list M of estimable terms. We know from classical factorial design (Chapter 7) that only some models are estimable for a given design and so any such theory must be intimately related to the problem of aliasing and we shall cover this in Section 11.6.

The set of all polynomials over a base field is a ring. Thus, given a base field K , we obtain the ring of polynomials, $K[x_1, \dots, x_k]$, which is the set of all polynomials with coefficients in the base field (see Cox et al. 2007). Our “ α ” notation involving exponents allows a generic polynomial to be written compactly as

$$f(x) = \sum_{\alpha \in M} \theta_\alpha x^\alpha,$$

where M is a finite set of distinct exponents. The element $\theta_\alpha \in K$ is the coefficient of term x^α ; thus clearly $f(x) \in K[x_1, \dots, x_k]$. The quantity θ_α has the same meaning of a parameter in a statistical polynomial model and usually θ_α will be allowed to be a real number so $K = \mathbb{R}$.

Given that we have launched into algebra, we need to introduce the first two essentials: *ideals* and *varieties*. In what follows we present only the basic ideas of the theory, pointing the reader to Cox et al. (2007) and Pistone et al. (2001) for further details. For a ring $K[x_1, \dots, x_k]$, we have special subsets of it called *ideals*.

Definition 11.2 *A subset $I \subset K[x_1, \dots, x_k]$ is an ideal if for any $f, g \in I$, we have $f + g \in I$, and for any $f \in I$ and $g \in K[x_1, \dots, x_k]$, we have $fg \in I$.*

The ideal generated by a finite set of polynomials $\{f_1, \dots, f_m\}$ is the set of all polynomial combinations of f_1, \dots, f_m , which we define as

$$\langle f_1, \dots, f_m \rangle = \{f_1g_1 + \dots + f_mg_m : g_1, \dots, g_m \in K[x_1, \dots, x_k]\}.$$

To have some immediate intuition, consider a single point $d \in \mathbb{R}^k$. If the point has components $d = (d_1, d_2, \dots, d_k)$, then $\langle x_1 - d_1, x_2 - d_2, \dots, x_k - d_k \rangle$ contains all polynomials that vanish over d . If say $d = (0, 1, -1) \in \mathbb{R}^3$, clearly $f(x_1, x_2, x_3) = x_1x_2 + x_2^2 - x_2x_3 - x_1 - x_2 + x_3 \in \langle x_1, x_2 - 1, x_3 + 1 \rangle$ because $f(x_1, x_2, x_3) = (x_2 - 1)(x_1 + x_2 - x_3)$. The ideal generated by a single point as previously mentioned is known as a *maximal ideal* (see Cox et al. 2007). Maximal ideals are the building blocks of design ideals as will be seen in Section 11.4.

The set of all polynomials $f(x) \in K[x_1, \dots, x_k]$ such that $f(x) = 0$ is an ideal, since for any polynomial $g(x)$ if $f(x) = 0$, we have $g(x)f(x) = 0$. The Hilbert basis theorem says that any (polynomial) ideal I is finitely generated, that is, for any ideal I , we can find a finite collection of polynomials $f_1, \dots, f_m \in K[x_1, \dots, x_k]$ such that $I = \langle f_1, \dots, f_m \rangle$, and for some integer m , recall the notation $\langle \cdot \rangle$ as defined earlier.

Readers will be familiar with linear varieties expressed by setting a linear polynomial function equal to zero. Thus a straight line can be written as the collection of points (x_1, x_2) in two dimensions such that

$$ax_1 + bx_2 + c = 0,$$

for constants a, b, c . An *affine variety* is the extension of this concept to simultaneous solutions of a set of polynomial equations, as stated in the following definition.

Definition 11.3 Let $f_1, \dots, f_m \in K[x_1, \dots, x_k]$ be a set of polynomials. The associated affine variety is the solution (also called the zero set) of a set of simultaneous equations defined by these polynomials:

$$V(f_1, \dots, f_m) = \{(a_1, \dots, a_k) \in K^k : f_i(a_1, \dots, a_k) = 0, i = 1, \dots, m\}.$$

Every affine variety V has an associated ideal which we write $I(V)$. It is the set of all polynomials which are zero on the variety:

$$I(V) = \{f \in K[x_1, \dots, x_k] : f(a_1, \dots, a_k) = 0, \text{ for all } (a_1, \dots, a_k) \in V\}.$$

What appears to be a straightforward relationship between ideals and varieties is actually very subtle. If we start with polynomials f_1, \dots, f_m and construct the corresponding variety V and form the ideal $I(V)$, is it true that $I(V) = \langle f_1, \dots, f_m \rangle$? For example, take the ideal $I = \langle x_1 - 1, x_1^2 + 1 \rangle$. If we work on the polynomial ring over the real field $K = \mathbb{R}$, the polynomial equation $x_1^2 + 1 = 0$ cannot be solved, and the variety $V(I)$ consists of the single point $\{1\}$. We can always claim that $\langle f_1, \dots, f_m \rangle \subset I(V)$, but the converse may not be true; refer to Cox et al. (2007) for a detailed discussion. Following the idea earlier, the ideal generated by the single point is $\langle x_1 - 1 \rangle$ which does not equal I , yet the following contention holds $\langle x_1 - 1, x_1^2 + 1 \rangle \subset \langle x_1 - 1 \rangle$. Fortunately, for a design with real points as will be seen in Section 11.4, the variety is collection of isolated single points and the equivalence holds and we may move freely between ideals and designs.

11.3 Gröbner Bases

At the core of this chapter, we have the concept of Gröbner bases (see Buchberger 2006). We will use them to define in a compact way the property of linear independence of model terms for a given design, and they will be used later on to study confounding in designs.

To start our study, we require the notion of *quotient*, which is perhaps the most important construction in abstract algebra. Given two polynomials $f, g \in K[x_1, \dots, x_k]$ and an ideal

$I \subset K[x_1, \dots, x_k]$, define the equivalence class $f \sim_I g$ if and only if $f - g \in I$. The members of the quotient $K[x_1, \dots, x_k]/I$ are the equivalence classes. Since $f_1 \sim_I f_2$ and $g_1 \sim_I g_2$ imply $f_1 + g_1 \sim_I f_2 + g_2$ and $f_1 g_1 \sim_I f_2 g_2$, then $K[x_1, \dots, x_k]/I$ is also a ring. Finding $K[x_1, \dots, x_k]/I$ in a particular case requires a *division algorithm*.

11.3.1 Term Orderings

We require term ordering to perform the division. The following simple example recalls division of polynomials in one dimension over the real field $K = \mathbb{R}$.

Example 11.2

If we divide $1 + 3x + 2x^2 + x^3$ by $2 + x$, we obtain the tableau

$$\begin{array}{r|l}
 & x^2 + 3 \\
 x + 2 & | \quad x^3 + 2x^2 + 3x + 1 \\
 & \quad \underline{-(x^3 + 2x^2)} \\
 & \qquad \qquad 3x + 1 \\
 & \qquad \qquad \underline{-(3x + 6)} \\
 & \qquad \qquad \qquad \qquad -5
 \end{array}$$

giving $x^3 + 2x^2 + 3x + 1 = (x^2 + 3)(x + 2) - 5$.

We give this example to remind ourselves that at each stage of the division, we need to use the leading terms of the intervening polynomials. The leading term is the largest monomial of a polynomial. To obtain the leading term, we need an ordering for monomials. In one dimension, the ordering is

$$1 < x < x^2 < \dots,$$

that is, we order by degree and division is unique. This is generalised to a special total ordering on the set of all monomials x^α where the exponent vector α has non-negative integer entries.

Definition 11.4 A monomial term ordering, $<$, is a total ordering of monomials such that (i) $1 < x^\alpha$ for all $\alpha \geq 0$, $\alpha \neq 0$ and (ii) for all $\gamma \geq 0$, $x^\alpha < x^\beta$ implies $x^{\alpha+\gamma} < x^{\beta+\gamma}$.

The notation $\alpha \geq 0$ means that all elements of α are simultaneously greater than or equal to zero. There are a number of standard monomial orderings:

1. *Lexicographic ordering, Lex*: $x^\alpha <_{Lex} x^\beta$ when the leftmost nonzero entry of $\beta - \alpha$ is positive.

Graded orderings are orderings in which the first comparison between monomials is determined by their degree. The aforementioned degree lexicographic and degree reverse lexicographic term orders fall in this class.

2. *Degree lexicographic ordering, DegLex*: $x^\alpha <_{DegLex} x^\beta$ if (i) the degree of x^α is less than that of x^β , that is, $|\alpha| < |\beta|$ or (ii) $|\alpha| = |\beta|$ and $\alpha <_{Lex} \beta$.
3. *Degree reverse lexicographic ordering, DegRevLex*: $x^\alpha <_{DegrevLex} x^\beta$ if (i) $|\alpha| < |\beta|$ or (ii) $|\alpha| = |\beta|$ and $\bar{\alpha} <_{Lex} \bar{\beta}$, where the overline means reverse the entries.

Example 11.3

Under a lexicographic term ordering, $x_2^3 x_3^4 <_{Lex} x_1$ because $\beta - \alpha = (1, -3, -4)$, while for a degree lexicographic ordering, $x_1 <_{DegLex} x_2^3 x_3^4$ since $|(1, 0, 0)| = 1 < |(0, 3, 4)| = 7$.

11.3.2 Matrix- and Vector-Based Term Orderings

Monomial term orderings can be defined using products with matrices and element-wise comparisons. If the exponents of monomials x^α, x^β are considered as row vectors, we say that $x^\alpha <_M x^\beta$ if $\alpha M' < \beta M'$, where M is a non-singular matrix, and the inequality is tested in a lexicographic manner, starting from the leftmost element. We refer to M as an ordering matrix. It satisfies certain conditions which are stated in the following theorem taken from Pistone et al. (2001).

Theorem 11.1 *Let M be a full rank matrix of size $k \times k$ such that the first non-zero entry in each column is positive. Then M defines a term ordering in the following sense:*

1. For every vector of non-negative integers α with $\alpha \neq (0, \dots, 0)$, $(0, \dots, 0) < \alpha M'$.
2. For any non-negative vectors α, β, γ such that $\alpha M' < \beta M'$, $(\alpha + \gamma) M' < (\beta + \gamma) M'$.

Example 11.4

The identity matrix of size k corresponds to the lexicographic term ordering. For $k = 3$ and monomials of Example 11.3, we test whether $x_2^3 x_3^4 < x_1$ by comparing row vectors $(0, 3, 4)$ and $(1, 0, 0)$ (after right multiplication by the identity). As the first entry is smaller on the left vector, then the ordering is true.

Note that the relationship between ordering matrices and term orderings is not one to one. A matrix M_2 defining the same ordering as M can be obtained by multiplying each row of M by a positive constant, so, for instance, the matrix with diagonal $1, 2, \dots, k$ and zeroes elsewhere also defines a lexicographic term ordering. Usually only integer entries are used for computations although the theory does not preclude using, for instance, matrices with rational or real entries (see Cox et al. 2007).

An important case of ordering matrices is that of matrices for graded orderings. Any full rank matrix M in which all elements of the first row are a positive constant defines a graded ordering. The degree lexicographic term ordering is built with a matrix M with all entries one in its first row, and the remaining rows are the top $k - 1$ rows of an identity matrix.

A more specialised and efficient instance of matrix orderings is produced by using a single row matrix with positive integer entries, in which case we say *ordering vector*. An ordering vector defines only a partial but not a total ordering over $K[x_1, \dots, x_k]$. For example, the vector $w = (1, 1, 1)$ naturally produces the ordering $xy^2 \succ_w xz$ because

$$(1, 2, 0)(1, 1, 1)' = 3 > 2 = (1, 0, 1)(1, 1, 1)',$$

yet it cannot distinguish between monomials of the same degree such as xy^2 and z^3 because $(1, 2, 0)(1, 1, 1)' = (0, 0, 3)(1, 1, 1)' = 3$.

The main utility of vector orderings comes when noting that Gröbner bases of Section 11.3.6 are computed over finite sets of monomials rather than over all monomials with exponents in $\mathbb{Z}_{\geq 0}^k$. This last fact, together with a careful selection of the ordering vector, is at the core of a special algorithm to compute Gröbner bases (see Babson et al. 2003; Maruri-Aguilar 2005).

11.3.3 Software Computations: CoCoA

The free software CoCoA is useful to perform computations in commutative algebra. The package has extensive documentation for functions and allows the user to program in its own language (see CoCoATeam 2009). Versions of the software for different platforms are downloadable from the url `cocoa.dima.unige.it`. Here and elsewhere in the chapter we give code examples which can be run either in the command line version or the graphical user interface (GUI).

We create the ring $\mathbb{Q}[x, y, z]$ and specify a degree lexicographic term ordering with the CoCoA command

```
Use T:=Q[x,y,z], DegLex;
```

This ring is the set of all polynomials with rational coefficients, that is, $K = \mathbb{Q}$. We can specify the same ring and ordering using matrix orderings; the command `Mat` is used to create a matrix which then defines a term order with the command `Ord` shown as follows:

```
M:=Mat([[1,1,1],[1,0,0],[0,1,0]]);
Use T:=Q[x,y,z], Ord(M);
```

Using this ordering, the query “ $xy^2 > x^2z$;” yields output `FALSE`, which means that $xy^2 < x^2z$ under the graded lexicographic order in which $x > y > z$. The standard ordering in the software system CoCoA is the degree reverse lexicographic (`DegRevLex`), which is implicit in the following ring definition:

```
Use T:=Q[x,y,z];
xy^2>x^2z;
```

The output of the query is now `TRUE`, and this is interpreted as $xy^2 > x^2z$ under a degree reverse lexicographic term ordering in which $x > y > z$. Note the reversal of the ordering between the two monomials obtained for the previous graded order.

11.3.4 Other Software Tools

Some popular computational algebra packages relevant for this chapter are CoCoA (introduced earlier), *macaulay2* (Grayson and Stillman 2009), *gfan* (Jensen 2009) and *Singular* (Decker et al. 2012). A rough list of capabilities relevant to this chapter is as follows:

1. Construction of monomial orderings: the standard ones are usually named and immediately available.
2. Ideal operations such as unions, intersections and elimination.
3. Buchberger algorithm and modern improvements, quotienting and computing normal forms.
4. Special algorithms for *ideals of points*.
5. Gröbner fan computations.

We use most of capabilities 1–4 in examples in Section 11.4. Buchberger’s algorithm is not used explicitly, yet it is at the core of Gröbner basis computations. The computations of Gröbner fan are used in Section 11.8.

11.3.5 Monomial Ideals and Hilbert Series

Now that we have defined a total ordering, any finite set of monomials has a leading term. In particular, since a polynomial, f , is based on a finite set of monomials, it has a unique leading term. We write it $LT_{<}(f)$, or, if $<$ is assumed, just $LT(f)$. A monomial ideal is an ideal generated by monomials.

Definition 11.5 *Let $f_1(x), \dots, f_m(x)$ be a collection of monomials. The monomial ideal generated by $f_1(x), \dots, f_m(x)$ is the set of all polynomials $g(x)$ that can be expressed as a sum*

$$g(x) = \sum_{i=1}^m g_i(x)f_i(x), \quad (11.3)$$

where $g_1(x), \dots, g_m(x)$ are polynomials in $K[x_1, \dots, x_k]$.

Multiplication of monomials is just achieved by adding exponents:

$$x^\alpha x^\beta = x^{\alpha+\beta}.$$

The set of all monomials in a monomial ideal is the union of all positive orthants whose corners are given by the exponent vectors of the generating monomials $f_1(x), \dots, f_m(x)$.

Example 11.5

Consider the monomial ideal $I \subset \mathbb{Q}[x, y]$ generated by monomials xy and y^3 , that is, $I = \langle xy, y^3 \rangle$. All monomial terms such as x^2y, x^3y, x^3y^3 or y^4 belong to I as they can be divided by at least one of xy or y^3 . Note that besides monomials, the monomial ideal I also contains polynomials, for example, $4xy + y^3$ and $x^2y + xy^3 - xy^2 + y^3$ are both in I . Note that this second polynomial equals $(1+x)y^3 + (x-y)xy$. Polynomials such as $x - xy, y - 1$ or $y^4 - 1$ do not belong to I as they cannot be written as in (11.3).

For a given monomial ideal, a complete degree-by-degree description of the monomials inside the ideal or, equivalently, those outside the ideal is given by the Hilbert function

and series. In an important paper, Hilbert studied the distribution of monomials outside an ideal to determine the dimension of the variety associated with the ideal (see Hilbert (1890)). Here we only give the basic idea, referring the reader to references Cox et al. (1998, 2007) for a full description.

Definition 11.6 Let I be a monomial ideal in $K[x_1, \dots, x_k]$.

1. For a non-negative integer j , the Hilbert function $HF_I(j)$ is the number of monomials not in I of degree j .
2. The Hilbert series of I is the formal series $HS_I(s) = \sum_{j=0}^{\infty} s^j HF_I(j)$.

The Hilbert series is the generating function of the Hilbert function and it is used to count monomials which are not in the monomial ideal I . In what follows, unless it is required, we omit the subindex referring to the monomial ideal.

Example 11.6

Consider the monomial ideal $I = \langle x^3, xy^2, y^4 \rangle \subset K[x, y]$. The monomials which do not belong to I are $1, x, x^2, y, xy, x^2y, y^2$ and y^3 , so the Hilbert function equals $1, 2, 3, 2$ for $j = 0, 1, 2, 3$ and zero for all $j \geq 4$. The Hilbert series is thus $HS_I(s) = 1 + 2s + 3s^2 + 2s^3$. The Hilbert function and Hilbert series are defined in an opposite manner in CoCoA in that we need to ask for the Hilbert function of the quotient $K[x_1, \dots, x_k]/I$. The following commands will produce the results named earlier:

```
Use T:=Q[x,y];
I:=Ideal(x^3,x*y^2,y^4);
Hilbert(T/I);
HilbertSeries(T/I);
```

The Hilbert function for monomials inside I is thus obtained by the command `Hilbert(I)`.

Example 11.7

For the monomial ideal I of Example 11.5, the Hilbert function equals $1, 2, 2$ for $j = 0, 1, 2$, corresponding to $1, x, y, x^2, y^2$, and takes value one for $j \geq 3$ which corresponds to x^j for $j \geq 3$. Therefore, the Hilbert series is

$$HS_I(s) = 1 + 2s + 2s^2 + \sum_{j=3}^{\infty} s^j = \frac{1 + s - s^3}{1 - s}.$$

As this example shows, the Hilbert series can have an infinite number of non-zero terms. Later in this chapter, we will work with monomial ideals stemming from zero-dimensional polynomial ideals which in practical terms implies that the Hilbert series will be finite. The total number of monomials in the first orthant is counted with the formal series

$$\sum_{j=0}^{\infty} \binom{j+k-1}{j} s^j = \frac{1}{(1-s)^k},$$

where k is the number of variables. Then, by subtracting the Hilbert series $HS_I(s)$ from the last expression, we have a generating function to count monomials inside I .

Example 11.8

Using $k = 2$ dimensions to continue Example 11.6, the formal series for monomials in the first orthant is

$$\sum_{j=0}^{\infty} \binom{j+1}{j} s^j = \sum_{j=0}^{\infty} (j+1) s^j = \frac{1}{(1-s)^2},$$

which counts the monomials in the first quadrant. The generating function for the number of monomials in I for each degree is found by subtraction:

$$\frac{1}{(1-s)^2} - (1 + 2s + 3s^2 + 2s^3) = \frac{2s^3 + s^4 - 2s^5}{(1-s)^2}.$$

The signs in the polynomial in the numerator are related to inclusion–exclusion rules. Those rules in monomial ideals have been applied to compute bounds for reliability of systems (see Sáenz-de Cabezón and Wynn 2009, 2011).

Example 11.9

Consider the monomial ideal in $K[x_1, \dots, x_7]$ generated by monomials x_1^2, \dots, x_7^2 and all pairs $x_i x_j, 1 \leq i < j \leq 7$. This ideal has Hilbert function with values 1 and 7 for $j = 0, 1$ and zero for $j \geq 2$ so its Hilbert series is $HS(s) = 1 + 7s$, that is, one monomial of degree zero and seven monomials of degree one outside the ideal. The monomials outside this ideal are $1, x_1, x_2, x_3, x_4, x_5, x_6, x_7$, which later will be understood as a model for the Plackett–Burman design in Examples 11.19 and 11.30 and is also covered in Chapter 9. See also first row of Table 11.7 in Section 11.8.

11.3.6 Gröbner Bases

Dickson’s Lemma (Cox et al. 2007) states that, even if we define a monomial ideal with an infinite set of monomials f_i , we can find a finite set of monomials h_1, \dots, h_l such that $I = \langle h_1, \dots, h_l \rangle$. But there are, in general, many ways to express an ideal I as being generated from a basis $I = \langle g_1, \dots, g_m \rangle$. We define now a special basis for an ideal called the Gröbner basis (see Buchberger 2006; Cox et al. 2007).

Definition 11.7 Given an ideal $I \subset K[x_1, \dots, x_k]$, a set of polynomials $\{g_1, \dots, g_m\} \subset I$ is called a Gröbner basis of I if

$$\langle LT(g_1), \dots, LT(g_m) \rangle = \langle LT(I) \rangle,$$

where $\langle LT(I) \rangle$ is the ideal generated by the leading terms of all the polynomials in I , that is, $\langle LT(I) \rangle := \langle LT(f) : f \in I \rangle$.

We refer to $\langle LT(I) \rangle$ as the *leading term ideal* of I . Although the aforementioned definition only requires specification of leading terms and thus not term orders, in order to compute a Gröbner basis, we do need to specify a term ordering.

Example 11.10

Consider the ideal $I \subset \mathbb{Q}[x_1, x_2]$ whose variety is the set of four points in \mathbb{Q}^2 with coordinates $(\pm 1, \pm 1)$. The ideal I is the set of all polynomials that vanish on those four points. For any term ordering, the set $\{x_1^2 - 1, x_2^2 - 1\}$ is a Gröbner basis for I . This computation can be performed using the following CoCoA commands:

```
Use T:=Q[x[1..2]];
I:=IdealOfPoints([[ -1, -1], [ -1, 1], [ 1, -1], [ 1, 1]]);
GBasis(I);
```

The last command retrieves the two polynomials $x_1^2 - 1, x_2^2 - 1$ which form the Gröbner basis of I . The list of points in the second line previously can be substituted by the following CoCoA instruction which generates all pairs of two numbers using values +1 and -1:

```
I:=IdealOfPoints(Tuples([-1,1],2));
```

The following result is from Cox et al. (2007).

Lemma 11.1 *Any ideal $I \subset K[x_1, \dots, x_k]$ has a Gröbner basis and any Gröbner basis in the ideal I is a basis of I .*

Note that Gröbner bases do not operate as minimal spanning sets in the sense of a linear basis of a vector space. Polynomials in the Gröbner basis are used to generate polynomials in the ideal and thus redundancy is allowed. We can augment a given Gröbner basis with the product of two elements of it, and this set will remain to be a Gröbner basis in the sense of Definition 11.7. As an example of the redundancy allowed in Gröbner bases, a Gröbner basis for the monomial ideal I of Example 11.6 is $\{x^3, xy^2, y^4\}$. The augmented set of monomials $\{x^3, xy^2, x^3y^4, y^4\}$ is also a Gröbner basis for this ideal I . We achieve uniqueness in the basis when computing a reduced Gröbner basis, which is defined next.

Definition 11.8 *Let $<$ be a term order. A reduced Gröbner basis of an ideal I is a Gröbner basis $G_{<}$ such that*

1. *The coefficient of $LT_{<}(g)$ is one for all $g \in G_{<}$*
2. *For all $g \in G_{<}$, no monomial of g lies in $\langle LT(f) : f \in G_{<} \setminus g \rangle$*

Given an ideal $I \subset K[x_1, \dots, x_k]$ and a term ordering $<$, we can compute the reduced Gröbner basis of I . If this set of polynomials is g_1, \dots, g_m , then by Lemma 11.1, this Gröbner basis generates I , that is, $I = \langle g_1, \dots, g_m \rangle$. Now consider a polynomial $f(x) \in K[x_1, \dots, x_k]$ and the same term ordering $<$. The polynomial $f(x)$ has a unique remainder, $r(x)$ with respect the quotient operation $K[x_1, \dots, x_k]/I$, that is,

$$f(x) = \sum_{i=1}^m s_i(x)g_i(x) + r(x). \quad (11.4)$$

We call the remainder $r(x)$ the *normal form* of $f(x)$ with respect to I and write $r(x) = NF_I(f)$. Or, to stress the fact that it may depend on \prec , we write $NF_I(f, \prec)$.

The division of a polynomial in (11.4) is the generalisation of simple polynomial division such as that of Example 11.2, where the result is $s_1(x) = x^2 + 3$ with remainder $r = -5$. In other words, the normal form of $1 + 3x + 2x^2 + x^3$ with respect to the ideal generated by $g_1(x) = 2 + x$ is -5 . The CoCoA commands are

```
Use T := Q[x];
NF(1+3*x+2*x^2+x^3, Ideal(x+2));
```

The normal form can be equivalently determined using polynomial division directly:

```
DivAlg(1+3*x+2*x^2+x^3, [x+2]);
```

In general, however, the results of polynomial division depend on the order of operations performed (see, e.g., Cox et al. 2007, Chapter 2.3, Exercise 5). Division by Gröbner bases avoids this lack of uniqueness.

Subtraction of the normal form from a given polynomial will cause this difference to belong to the ideal. This is the basis of congruence between a polynomial and its normal form. The following lemma formalises this relation.

Lemma 11.2 *Given an ideal I and a monomial ordering \prec , for every $f(x) \in K[x_1, \dots, x_k]$, there is a unique normal form $NF(f)$ such that $f - NF(f) \in I$.*

In what follows, we collect the main facts relating (1) Grobner bases, (2) the division algorithm and (3) the nature of the normal form and the congruence generated by an ideal. The basic reference for this material is Cox et al. (2007).

1. Given an ideal, I , a monomial ordering \prec and a polynomial $f(x)$, there are algorithms which deliver the remainder $r(x)$ by successively dividing by the Gröbner basis terms g_1, \dots, g_m . The best known is the Buchberger algorithm (see Buchberger 2006; Cox et al. 2007).
2. The remainder of $f(x)$ divided by I takes the form $r(x) = NF_I(f) = \sum_{\alpha \in L} \theta_\alpha x^\alpha$. The set $\{x^\alpha, \alpha \in L\}$ is precisely the set of monomials *not* divisible by any of the leading terms of the Gröbner basis of I , namely, $\{LT(g_i), i = 1, \dots, m\}$.
3. The remainder $r(x) = NF(f)$ depends on the term ordering used, but does not depend on which order the Gröbner basis terms $g_i(x)$ are used in the division algorithm.
4. The set $\{x^\alpha, \alpha \in L\}$, which appears in remainder $r(x)$, is a basis of the quotient ring considered as a vector space of functions over $K[x_1, \dots, x_k]/I$. This set of monomials is also referred to as quotient basis. The terms are linearly independent over I , that is, the congruence

$$\sum_{\alpha \in L} \theta_\alpha x^\alpha \sim_I 0, \quad (11.5)$$

implies $\theta_\alpha = 0$ for all $\alpha \in L$ (see Cox et al. 2007, Chapter 5.3).

11.4 Experimental Design

We have indicated already that for applications to design, we should think of design as a list of points,

$$D = \{x^{(1)}, \dots, x^{(n)}\},$$

with each point $x^{(i)} \in K^k$ for $i = 1, \dots, n$. The design ideal is the set of all polynomials that vanish over the design points

$$I(D) = \{f(x) : f(x) = 0, x \in D\}.$$

In most cases, we use design points that have rational coordinates and thus set $K = \mathbb{Q}$. If designs are meant to have real coordinates, then the coefficient field should be set to $K = \mathbb{R}$. Even when most packages, for instance, CoCoA, work with rationals or modular integer fields, we can still define the real objects we want. For example, the variety associated with the ideal $\langle x_1, x_2^2 - 2 \rangle$ achieves the desired result of adding star points $(0, \pm\sqrt{2})$; see Example 11.13. See also Example 11.18 for a case of working in modular integer field.

The use of polynomials to define design is clearly not new. For example, a 2^k full factorial design with points $\{\pm 1, \dots, \pm 1\}$ is expressed as the solution of the simultaneous equations:

$$\{x_i^2 - 1 = 0, i = 1, \dots, k\}.$$

To obtain fractions we impose additional equations the points in the fraction must satisfy, for example, $x_1 x_2 \dots x_k = 1$.

Operations between varieties as geometric objects have a counterpart with ideals, which are algebraic objects. For example, the ideal of a union of points is the intersection of the ideals of single points; the ideal of an intersection of varieties is the *radical* of the sum of the ideals for individual varieties. These relations form the algebra-geometry dictionary (see Cox et al. 2007). A practical implication of such results is the direct generation of ideals for regular fractions of factorial designs (see Examples 11.16 through 11.18). A further reference for the generation of design ideals is Abbott et al. (2000).

Example 11.11

The generation of the design ideal for a given design can be achieved in CoCoA simply by using the command "IdealOfPoints" as is done in Example 11.10 for the design with coordinates $(\pm 1, \pm 1)$. This command implements the intersection of ideals corresponding to single points. For this same design $(\pm 1, \pm 1)$, the design ideal can also be written as the intersection of four ideals $I(D) = \langle x_1 + 1, x_2 + 1 \rangle \cap \langle x_1 + 1, x_2 - 1 \rangle \cap \langle x_1 - 1, x_2 + 1 \rangle \cap \langle x_1 - 1, x_2 - 1 \rangle$. The intersection gives precisely what we would intuitively expect from Example 11.10, that is, $I(D) = \langle x_1^2 - 1, x_2^2 - 1 \rangle$. The following CoCoA commands implement the intersection:

```
Use T:=Q[x[1..2]];
I:=Intersection(Ideal(x[1]+1,x[2]+1),Ideal(x[1]+1,x[2]-1),
Ideal(x[1]-1,x[2]+1),Ideal(x[1]-1,x[2]-1));
```

Alternatively, the design ideal can be specified by a system of equations whose solution give the design points. For the same example, noting that the design points are the intersection of the circle $x_1^2 + x_2^2 = 2$ and the cross created by union of lines $x_1 + x_2 = 0$ and $x_1 - x_2 = 0$ suggests the following design ideal, which can be verified to coincide with ideal $I(D)$ earlier:

$$\text{Ideal}(x[1]^2+x[2]^2-2)+\text{Ideal}((x[1]+x[2])*(x[1]-x[2]));$$

We now give what can loosely be described as the *algebraic method* in the title of this chapter. The goal is to identify estimable model(s) for a given design D and study confounding induced by the design. The estimable model is given as a list of polynomial terms.

1. Choose a design D of size n with no replications.
2. Select a monomial term ordering, \prec .
3. Build the design ideal $I(D)$ and compute the reduced Gröbner basis for $I(D)$ for the given monomial ordering \prec .
4. The quotient ring

$$K[x_1, \dots, x_k]/I(D)$$

of the ring of polynomials $K[x_1, \dots, x_k]$ in x_1, \dots, x_k is a vector space spanned by a special set of monomials: $\{x^\alpha, \alpha \in L\}$. These are all the monomials not divisible by the leading terms of the Gröbner basis. There are n such monomials, which is the size of the list of vectors L and the number of points in D .

5. The set of integer vectors L has the "order ideal" property: $\alpha \in L$ implies $\beta \in L$ for any $0 \leq \beta \leq \alpha$. For example, if exponent for $x_1^2 x_2$ is in L so are exponents for its divisors $1, x_1, x_2, x_1 x_2$ and x_1^2 .
6. Any function $y(x)$ on D has a unique polynomial interpolator given by

$$f(x) = \sum_{\alpha \in L} \theta_\alpha x^\alpha,$$

such that $y(x) = f(x), x \in D$. In other words, the polynomial model in the right-hand side previously is estimable for the given design D .

7. The design-model matrix X for design D and model with basis terms $\{x^\alpha : \alpha \in L\}$ has rank n . This matrix is of size $n \times n$ with rows indexed by the design points in D and columns indexed by terms in the basis:

$$X = \{x^\alpha\}_{x \in D, \alpha \in L}. \tag{11.6}$$

For further explanation of the relation between the representation of linear independence as congruence in (11.5) and as linear independence of columns in the design-model matrix (11.6), see Babson et al. (2003) and Cox et al. (2007).

When doing computations for items 1–3 previously, there are several possibilities. If design points are known and have rational coordinates, the CoCoA function

IdealOfPoints can be used to generate the design ideal $I(D)$ from the point coordinates. Alternatively, $I(D)$ can be built with the function `Ideal` and the system of equations whose solution gives the design points. Ideal operations can be used as well to construct $I(D)$, as will be described in examples as follows. The support of the model in items 4–5 is given by the command `QuotientBasis`. Interpolation of a function in item 6 is not necessarily required, but if so, the `CoCoA` command `Interpolate` can be used. The design-model matrix in item 7 is not explicitly computed as part of `CoCoA` analysis, but it is a consequence of the methodology developed. The following simple example intends to present the application of the algebraic method in a simple example.

Example 11.12

Consider the ideal I generated by polynomials $x_1^2 - 1$ and $x_2^2 - 1$ of Example 11.10. This is the design ideal for a factorial design 2^2 with point coordinates $(\pm 1, \pm 1)$, that is, $I = I(D) = \langle x_1^2 - 1, x_2^2 - 1 \rangle$. For any term ordering \prec , the set of polynomials $\{x_1^2 - 1, x_2^2 - 1\}$ is a reduced Gröbner basis for I ; see `CoCoA` commands in Example 11.12. The leading term ideal for this design is $\langle LT(I) \rangle = \langle x_1^2, x_2^2 \rangle$, so the set of monomials that do not belong to the ideal of leading terms is $\{1, x_1, x_2, x_1x_2\}$. These monomials cannot be divided by the leading terms of the Gröbner basis of I and are obtained with the command

```
QuotientBasis(I);
```

See Figure 11.1 for a depiction of exponents of monomials in the model $\{1, x_1, x_2, x_1x_2\}$ and monomials in the ideal of leading terms $\langle LT(I) \rangle$. As stated previously, the set of monomials in the model is linearly independent modulo I , and this property translates in the fact that the design-model matrix X for points $(\pm 1, \pm 1)$ and monomial columns $\{1, x_1, x_2, x_1x_2\}$ is full rank. We give the explicit version of this well-known design-model matrix as follows:

Points	Monomials
	1 x_1 x_2 x_1x_2
$(-1, -1)$	(1 -1 -1 1)
$(-1, 1)$	(1 -1 1 -1)
$(1, -1)$	(1 1 -1 -1)
$(1, 1)$	(1 1 1 1)




FIGURE 11.1

Staircase diagrams for (a) Example 11.12 and (b) Example 11.13. Black dots represent exponents of model terms and gray dots are exponents of monomials in the ideal of leading terms.

The implications of the method are considerable. But at its most basic it says that *we can always find a saturated polynomial $f(x)$ interpolating data over an arbitrary design D* . Indeed, the monomials $\{x^\alpha : \alpha \in L\}$ form a vector basis for K -valued functions defined over D , and thus the method retrieves polynomial functions which form a very flexible class (see Cox et al. 2007).

For example, suppose the vector \mathbf{y} has hypothetical values $\mathbf{y}' = (2, 2.8, -0.5, 3.3)$, where the values are listed in the same order as design points D of Example 11.12. Those values could be observations at design points in D or values from a function $y(x)$ taken at design points. For both cases, the objective is to produce an interpolator with a simple hierarchical structure. Following the basis obtained in the same example, we build the saturated model to fit data values:

$$y(x_1, x_2) = \theta_{00} + \theta_{10}x_1 + \theta_{01}x_2 + \theta_{11}x_1x_2.$$

This model has no degrees of freedom left for error and we obtain the matrix system $\mathbf{y} = \mathbf{X}\theta$, where \mathbf{y} is the vector of hypothetical response values, \mathbf{X} is the design-model matrix of Example 11.12 and $\theta = (\theta_{00}, \theta_{10}, \theta_{01}, \theta_{11})'$. Solving the system $\theta = \mathbf{X}^{-1}\mathbf{y}$ yields the fitted model $\hat{y}(x_1, x_2) = 1.9 - 0.5x_1 + 1.15x_2 + 0.75x_1x_2$ that agrees with the response values over the design. Recall that the main interest is to study potential estimable models for a design rather than suggesting the use of interpolator models in data analysis.

The structure of the model index set L arising from the order ideal property is important. This property of item 5 earlier gives exactly the shape which in statistical literature has been called variously: *staircase models*, *hierarchical models*, *well-formulated models*, or *marginality condition* (see Nelder 1977; Peixoto 1990). It can be seen easily from the fact that the multi-index terms given by L are the complement in the non-negative integer orthant of those given by the monomials in the monomial ideal of leading terms: the complement of a union of orthants has the staircase property. Figure 11.1 depicts two bidimensional examples with exponents of monomials in the model and in the complement of the model. Figures 11.2 and 11.3 also exhibit the staircase property in three dimensions.

11.5 Examples

An estimable model for a simple factorial design 2^2 was already identified in Example 11.12. In what follows, we give survey of applications of the algebraic method for different designs. We start in Section 11.5.1 with response surface designs (see also Chapter 5), then cover regular and non-regular fractional factorial designs in Section 11.5.2 (see also Chapters 7 and 9 for background). In Section 11.5.3, we study other design structures such as block designs (see also Chapter 3) and latin hypercube sampling (see Chapter 17). In the first example, we highlight the steps of the algebraic method giving full CoCoA commands, while in other examples, we concentrate on description of models identified.

11.5.1 Response Surface Designs

Example 11.13

Central composite design. Consider the central composite design of Example 11.1. This design is the union of four factorial points, four axial points and the origin; therefore,




FIGURE 11.2
Algebraic fan of design L_1 of Example 11.23.

its design ideal is the intersection of ideals (cf. Example 11.11) corresponding to each subdesign:

$$I(D) = \langle x_1^2 - 1, x_2^2 - 1 \rangle \cap \langle x_1^2 - 2, x_2 \rangle \cap \langle x_1, x_2^2 - 2 \rangle \cap \langle x_1, x_2 \rangle.$$

For example, the variety associated with the ideal $\langle x_1^2 - 2, x_2 \rangle$ is the simultaneous solution of $x_1^2 - 2 = 0$ and $x_2 = 0$ which yields the axial points $(\pm\sqrt{2}, 0)$. The CoCoA commands are

```
Use T:=Q[x[1..2]];
I:=Intersection(Ideal(x[1]^2-2,x[2]), Ideal(x[1],x[2]^2-2)); -- axial
I:=Intersection(I,Ideal(x[1]^2-1,x[2]^2-1)); -- factorial
I:=Intersection(I,Ideal(x[1],x[2])); -- origin
I;
GBasis(I);
```




FIGURE 11.3
Algebraic fan of design L_2 of Example 11.23.

```
LT(I);
QuotientBasis(I);
```

Leaving term ordering unspecified in the first line earlier automatically selects degree reverse lexicographic. The reduced Gröbner basis is $\{x_1^3 + x_1x_2^2 - 2x_1, x_1^2x_2 + x_2^3 - 2x_2, x_1x_2^3 - x_1x_2, x_2^5 - 3x_2^3 + 2x_2\}$ with leading terms $x_1^3, x_1^2x_2, x_1x_2^3, x_2^5$ so that the ideal of leading terms is $\langle x_1^3, x_1^2x_2, x_1x_2^3, x_2^5 \rangle$. The monomials which cannot be divided by the leading terms are $1, x_1, x_2, x_1^2, x_1x_2, x_2^2, x_2^3, x_2^4, x_1x_2^2$. They are precisely one of the bases detected in Example 11.1. The model cannot contain x_1^3 under this ordering. However, by rearranging the first polynomial in the Gröbner basis, we obtain the equality $x_1^3 = -x_1x_2^2 + 2x_1$, which holds over the design and the terms $x_1x_2^2$ and x_1 are included in the model. This type of aliasing or confounding between terms will be further covered in Section 11.6.

The ideal of leading terms can be used to count distribution of terms by total degree, as counted by the CoCoA command “`Hilbert(T/LT(I));`” which retrieves the Hilbert function of the complement of $LT(I)$ in the first orthant. It takes values 1, 2, 3, 2, 1 for $j = 0, 1, 2, 3, 4$ and zero for $j \geq 5$. All the aforementioned results depend on the assumed term ordering. In Section 11.8, we study results under different term orderings.

Example 11.14

Screening designs. A class of designs for main effect estimation while simultaneously avoiding biases caused by the presence of second-order effects and avoiding confounding of any pair of second-order effects was recently proposed by Jones and Nachtshiem (2011). The authors produced designs of size $n = 2k + 1$ for different dimensions ranging from $k = 4$ up to $k = 30$, and their construction is based on foldover (see Chapter 7) of a certain small fraction of size k of a 3^k design with levels $-1, 0, 1$ and then adding the center point. Naturally, after foldover and adding the origin, the screening design still remains a special fraction of a 3^k design. Points to be folded over can be related to the rows of a matrix called a *conference matrix*, (see Xiao et al. 2012).

Here we consider the designs for $k = 4, 7$ and 10. For $k = 4$, points in the screening design are creating by the foldover of the matrix

$$\begin{pmatrix} 0 & 1 & -1 & -1 \\ -1 & 0 & -1 & 1 \\ -1 & -1 & 0 & -1 \\ -1 & 1 & 1 & 0 \end{pmatrix}$$

together with center point (0,0,0,0) for a total of 9 points. For a lexicographic term ordering, we retrieve the model $1, x_3, x_4, x_4^2, x_3x_4, x_3x_4^2, x_3^2, x_3^2x_4, x_3^2x_4^2$, while with a degree reverse lexicographic term ordering, we retrieve the model $1, x_1, x_2, x_3, x_4, x_3^2, x_3x_4, x_4^2, x_2x_4$. Note the flexibility of the design which can identify either a full product model in two variables or a more balanced model with all four variables.

If we instead set a term order which eliminates x_1 and then orders x_2, x_3, x_4 using a graded order, we retrieve the model $1, x_2, x_3, x_4, x_2x_4x_3x_4, x_3^2, x_4^2, x_3x_4^2$. This is achieved in CoCoA using the following ring definition:

```
W:=Mat([[1,0,0,0],[0,1,1,1],[0,1,0,0],[0,0,1,0]]);
Use T:=Q[x[1..4]], Ord(W);
```

For $k = 7$, the design is obtained by adding the center point to a foldover of the points shown in the left panel of Table 11.2 to give a total of 15 points. Under the usual degree reverse lexicographic ordering in CoCoA, we retrieve the estimable model with terms $1, x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_5^2, x_6^2, x_7^2, x_2x_7, x_3x_7, x_4x_7, x_5x_7$ and x_6x_7 . We note that use of a graded order allows for the inclusion of all terms of degree one before the addition of terms of second degree, and the total degree of this model (addition of all degrees of exponents in the model) is 21. If a degree lexicographic order is used, the model remains with the same total degree, but it interchanges one interaction for a quadratic term $1, x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_5^2, x_6^2, x_7^2, x_5x_6, x_4x_7, x_5x_7, x_6x_7$. Lexicographic term orderings work in rather the opposite manner than graded orderings. For a lexicographic ordering, all terms with x_7 are included first. As this cannot go further than $1, x_7, x_7^2$ because $x_7^3 = x_7$, then term inclusion adds terms and interactions with x_6 . Once these are exhausted, terms involving x_5 are included. This process ends with a model with terms $1, x_7, x_7^2, x_6, x_6x_7, x_6x_7^2, x_6^2, x_6^2x_7, x_6^2x_7^2, x_5, x_5x_7, x_5x_6, x_5x_6x_7, x_5^2, x_5^2x_7$ and total degree 31. Looking at all possible term orderings is in general a complex and expensive task. In Section 11.8, we discuss and comment on the whole set of models identified by the four and seven factor designs, when considering all possible term orderings.

TABLE 11.2

Points to Be Folded Over for Two Screening Designs in Example 11.14

x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_1	x_2	x_3	x_4	x_5	x_6	x_7	x_8	x_9	x_{10}
0	1	-1	1	-1	1	-1	0	1	1	1	1	1	1	1	1	1
-1	0	1	-1	1	1	-1	1	0	-1	-1	-1	-1	1	1	1	1
1	-1	0	1	1	1	1	1	-1	0	-1	1	1	-1	-1	1	1
1	-1	-1	0	1	-1	-1	1	-1	-1	0	1	1	1	1	-1	-1
-1	-1	1	1	0	-1	-1	1	-1	1	1	0	-1	-1	1	-1	1
-1	1	-1	1	1	0	1	1	-1	1	1	-1	0	1	-1	1	-1
1	1	1	1	1	-1	0	1	1	-1	1	-1	1	0	-1	-1	1
							1	1	-1	1	1	-1	-1	0	1	-1
							1	1	1	-1	-1	1	-1	1	0	-1
							1	1	1	-1	1	-1	1	-1	-1	0

For $k=10$, the points to be folded are given in the right panel of Table 11.2. The standard term ordering in CoCoA was used to identify with this design a model of total degree 30 with constant, all ten linear terms x_1, \dots, x_{10} , two quadratic terms x_9^2, x_{10}^2 and eight two-factor interactions between x_{10} and one of x_2, \dots, x_9 . A lexicographic term ordering produces a model of much higher total degree (44) which contains monomials in four variables only (x_7, x_8, x_9 and x_{10}): $1, x_7, x_8, x_9, x_{10}, x_8^2, x_9^2, x_{10}^2, x_7x_8, x_7x_9, x_7x_{10}, x_8x_9, x_8x_{10}, x_9x_{10}, x_7x_8x_{10}, x_7x_9x_{10}, x_8x_9x_{10}, x_8^2x_{10}, x_9x_{10}^2, x_9^2x_{10}$ and $x_9^2x_{10}^2$. The distribution of terms by degree in this case is counted by the Hilbert function of the leading term ideal which equals 1, 4, 9, 6, 1 for $j = 0, 1, 2, 3, 4$ and zero for $j \geq 5$.

Example 11.15

Response surface design, non-standard. We take a 16-point design which is a 5^2 factorial with variable levels $-2, -1, 0, 1, 2$ with all internal points (forming a 3^2 design with levels $-1, 0, 1$) removed:

(2, 0)	(2, 1)	(2, 2)	(1, 2)	(0, 2)	(-1, 2)	(-2, 2)	(-2, 1)
(-2, 0)	(-2, -1)	(-2, -2)	(-1, -2)	(0, -2)	(1, -2)	(2, -2)	(2, -1)

If I_1 is the ideal of the full 5^2 factorial and I_2 is the ideal of the 3^2 factorial, then the design ideal corresponding to removal of internal points is given by the ideal quotient $I_1 : I_2$ (see the algebra–geometry dictionary in Cox et al. 2007). This ideal quotient is generated by the polynomials $x_2^5 - 5x_2^3 + 4x_2, x_1^5 - 5x_1^3 + 4x_1$ and $x_1^2x_2^2 - 4x_1^2 - 4x_2^2 + 16$. It can be shown that, for any term ordering, these polynomials form a reduced Gröbner basis, and thus the design identifies a single model with terms $1, x_2, x_2^2, x_2^3, x_2^4, x_1, x_1x_2, x_1x_2^2, x_1x_2^3, x_1x_2^4, x_1^2, x_1^2x_2, x_1^3, x_1^3x_2, x_1^4, x_1^4x_2$. This is an example of a design which is a complement of points with respect to a factorial grid (see Maruri-Aguilar et al. 2013).

11.5.2 Two-Level Designs

Example 11.16

Regular fraction. Let us take a 2^{6-2} fractional factorial design in six variables with resolution III (all main effects estimated independently of interaction) (see Chapters 1 and 7). In classical notation, this has defining contrasts: $\{ABCD, CDEF\}$. We select the fraction with $ABCD = CDEF = 1$ as generators of the defining contrast subgroup, and instead of A, \dots, F , we use variables x_1, \dots, x_6 . The design ideal is

$$I(D) = \langle x_1^2 - 1, x_2^2 - 1, x_3^2 - 1, x_4^2 - 1, x_5^2 - 1, x_6^2 - 1 \rangle + \langle x_1x_2x_3x_4 - 1, x_3x_4x_5x_6 - 1 \rangle.$$

This is another instance of using the algebra–geometry dictionary in Cox et al. (2007). The variety corresponding to the left ideal previously is the full factorial design 2^6 with 64 points. Adding the left ideal previously corresponds to the geometric intersection of varieties, with the effect that we only keep those points that satisfy the two generators. The design ideal is created in the following CoCoA code as the sum of the ideal defining the full factorial design and the ideal with defining equations of the desired fraction:

```
Use T:=Q[x[1..6]];
I:=Ideal([A^2-1|A In Indets()])
+Ideal(x[1]*x[2]*x[3]*x[4]-1, x[3]*x[4]*x[5]*x[6]-1);
```


TABLE 11.3

Aliasing Table for Example 11.16

1	$x_1x_2x_3x_4$	$x_3x_4x_5x_6$	$x_1x_2x_5x_6$
x_1	$x_2x_3x_4$	$x_1x_3x_4x_5x_6$	$x_2x_5x_6$
x_2	$x_1x_3x_4$	$x_2x_3x_4x_5x_6$	$x_1x_5x_6$
x_3	$x_1x_2x_4$	$x_4x_5x_6$	$x_1x_2x_3x_5x_6$
x_4	$x_1x_2x_3$	$x_3x_5x_6$	$x_1x_2x_4x_5x_6$
x_5	$x_1x_2x_3x_4x_5$	$x_3x_4x_6$	$x_1x_2x_6$
x_6	$x_1x_2x_3x_4x_6$	$x_3x_4x_5$	$x_1x_2x_5$
x_1x_4	x_2x_3	$x_1x_3x_5x_6$	$x_2x_4x_5x_6$
x_1x_6	$x_2x_3x_4x_6$	$x_1x_3x_4x_5$	x_2x_5
x_2x_4	x_1x_3	$x_2x_3x_5x_6$	$x_1x_4x_5x_6$
x_2x_6	$x_1x_3x_4x_6$	$x_2x_3x_4x_5$	x_1x_5
x_3x_6	$x_1x_2x_4x_6$	x_4x_5	$x_1x_2x_3x_5$
x_4x_6	$x_1x_2x_3x_6$	x_3x_5	$x_1x_2x_4x_5$
x_5x_6	$x_1x_2x_3x_4x_5x_6$	x_3x_4	x_1x_2
$x_1x_4x_6$	$x_2x_3x_6$	$x_1x_3x_5$	$x_2x_4x_5$
$x_2x_4x_6$	$x_1x_3x_6$	$x_2x_3x_5$	$x_1x_4x_5$

The CoCoA command `QuotientBasis(I);` gives the basis of the quotient ring:

```
[1, x[6], x[5], x[5]x[6], x[4], x[4]x[6], x[3],
  x[3]x[6], x[2], x[2]x[6],
  x[2]x[4], x[2]x[4]x[6], x[1], x[1]x[6], x[1]x[4], x[1]x[4]x[6]]
```

If the aliasing relation is desired for a given monomial, this is computed using the normal form. For example, `NF(x[2]*x[3]*x[6], I);` with output `x[1]x[4]x[6]` shows that over the design, the term $x_2x_3x_6$ is aliased with $x_1x_4x_6$, equivalently $x_2x_3x_6 - x_1x_4x_6 \in I(D)$, and thus both terms appear in the same row of the aliasing (Table 11.3). The aliasing table is read row-wise, for example, the first row implies that over the design $1 = x_1x_2x_3x_4 = x_3x_4x_5x_6 = x_1x_2x_5x_6$. Note that the first column of Table 11.3 contains the monomials in the quotient basis computed earlier and that the row containing the monomial 1 has the generators of the defining contrast subgroup.

For regular fractions like this case, the effect of different term orderings in the model means selecting (possibly) a different representative per row of the aliasing table. If the command `"Use T:=Q[x[1..6]], Lex;"` is used instead of the first line in the previous CoCoA code, the model is now selected using a lexicographic term ordering. Ten terms $1, x_2, x_4, x_5, x_6, x_2x_4, x_2x_6, x_4x_6, x_5x_6, x_2x_4x_6$ of the model coincide with the model identified earlier, and six terms $x_1, x_3, x_1x_4, x_1x_6, x_3x_6, x_1x_4x_6$ are replaced by $x_2x_5x_6, x_4x_5x_6, x_2x_4x_5x_6, x_2x_5, x_4x_5, x_2x_4x_5$. In each case, the replacement monomial is taken from the same row.

Example 11.17

Regular fraction. Consider the 2^{9-4} regular fraction of Chapter 7, Section 7.2.3. This fraction is defined by generating words **1236**, **1247**, **1258** and **13459**. Sixteen fractions are possible depending on the signs allocated to the generating words. Here we generate one possible fraction by adding $\langle x_1x_2x_3x_6 - 1, x_1x_2x_4x_7 - 1, x_1x_2x_5x_8 - 1, x_1x_3x_4x_5x_9 - 1 \rangle$ to the ideal of the full factorial design with levels ± 1 . This is achieved by the following CoCoA commands:

TABLE 11.4
Hilbert Function for Regular Fraction 2^{9-4} of Example 11.17

Ordering	Degree					
	0	1	2	3	4	5
Lex	1	5	10	10	5	1
DegRevLex	1	9	21	1	0	0

```
Use T::=Q[x[1..9]],Lex;
I1:=Ideal([A^2-1|A In Indets()]); -- factorial
J:=Ideal(x[1]*x[2]*x[3]*x[6]-1,
         x[1]*x[2]*x[4]*x[7]-1,
         x[1]*x[2]*x[5]*x[8]-1,
         x[1]*x[3]*x[4]*x[5]*x[9]-1); -- generator
I:=I1+J; -- design ideal
Hilbert(T/LT(I)); -- count of monomials in model
```

The Hilbert function for the lexicographic order selected is shown in the first row of Table 11.4. We observe that the model identified has one monomial term of total degree five. Due to the square-free nature of terms in a two-level design, this means the model has one-five factor interaction. If instead a degree reverse lexicographic order is selected, the model contains 21 two-factor interactions and only one three-factor interaction. This is shown in the second row of the same table.

Given that this design is a regular fraction, selecting models is equivalent to selecting representatives from each row in the alias table, according to the term ordering selected. For example, the term five-factor interaction $x_5x_6x_7x_8x_9$ that appeared under a Lex term ordering is replaced by term x_1x_8 using a degree reverse lexicographic ordering. All monomials whose normal form equals one coincide precisely with the entire defining contrast subgroup given in (7.8) of Chapter 7. The following are those elements:

$$\begin{array}{ccccccc}
 1 & x_4x_5x_7x_8 & x_3x_5x_6x_8 & & x_3x_4x_6x_7 & x_2x_6x_7x_8x_9 & x_2x_4x_5x_6x_9 \\
 x_2x_3x_5x_7x_9 & x_2x_3x_4x_8x_9 & x_1x_5x_6x_7x_9 & & x_1x_4x_6x_8x_9 & x_1x_3x_7x_8x_9 & x_1x_3x_4x_5x_9 \\
 x_1x_2x_5x_8 & x_1x_2x_4x_7 & x_1x_2x_3x_6 & x_1x_2x_3x_4x_5x_6x_7x_8 & & &
 \end{array}$$

These monomials are equal over the fraction and we say that they are aliased. As for every term ordering, the monomial 1 satisfies $1 < x^\alpha$ for $\alpha \neq 0$; then for the aforementioned elements, the representative obtained by the algebraic techniques will always be the monomial 1. In Section 11.6, we further discuss aliasing under the perspective of ideals.

Example 11.18

Consider the design $D = \{(0, 0, 1), (0, 1, 0), (1, 0, 0), (1, 1, 1)\}$. This design is a half fraction of a 2^3 design with factor levels 0, 1. For a degree reverse lexicographic ordering and working in the ring $\mathbb{Q}[x_1, x_2, x_3]$, its design ideal has reduced Gröbner basis $\{x_1^2 - x_1x_2^2 - x_2, x_3^2 - x_3, x_1x_2 - 0.5x_1 - 0.5x_2 - 0.5x_3 + 0.5, x_1x_3 - 0.5x_1 - 0.5x_2 - 0.5x_3 + 0.5, x_2x_3 - 0.5x_1 - 0.5x_2 - 0.5x_3 + 0.5\}$. The first three polynomials form the full factorial design, while the latter three select points and express the confounding relation between the three two-factor interactions. For example, the polynomial $x_1x_2 - 0.5x_1 - 0.5x_2 - 0.5x_3 + 0.5 = 0$ earlier leads into the confounding $x_1x_2 = 0.5x_1 + 0.5x_2 + 0.5x_3 - 0.5$. The model for this design has terms $\{1, x_1, x_2, x_3\}$.

Now consider the same design and term ordering, but now work over the ring $\mathbb{Z}_2[x_1, x_2, x_3]$, which implies that operations are performed with arithmetic modulo 2. In this setting the design is a regular fraction with generator $x_1 + x_2 + x_3 = 1 \pmod{2}$. The reduced Gröbner basis is $\{x_2^2 + x_2, x_3^2 + x_3, x_1 + x_2 + x_3 + 1\}$ and the model identified for this design is $\{1, x_3, x_2, x_2x_3\}$. The effect of modulo 2 arithmetic is evident in the first two polynomials of the Gröbner basis corresponding to the full factorial part of D . The polynomial $x_1^2 + x_1$ would complete the full factorial part of D but is redundant and thus it is not included in the Gröbner basis. The last polynomial in the basis is the expected generator which prevents the inclusion of term x_1 in the model. The following CoCoA commands perform the required operations:

```
Use T:=Z/(2)[x[1..3]];
I:=IdealOfPoints(Tuples([0,1],3));
J:=I+Ideal(Sum(Indets())+1);
GBasis(J);
QuotientBasis(J);
```

Example 11.19

Plackett–Burman, PB(8). Consider the Plackett–Burman design (Plackett and Burman 1946) with 8 points in $k=7$ dimensions generated by circular shifts of $(+1, +1, +1, -1, +1, -1, -1)$ together with the point $(-1, -1, -1, -1, -1, -1, -1)$. Using the standard ordering in CoCoA (degree reverse lexicographic), we retrieve the usual first order model for this design: $1, x_1, x_2, x_3, x_4, x_5, x_6, x_7$. If a lexicographic term ordering in which $x_1 > \dots > x_7$ is used, the model retrieved is a “slack” model in only four variables with terms $1, x_5, x_6, x_7, x_5x_7, x_5x_6, x_6x_7, x_5x_6x_7$.

11.5.3 Other Designs

Example 11.20

Græco-Latin square. It is a straightforward exercise to code up combinatorial designs using indicator variables that associate treatment (Greek, Latin) and plot (row, column) factors with vector spaces. Let us take as an example the 4×4 Graeco-Latin square derived via the standard Galois field method (see Bailey 2008). The square is

$A\alpha$	$B\beta$	$C\gamma$	$D\delta$
$B\gamma$	$A\delta$	$D\alpha$	$C\beta$
$C\delta$	$D\gamma$	$A\beta$	$B\alpha$
$D\beta$	$C\alpha$	$B\delta$	$A\gamma$

The first step is to code the design with indicator functions: $x_{ij} = 1$ for the j th level of factor i (and zero otherwise) for $i, j = 1, \dots, 4$. Here i indexes the factors: rows, columns, Latin and Greek letters, respectively. This results in an array of 16 rows (total of cells) and 16 columns (4 columns by each of factors row, column, Latin and Greek). The design points are shown as rows in Table 11.5, where vertical lines are only added for clarity. For example, the top-left cell in the Græco-Latin square has row and column equal to one and treatment factors $A\alpha$ so it corresponds to point $(1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0)$ in the first row of Table 11.5. The design ideal $I(D)$ is generated for the aforementioned point configuration. Using the degree lexicographic term ordering in CoCoA, the model identified for the design has terms

$$1, u_2, u_3, u_4, t_2, t_3, t_4, r_2, r_3, r_4, c_2, c_3, c_4, t_4u_2, t_4u_3, t_4u_4,$$

TABLE 11.5
Design Points for Graeco-Latin Design of Example 11.20

Row r				Column c				Factor u				Factor t			
1	2	3	4	1	2	3	4	A	B	C	D	α	β	γ	δ
1	0	0	0	1	0	0	0	1	0	0	0	1	0	0	0
1	0	0	0	0	1	0	0	0	1	0	0	0	1	0	0
1	0	0	0	0	0	1	0	0	0	1	0	0	0	1	0
1	0	0	0	0	0	0	1	0	0	0	1	0	0	0	1
0	1	0	0	1	0	0	0	0	1	0	0	0	0	1	0
0	1	0	0	0	1	0	0	1	0	0	0	0	0	0	1
0	1	0	0	0	0	1	0	0	0	0	1	1	0	0	0
0	1	0	0	0	0	0	1	0	0	1	0	0	1	0	0
0	0	1	0	1	0	0	0	0	0	1	0	0	0	0	1
0	0	1	0	0	1	0	0	0	0	0	1	0	0	1	0
0	0	1	0	0	0	1	0	1	0	0	0	0	1	0	0
0	0	1	0	0	0	0	1	0	1	0	0	1	0	0	0
0	0	0	1	1	0	0	0	0	0	0	1	0	1	0	0
0	0	0	1	0	1	0	0	0	0	1	0	1	0	0	0
0	0	0	1	0	0	1	0	0	1	0	0	0	0	0	1
0	0	0	1	0	0	0	1	1	0	0	0	0	0	1	0

where the variables u, t identify treatments (Latin and Greek letters) and the variables r, c identify rows and columns of the design (see heading of Table 11.5). Note in the following exhibit the neat decomposition of model terms that coincides with the standard analysis of variance for this orthogonal design (with $df =$ degrees of freedom), where the interaction between treatment factors (three terms involving t and u earlier) is often allocated to the residual:

Source	df
Mean	1
u (treatment factor 1)	3
t (treatment factor 2)	3
r (row factor)	3
c (column factor)	3
Interactions	3
Total	16

Example 11.21

Balanced incomplete block design (BIBD). Consider the following BIBD: $\{1, 2, 4\}, \{2, 3, 5\}, \{3, 4, 6\}, \{4, 5, 0\}, \{5, 6, 1\}, \{6, 0, 2\}$ and $\{0, 1, 3\}$, that is, the first block b_0 contains treatments t_1, t_2 and t_4 . This design has $n = 21$ runs, $t = 7$ treatments t_0, \dots, t_6 and $b = 7$ blocks b_0, \dots, b_6 . Using a degree reverse lexicographic ordering, we retrieve the following model terms:

$$1, t_1, t_2, t_3, t_4, t_5, t_6, b_1, b_2, b_3, b_4, b_5, b_6, t_1b_6, t_2b_5, t_3b_6, t_4b_2, t_5b_3, t_5b_4, t_6b_4, t_6b_5,$$

while for a lexicographic term ordering, we retrieve

$$1, t_0, t_1, t_2, t_3, t_4, t_5, b_0, b_1, b_2, b_3, b_4, b_5, t_0b_3, t_1b_0, t_1b_4, t_2b_0, t_2b_1, t_3b_1, t_3b_2, t_4b_2.$$

In both cases, we retrieve $t - 1 = 6$ treatment terms and $b - 1 = 6$ block terms while having 8 interaction terms.

Example 11.22

Incomplete block design (Chapter 3). Consider an incomplete block design with $n = 12$ runs and $t = 6$ treatments t_1, \dots, t_6 arranged in $b = 6$ blocks of size two (t_i, t_j) for the following pairs (i, j) : $(1, 2), (2, 3), (3, 4), (4, 5), (5, 6), (1, 6)$. Using the standard term ordering in CoCoA gives the following model:

$$1, t_6, t_5, t_4, t_3, t_2, b_6, b_6t_6, b_5, b_4, b_3, b_2.$$

This model has $t - 1 = 5$ degrees of freedom for treatments, $b - 1 = 5$ for blocks plus the interaction b_6t_6 . A possibility for analysis would be to allocate the interaction b_6t_6 to the residual error, with only one degree of freedom. Under a lexicographic ordering, we retrieve the same model as the previous. This result is not extremely surprising given the highly restricted range of monomial terms for the model for this design.

Example 11.23

Latin hypercube sample (Chapter 17). Latin hypercubes (McKay et al. 1979) are widely used schemes in the design and analysis of computer experiments. The design region is often the hypercube $[0, 1]^k$ and designs of interest are often those that efficiently cover the design region. Latin hypercubes have at least two clear advantages over a random selection of points: univariate projections of the design are uniform and they are simple to generate.

The design L_1 with points $(0, 0, 0), (1/5, 1, 4/5), (2/5, 3/5, 2/5), (3/5, 4/5, 1/5), (4/5, 1/5, 1)$ and $(1, 2/5, 3/5)$ is an example of randomly generated latin hypercube in $k = 3$ dimensions and $n = 6$ runs. Under the standard term ordering in CoCoA, the design L_1 identifies the model $1, x_1, x_2, x_3, x_2x_3, x_3^2$. Experimentally, some latin hypercubes have been found to identify certain types of models which are of minimal degree called *corner cut models*, (see Onn and Sturmfels 1999 and also Berstein et al. 2010). The design L_1 belongs to such a class and will be discussed further in Section 11.8.

A second example of latin hypercube is L_2 with points $(0, 0, 4/5), (1/5, 1/5, 2/5), (2/5, 2/5, 1), (3/5, 3/5, 0), (4/5, 4/5, 3/5)$ and $(1, 1, 1/5)$. Under the same ordering as previous, L_2 identifies the model $1, x_2, x_3, x_2x_3, x_3^2, x_3^3$.

11.6 Understanding Aliasing

The algebraic method is not only a way of obtaining candidate models for a given design, but it does, we claim, deliver considerable understanding of the notion of aliasing. Aliasing is close to the idea of equivalence used in Section 11.3.6 to define the quotient operation and is not restricted to factorial designs. Let $I(D)$ be the design ideal and for two polynomials $f, g \in K[x_1, \dots, x_k]$, define

$$f(\mathbf{x}) \sim_D g(\mathbf{x}),$$

to mean the polynomials have the same value at the design points; that is,

$$f(x) = g(x), x \in D.$$

This is equivalent to $f(x) - g(x) \in I(D)$ (see Section 11.4). Equivalently, we have, with respect to any monomial ordering,

$$NF(f(x)) = NF(g(x)),$$

for $x \in D$ since by Lemma 11.2, $f(x) = NF(f(x))$ for $x \in D$ and similarly for $g(x)$. We call this *algebraic aliasing*. However, this is not quite the same as the statistical idea of aliasing that $f(x) = cg(x)$ over the design for some non-zero constant c , because scalar multiples can be absorbed into constants in a regression model. That is, both $f(x)$ and $g(x)$ should not both be in the same regression model.

We want to link this notion with vectors which span the same space. The values of a polynomial $f(x)$ on the design expressed as a (column) vector are defined as $\text{supp}_D(f(x)) := (f(x) : x \in D)'$. Then $f(x) \sim_D g(x)$ is equivalent to $\text{supp}_D(f(x)) = \text{supp}_D(g(x))$. For example, consider the design $(\pm 1, \pm 1, \pm 1)$ and the standard ordering of design points $(-1, -1, -1), (-1, -1, 1), \dots, (1, 1, 1)$, we have $\text{supp}(x_1x_2x_3) = (-1, 1, 1, -1, 1, -1, -1, 1)'$ that is, the column for $x_1x_2x_3$ which would appear in the X matrix.

Definition 11.9 *Collections of polynomials F and G are said to be statistically aliased if*

$$\text{span}\{\text{supp}(f), f \in F\} = \text{span}\{\text{supp}(g), g \in G\}. \tag{11.7}$$

This aliasing is written as $F \approx_D G$. Here the span of a set of vectors is the collection of all linear combinations of them with real coefficients. This is a standard vector space.

Example 11.24

Consider the design in $k = 2$ dimensions with points $(1, 1/2), (-1/2, 1), (-1, -1/2)$ and $(1/2, -1)$. Two models are $\{1, x_1, x_1^2, x_1^3\}$ and $\{1, x_1, x_2, x_1x_2\}$: the first was obtained with a lexicographic term order, and the second was not retrieved by algebra but only by checking that its design-model matrix is full rank. Notice that the first two terms in both models are the same. We now convert the design-model matrix for the first model to an orthogonal matrix using Gram–Schmidt orthogonalisation; see, for example, the package `far` in `R` (R Core Team 2014). The original matrix and the orthogonal matrix are given in Table 11.6, where the headings indicate (a) the monomials involved and (b) the linear transformations applied to achieve orthogonality. The orthogonal version of the first model is $\{1, x_1, x_1^2 - 5/8, x_1^3 - 17x_1/20\}$, while the second model is already orthogonal. Since the first two terms of the models are the same, the last two terms must span the same bidimensional subspace, so we can write

$$\{x_2, x_1x_2\} \approx_D \{x_1^2 - 5/8, x_1^3 - 17x_1/20\}.$$

We refine these statements by careful observation of the columns of the orthogonal version of the design-model matrix for the aforementioned first model and that for the second model (not shown). We see that $\{x_2\} \approx_D \{x_1^3 - 17x_1/20\}$ and $\{x_1x_2\} \approx_D \{x_1^2 - 5/8\}$. The orthogonality is needed: it is not true that $\{x_2, x_1x_2\} \approx_D \{x_1^2, x_1^3\}$.

TABLE 11.6

(a) Design-Model Matrix for Model $\{1, x_1, x_1^2, x_1^3\}$ and (b) Its Orthogonal Version

(a)				(b)			
1	x_1	x_1^2	x_1^3	1	x_1	$x_1^2 - 5/8$	$x_1^3 - 17x_1/20$
$\left(\begin{array}{cccc} 1 & 1 & 1 & 1 \\ 1 & -0.5 & 0.25 & -0.125 \\ 1 & -1 & 1 & -1 \\ 1 & 0.5 & 0.25 & 0.125 \end{array} \right)$				$\left(\begin{array}{cccc} 1 & 1 & 0.375 & 0.15 \\ 1 & -0.5 & -0.375 & 0.3 \\ 1 & -1 & 0.375 & -0.15 \\ 1 & 0.5 & -0.375 & -0.3 \end{array} \right)$			

Given that any polynomial and its normal form satisfy $f(x) = NF(f(x))$ for $x \in D$, we can rewrite (11.7) as

$$\text{span}\{\text{supp}(NF(f)), f \in F\} = \text{span}\{\text{supp}(NF(g)), g \in G\}.$$

This means that any aliasing statement between collections of polynomials is equivalent to one for the normal forms (Section 11.3.6). For $f \in F$, let

$$f = \sum_{\alpha \in L} \theta_{\alpha, f} x^\alpha,$$

where L is as defined in Section 11.4 and depends on the design D and the term ordering. Let θ_f be the vector of $\theta_{\alpha, f}$ and define θ_g similarly. Then since the design-model matrix X is non-singular, by construction, we have

$$F \approx_D G \Leftrightarrow \text{span}\{\theta_f, f \in F\} = \text{span}\{\theta_g, g \in G\}.$$

Thus, statistical aliasing can be thought of in two stages: (1) First reduce to expressing each polynomial in F and G to its normal form using the algebra, and then (2) compare the coefficient subspaces. In the regular factorial fraction case, the normal form of a monomial is itself a monomial, which makes the interpretation easier, but in the general case, it is a polynomial.

We can often find the alias classes by inspection, once we have the normal form. Consider Example 11.15 and the monomials $\{x_1^2 x_2^2, x_1^4 x_2^4, x_1^6 x_2^6, x_1^8 x_2^8\}$. The normal forms are, respectively,

$$\begin{aligned} &4x_1^2 + 4x_2^2 - 16, \quad 16x_1^4 + 16x_2^4 - 256, \\ &320x_1^4 + 320x_2^4 - 256x_1^2 - 256x_2^2 - 4096, \\ &5376x_1^4 + 5376x_2^4 - 5120x_1^2 - 5120x_2^2 - 65536. \end{aligned}$$

Because of linear transformations, the span of $1, x_1^2 + x_2^2, x_1^4 + x_2^4$ can be transformed to become the span of $1, x_1^2 x_2^2, x_1^4 x_2^4$. This is because over the design, the following two statements are true:

$$x_1^2 x_2^2 = 4x_1^2 + 4x_2^2 - 16 = 4(x_1^2 + x_2^2) - 16(1) \text{ and}$$

$$x_1^4 x_2^4 = 16x_1^4 + 16x_2^4 - 256 = 16(x_1^4 + x_2^4) - 256(1).$$

The statements are built using Lemma 11.11 in Section 11.3.6 and the normal forms of monomials previously. A similar operation can be performed to show that the span of monomials $1, x_1^6 x_2^6, x_1^8 x_2^8$ is also equivalent to the span of $1, x_1^2 + x_2^2, x_1^4 + x_2^4$. Therefore, we see that

$$\{1, x_1^2 x_2^2, x_1^4 x_2^4\} \approx_D \{1, x_1^6 x_2^6, x_1^8 x_2^8\}.$$

The equivalence continues to all $\{1, x_1^{2k} x_2^{2k}, x_1^{2(k+1)} x_2^{2(k+1)}\}$. To retain the link to classical notation such as that in Chapter 7, Section 7.10, we might say that the collection $\{I, A^2 B^2, A^4 B^4\}$ is aliased with the collection $\{I, A^6 B^6, A^8 B^8\}$, and we might write

$$\{I, A^2 B^2, A^4 B^4\} \approx_D \{I, A^6 B^6, A^8 B^8\}.$$

This arises because $A^2 B^2 \approx A^2 + B^2 - 4I$ and $A^4 B^4 \approx A^4 + B^4 - 16$, and both the reduced forms are estimable. In this example, odd terms also pair up. The normal forms of $\{x_1^3 x_2^3, x_1^5 x_2^5, x_1^7 x_2^7, x_1^9 x_2^9\}$ are, respectively, $4x_1^3 x_2 + 4x_1 x_2^3 - 16x_1 x_2$, $80x_1^3 x_2 + 80x_1 x_2^3 - 384x_1 x_2$, $1344x_1^3 x_2 + 1344x_1 x_2^3 - 6656x_1 x_2$ and $21760x_1^3 x_2 + 21760x_1 x_2^3 - 108544x_1 x_2$, so that $\{1, x_1^3 x_2^3, x_1^5 x_2^5\} \approx_D \{1, x_1^7 x_2^7, x_1^9 x_2^9\}$ and so on. In classical notation, $\{I, A^3 B^3, A^5 B^5\} \approx \{I, A^7 B^7, A^9 B^9\}$.

11.7 Indicator Functions and Orthogonality

At times it is convenient to see the design D as a subset of a full factorial design \mathcal{N} . This is most usual when we start with some basic design, such as a full factorial, and consider a fraction as in Example 11.16. In such case, an algebraic description of the fraction can be made via an indicator function F_D , rather than via a Gröbner basis. The design ideal of D is unique, but the objects that change are the generating equations we choose to describe $I(D)$. These encode different information on D . An indicator function is a *single* additional polynomial function which we add to the generators of the ideal of the full factorial design to form the ideal of D .

Example 11.25

Consider the design ideal generated by regular fractional factorial 2^{3-1} design, that is, $I(D) = \langle x_1^2 - 1, x_2^2 - 1, x_3^2 - 1, x_1 x_2 x_3 + 1 \rangle$. The first three terms form the Gröbner basis of the full factorial $\{(\pm 1, \pm 1, \pm 1)\}$. From the equation $x_1 x_2 x_3 + 1 = 0$, we can deduce the indicator functions of D in \mathcal{N} as follows. We know that $x_1 x_2 x_3 = -1$ over points in D . By the regularity of the fraction, we have

$$x_1 x_2 x_3 = \begin{cases} -1 & \text{for } x \in D \\ 1 & \text{for } x \in \mathcal{N} \setminus D \end{cases}.$$

A straightforward transformation of this is used to build the indicator

$$g(x_1, x_2, x_3) = \frac{1 - x_1x_2x_3}{2} = \begin{cases} 1 & \text{for } x \in D \\ 0 & \text{for } x \in \mathcal{N} \setminus D. \end{cases}$$

Then, on the full factorial,

$$x_1x_2x_3 + 1 = 0 \Leftrightarrow g(x_1, x_2, x_3) = 1.$$

More generally, let \mathcal{N} be the basic design which is not necessarily a full factorial design, and let $D \subset \mathcal{N}$ be a fraction. Sometimes \mathcal{N} is called a *candidate set*. Fix a monomial order and, via the $I(\mathcal{N})$, construct a vector space basis for interpolation over \mathcal{N} . Then the indicator function of D interpolates the 0, 1 values as required:

$$g(x) = \begin{cases} 1, & x \in D \\ 0, & x \in \mathcal{N} \setminus D. \end{cases}$$

In this example, since \mathcal{N} is a full factorial design, there is only one basis for interpolation, $\{x_1^{\alpha_1}x_2^{\alpha_2}x_3^{\alpha_3} : \alpha_i \in \{0, 1\} \text{ for } i = 1, 2, 3\}$, and the indicator involves only terms for $\alpha = (0, 0, 0)$ and $\alpha = (1, 1, 1)$. This last fact is a consequence that the design is a regular half fraction.

Example 11.26

The indicator function for the regular fraction 2^{6-2} of Example 11.16 is

$$g(x_1, x_2, x_3, x_4, x_5, x_6) = \frac{1 + x_3x_4x_5x_6 + x_1x_2x_5x_6 + x_1x_2x_3x_4}{4},$$

which takes the value one at points in the fraction and zero in the rest of the 2^6 design. Thus adding the ideal $\langle g(x_1, x_2, x_3, x_4, x_5, x_6) - 1 \rangle$ to the ideal of the full factorial yields the ideal of the same fraction as described in Example 11.16. Conversely, to retrieve the 48 points which are the complement of the regular fraction, we would add $\langle g(x_1, x_2, x_3, x_4, x_5, x_6) \rangle$ to the ideal of the full factorial design.

The coefficients of the indicator functions expressed over the interpolation basis embed information on the geometric/combinatoric properties of the fraction. We exemplify this as follows in the binary case where \mathcal{N} is the full factorial design 2^d with coding $\{-1, 1\}$ (see Fontana et al. 2000). For factors with mixed levels, coding with complex numbers can be used (see Pistone and Rogantin 2008). Two monomials x^α, x^β are said to be *orthogonal* over $D \subset \mathcal{N}$ if the corresponding columns in the design-model matrix X are orthogonal:

$$\sum_{x \in D} x^\alpha x^\beta = \sum_{x \in D} x^{\alpha+\beta} = 0.$$

We can express this in terms of the indicator function $g(x)$ for the design \mathcal{N} and write

$$\sum_{x \in \mathcal{N}} x^{\alpha+\beta} g(x) = 0$$

because $g(x) = 0$ for points $x \in \mathcal{N} \setminus D$ and $g(x) = 1$ over the design, that is, for $x \in D$.

Following Example 11.25, we want to check whether two-way-factor interactions are orthogonal or not to the factor missing from the interaction. Setting $\alpha = (1, 0, 0)$ and $\beta = (0, 1, 1)$, we are checking orthogonality between x_1 and x_2x_3 . We use the previous indicator $g(x_1, x_2, x_3)$ and obtain

$$\sum_{x \in \mathcal{N}} x_1 x_2 x_3 \frac{1 - x_1 x_2 x_3}{2} = \sum_{x \in \mathcal{N}} \frac{x_1 x_2 x_3 - 1}{2} = \sum_{x \in \mathcal{N}} \frac{1}{2} x_1 x_2 x_3 - 4 = -4 \neq 0,$$

thus showing that there is no orthogonality between x_1 and x_2x_3 . Note that the calculations and result would be the same if we checked instead orthogonality between x_2 and x_1x_3 or between x_3 and x_1x_2 . In the previous text, we used the fact that the full factorial design \mathcal{N} has 8 points. In the following CoCoA code, the normal form of $x_1x_2x_3$ is a constant which when summed over the four design points does not equal zero, that is, x_1 and x_2x_3 are non-orthogonal:

```
Use T::=Q[x[1..3]];
G:=(1-x[1]*x[2]*x[3])/2; -- indicator function
J:=Ideal([A^2-1|A In Indets()])
  + Ideal(G-1); -- ideal of fraction D
NF(x[1]*x[2]*x[3],J);
```

It is no coincidence of the aforementioned that the coefficient of $x_1x_2x_3$ in the indicator function is not zero. This idea generalises for designs over $\{-1, 1\}^k$, all non-zero square-free monomials sum to zero, so orthogonality holds if and only if the constant term is zero.

A practical advantage of the indicator function is that we can take union and intersections of design by considering $g(x)$ as a Boolean function over the basic design \mathcal{N} :

$$g_{D_1 \cap D_2} = g_{D_1} g_{D_2}, \quad g_{D_1 \cup D_2} = g_{D_1} + g_{D_2} - g_{D_1} g_{D_2}.$$

Again the zero coefficients of the normal form of $g_{D_1 \cap D_2}$ and $g_{D_1 \cup D_2}$ over the interpolation monomial basis of \mathcal{N} are informative of the geometry of the intersection and union designs. This is an alternative to operations with ideals as suggested from the algebra-geometry dictionary in Cox et al. (2007) (see also Examples 11.11, 11.13 and 11.16).

Example 11.27

Take \mathcal{N} to be the full factorial design 2^3 with levels ± 1 . The indicator function for the half fraction $D_1 \subset \mathcal{N}$ defined by the design ideal $I(D_1) = \langle x_1^2 - 1, x_2^2 - 1, x_3^2 - 1 \rangle + \langle x_1 x_2 x_3 - 1 \rangle$ is $g_{D_1} = (1 + x_1 x_2 x_3)/2$. Set $D_2 \subset \mathcal{N}$ to be the design with points $(-1, -1, -1)$ and $(1, 1, -1)$. The indicator function for D_2 is $g_{D_2} = (1 + x_1 x_2 - x_3 - x_1 x_2 x_3)/4$. We now build the indicator of the union $D_1 \cup D_2$. This design is a $3/4$ fraction of the full design \mathcal{N} , and after simplification, the indicator is

$$g_{D_1 \cup D_2} = \frac{3 + x_1 x_2 - x_3 + x_1 x_2 x_3}{4}.$$

Orthogonality of factorial effects can be checked as illustrated.

11.8 Fans, State Polytopes and Linear Aberration

The computations of Gröbner bases and model identification with Gröbner bases described in Sections 11.3 and 11.4 depend upon the term ordering selected. Recall that for a given term order $<$, the Gröbner basis algorithm produces a basis for the quotient ring $K[x_1, \dots, x_k]/I(D)$ which consists precisely of the smallest monomials under $<$ that are linearly independent under the congruence (11.5). This set of monomials forms the support of a saturated model we search for. For this reason, setting a particular term order $<$ allows the analyst to put preference over terms which will be identified by the model, for instance, a graded ordering will include as many terms of order one as possible in all variables before adding terms of second degree in the model. On the other end of the spectrum, a lexicographic order selects as many terms in the first variable in the selected order before adding terms in the second variable and so on. In other instances, the experimenter might be interested in exploring the range of all models identifiable by the design using algebraic techniques and different term orders.

This collection of models obtained by varying term orders is termed the *fan of the design* and allows assessment of design properties like estimation capacity in Chen and Cheng (2004) and Cheng and Mukerjee (1998) or the minimal linear aberration of the design in Berstein et al. (2010) and its general case of non-linear aberration (see Berstein et al. 2008). Fan computations have been applied, among others, to industrial experiments (see Holliday et al. 1999 and systems biology in Dimitrova et al. 2007).

11.8.1 Algebraic Fan of a Design

Given a design ideal $I(D)$ and ranging over all possible term orderings, we have a collection of reduced Gröbner bases for $I(D)$. A crucial fact is that despite the collection of different term orderings being an infinitely countable set (excluding the trivial case of one dimension), the collection of reduced Gröbner bases has always a finite number of distinct elements (see Mora and Robbiano 1988). Associated to this collection of Gröbner bases, there is a collection of polyhedral cones, called the Gröbner fan (see Mora and Robbiano 1988). We use the term *algebraic fan* of the design for the collection of different bases for the quotient ring $K[x_1, \dots, x_k]/I(D)$. In other words, the algebraic fan is a collection of saturated models.

For some relatively simple designs, such as factorial designs, the algebraic fan has only a single model. The general class of designs with a single model is called the class of *echelon* designs (see Pistone et al. 2001). However, in general, computing the algebraic fan of a design is an expensive computation. Reverse search techniques are at the core of state-of-the-art software *gfan* (Jensen 2009), while other approaches remain under investigation, such as the polynomial-time approach based on partial orderings, matrix operations and zonotopes (see Babson et al. 2003; Maruri-Aguilar 2005). The well-known link between Gröbner basis calculations and matrix operations for design ideals allows these methodologies to be efficient (see De Loera et al. 2009; Lundqvist 2010).

Example 11.28

For the central composite design in two dimensions and axial points at $\sqrt{2}$, the algebraic fan has only two models. The models are listed in Example 11.1.

Example 11.29 (Continuation of Example 11.23)

The collection of all models identifiable by the design L_1 (algebraic fan of L_1) is computed. Design L_1 identifies 27 different models which can be classified into six types of models, up to permutations of variables: $1, x_1, x_1^2, x_1^3, x_1^4, x_1^5$ (3 models); $1, x_1, x_1^2, x_1^3, x_1^4, x_2$ (6 models); $1, x_1, x_1^2, x_1^3, x_2, x_1x_2$ (6 models); $1, x_1, x_2, x_1^2, x_1x_2, x_2^2$ (3 models); $1, x_1, x_2, x_3, x_1^2, x_1^3$ (3 models), and $1, x_1, x_2, x_3, x_1^2, x_1x_2$ (6 models). We say that this fan has a complete combinatorial structure, meaning that each class of models is closed under permutations of variables, for example, if the model $1, x_1, x_1^2, x_1^3, x_1^4, x_1^5$ is in the class, so are models $1, x_2, x_2^2, x_2^3, x_2^4, x_2^5$ and $1, x_3, x_3^2, x_3^3, x_3^4, x_3^5$. The algebraic fan of L_1 is depicted in Figure 11.2, where each model is represented as a staircase diagram using one small box for each monomial term, and axes for variables are x_1, x_2, x_3 in counterclockwise direction starting from bottom left. The figures and Gröbner basis computations were performed with the software `gf` an Jensen (2009). The models are presented by classes following the order described previously (row-wise from top left). For instance, the first diagram shows the model $1, x_3, x_3^2, x_3^3, x_3^4, x_3^5$, the second is $1, x_1, x_1^2, x_1^3, x_1^4, x_1^5$ and the third is $1, x_2, x_2^2, x_2^3, x_2^4, x_2^5$. Now we turn our attention to the other latin hypercube L_2 in Example 11.23. From the design coordinates we note that this design has complete confounding between x_1 and x_2 , and we should expect a much more limited collection of models. Indeed, this design identifies only 11 models which are depicted in Figure 11.3. Only one of the models contains terms with x_1 (first from left in second row), while the rest of the models have monomials in x_2 and x_3 . The models can be classified in three classes, only one of which is closed under permutation of variables (shown in the left column in Figure 11.3).

Example 11.30 (Continuation of Example 11.19)

In total there are 218 different hierarchical models identifiable by the Plackett–Burman design. Those models belong to 6 different classes obtained by permutations of variables. As the design has only two levels in each factor, the models identified by this design are all multilinear with total degree of models ranging between 7 and 12. See Table 11.7 for

TABLE 11.7

Summary of the Algebraic Fan of the Plackett–Burman Design

Simplicial Complex	Degree	$HS(s)$	Class Size	Example Vertex/Model
	7	$1 + 7s$	1	$(1, 1, 1, 1, 1, 1, 1)$ $1, x_1, x_2, x_3, x_4, x_5, x_6, x_7$
	8	$1 + 6s + s^2$	21	$(1, 1, 1, 2, 1, 2, 0)$ $1, x_1, x_2, x_3, x_4, x_5, x_6, x_4x_6$
	9	$1 + 5s + 2s^2$	84	$(1, 0, 2, 2, 1, 0, 3)$ $1, x_1, x_3, x_4, x_5, x_7, x_3x_7, x_4x_7$
	10	$1 + 4s + 3s^2$	56	$(0, 0, 0, 2, 4, 2, 2)$ $1, x_4, x_5, x_6, x_7, x_4x_5, x_5x_6, x_5x_7$
			28	$(0, 0, 1, 3, 3, 0, 3)$ $1, x_3, x_4, x_5, x_7, x_4x_5, x_4x_7, x_5x_7$
	12	$1 + 3s + 3s^2 + s^3$	28	$(0, 0, 0, 0, 4, 4, 4)$ $1, x_5, x_6, x_7, x_5x_6, x_5x_7, x_6x_7, x_5x_6x_7$

details and examples for each class. The Hilbert series $HS(s)$ was included to count the number of terms by degree in each class.

Example 11.31 (Continuation of Example 11.16)

The fan of the regular 2^{6-2} fraction with generators $\{ABCD, CDEF\}$ is of relatively modest size: 132 models which belong to six equivalence classes whose sizes range from 12 to 24. Models range from total degree 26 to 32 and none of the equivalence classes is closed under permutation of variables.

Despite this apparent fan simplicity, these six classes share only three different total degrees and Hilbert functions. For instance, three different model classes share the same total degree 26, while other two different model classes have total degree 28. Table 11.8 shows a summary of the fan computations for this design, and Figure 11.4 shows simplicial representation of models in each class (vertices refer to single variables, edges to two-factor interactions and so on).

TABLE 11.8

Summary of the Algebraic Fan of Regular Fraction 2^{6-2}

Class	Total Degree	$HS(s)$	Class Size	Example Vertex/Model
I	26	$1 + 6s + 7s^2 + 2s^3$	24	(5, 5, 6, 2, 1, 7) $1, x_1, x_2, x_3, x_4, x_5, x_6, x_1x_2, x_1x_3, x_1x_6,$ $x_2x_3, x_2x_6, x_3x_6, x_4x_6, x_1x_3x_6, x_2x_3x_6$
II			24	(7, 1, 4, 4, 3, 7) $1, x_1, x_2, x_3, x_4, x_5, x_6, x_1x_3, x_1x_4, x_1x_5,$ $x_1x_6, x_3x_6, x_4x_6, x_5x_6, x_1x_3x_6, x_1x_4x_6$
III			24	(4, 4, 6, 2, 2, 8) $1, x_1, x_2, x_3, x_4, x_5, x_6, x_1x_3, x_1x_6, x_2x_3,$ $x_2x_6, x_3x_6, x_4x_6, x_5x_6, x_1x_3x_6, x_2x_3x_6$
IV	28	$1 + 5s + 7s^2 + 3s^3$	24	(6, 6, 2, 6, 0, 8) $1, x_1, x_2, x_3, x_4, x_6, x_1x_2, x_1x_4, x_1x_6, x_2x_4,$ $x_2x_6, x_3x_6, x_4x_6, x_1x_2x_6, x_1x_4x_6, x_2x_4x_6$
V			24	(4, 8, 0, 8, 4, 4) $1, x_1, x_2, x_4, x_5, x_6, x_1x_2, x_1x_4, x_2x_4, x_2x_5,$ $x_2x_6, x_4x_5, x_4x_6, x_1x_2x_4, x_2x_4x_5, x_2x_4x_6$
VI	32	$1 + 4s + 6s^2 + 4s^3 + s^4$	12	(8, 0, 8, 8, 0, 8) $1, x_1, x_3, x_4, x_6, x_1x_3, x_1x_4, x_1x_6, x_3x_4, x_3x_6,$ $x_4x_6, x_1x_3x_4, x_1x_4x_6, x_1x_3x_6, x_3x_4x_6, x_1x_3x_4x_6$




FIGURE 11.4 Depiction of simplicial models for fan classes I-VI (left to right), design 2^{6-2} .

Example 11.32 (Continuation of Example 11.14)

For $k = 4$ dimensions, the algebraic fan of this screening design has 54 models. Models in the fan belong to only three equivalence classes, each of which is closed under permutation of variables. Representatives of each class are shown in Example 11.14: the model obtained with a lexicographic order belongs to a class of size 6; the model obtained under degree lexicographic belongs to a class of size 24, while the class of models obtained with an elimination ordering has also size 24.

The algebraic fan of the screening design for seven factors $k = 7$ and $n = 15$ runs is a complicated and large object which nevertheless exhibits in some instances combinatorial symmetry. The design identifies 18,368 staircase models which can be classified in 25 equivalence classes. The class sizes range from 7 to 2520, while the total degree of models ranges from 21 to 31. Six equivalence classes are closed under permutation of variables, and this includes the classes of models identified by degree lexicographic (420 models) and by degree reverse lexicographic (210 models); examples of models for each ordering are computed in Example 11.14. Two other equivalence classes are almost closed under permutations of variables; each can be paired with another small equivalence class.

11.8.2 State Polytope and Linear Aberration

The *state polytope* of $I(D)$ is a geometric object which is associated with the Gröbner fan of $I(D)$ (see Bayer and Morrison 1988; Mora and Robbiano 1988). The state polytope is constructed as the convex hull of state vectors, and each state vector is built from a model in the algebraic fan by simply adding the exponents of the model. Aside from a proportionality constant, each state vector is indeed the centroid of the staircase diagram represented by the model and thus the state polytope is the convex hull of all those centroids (see the following example).

Example 11.33

For the central composite design of Examples 11.1, 11.13 and 11.28, the state vector of the model with terms $1, x_1, x_2, x_1^2, x_1x_2, x_2^2, x_1^3, x_1^4, x_1^2x_2$ is $(0, 0) + (1, 0) + (0, 1) + (2, 0) + (1, 1) + (0, 2) + (3, 0) + (4, 0) + (2, 1) = (13, 5)$. The other model in the fan is the conjugate of the earlier, that is, it is obtained by permuting the variables thus its state vector is $(5, 13)$. Therefore, the state polytope of the central composite design is the convex hull $\text{conv}(\{(13, 5), (5, 13)\})$.

The state polytope of $I(D)$ encodes information in terms of each variable about the weighted total degree of each model in the fan of design D . A simple argument of linear programming shows that models in the algebraic fan are those that minimise a simple linear cost function on the weighted degree of the model. This is the idea of *linear aberration* defined in Berstein et al. (2010). This concept of associating cost functions to staircase polynomial models has been generalised to nonlinear cost functions in Berstein et al. (2008) and De Loera et al. (2009).

Example 11.34

For the Latin hypercube design L_1 of Example 11.23 the state polytope of its design ideal is built with state vectors for each of the 27 models enumerated in Example 11.29.

For instance, the model $1, x_1, x_1^2, x_1^3, x_1^4, x_1^5$ has state vector $(15, 0, 0)$, and as the other two models in this class are created by permutations of variables, the same action is performed on the state vectors, so for this class, we have three vectors: $(15, 0, 0)$, $(0, 15, 0)$ and $(0, 0, 15)$. A similar construction and arguments are used for each model in the fan of L_1 , and we have 6 vectors for each of the permutations of $(10, 1, 0)$, $(7, 2, 0)$ and $(4, 2, 1)$, three permutations for each of $(4, 4, 0)$ and $(6, 1, 1)$.

There is a special type of polynomial model which is of minimal weighted degree. Such models are termed corner cut staircases (see Onn and Sturmfels 1999), as their exponents can be separated by their complement by a single hyperplane. The properties of corner cut staircases and their cardinality have been studied in literature (see Corteel et al. 1999; Wagner 2002).

A design that identifies all corner cut models is termed a *generic design* (see Onn and Sturmfels 1999), and automatically a generic design is of minimal linear aberration (see Berstein et al. 2010). The collection of models identified by design L_1 (of Examples 11.23, 11.29 and 11.34) is the set of all corner cut staircases for $k = 3, n = 6$, and thus L_1 is a generic design. State polytopes associated with corner cuts and generic designs are described in Müller (2003).

In addition to information about degrees of models in the fan, the state polytope also encodes information to compute Gröbner bases. To each vertex of the state polytope, a normal cone is associated (see Ziegler 1995). The collection of all those cones is precisely the Gröbner fan of $I(D)$, in the sense that the interior of each full dimensional cone contains ordering vectors necessary to compute the Gröbner basis (and identify the model) for the corresponding vertex.

In Figure 11.5, cones in the fan of state polytopes for designs L_1 and L_2 are depicted. As in each case the tridimensional cones form a partition of the first orthant, the figures show a slice of the cones when intersected with the standard simplex. The diagram for design L_1 (a) shows 27 cells, one for each model. The central symmetry of the diagram corresponds to symmetry of models under permutation of variables. Ordering vectors taken from the same cell will yield the same vertex (and corresponding model). Design L_2 produced the (b) in Figure 11.5. The diagram shows still some symmetry, but not central symmetry.




FIGURE 11.5 Gröbner fan for designs L_1 (a) and L_2 (b) of Example 11.23.

This symmetry reflects the range of models computed for L_2 in Example 11.29, where only 11 models are identifiable by L_2 , and 10 models are in terms of x_2 and x_3 .

11.9 Other Topics and References

The algebraic method in the form discussed here started with Pistone and Wynn (1996), Diaconis and Sturmfels (1998) and Dinwoodie (1998), and the basic ideas are presented in the monograph by Pistone et al. (2001). A short review is in Riccomagno (2009). More extensive work on the computation of universal Gröbner bases with zonotopes appears in Babson et al. (2003). Recent work showed the relation between minimal aberration models and the border description of models in terms of Betti numbers of monomial ideals (see Maruri-Aguilar et al. 2012). For the link between Gröbner basis methods and projection properties of factorial designs, see Evangelaras and Koukouvinos (2006). Applications to mixture designs are found in Giglio et al. (2001) and Maruri-Aguilar et al. (2007). Often, design coordinates are not known or perhaps are measured with some error in which case numerically stable computations have been studied using *border bases* rather than Gröbner bases (see Heldt et al. (2009); Robbiano and Torrente (2013)). Industrial applications were performed, perhaps surprisingly early (Holliday et al. 1999; Pistone et al. 2000).

For an excellent summary of the wider work in the field of algebraic statistics, see Drton et al. (2009). One topic omitted from this chapter, but important for conducting exact conditional test for contingency tables via Markov chain Monte Carlo is, the construction of Markov bases; see Hara et al. (2010); Hara et al. (2012); Hara and Takemura (2010) and Carlini and Rapallo (2011). Important applications to biology, which continue, are covered in Pachter and Sturmfels (2005). Related and of considerable recent interest is the algebraic study of boundary exponential models (Rauh et al. 2011; Cena and Pistone 2007; Drton and Sullivant 2007).

In this chapter we study direct problems in design of experiments, that is, given a design, the properties it has. Inverse problems in design have also been studied from the perspective of algebra. From a starting configuration of points and a given set of monomial terms of interest (model), algorithms have been developed to retrieve subsets of the original configuration that are able to identify the given model. The starting configuration is typically a full factorial equally spaced grid. Here we only give an example, adapted from Fontana et al. (2013).

Example 11.35

Consider the full factorial design 2^4 with levels 0 and 1 and the model with all linear terms and all two-factor interactions. This model has 11 terms: $1, x_1, x_2, x_3, x_4, x_1x_2, x_1x_3, x_1x_4, x_2x_3, x_2x_4$ and x_3x_4 . Simple but expensive enumeration techniques would allow to detect that out of $\binom{16}{11} = 4368$ possible subsets of the full design; there are 3008 fractions that identify the model, while the remaining 1360 cannot identify it.

The techniques using *circuits* of matrices in Fontana et al. (2013) avoid computing ranks of design-model matrices thus making the search very efficient. Results for the inverse problem have also been given using border bases (see Caboara and Robbiano 2001; Robbiano 2006). Recall that the emphasis of these results is in obtaining design-model matrices that are full rank.

References

- Abbott, J., Bigatti, A., Kreuzer, M., and Robbiano, L. (2000), Computing ideals of points, *Journal of Symbolic Computation*, 30, 341–356.
- Babson, E., Onn, S., and Thomas, R. (2003), The Hilbert zonotope and a polynomial time algorithm for universal Gröbner bases, *Advances in Applied Mathematics*, 30, 529–544.
- Bailey, R. A. (2008), *Design of Comparative Experiments*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge, U.K.: Cambridge University Press.
- Bayer, D. and Morrison, I. (1988), Standard bases and geometric invariant theory. I. Initial ideals and state polytopes, *Journal of Symbolic Computation*, 6, 209–217, Computational Aspects of Commutative Algebra.
- Berstein, Y., Lee, J., Maruri-Aguilar, H., Onn, S., Riccomagno, E., Weismantel, R., and Wynn, H. (2008), Nonlinear matroid optimization and experimental design, *SIAM Journal on Discrete Mathematics*, 22, 901–919.
- Berstein, Y., Maruri-Aguilar, H., Onn, S., Riccomagno, E., and Wynn, H. (2010), Minimal average degree aberration and the state polytope for experimental designs, *Annals of the Institute of Statistical Mathematics*, 62, 673–698.
- Buchberger, B. (2006), An algorithm for finding the basis elements of the residue class ring of a zero dimensional polynomial ideal, *Journal of Symbolic Computation*, 41, 475–511, translated from the 1965 German original by Michael P. Abramson.
- Caboara, M. and Robbiano, L. (2001), Families of estimable terms, in *Proceedings of the 2001 International Symposium on Symbolic and Algebraic Computation*, New York: ACM, pp. 56–63 (electronic).
- Carlini, E. and Rapallo, F. (2011), A class of statistical models to weaken independence in two-way contingency tables, *Metrika: International Journal for Theoretical and Applied Statistics*, 73, 1–22.
- Cena, A. and Pistone, G. (2007), Exponential statistical manifold, *Annals of the Institute of Statistical Mathematics*, 59, 27–56.
- Chen, H. H. and Cheng, C.-S. (2004), Aberration, estimation capacity and estimation index, *Statistica Sinica*, 14, 203–215.
- Cheng, C.-S. and Mukerjee, R. (1998), Regular fractional factorial designs with minimum aberration and maximum estimation capacity, *The Annals of Statistics*, 26, 2289–2300.
- CoCoATeam (2009), CoCoA: A system for doing computations in commutative algebra, Available at cocoa.dima.unige.it.
- Cortee, S., Rémond, G., Schaeffer, G., and Thomas, H. (1999), The number of plane corner cuts, *Advances in Applied Mathematics*, 23, 49–53.
- Cox, D., Little, J., and O’Shea, D. (1998), *Using Algebraic Geometry*, vol. 185 of Graduate Texts in Mathematics, New York: Springer-Verlag.
- Cox, D., Little, J., and O’Shea, D. (2007), *Ideals, Varieties, and Algorithms: An Introduction to computational Algebraic Geometry and Commutative Algebra*, 3rd edn. Undergraduate Texts in Mathematics, New York: Springer.
- De Loera, J. A., Haws, D. C., Lee, J., and O’Hair, A. (2009), Computation in multicriteria matroid optimization, *ACM Journal of Experimental Algorithmics*, 14, Paper 1.8, 33.
- Decker, W., Greuel, G.-M., Pfister, G., and Schönemann, H. (2012), SINGULAR 3-1-6 — A computer algebra system for polynomial computations, Available at www.singular.uni-kl.de.
- Diaconis, P. and Sturmfels, B. (1998), Algebraic algorithms for sampling from conditional distributions, *The Annals of Statistics*, 26, 363–397.
- Dimitrova, E. S., Jarrah, A. S., Laubenbacher, R., and Stigler, B. (2007), A Gröbner fan method for biochemical network modeling, in *ISSAC 2007*, New York: ACM, pp. 122–126.
- Dinwoodie, I. H. (1998), The Diaconis-Sturmfels algorithm and rules of succession, *Bernoulli: Official Journal of the Bernoulli Society for Mathematical Statistics and Probability*, 4, 401–410.

- Drton, M., Sturmfels, B., and Sullivant, S. (2009), *Lectures on Algebraic Statistics*, vol. 39 of Oberwolfach Seminars, Basel, Switzerland: Birkhäuser Verlag.
- Drton, M. and Sullivant, S. (2007), Algebraic statistical models, *Statistica Sinica*, 17, 1273–1297.
- Evangelaras, H. and Koukouvinos, C. (2006), A comparison between the Gröbner bases approach and hidden projection properties in factorial designs, *Computational Statistics & Data Analysis*, 50, 77–88.
- Fontana, R., Pistone, G., and Rogantin, M. P. (2000), Classification of two-level factorial fractions, *Journal of Statistical Planning and Inference*, 87, 149–172.
- Fontana, R., Rapallo, F., and Rogantin, M. P. (2014), A characterization of saturated designs for factorial experiments, *Journal of Statistical Planning and Inference*, 147, 204–211.
- Giglio, B., Wynn, H. P., and Riccomagno, E. (2001), Gröbner basis methods in mixture experiments and generalisations, in *Optimum Design 2000 (Cardiff)*, vol. 51 of Nonconvex Optimization and Applications, Dordrecht, the Netherland: Kluwer Academic Publisher, pp. 33–44.
- Grayson, D. R. and Stillman, M. E. (2009), Macaulay2, a software system for research in algebraic geometry, Available at www.math.uiuc.edu/Macaulay2/.
- Hara, H., Aoki, S., and Takemura, A. (2010), Minimal and minimal invariant Markov bases of decomposable models for contingency tables, *Bernoulli: Official Journal of the Bernoulli Society for Mathematical Statistics and Probability*, 16, 208–233.
- Hara, H., Sei, T., and Takemura, A. (2012), Hierarchical subspace models for contingency tables, *Journal of Multivariate Analysis*, 103, 19–34.
- Hara, H. and Takemura, A. (2010), Connecting tables with zero-one entries by a subset of a Markov basis, in *Algebraic Methods in Statistics and Probability II*, vol. 516 of Contemporary Mathematics, Providence, RI: American Mathematical Society, pp. 199–213.
- Heldt, D., Kreuzer, M., Pokutta, S., and Poulisse, H. (2009), Approximate computation of zero-dimensional polynomial ideals, *Journal of Symbolic Computation*, 44, 1566–1591.
- Hilbert, D. (1890), Über die Theorie der algebraischen Formen, *Mathematische Annalen*, 36, 473–534, reprinted in *Gesammelte Abhandlungen*, Volume II, Chelsea, New York, 1965.
- Holliday, T., Pistone, G., Riccomagno, E., and Wynn, H. P. (1999), The application of computational algebraic geometry to the analysis of designed experiments: A case study, *Computational Statistics*, 14, 213–231.
- Jensen, A. N. (2009), Gfan, a software system for Gröbner fans and tropical varieties, Available at home.imf.au.dk/jensen/software/gfan/gfan.html.
- Jones, B. and Nachtsheim, C. (2011), A class of three level designs for definitive screening in the presence of second order effects, *Technometrics*, 43, 1–15.
- Lundqvist, S. (2010), Vector space bases associated to vanishing ideals of points, *Journal of Pure and Applied Algebra*, 214, 309–321.
- Maruri-Aguilar, H. (2005), Universal Gröbner bases for designs of experiments, *Rendiconti dell'Istituto di Matematica dell'Università di Trieste*, 37, 95–119 (2006).
- Maruri-Aguilar, H., Notari, R., and Riccomagno, E. (2007), On the description and identifiability analysis of experiments with mixtures, *Statistica Sinica*, 17, 1417–1440.
- Maruri-Aguilar, H., Sáenz-de Cabezón, E., and Wynn, H. P. (2012), Betti numbers on polynomial hierarchical models for experimental designs, *Annals of Mathematics and Artificial Intelligence*, 64, 411–426.
- Maruri-Aguilar, H., Sáenz-de Cabezón, E., and Wynn, H.P. (2013), Alexander duality in experimental designs, *Annals of the Institute of Statistical Mathematics*, 65, 667–686.
- McKay, M. D., Beckman, R. J., and Conover, W. J. (1979), A comparison of three methods for selecting values of input variables in the analysis of output from a computer code, *Technometrics*, 21, 239–245.
- Mora, T. and Robbiano, L. (1988), The Gröbner fan of an ideal, *Journal of Symbolic Computation*, 6, 183–208, computational aspects of commutative algebra.
- Müller, I. (2003), Corner cuts and their polytopes, *Beiträge zur Algebra und Geometrie. Contributions to Algebra and Geometry*, 44, 323–333.

- Nelder, J. A. (1977), A reformulation of linear models, *Journal of the Royal Statistical Society. Series A*: 140, 48–76, with discussion.
- Onn, S. and Sturmfels, B. (1999), Cutting corners, *Advances in Applied Mathematics*, 23, 29–48.
- Pachter, L. and Sturmfels, B. (2005), Statistics, in *Algebraic Statistics for Computational Biology*, New York: Cambridge University Press, pp. 3–42.
- Peixoto, J. L. (1990), A property of well-formulated polynomial regression models, *The American Statistician*, 44, 26–30.
- Pistone, G., Riccomagno, E., and Wynn, H. (2000), Gröbner basis methods for structuring and analysing complex industrial experiments, *International Journal of Reliability, Quality and Safety Engineering*, 7, 285–300.
- Pistone, G., Riccomagno, E., and Wynn, H. P. (2001), *Algebraic Statistics*, Boca Raton, FL: Chapman & Hall/CRC, Computational Commutative Algebra in Statistics.
- Pistone, G. and Rogantin, M.-P. (2008), Indicator function and complex coding for mixed fractional factorial designs, *Journal of Statistical Planning and Inference*, 138, 787–802.
- Pistone, G. and Wynn, H. P. (1996), Generalised confounding with Gröbner bases, *Biometrika*, 83, 653–666.
- Plackett, R. L. and Burman, J. P. (1946), The design of optimum multifactorial experiments, *Biometrika*, 33, 305–325.
- R Core Team (2014), *R: A Language and Environment for Statistical Computing*, Vienna, Austria: R Foundation for Statistical Computing.
- Rauh, J., Kahle, T., and Ay, N. (2011), Support sets in exponential families and oriented matroid theory, *International Journal of Approximate Reasoning*, 52, 613–626.
- Riccomagno, E. (2009), A short history of algebraic statistics, *Metrika: International Journal for Theoretical and Applied Statistics*, 69, 397–418.
- Robbiano, L. (2006), Zero-dimensional ideals or the inestimable value of estimable terms, in *Constructive Algebra and Systems Theory*, vol. 53 of *Verhandelingen, Afdeling Natuurkunde. Eerste Reeks*. Koninklijke Nederlandse Akademie van Wetenschappen, Amsterdam, the Netherlands: Royal Netherlands Academy of Arts and Science, pp. 95–114.
- Robbiano, L. and Torrente, M.-L. (2013), Zero-dimensional families of polynomial systems, *Le Matematiche*, 68, 137–164.
- Sáenz-de Cabezón, E. and Wynn, H. P. (2009), Betti numbers and minimal free resolutions for multi-state system reliability bounds, *Journal of Symbolic Computation*, 44, 1311–1325.
- Sáenz-de Cabezón, E. and Wynn, H. P. (2011), Computational algebraic algorithms for the reliability of generalized k -out-of- n and related systems, *Mathematics and Computers in Simulation*, 82, 68–78.
- Wagner, U. (2002), On the number of corner cuts, *Advances in Applied Mathematics*, 29, 152–161.
- Xiao, L., Lin, D. K. J., and Bai, F. (2012), Constructing definitive screening designs using conference matrices, *Journal of Quality Technology*, 44, 2–8.
- Ziegler, G. M. (1995), *Lectures on Polytopes*, vol. 152 of *Graduate Texts in Mathematics*, New York: Springer-Verlag.

Section IV

Optimal Design for Nonlinear and Spatial Models

12

Optimal Design for Nonlinear and Spatial Models: Introduction and Historical Overview

Douglas P. Wiens

CONTENTS

12.1 Introduction.....	457
12.2 Generalized Linear Models.....	457
12.3 Selected Nonlinear Models.....	460
12.4 Spatial Models.....	465
References.....	467

12.1 Introduction

The topic of this part of the handbook—optimal design for *nonlinear* and spatial models—allows for a very broad range of subtopics. We should first distinguish these from those formulated for *linear* models. A salient feature of design problems for linear models is that the common functions expressing the experimenter’s loss, when estimating the mean response, do not depend on the unknown parameters being estimated. In this chapter, a number of design problems are introduced in which this very convenient feature is absent, and ways of dealing with its absence are discussed in general terms. Thus, although we treat classical nonlinear regression models in which a response variable y is measured with additive error and $E[y|x]$ is a nonlinear function of parameters θ to be estimated after the experiment is conducted, there is a multitude of other applications. In this chapter, these subjects will be introduced in broad generality only, and some historical context provided; precise details and examples are given in the three chapters which follow:

- Designs for Generalized Linear Models (Chapter 13)
- Designs for Selected Nonlinear Models (Chapter 14)
- Optimal Design for Spatial Models (Chapter 15)

Chapters 22, 24 and 25 deal with special applications that use nonlinear models.

12.2 Generalized Linear Models

For a book-length treatment of generalized linear models (GLMs), we refer the reader to the now classic text McCullagh and Nelder (1989). Briefly, the response variable y , given a

covariate vector x chosen by the experimenter, follows a distribution from the exponential family, with (canonical) density

$$p(y|\theta, \phi, x) = \exp \left\{ \frac{y\theta - b(\theta)}{a(\phi)} + c(y, \phi) \right\},$$

for scalar functions $a(\cdot)$, $b(\cdot)$ and $c(\cdot, \cdot)$. The canonical parameter θ relates the systematic linear component $\eta(x) = f'(x)\beta$, with regressors $f(x)$ and regression parameters β , to the mean $\mu = db(\theta)/d\theta$ via an invertible *link function* g , namely, $\eta = g(\mu)$. We write $h = g^{-1}$; $h^{(1)}$ and $h^{(2)}$ are the first and second derivatives with respect to η .

The parameters are typically estimated by maximum likelihood, computed from observations $\{y_i\}_{i=1}^n$ made at points $\{x_i\}_{i=1}^n$ chosen from a design space χ . The asymptotic variance of $\hat{\beta}$ is the inverse $I^{-1}(\beta)$ of the information matrix:

$$I(\beta) = X'UX,$$

where X is the model matrix, with i th row $f'(x_i)$ ($i = 1, \dots, n$) and U is the diagonal matrix of *weights*, with i th diagonal element $(h^{(1)}(\eta(x_i)))^2 / \text{Var}[y|x_i]$.

If the designer is primarily interested in precise estimation of β , then he or she will aim to maximize, in some sense, $I(\beta)$; this leads to the adoption of classical alphabetic optimality criteria—notably D-optimality, in which the goal is maximization of $\det(I(\beta))$.

The mean is estimated by

$$\hat{\mu}(x) = h(f'(x)\hat{\beta}),$$

with asymptotic variance and asymptotic bias given by (Robinson and Khuri 2003)

$$\text{Var}[\hat{\mu}(x)] = (h^{(1)}(\eta(x)))^2 f'(x)I^{-1}(\beta)f(x),$$

$$\text{Bias}[\hat{\mu}(x)] = h^{(1)}(\eta(x))f'(x)I^{-1}(\beta)X'U\psi + \frac{1}{2}h^{(2)}(\eta(x))f'(x)I^{-1}(\beta)f(x),$$

where $\psi_{n \times 1}$ has elements

$$\psi_i = -\frac{h^{(2)}(\eta(x_i))}{2}f'(x_i)I^{-1}(\beta)f(x_i).$$

If interest focusses on prediction of mean values, then the designer will aim to minimize some function of the mean squared errors (MSEs)

$$\text{MSE}[\hat{\mu}(x)] = \text{Var}[\hat{\mu}(x)] + \text{Bias}^2[\hat{\mu}(x)],$$

an obvious choice is the integral or average of $\text{MSE}[\hat{\mu}(x)]$ over the design space χ .

The class of GLMs spawns a wealth of particular applications and related design issues. Prominent among these is logistic regression, in which a binary response y has

$$P(y = 1) \stackrel{\text{def}}{=} \pi = L(\alpha + \beta x),$$

for $L(\eta) = 1/(1 + e^{-\eta})$, the logistic distribution. Here $\mu = \pi$ and $\eta = g(\mu) = \ln(\pi/(1 - \pi))$, the *logit*. One might seek a design—a choice of values of x and a specification of the frequencies with which y is to be observed at these values—in order to estimate the linear parameters efficiently, or to study functions of these parameters. For instance, in bioassay and dose–response problems, interest often focusses on the covariate value

$$x_{\pi_0} = \frac{L^{-1}(\pi_0) - \alpha}{\beta},$$

required to attain a response $y = 1$ in a specified proportion π_0 of the population.

The role of the logistic distribution in the aforementioned may of course be played by other distributions; if L is replaced by the Gaussian distribution function, then one is dealing with *probit regression*, and similar design problems are of interest. One of the earliest instances of nonlinear regression design is for the exponential regression GLM—Fisher (1922) considered a *dilution-series* problem, with $P(y = 1) = \exp(-\theta x)$ with $\theta, x > 0$. This problem is also the subject of an example by Fedorov (1972, pp. 121–122), who notes that the information *matrix* for θ is a scalar, maximized by placing all observations at the solution x ($\approx 1.6/\theta$) of the equation $2e^{-\theta x} + \theta x = 2$.

In the Poisson count model, y follows a Poisson distribution with mean $\mu = f'(x)\beta$, and the experimenter is typically interested in efficiently estimating functions of β . The optimal design will of course depend on which such function is of interest. Particular examples are discussed, for response surface exploration in an environmetric setting, by Myers (1999).

In these problems, and indeed in virtually all design problems for GLMs, one begins by determining an optimal design under the assumption that certain parameters—even those to be estimated from the experimental data—are known beforehand. This, clearly untenable, assumption might then be dropped in a number of ways, all discussed in detail in the chapters which follow. One can content oneself with a *locally optimal* design, in which optimality is sought only at, or in a small neighbourhood of, these assumed parameter values. Alternatively one might design so as to minimize the *maximum* loss, with the maximum evaluated over a set of parameters—a mild robustness criterion which is also discussed in Chapter 20. Another approach is to choose the design points sequentially, at each stage using parameter estimates derived from the preceding observations.

A further possibility, when the loss function being minimized depends on unknown parameters, is to integrate them out, with respect to a *prior*, and to then minimize the average loss so obtained. This ‘*pseudo*’ Bayesian criterion is discussed in Chapter 13 and is a topic to which we return in Section 12.3.

The field of optimal design for GLMs seems to have blossomed in the 1980s, and many contributors acknowledge a debt to Ford et al. (1989), who surveyed the then current state of research in a more general context of nonlinear design. Burrige and Sebastiani (1992) obtained locally D-optimal designs, that is, designs maximizing $\det(\mathbf{I}(\beta))$ for fixed values of β . For this, they pointed out that if the parameters are known, then the problem can be transformed to the D-optimality problem for a linear model with model matrix $\mathbf{U}^{1/2}\mathbf{X}$; they applied methods developed for linear design theory to derive optimal designs in this transformed problem and then translated these back to the original problem. In a small simulation study, with a bivariate linear predictor η and canonical link $\eta = \mu^{1/k}$ for various values of k , the efficiencies turned out to be relatively insensitive to the settings of the parameter values.

Ford et al. (1992) refer to this transformed problem, in terms of $\mathbf{U}^{1/2}\mathbf{X}$, as the *canonical form* of the problem. They consider the structure of the induced design space in some depth and use methods of Elfving (1952) to obtain locally D-optimal and c-optimal designs; the latter are designs minimizing $\mathbf{c}'\mathbf{I}^{-1}(\boldsymbol{\beta})\mathbf{c}$ for fixed \mathbf{c} . As do Burrigge and Sebastiani (1992), they concentrate on examples with two linear parameters (β_0, β_1) ; in both of these papers, the optimal designs turn out to be concentrated on one, two or three points. Atkinson and Haines (1996) apply this canonical approach to, among others, examples of multifactor experiments.

A class of attractive alternatives to local optimality is given by sequential designs. The asymptotic theory related to this is most well developed for the case of D-optimality. Here it is supposed that one will obtain n_1 observations from an initial, static design. These are used to give initial estimates of the parameters, following which the remaining $n - n_1$ observations are made sequentially, at each stage choosing the next design point so as to maximize the determinant of the information matrix evaluated at the current parameter estimates. Chaudhuri and Mykland (1993) show that, under certain conditions, the sequence of designs so obtained converges to the D-optimal design for the true parameters. These conditions include the requirement that $n_1/n \rightarrow 0$ as both n and n_1 tend to infinity and an assumption that the parameter estimates be consistent. A consequence is that inferences made from a sequentially constructed design have the same asymptotic properties as if they were made following a static design—an observation previously made by Wu (1985) in a related context. Sinha and Wiens (2002) extend the ideas of Chaudhuri and Mykland, and incorporate some uncertainty as to the nature of the parametric model. Dror and Steinberg (2008) introduce significant improvements to these methods; in particular their sequential procedure for design construction is easily adapted to multifactor experiments and to a range of possible models.

One likely reason for the popularity of the D-optimality criterion in these problems is its invariance under non-singular transformations of the design space, leading to the possibility of transforming to the aforementioned canonical form of the problem. Failing this, other methods are available. Yang (2008) takes a direct algebraic approach to obtain A-optimal designs (minimizing $\text{trace}(\mathbf{I}^{-1}(\boldsymbol{\beta}))$) for logistic, probit and Laplace models with two linear parameters. Other criteria—minimizing the integrated $MSE[\hat{\mu}(x)]$, for instance—rely more heavily on numerical methods of design construction. One sequential approach of considerable interest involves *stochastic approximation*—see the discussion in Khuri et al. (2006) and, in a dose-finding framework, Cheung (2010). Once a design is constructed—by this or another method—it is of obvious interest to compare its performance with other candidate designs; the quantile dispersion graphs of Robinson and Khuri (2003) provide a possible means for doing this.

Here, but a few of the many facets of design for GLMs have been touched upon; these, and the broad spectrum of topics discussed in Chapter 13, illustrate that design theory for GLMs continues to be an active and exciting area of research.

12.3 Selected Nonlinear Models

Part of the richness of the theory of design for nonlinear models stems from the physical settings in which the various models arise, each resulting in unique approaches to the design problems. Some particular nonlinear regression models, of the form

$$y = \eta(x; \theta) + \varepsilon,$$

where ε is random error and $\eta(x; \theta)$ is an at least partially nonlinear function of a p -dimensional parameter vector θ , correspond to the following response functions:

- The response

$$\eta(x; \theta) = \frac{\theta_1}{\theta_1 - \theta_2} \left(e^{-\theta_2 x} - e^{-\theta_1 x} \right), \quad \theta_1, \theta_2 > 0, x > 0,$$

for which the design problem was studied by Box and Lucas (1959), is used in chemometrics to model reactions in which a substance decomposes from a state A to a state B and finally to a state C. The parameters θ_1 and θ_2 measure the rates of these two decompositions, and η is the mean yield in state B. The design variable x represents time; a consequence is that, in contrast to many other design problems, there is no possibility of replication—only one observation can be made at a specific value of x .

Here and elsewhere, we define $f(x; \theta)$ to be the gradient

$$f(x; \theta) = \left(\frac{\partial \eta(x; \theta)}{\partial \theta_1}, \dots, \frac{\partial \eta(x; \theta)}{\partial \theta_p} \right)', \quad (12.1)$$

and $F(\theta)$ to be the $n \times p$ matrix with i th row $f'(x_i; \theta)$, where x_i denotes the settings of the variables in the i th run of the experiment. Box and Lucas make *preliminary guesses* $\theta^* = (\theta_1^*, \theta_2^*)$ and adopt the local D-optimality criterion, which aims to maximize the determinant $|F'(\theta^*) F(\theta^*)|$. A motivation is that when the asymptotic distribution of the parameter estimates is employed and if the initial guesses are correct, then such a design results in confidence ellipsoids of minimum volume. When the *de la Garza phenomenon* holds or is assumed—this is expanded upon and exploited in Chapter 14—an optimal design will have only p support points and thus $|F'(\theta^*) F(\theta^*)| = |F(\theta^*)|^2$; this simplifies the search for the optimal points, at least when p is small and when analytic rather than numerical, methods are being used. Box and Lucas obtained optimal points (x_1, x_2) through a combination of geometric and analytic arguments, and used this example to illustrate a stepwise journey to the optimum, through fitting a sequence of quadratic models, in x_1 and x_2 , very similar to common practice in response surface exploration.

- The Michaelis–Menten enzyme kinetic function is

$$\eta(x; \theta) = \frac{\theta_1 x}{\theta_2 + x}, \quad \theta_1, \theta_2 > 0, x > 0,$$

where x is the concentration of substrate, θ_1 the maximum reaction velocity (i.e., the horizontal asymptote as $x \rightarrow \infty$), and θ_2 is the *half-saturation* constant, that is, the value of x at which the mean velocity η attains one-half of its asymptotic value.

An important feature of this model from a design standpoint is that it is nonlinear in θ_2 but not in θ_1 , and then the loss function for D-optimality depends (up to a constant of proportionality) only on θ_2 . Currie (1982) discusses various designs for

this model. Assuming homoscedastic normal errors, the maximum likelihood estimates are obtained by least squares, leading to the D-optimal design which places half of the observations at θ_2 and the other half at as large a value of x as possible. (Bates and Watts 1988, p. 126) state instead that the two locations are $x_1 = x_{\max}$, the maximum allowable value, and $x_2 = \theta_2/(1 + 2(\theta_2/x_{\max}))$, in agreement with Currie if x_2 is evaluated at $x_{\max} = \infty$.)

An obvious drawback to this design, shared by others in which the number of distinct locations of the explanatory variables is no larger than p , is that there is no possibility to check the validity of the model—a point which is also discussed in Chapter 20. Thus Currie discusses as well more *ad hoc*, but sensible, designs in which the majority of the design points are spread out over the low range of concentration, with the rest distributed throughout the higher range. He finds that the value of $|F'(\theta)F(\theta)|$ (evaluated at the assumed value of θ_2) can be substantially smaller than that for the locally D-optimal design, but that the performance of this latter design can itself deteriorate markedly if the experimenter’s guess at the value of θ_2 is inaccurate. An obvious remedy, if conditions permit, is to design sequentially, with past observations used to give improved estimates of θ_2 .

The Michaelis–Menten model is used throughout Chapter 14 for illustration of the concepts there.

- The rational function response

$$\eta(x; \theta) = \frac{\theta_1 \theta_3 x_1}{1 + \theta_1 x_1 + \theta_2 x_2}, \quad x_1, x_2 > 0,$$

models chemical reactions of the type $R \rightarrow P_1 + P$, with η representing the speed of the reaction, x_1 the partial pressure of the sought product P , x_2 the partial pressure of the product P_1 , θ_2 the absorption equilibrium constant for P_1 , θ_3 the effective constant of the speed of reaction (appearing linearly) and θ_1 the absorption equilibrium constant for the reagent R (Fedorov 1972, pp. 193–199). Box and Hunter (1965) propose a sequential approach with, at each stage, new locations $x = (x_1, x_2)'$ chosen to maximize the resulting value of $|F'(\hat{\theta})F(\hat{\theta})|$ evaluated at the current estimates $\hat{\theta}$. Fedorov (1972) discusses this example in detail. Initial estimates θ^* of the parameter values are obtained from a preliminary experiment, with observations made at the four combinations of $x_1, x_2 \in \{1, 2\}$. Given a design specifying n observations, and resulting in parameter estimates $\hat{\theta}_{(n)}$, the next location is given by

$$x_{n+1} = \arg \max_x f'(x; \hat{\theta}_{(n)}) [F'(\hat{\theta}_{(n)})F(\hat{\theta}_{(n)})]^{-1} f(x; \hat{\theta}_{(n)}),$$

stopping once the changes in the parameter estimates become insignificant. The asymptotic optimality results of Chaudhuri and Mykland (1993) and Wu (1985), mentioned in Section 12.2, apply.

Recall that the volume of a confidence ellipsoid on the parameters is proportional to $|F'F|^{-1/2}$. Even under exact normality, the coverage probability of such regions equals the nominal value only for linear response surfaces. Hamilton et al. (1982) obtain corrected, *second-order* expressions for the volume of such regions, with the

correction term, which is $O_p(n^{-1})$, depending on the degree of nonlinearity of the response. Hamilton and Watts (1985) then reconsider the sequential design procedure for this rational function example as an illustration of their *quadratic design criterion*, which aims to minimize the corrected value of the volume. They find that each subsequent observation diminishes the effect of the nonlinearity and also that the designs can be quite different from those of Box and Hunter.

As in these examples, a preliminary goal of the experimenter might be to design for efficient estimation of the parameters; in this case, the same alphabetic optimality criteria as in linear regression are available. Or the experimenter might seek a design which aids in the selection of an appropriate model. When this is phrased as a discrimination problem, the mathematical goal could be the maximization of the power of a test of a hypothesis $\eta = \eta_0$ versus $\eta = \eta_1$, each specified up to its parameter values. If the densities $p_0(y; \eta_0(x))$ and $p_1(y; \eta_1(x))$ of y under the two models are both Gaussian, this leads to the notion of T-optimality (Atkinson and Fedorov 1975a,b). More generally (López-Fidalgo et al. 2007), it leads to KL-optimality, in which the goal is to find a design ξ maximizing

$$\inf_{\theta_0} \int \mathcal{I}(\eta_0(x|\theta_0), \eta_1(x|\theta_1)) \xi(dx), \tag{12.2}$$

here

$$\mathcal{I}(\eta_0(x), \eta_1(x)) = \int_{-\infty}^{\infty} p_1(y; \eta_1(x)) \log \left\{ \frac{p_1(y; \eta_1(x))}{p_0(y; \eta_0(x))} \right\} dy,$$

is the Kullback–Leibler divergence, measuring the information which is lost when p_0 is used to approximate p_1 . In (12.2), θ_1 is assumed known. Both static and sequential approaches are available; robustifications of this approach are discussed in Chapter 20.

Whatever might be the parameter-dependent loss function, a possibility is to seek a design minimizing the *average loss*; namely,

$$\xi_0 = \arg \min_{\xi} \int_{\Theta} \mathcal{L}(\xi; \theta) \pi(\theta) d\theta, \tag{12.3}$$

where $\mathcal{L}(\xi; \theta)$ is the loss corresponding to a design ξ when the true model is parameterized by θ and $\pi(\cdot)$ is a user-chosen function assigning greater weight to parameter values thought to be most plausible or perhaps values against which one desires greater protection.

For instance, the choice $\mathcal{L}(\xi; \theta) = -\log |\mathbf{M}(\xi; \theta)|$, where

$$\mathbf{M}(\xi; \theta)_{p \times p} = \int_{\mathcal{X}} \mathbf{f}(x; \theta) \mathbf{f}'(x; \theta) \xi(dx),$$

gives an analogue of classical D-optimality. For this choice, an equivalence theorem (Läuter 1974; see also Section 7.3 of Cox and Reid 2000) applies and states that, under mild conditions, ξ_0 satisfies (12.3) if and only if

$$d(x; \xi_0) = \int_{\Theta} \mathbf{f}'(x; \theta) \mathbf{M}^{-1}(\xi_0; \theta) \mathbf{f}(x; \theta) \pi(\theta) d\theta \leq p,$$

at all points x in the design space, with equality at the support points of ξ_0 . As a simple yet instructive example, suppose that one intends to fit an exponential response

$$\eta(x; \theta) = e^{-\theta x}, \quad (12.4)$$

with additive error, by least squares. Then in (12.1), $p = 1$, $f(x; \theta) = -xe^{-\theta x}$ and the requirement becomes, in an obvious notation,

$$E_{\pi} \left[\frac{x^2 e^{-2\theta x}}{E_{\xi_0} [x^2 e^{-2\theta x}]} \right] \leq 1, \quad (12.5)$$

with equality at the support points. With a design region $\chi = (0, 1]$, (12.5) applied to a one-point design with all mass at $x_0 \in \chi$ becomes

$$E_{\pi} [e^{-2\theta(x-x_0)}] \leq (x_0/x)^2.$$

Some calculus yields $x_0 = \min \{1, 1/E_{\pi} [\theta]\}$, as given in Chaloner (1993) and restated in Dette and Neugebauer (1997), where, as well, conditions on π are given under which this one-point design is optimal, that is, satisfies (12.3), within the class of *all* designs. These conditions fail if, for instance, π is uniform on $\Theta = [1, \theta_{\max}]$, for θ_{\max} sufficiently large. Then numerical methods must be used to obtain the maximizer in (12.3) directly, with (12.5) checked for verification of the optimality.

An overview of this approach to design, in which the weight functions $\pi(\cdot)$ are chosen according to a Bayesian paradigm, is given in Chaloner and Verdinelli (1995). For multi-parameter models and priors, the integrations in (12.3) can become a significant part of the problem, requiring methods such as Markov chain Monte Carlo—see, for instance, the discussions in Atkinson and Haines (1996) and Atkinson et al. (1995).

Another possibility is to design so as to test the assumed response function for lack of fit (O'Brien 1995). Designs optimal for discrimination or for lack of fit testing are typically not very efficient for estimating the parameters of the final model; this leads to designs which optimize some mixture of these goals—see Hill et al. (1968) and the discussion in Chapter 14 of the approach of Dette et al. (2005).

Similar in nature to calibration problems in linear regression are dose finding studies, which are also discussed in Chapter 24. Here one seeks the value of x resulting in a specified mean response $\eta(x; \theta)$. If η is explicitly invertible—in particular, if it is linear in the parameters—then estimates of x may be obtained from those of θ , and so the design problem is concentrated on efficient estimation of a function of the parameters. Otherwise, a possible approach is to design sequentially, guided by stochastic approximation (Cheung 2010).

A class of design problems, apparently first studied by Chernoff (1962), arises in quality control and concerns accelerated life testing. One assumes a, typically nonlinear, response relating the lifetime (y) of a product to stress levels (x) and possibly to other covariates. The experimenter can usually not wait for a product on test to fail under normal stress levels, and so attempts to obtain inferences upon subjecting the product to abnormally high stresses. The goal is accurate prediction of product lifetime at normal stress levels, so that there is a natural link here to the more general problem of designing experiments for purposes of extrapolation.

The list of design problems and applications goes on; these and others are expanded upon in Chapter 14, where as well the mathematical theory is outlined. Other useful references include Bates and Watts (1988) and Seber and Wild (1989), each of which discusses modelling, inference, computations and to some extent design, in a comprehensive manner.

12.4 Spatial Models

Spatial models pose some unique problems, both in inference and in design. Cressie (1993, p. 313) distinguishes between *spatial experimental design*, in which locations are fixed and the design consists of an allocation of treatments to these locations, and *spatial sampling*, in which the designer is faced with a spatial stochastic process (a *random field*), from which he or she is to choose locations at which to make observations.

Much of the impetus for spatial experimental design derived from agricultural experimentation, and hence a large debt is owed to R. A. Fisher, who introduced in the 1920s and 1930s the now common notions of replication, randomization, blocking, etc.; see Martin (1996). Randomized designs came to be replaced by more systematic layouts, the analysis of which led to particular requirements in accounting for the spatial dependence. One of such is *neighbour balance*—the requirement that, for instance, each treatment occurs the same number of times next to each other treatment. This might arise because of competition or interference between treatments.

The achievement of neighbour balance in a design can lead to interesting combinatorial problems; see, for instance, Druilhet and Walter (2012). Typically, efficiency of estimation of model parameters is not a particularly important goal in spatial studies; this is however the aim of many designs which take account of spatial information by instead adopting a particular structure of dependence between nearby observations. Commonly, the ensuing analysis utilizes generalized least squares estimates, tailored for the particular dependence structure assumed. An optimal design then might be one which minimizes a particular loss function associated with these estimates or predictions.

In all these cases, there might be dependence on covariates besides location; a possible model of the mean response at location t , with treatment covariates x , might be

$$E[y|x,t] = f'(x)\theta_1 + g'(t)\theta_2.$$

In this case the locations are fixed but the covariates x are to be chosen by the designer. That this is a nonlinear model arises from the spatial dependence between observations, hence the dependence of the loss on the unknown parameters of the correlation structure.

In spatial sampling, as in spatial experimental design, efficiency might take a back seat to other goals dictated by the physical setting of the problem; see Thompson (1997) and Müller (2005). *Geometry-based* designs, often intended for exploratory purposes, might aim to be *space filling*. If model-free imputation of missing observations is the primary goal, then the designer might use *probability sampling* (Matérn 1960). When the probabilistic structure is known and prediction is the goal, then an information theoretic approach might be apt—see Caselton and Zidek (1984), who propose the maximization of mutual information based on Shannon's entropy, and the environmental application in Zidek et al. (2000).

On the other hand, when efficient parameter estimation and parametric inference is the aim, we are in the realm of optimal sampling design. The first step is often the choice of a correlation function specifying the nature and degree of the dependencies between observations made at various locations. This function plays a central role in the prediction of the response at unsampled locations—typically through *kriging* (i.e. best linear unbiased prediction)—and hence on the construction of designs. The choice of a particular spatial model is discussed at length in the companion handbook Gelfand et al. (2010), and so we do not discuss this here.

A common aim of the designer is to minimize the integral (or sum, if the set of locations is discrete) of the MSEs of the predictions over all locations in the region of interest. Minimizing the maximum MSE is another possibility. This MSE might arise from the spatial variation and its estimation; another contributing factor might be the estimation of the mean response $E[y|x,t]$, modelled parametrically. When a regression response is modelled, the usual alphabetic optimality criteria become germane. In some applications, physical interpretations of covariance function parameters are also important and can become the objective of the design.

To give some idea of the flavour of the techniques, consider the following design problem studied by Müller (2005). A region in the Danube river basin in Austria currently has a network of 36 water quality monitoring stations. The locations are labelled relative to a grid overlying the region. To predict chloride concentration (y) at location x , the experimenter fits a regression model with spatially correlated errors and a parametric covariance function:

$$y(x) = f'(x) \beta + \varepsilon(x),$$

$$Cov[\varepsilon(x), \varepsilon(x')] = c(x, x'; \theta).$$

For illustrative purposes, Müller redesigns this network of 36 stations in several ways. In all cases, an important feature is that there is no notion of replication—only one monitoring station may be placed at a particular location. The first design illustrated is D-optimal, maximizing the determinant of the information matrix for β (with $f(x) = (1, x')$); this matrix of course depends on the covariance function, taken to be

$$c(x, x'; \theta) = \begin{cases} \theta_1 + \theta_2, & x = x', \\ \theta_2 \left\{ 1 - \frac{3}{2} \left(\frac{\|x-x'\|}{\theta_3} \right) + \frac{1}{2} \left(\frac{\|x-x'\|}{\theta_3} \right)^3 \right\}, & 0 < \|x-x'\| \leq \theta_3, \\ 0, & \|x-x'\| > \theta_3. \end{cases}$$

Exchange algorithms are introduced to carry out the optimization. The resulting design is in Figure 12.1; a notable feature is that the design calls for all stations to be concentrated at





FIGURE 12.1 D-optimal network of chlorine monitoring stations. (From Muller, W.G., *Environmetrics*, 16, 495, 2005.)

**FIGURE 12.2**

Network of chlorine monitoring stations obtained via an expansion of the covariance kernel, followed by D-optimality. (From Muller, W.G., *Environmetrics*, 16, 495, 2005.)

the boundary of the region, but to be somewhat evenly distributed on this boundary. Presumably the managers of such a network would be asked if they were perhaps duplicating efforts of others immediately across the geographic boundary of their region.

Another method of D-optimal design construction in Müller (2005) relies on an expansion of the covariance function in terms of eigenfunctions $\{\phi_l(x)\}$, resulting in an approximation of the process as

$$y(x) = f'(x)\beta + \sum_{l=1}^p \gamma_l \phi_l(x) + e(x),$$

with uncorrelated errors $\{e(x)\}$. Here the $\{\gamma_l\}$ are uncorrelated random variables with variances given by the eigenvalues corresponding to the ϕ_l . This representation allows for an analysis by random coefficient regression, leading to the design in Figure 12.2, exhibiting a greater coverage of the region than that of Figure 12.1.

There is a close relationship between spatial sampling and the design of computer experiments. Although there is no *random error*, in the usual sense, in such experiments, it is common to model the dependencies between the outputs of experiments, with distinct inputs, via spatial correlation structures. This then engenders a certain similarity in the design problems—the inputs to the computer experiment, to be chosen by the designer, play much the same role as do the locations in spatial sampling. Designs for computer experiments are discussed in Section V.

The computational demands involved in constructing spatial designs can be immense. Some techniques which have been attempted, with varying measures of success, are exchange algorithms, simulated annealing and genetic algorithms. These, and many of the topics touched on previously, are discussed at length in Chapter 15.

References

- Atkinson, A. C., Demetrio, C. G. B., and Zocchi, S. S. (1995), Optimum dose levels when males and females differ in response, *Applied Statistics*, 44, 213–226.
- Atkinson, A. C. and Fedorov, V. V. (1975a), The design of experiments for discriminating between two rival models, *Biometrika*, 62, 57–70.
- Atkinson, A. C. and Fedorov, V. V. (1975b), Optimal design: Experiments for discriminating between several models, *Biometrika*, 62, 289–303.

- Atkinson, A. C. and Haines, L. M. (1996), Designs for nonlinear and generalized linear models, in: *Design and Analysis of Experiments*, Handbook of Statistics, Vol. 13, pp. 437–475; eds. Ghosh, S. and Rao, C. R., Elsevier/North-Holland.
- Bates, D. M. and Watts, D. G. (1988), *Nonlinear Regression Analysis and Its Applications*, Wiley, New York.
- Box, G. E. P. and Hunter, W. G. (1959), Design of experiments in non-linear situations, *Biometrika*, 46, 77–90.
- Box, G. E. P. and Lucas, H. L. (1965), The experimental study of physical mechanisms, *Technometrics*, 7, 23–42.
- Burrige, J. and Sebastiani, P. (1992), Optimal designs for generalized linear models, *Journal of the Italian Statistical Society*, 1, 183–202.
- Caselton, W. F. and Zidek, J. V. (1984), Optimal monitoring network designs, *Statistics and Probability Letters*, 2, 223–227.
- Chaudhuri, P. and Mykland, P. A. (1993), Nonlinear experiments: Optimal design and inference based on likelihood, *Journal of the American Statistical Association*, 88, 538–546.
- Chaloner, K. (1993), A note on optimal bayesian design in nonlinear problems, *Journal of Statistical Planning and Inference*, 37, 229–236.
- Chaloner, K. and Verdinelli, I. (1995), Bayesian experimental design: A review, *Statistical Science*, 10, 273–304.
- Chernoff, H. (1962), Optimal accelerated life designs for estimation, *Technometrics*, 4, 381–408.
- Cheung, Y. K. (2010), Stochastic approximation and modern model-based designs for dose-finding clinical trials, *Statistical Science*, 25, 191–201.
- Cox, D. R. and Reid, N. (2000), *The Theory of the Design of Experiments*, Chapman & Hall.
- Currie, D. J. (1982), Estimating Michaelis-Menten parameters: Bias, variance and experimental design, *Biometrics*, 38, 907–919.
- Cressie, N. (1993), *Statistics for Spatial Data*, Wiley, New York.
- Dette, H., Melas, V. B., and Wong, W.-K. (2005), Optimal design for goodness-of-fit of the Michaelis-Menten enzyme kinetic function, *Journal of the American Statistical Association*, 100, 1370–1381.
- Dette, H. and Neugebauer, H.-M. (1997), Bayesian D-optimal designs for exponential regression models, *Journal of Statistical Planning and Inference*, 60, 331–349.
- Dror, H. A. and Steinberg, D. M. (2008), Sequential experimental designs for generalized linear models, *Journal of the American Statistical Association*, 103, 288–298.
- Druilhet, P. and Walter, T. (2012), Efficient circular neighbour designs for spatial interference model, *Journal of Statistical Planning and Inference*, 142, 1161–1169.
- Elfving, G. (1952), Optimal allocation in linear regression theory, *Annals of Mathematical Statistics*, 23, 255–262.
- Fedorov, V. V. (1972), *Theory of Optimal Experiments*, Academic Press, New York.
- Fisher, R. A. (1922), On the mathematical foundations of theoretical statistics, *Philosophical Transactions of the Royal Society of London, Series A*, 22, 309–368.
- Ford, I., Titterton, D. M., and Kitsos, C. P. (1989), Recent advances in nonlinear experimental design, *Technometrics*, 31, 49–60.
- Gelfand, A. E., Diggle, P., Fuentes, M., and Guttorp, P. (2010), *Handbook of Spatial Statistics*, Chapman & Hall, New York.
- Hamilton, D. C. and Watts, D. G. (1985), A quadratic design criterion for precise estimation in nonlinear regression models, *Technometrics*, 27, 241–250.
- Hamilton, D. C., Watts, D. G., and Bates, D. C. (1982), Accounting for intrinsic nonlinearity in nonlinear regression parameter inference regions, *Annals of Statistics*, 10, 386–393.
- Hill, W. J., Hunter, W. G., and Wichern, D. W. (1968), A joint design criterion for the dual problem of model discrimination and parameter estimation, *Technometrics*, 10, 145–160.
- Khuri, A. I., Mukherjee, B., Sinha, B. K., and Ghosh, M. (2006), Design issues for generalized linear models: A review, *Statistical Science*, 21, 376–399.
- Läuter, E. (1974), Experimental design in a class of models, *Mathematische Operations-forschung Statistik*, 5, 379–396.

- López-Fidalgo, J., Tommasi, C., and Trandafir, P. C. (2007), An optimal experimental design criterion for discriminating between non-normal models, *Journal of the Royal Statistical Society B*, 69, 231–242.
- Martin, R. J. (1996), Spatial experimental design, in: *Design and Analysis of Experiments*, Handbook of Statistics, Vol. 13, pp. 437–475; eds. Ghosh, S. and Rao, C. R., Elsevier/North-Holland.
- Matérn, B. (1960), *Spatial Variation*, Springer-Verlag, Berlin.
- McCullagh, P. and Nelder, J. A. (1989), *Generalized Linear Models*, Wiley, New York.
- Müller, W. G. (2005), A comparison of spatial design methods for correlated observations, *Environmetrics*, 16, 495–505.
- Myers, R. H. (1999), Response surface methodology—Current status and future directions, *Journal of Quality Technology*, 31, 30–44.
- O'Brien, T. E. (1995), Optimal design and lack of fit in nonlinear regression models, in: *Proceedings of the 10th International Workshop on Statistical Modelling, Lecture Notes in Statistics*, Springer-Verlag, New York, pp. 201–206.
- Robinson, K. S. and Khuri, A. I. (2003), Quantile dispersion graphs for evaluating and comparing designs for logistic regression models, *Computational Statistics and Data Analysis*, 43, 47–62.
- Seber, G. A. F. and Wild, C. J. (1989), *Nonlinear Regression*, Wiley, New York.
- Sinha, S. and Wiens, D. P. (2002), Robust sequential designs for nonlinear regression, *The Canadian Journal of Statistics*, 30, 601–618.
- Thompson, S. K. (1997), Effective sampling strategies for spatial studies, *Metron*, 55, 1–21.
- Wu, C. F. J. (1985), Asymptotic inference from sequential design in a nonlinear situation, *Biometrika*, 72, 553–558.
- Yang, M. (2008), A-optimal designs for generalized linear models with two parameters, *Journal of Statistical Planning and Inference*, 138, 624–641.
- Zidek, J. V., Sun, W., and Le, N. D. (2000), Designing and integrating composite networks for monitoring multivariate gaussian pollution fields, *Applied Statistics*, 49, 63–79.

13

Designs for Generalized Linear Models

Anthony C. Atkinson and David C. Woods

CONTENTS

13.1	Introduction.....	472
13.2	Generalized Linear Models.....	473
13.2.1	Family of Models.....	473
13.2.2	Normal Distribution.....	473
13.2.3	Binomial Data.....	474
13.2.4	Poisson Data.....	475
13.2.5	Gamma Data.....	475
13.3	Optimal Experimental Designs.....	476
13.3.1	Theory.....	476
13.3.2	D-Optimal Designs.....	478
13.3.3	Design Efficiency.....	479
13.3.4	Parameter Uncertainty.....	479
13.3.4.1	Bayesian Designs.....	480
13.3.4.2	Sequential Designs.....	480
13.3.4.3	Minimax and Maximin Designs.....	480
13.3.4.4	Cluster Designs.....	481
13.3.5	Small Effects.....	481
13.4	Locally Optimal Designs.....	481
13.4.1	Binomial Data: Logistic Regression in One Variable.....	481
13.5	Optimal Designs for Two or More Explanatory Variables.....	484
13.5.1	Binomial Data.....	484
13.5.1.1	Induced Design Region.....	484
13.5.1.2	First-Order Models.....	485
13.5.1.3	Induced Design Region for the Logistic Link.....	487
13.5.1.4	Theoretical Results for a First-Order Predictor.....	488
13.5.1.5	Second-Order Response Surface.....	490
13.5.1.6	Bayesian D-Optimal Designs.....	492
13.5.2	Gamma Data.....	494
13.5.2.1	Theoretical Results for First-Order Designs with the Power Link.....	494
13.5.2.2	Examples of First-Order Designs with the Power Link.....	495
13.5.2.3	Second-Order Response Surface with the Power Link.....	496
13.5.2.4	Efficient Standard Designs for Gamma Models.....	498
13.5.2.5	Designs for the Power Link and for the Box-Cox Link.....	498
13.5.3	Poisson Data.....	499

13.6 Designs with Dependent Data.....	503
13.6.1 Random Intercept Model.....	504
13.6.2 Continuous Block Designs.....	505
13.6.3 Information Matrix for a Random Intercept Model.....	505
13.6.4 Approximating the Information Matrix Using Estimating Equations.....	506
13.6.5 Comparison of Approximations.....	507
13.7 Extensions and Further Reading.....	509
Acknowledgments.....	510
References.....	510

13.1 Introduction

The methods of experimental design described in the majority of the earlier chapters are appropriate if the continuous response, perhaps after transformation, has independent errors with a variance that is known up to a multiplicative constant. (An exception is Chapter 6, which describes designs for correlated errors.) However, this is not a characteristic of the Poisson and binomial distributions, where there is a strong relationship for these discrete random variables between mean and variance. The main emphasis of this chapter is on designs for generalized linear models (GLMs) appropriate for data from these and other distributions.

The classic account of GLMs is McCullagh and Nelder (1989). Issues in the design of experiments for these models are reviewed by Khuri et al. (2006); in addition to the methods of optimal experimental design, they consider stochastic approximation (Robbins and Monro 1951) and adaptations of response surface methodology (Box and Draper 1963 and Chapter 5). Their emphasis is mainly on models with a single explanatory variable. On the other hand, the review of Atkinson (2006) focuses on optimal design and models with several explanatory variables as, to some extent, does the more recent review of Stufken and Yang (2012b), where the emphasis is towards analytical results. Here, we follow the approach of Atkinson but focus on more recent results and on computational methods.

The assumptions of normality and constancy of variance for regression models enter the criteria of optimal design through the form of the information matrix $\mathbf{X}'\mathbf{X}$, where, as in other chapters, \mathbf{X} is the $n \times p$ model, or extended design, matrix. Other forms of information matrix arise from other distributions (see Atkinson et al. 2014). Given the appropriate information matrix, the principles of optimal design are the same as those described in earlier chapters. In designs for GLMs, the asymptotic covariance matrix of the parameters of the linear model is of the form $\mathbf{X}'\mathbf{U}\mathbf{X}$, where the $n \times n$ diagonal matrix of weights \mathbf{U} depends on the parameters of the linear predictor, on the error distribution and on the link between them. The dependence of the designs on parameters whose values are unknown prior to experimentation means that, in general, designs for GLMs require some specification of prior information.

In Section 13.2, we briefly review the class of GLMs, with particular emphasis on models for binomial, Poisson and gamma data. Some fundamental ideas in optimal design of experiments are reviewed in Section 13.3, and the optimality criteria employed in this chapter are introduced. We emphasize the reliance of optimal designs on the unknown parameter values and discuss methods of overcoming this dependence. Locally optimal designs are introduced in Section 13.4 for logistic models for binomial data with a single

explanatory variable. In Section 13.5, we move on to designs for binomial, Poisson and gamma data with several explanatory variables. This latter section also includes results on Bayesian designs. In Section 13.6, we discuss designs for dependent non-normal data, for example, arising from blocked experiments, and demonstrate optimal design for generalized linear mixed models (GLMMs) through an example. In Section 13.7, we give some extensions and suggestions for further reading.

13.2 Generalized Linear Models

13.2.1 Family of Models

The family of GLMs extends normal theory regression to any distribution belonging to the one-parameter exponential family. As well as the normal (with known variance), this includes the gamma, Poisson and binomial distributions, all of which are important in the analysis of data. The three components of a GLM are

1. A distribution for the univariate response y with mean μ
2. A linear predictor $\eta = f'(x)\theta$ where $f(x)$ is a p -vector of known functions of the k explanatory variables, x , and θ is a p -vector of unknown model parameters
3. A link function $g(\mu) = \eta$, relating x to the mean μ

The distribution of y determines the relationship between the mean and the variance of the observations. The variance is of the form

$$\text{Var}(y) = \phi V(\mu), \tag{13.1}$$

where ϕ is a *dispersion parameter*, equal to σ^2 for the normal distribution and equal to one for the binomial and Poisson distributions. The variance function $V(\mu)$ is specific to the error distribution.

The information matrix for a single observation at a point x is

$$M(x; \theta) = u(x)f'(x)f'(x), \tag{13.2}$$

with the weights for individual observations given by

$$u(x) = V^{-1}(\mu) \left(\frac{d\mu}{d\eta} \right)^2. \tag{13.3}$$

These weights depend both on the distribution of y and on the link function.

13.2.2 Normal Distribution

The linear multiple regression model can be written as

$$E(y) = \mu = \eta = f'(x)\theta, \tag{13.4}$$

where μ , the mean of y for given x , is equal to the linear predictor η . In (13.1), $V(\mu) = u(x) = 1$, and in this simple case, $\phi = \sigma^2$.

Important design problems arise with extensions to this model, particularly those in which the variance is parameterized through a link function and linear predictor, that may include parameters in common with the linear predictor for the mean (Muirhead 1982; Magnus and Neudecker 1988; Atkinson and Cook 1995; Fedorov and Leonov 2014, Section 6.3.1). Some references for designs related to the extension of the model to include transformation of the response and to transformation of both sides of the model are in Section 13.7.

13.2.3 Binomial Data

For the binomial distribution, the variance function (13.1) is

$$V(\mu) = \mu(1 - \mu). \quad (13.5)$$

In models for binomial data with R successes in n observations, the response y is defined to be R/n . The link function should be such that, however, the values of x and θ vary, the mean μ satisfies the physically meaningful constraint that $0 \leq \mu \leq 1$. We list four link functions that have been found useful in the analysis of data.

1. Logistic

$$\eta = \log\left(\frac{\mu}{1 - \mu}\right). \quad (13.6)$$

The ratio $\mu/(1 - \mu)$ is the odds that $y = 1$. In the logistic model, the log odds is therefore equal to the linear predictor. Apart from a change of sign, the model is unaltered if *success* is replaced with *failure*. For this canonical link, it follows from calculation of $d\eta/d\mu$ that

$$u(x) = \mu(1 - \mu), \quad (13.7)$$

a simpler form than for the other three link functions we shall discuss.

2. Probit

$$\eta = \Phi^{-1}(\mu), \quad (13.8)$$

where Φ is the normal cumulative distribution function. This link has very similar properties to those of the logistic link. In this case,

$$u(x) = \frac{\phi^2(\eta)}{\Phi(\eta)\{1 - \Phi(\eta)\}}. \quad (13.9)$$

Here, ϕ is the standard normal density.

3. Complementary log-log

$$\eta = \log\{-\log(1 - \mu)\}, \quad (13.10)$$

for which

$$u(x) = \frac{(1 - \mu)}{\mu} \{\log(1 - \mu)\}^2. \tag{13.11}$$

The complementary log-log link is not symmetric in success and failure, so providing a model with properties distinct from those of the logistic and probit links. Interchanging success and failure gives the following log-log link.

4. *Log-log*

$$\eta = \log(-\log \mu). \tag{13.12}$$

A plot of these four link functions is in Figure 4.1 of McCullagh and Nelder (1989). Atkinson and Riani (2000, Section 6.18) describe a fifth link, the arcsine link, which has some desirable robustness properties for binary data. In our examples, we only calculate designs for the logistic link. Atkinson et al. (2007, Section 22.4.3) compare designs for logistic and complementary log-log links when there is a single explanatory variable.

13.2.4 Poisson Data

For Poisson data, where $V(\mu) = \mu$, we require that $\mu > 0$. The log link, $\log \mu = \eta$, is standard for the analysis of Poisson data in areas such as medicine, social science and economics; see, for example, Agresti (2002, Chapter 9), Winkelmann (2008) and von Eye et al. (2011). This link leads to models with $\mu = \exp \eta$, which satisfy the constraint on values of μ , and weights $u(x) = \mu$.

13.2.5 Gamma Data

The gamma family is one in which the correct link is often in doubt. The physical requirement is again that $\mu > 0$. A useful, flexible family of links that obeys this constraint is the Box and Cox family, in which

$$g(\mu) = \begin{cases} (\mu^\lambda - 1)/\lambda & (\lambda \neq 0) \\ \log \mu & (\lambda = 0). \end{cases} \tag{13.13}$$

See Box and Cox (1964) for the use of this function in the transformation of data.

Differentiation of (13.13) combined with (13.1) shows that the weight for the gamma distribution with this link family is

$$u(x) = \mu^{-2\lambda}. \tag{13.14}$$

When $\lambda = 0$, that is, for the log link, the weights in (13.14) are equal to one. It therefore follows that optimal designs for gamma models with this link and $\lambda = 0$ are identical to optimal designs for regression models with the same linear predictors. Unlike designs for binomial and Poisson GLMs, the designs when $\lambda = 0$ do not depend on the parameter θ .

This link is seemingly equivalent to the power family of links

$$g(\mu) = \mu^\kappa, \tag{13.15}$$

for which differentiation shows that

$$u(x) = \frac{1}{\kappa^2} \frac{1}{\mu^{2\kappa}} = \frac{1}{\kappa^2} \frac{1}{\eta^2}. \quad (13.16)$$

Since $1/\kappa^2$ is constant over \mathcal{X} , the optimal design depends on η , but not on κ . The situation is similar to that for normal theory regression, in which D-optimal designs are independent of the value of the variance σ^2 , although the information matrix is a function of that parameter. In Section 13.5.2, we discuss the relationship of the designs produced by these two links.

The gamma model is often an alternative to response transformation (Section 13.7). In particular, with a log link, it may be hard to distinguish the gamma from a linear regression model with logged response. A discussion is in Sections 8.1 and 8.3.4 of McCullagh and Nelder (1989) with examples of data analyses in Section 7.5 of Myers et al. (2010).

13.3 Optimal Experimental Designs

13.3.1 Theory

The first exposition of optimal design in its modern form is Kiefer (1959), although the subject goes back to Smith (1918) (see Chapter 1, especially Section 1.9.3). Book length treatments include Fedorov (1972), Pázman (1986) and Pukelsheim (1993). The focus of Silvey (1980) and Fedorov and Hackl (1997) is on the mathematical theory; Atkinson et al. (2007) and Berger and Wong (2009) are more concerned with applications, whereas Goos and Jones (2011) introduce theory and examples via the JMP software. Pronzato and Pázman (2013) present the theory of optimal design for nonlinear models, whilst Fedorov and Leonov (2014) illustrate their theory with pharmaceutical applications, particularly dose finding.

As in Chapter 2, Section 2.3.1.2, where interest was in linear models, we let an experimental design ξ place a fraction w_i of the experimental trials at the conditions x_i . A design with t points of support is written as

$$\xi = \left\{ \begin{array}{cccc} x_1 & x_2 & \dots & x_t \\ w_1 & w_2 & \dots & w_t \end{array} \right\}, \quad (13.17)$$

where $w_i > 0$ and $\sum_{i=1}^t w_i = 1$. There are thus two sets of weights: w_i which determine the proportion of experimentation at x_i and $u(x_i)$ which, from (13.2), partially determine the amount of information from observations at that point. Any realizable experimental design for a total of n trials will require that the weights are ratios of integers, that is, $w_i = r_i/n$, where r_i is the number of replicates at condition x_i . Such designs are called *exact* and are labelled d_n . The mathematics of finding optimal experimental designs and demonstrating their properties is greatly simplified by consideration of *continuous* designs in which the integer restriction is ignored.

The resulting design is a list of conditions under which observations are to be taken. The order in which the observations are made is usually also important; for example, the order should be randomized subject to any restrictions imposed by, for example, blocking factors (Section 13.6.2).

Optimal experimental designs minimize some measure of uncertainty of the parameter estimators, typically a function of the information matrix. They require the specification of the following:

1. A model, or set of models, of interest. For GLMs, the specifications will include a set of parameter values $\theta \in \Theta$, perhaps with an accompanying prior distribution $p(\theta)$.
2. A design criterion, for example, the minimization of a function of one or more information matrices.
3. A design region $\mathcal{X} \subseteq \mathfrak{R}^k$ to which the x_i belong.

The information matrix for the design ξ with t support points is, from (13.2),

$$M(\xi; \theta) = \sum_{i=1}^t w_i M(x_i; \theta) = \sum_{i=1}^t w_i u(x_i) f(x_i) f'(x_i). \tag{13.18}$$

As we are concerned with GLMs, the parameter values enter only through the GLM weights $u(x_i)$.

For continuous designs, we consider minimization of the general measure of imprecision $\Psi\{M(\xi; \theta)\}$. Under very mild assumptions, the most important of which are the compactness of \mathcal{X} and the convexity and differentiability of Ψ , designs that minimize Ψ also satisfy a second criterion. The relationship between these two provides a *general equivalence theorem*, one form of which was introduced by Kiefer and Wolfowitz (1960). See Chapter 2, Section 2.3 for a discussion of such theorems for linear models.

Let the measure $\bar{\xi}$ put unit mass at point x and let the measure ξ^α be given by

$$\xi^\alpha = (1 - \alpha)\xi + \alpha\bar{\xi} \quad (0 \leq \alpha \leq 1). \tag{13.19}$$

Then, from (13.18),

$$M(\xi^\alpha; \theta) = (1 - \alpha)M(\xi; \theta) + \alpha M(\bar{\xi}; \theta). \tag{13.20}$$

Accordingly, the derivative of Ψ in the direction $\bar{\xi}$ is

$$\psi(x, \xi; \theta) = \lim_{\alpha \rightarrow 0^+} \frac{1}{\alpha} [\Psi\{(1 - \alpha)M(\xi; \theta) + \alpha M(\bar{\xi}; \theta)\} - \Psi\{M(\xi; \theta)\}]. \tag{13.21}$$

The values of $M(\xi; \theta)$ and of the Fréchet derivative $\psi(x, \xi; \theta)$ depend on the parameter value θ , as may the design ξ^* minimizing $\Psi\{M(\xi; \theta)\}$. Variation of ξ^* with θ is one of the themes of this chapter. In order to incorporate uncertainty in θ , we define (pseudo) Bayesian criteria and state the equivalence theorem in a form that incorporates the prior distribution $p(\theta)$ through use of

$$\Psi\{\xi\} = E_\theta \Psi\{M(\xi; \theta)\} = \int_\Theta \Psi\{M(\xi; \theta)\} p(\theta) d\theta, \tag{13.22}$$

and

$$\psi(\mathbf{x}, \xi) = E_{\theta} \psi(\mathbf{x}, \xi; \theta). \quad (13.23)$$

The general equivalence theorem states the equivalence of the following three conditions on ξ^* :

1. The design ξ^* minimizes $\Psi(\xi)$.
2. The design ξ^* maximizes the minimum over \mathcal{X} of $\psi(\mathbf{x}, \xi)$.
3. The minimum over \mathcal{X} of $\psi(\mathbf{x}, \xi^*)$ is equal to zero, this minimum occurring at the points of support of the design.

As a consequence of 3, we obtain the further condition:

4. For any non-optimal design ξ , the minimum over \mathcal{X} of $\psi(\mathbf{x}, \xi) < 0$.

The proof of this theorem follows Whittle (1973). See Chaloner and Larntz (1989) and Woods and Lewis (2011) for general equivalence theorems developed for, and applied to, optimal designs for GLMs.

As we illustrate later, the theorem provides methods for the construction and checking of designs. However, it says nothing about t , the number of support points of the design. If $p(\theta)$ is a point prior, putting all mass at a single value θ_0 , the designs are called *locally optimal*. A bound on t can then be obtained from the nature of $M(\xi; \theta_0)$, which is a symmetric $p \times p$ matrix. Due to the additive nature of information matrices (13.18), it follows from Carathéodory's theorem that the information matrix of a design can be represented as a weighted sum of, at most, $p(p + 1)/2 + 1$ rank-one information matrices. The maximum number of support points is therefore $p(p + 1)/2 + 1$. A careful discussion is given by Pukelsheim (1993, Section 8.3), with a shorter treatment by Fedorov and Leonov (2014, Section 2.4.1) who state a simple and usually satisfied condition under which the maximum number of support points reduces to $p(p + 1)/2$. In the examples that follow, the number of support points of optimal designs is usually appreciably less than this; often as few as p is optimal. Of course, such designs with minimum support provide no means of model checking (see Section 13.7). For more general priors $p(\theta)$, the number of support points may be larger, increasing with the variance of the prior. Atkinson et al. (2007, p. 300) give examples for one-parameter nonlinear normal models in which the optimal designs have up to five support points.

13.3.2 D-Optimal Designs

The most widely used design criterion is that of D-optimality (see Chapter 2) in which $\Psi\{M(\xi; \theta)\} = -\log |M(\xi; \theta)|$, so that the log determinant of the information matrix is to be maximized. Then (13.22) becomes

$$\Psi(\xi) = -E_{\theta} \log |M(\xi; \theta)|. \quad (13.24)$$

This (pseudo) Bayesian D-optimality criterion has been used to find designs for GLMs by Chaloner and Larntz (1989) and Woods et al. (2006), amongst others. See Section 13.3.4

for a comment on the distinction between designs under such criteria and truly Bayesian designs.

Fedorov and Leonov (2014, p. 68 and Section 10.1) provide the mathematical results, including matrix differentiation, required for calculation of the Fréchet derivative (13.21). The expected derivative (13.23) then becomes

$$\begin{aligned} \psi(x, \xi) &= E_{\theta} \left\{ p - u(x)f'(x)M^{-1}(\xi; \theta)f(x) \right\}, \\ &= p - E_{\theta} \left\{ u(x)f'(x)M^{-1}(\xi; \theta)f(x) \right\}. \end{aligned} \tag{13.25}$$

For locally D-optimal designs, the number of support points t may be between p and $p(p + 1)/2 + 1$. If $t = p$, the optimal design weights are $w_i = 1/p$.

13.3.3 Design Efficiency

Efficiencies of designs can be compared through the values of the objective function (13.22). If ξ_0^* is the optimal design for the prior $p_0(\theta)$, comparison is of the values of $\Psi\{\xi_0^*\}$ and of $\Psi\{\xi\}$, where ξ is some other design to be compared, and both expectations are taken over the prior $p_0(\theta)$.

There is a particularly satisfying form of efficiency for D-optimality when $p(\theta)$ is a point prior. Then from (13.24), the locally D-optimal design maximizes $|M(\xi; \theta_0)|$, and the efficiency of the design ξ is

$$\text{Eff}_D = \left\{ \frac{|M(\xi; \theta_0)|}{|M(\xi_0^*; \theta_0)|} \right\}^{1/p}. \tag{13.26}$$

Raising the ratio of determinants to the power $1/p$ results in an efficiency measure which has the dimension of variance; a design with an efficiency of 50% requires twice as many trials as the D-optimal design to give the same precision of estimation. If only one parameter is of interest, (13.26) reduces to comparison of the variances of the estimated parameters under different designs (see Chapter 1, Sections 1.9.1 and 1.9.2 for a discussion of efficiency for wider classes of designs and for functions of parameter estimates including contrasts).

If the prior is not concentrated on a single point, the optimal design has to be found by taking the expectation of the determinant over the prior distribution. Usually this requires a numerical approximation to the value. The efficiency in (13.26) in addition requires calculation of the expectation of the determinant for the design ξ . An informative alternative, which we illustrate in Section 13.5.1 is to look instead at the distribution of efficiencies found by simulation from the prior distribution. This procedure avoids taking expectations, since we calculate (13.26) for each sampled value of θ_0 .

13.3.4 Parameter Uncertainty

In Section 13.4.1, we illustrate the dependence of the locally optimal design on the value of θ for logistic regression with a single explanatory variable. An example for two-variable logistic regression is in Section 13.5.1.5 and for two-variable gamma regression in Section 13.5.2.3. In addition to the Bayesian design of Section 13.3.4.1, which we exemplify in Section 13.5.1.6, we here list some other approaches to parameter uncertainty.

13.3.4.1 Bayesian Designs

Bayesian designs are found to maximize expectations such as (13.22). The ease of calculation depends on the form of the prior $p(\theta)$ and of $\Psi(\cdot)$ as well as, often crucially, on the region of integration Θ . The easiest case is when the θ has a multivariate normal distribution over \mathbb{R}^p , although numerical approximation is still needed. Sometimes a transformation of the parameters is required to achieve this simple structure for multivariate priors. For less amenable cases, a standard solution is to sample from the prior distribution and to use an average objective function. An example for a nonlinear model is in Section 18.5 of Atkinson et al. (2007).

Designs maximizing expectations such as (13.22) ignore the additional effect of the prior information about θ on the information matrix and make no assumption of a Bayesian analysis. The designs are accordingly sometimes called pseudo-Bayesian. A discussion of Bayesian experimental design is given by Chaloner and Verdinelli (1995).

13.3.4.2 Sequential Designs

Where possible, sequential experimentation can provide an efficient strategy in the absence of knowledge of plausible parameter values. The usual steps are as follows:

1. Start with some preliminary information providing an estimate, or guess, of the parameter values. This may lead to either a point prior θ_0 or a prior distribution $p(\theta)$.
2. One or a few trials of the optimal design are executed and analysed. If the new estimate of θ is sufficiently accurate, the process stops. Otherwise, step 2 is repeated for the new estimate, and the process continues until sufficient accuracy is obtained or the experimental resources are exhausted.

An early example, for nonlinear regression, is Box and Hunter (1965), extended by Atkinson et al. (2007, Section 17.7). Dror and Steinberg (2008) developed a Bayesian sequential design methodology for GLMs.

13.3.4.3 Minimax and Maximin Designs

The minimax approach overcomes dependence of designs on the unknown value of θ by finding the best design for the worst case when the parameter θ belongs to a set Θ . A design ξ^* is found for which

$$\Psi\{M(\xi^*)\} = \min_{\xi} \max_{\theta \in \Theta} \Psi\{M(\xi; \theta)\}. \quad (13.27)$$

In (13.27), the criterion $\Psi(\cdot)$ needs to be chosen with care. Suppose interest is in D-optimality, with $\xi^*(\theta')$ the locally D-optimal design for parameter value θ' . Unlike with linear models, the value of $|M\{\xi^*(\theta'); \theta'\}|$ depends on the value of θ' . Accordingly, maximization of the related design efficiencies is often used as a criterion, when the maximin design ξ^* is found to maximize the minimum efficiency:

$$\text{Eff}_D(\xi^*) = \max_{\xi} \min_{\theta' \in \Theta} \text{Eff}_D(\xi; \theta').$$

A potential objection to these designs is that the minimax or maximin design is often close to a combination of locally optimal designs for values of θ at the edges of the parameter space. If a prior distribution is available, such points may have a very low probability; their importance in the design criterion may therefore be overemphasized by the minimax criterion. Providing adequate performance in these unlikely worst-case scenarios may greatly affect overall design performance.

A computational difficulty is that such designs can be hard to find. Numerical procedures are described by Nyquist (2013), Fedorov and Leonov (2014, p. 82 and p.130) and, in greatest detail, by Pronzato and Pázman (2013, Section 9.3). Minimax designs for GLMs have been found by Sitter (1992) and King and Wong (2000). A practical difficulty is that the designs may have an appreciable number of support points, some with very low weights. Of course, approximations to the optimal design with fewer support points can always be evaluated, provided the optimal design has been found.

13.3.4.4 Cluster Designs

Some of the computational issues associated with finding Bayesian or minimax designs can be alleviated through the use of clustering of design points, or other less formal techniques, to find designs that incorporate the overall structure of the set of locally optimal designs. Cluster designs are found by (1) sampling parameter vectors θ from a prior distribution, (2) finding a locally optimal design for each sampled vector, (3) clustering the support points of these designs and (4) forming a new, robust design having equally weighted support points that are the cluster centroids. See Dror and Steinberg (2006). Such methods are particularly effective in reducing computation when the locally optimal designs are available analytically (Russell et al. 2009).

13.3.5 Small Effects

If the effects of the factors are slight, the means of observations at different x_i will be similar, and so will each corresponding model weight $u(x_i)$. The information matrix will then, apart from a scaling factor, be close to the unweighted information matrix $X'X$, and the optimal designs for the weighted and unweighted matrices will be close. Since the designs minimizing functions of $X'X$ are the optimal regression designs, these will be optimal, or near optimal, for GLMs with small effects (Cox 1988). This is in addition to the result of Section 13.2.5 that regression designs are optimal for the gamma model with log link.

13.4 Locally Optimal Designs

13.4.1 Binomial Data: Logistic Regression in One Variable

The logistic model is widely used for dose–response data when the response is binomial. For example, Bliss (1935) gives data from subjecting groups of around 60 beetles to eight different levels of an insecticide. The response is the number of beetles killed. The data are reported in numerous textbooks, including Agresti (2002), Collett (2002), Dobson (2001) and Morgan (1992). The original analysis used a probit model, with more recent analyses preferring a logistic link. Atkinson and Riani (2000, Section 6.14) use a goodness of link test




FIGURE 13.1

Logistic regression on a single explanatory variable. Response functions for $\theta_1 = 0.5, 1,$ and 2 (the steepest curve). The D-optimal design points \bullet are clearly highly dependent on the value of θ_1 .

to argue that a complementary log–log link is preferable; Atkinson et al. (2007, Section 22.4) present locally optimal designs for both models.

With a single explanatory variable and the logistic link,

$$\log \left\{ \frac{\mu}{(1 - \mu)} \right\} = \eta = \theta_0 + \theta_1 x. \tag{13.28}$$

Figure 13.1 shows these response curves for $\theta_0 = 0$ and $\theta_1 = 0.5, 1$ and 2 . As θ_1 increases, so does the rate of increase of the response at $x = 0$.

It is clear that optimal designs for the three sets of parameter values will be rather different: experiments for values of x for which the responses are virtually all 0 or 1 will be uninformative about the values of the parameters. This intuitive result also follows from the weight $u(x)$ in (13.7) which goes to zero for extreme values of μ . Further, this result is related to the phenomenon of separation in the analysis of binary data (see Firth 1993, Woods et al. 2006), where a hyperplane of the design space separates the observed data into zeroes and ones.

We start by finding the locally D-optimal design for the *canonical* case of $\theta = (0, 1)'$. From Section 13.3.2, we know that the D-optimal design will have either two or three support points. A standard way to proceed is to assume that there are two, when, again from Section 13.3.2, each $w_i = 0.5$. Further, with a response symmetrical about zero, the design can also be assumed symmetrical about zero. It is thus simple to calculate the D-optimal design within this class. The equivalence theorem is then used to check whether the design is indeed optimal. In this case, see Figure 13.2, this is the correct form and the D-optimal design for a sufficiently large design region \mathcal{X} is

$$\xi^* = \left\{ \begin{array}{cc} -1.5434 & 1.5434 \\ 0.5 & 0.5 \end{array} \right\}, \tag{13.29}$$

that is, equal weights at two points symmetrical about $x = 0$. At these support points, $\mu = 0.176$ and $1 - 0.176 = 0.824$.

Although designs for other values of the parameters can likewise be found numerically, design problems for a single x can often be solved in a canonical form, yielding a structure




FIGURE 13.2 Logistic regression on a single explanatory variable. Equivalence theorem: derivative functions for three D-optimal designs for different values of θ_1 assessed for $\theta_1 = 1$. Curves labelled by value of θ_1 ; • design points.

for the designs independent of the particular parameter values (see Section 13.5.1.1). The translation into a design for other parameter values depends, of course, on the particular θ .

For the upper design point in (13.29), the linear predictor $\eta = 0 + 1 \times x$ has the value 1.5434, which is the value we need for the optimal design whatever the parameterization. If we solve (13.28) for the η giving this value, the upper support point of the design is given by

$$x^* = \frac{(1.5434 - \theta_0)}{\theta_1}. \tag{13.30}$$

For linear regression, the D-optimal design puts half of the design weight at the two extreme values of \mathcal{X} , whereas, for logistic regression, the design does not span \mathcal{X} , provided the region is sufficiently large. Note that as $\theta_1 \rightarrow 0$, the value of x^* increases without limit. This is an example of the result of Cox (1988) mentioned earlier that for small effects of the variables, the design tends to that for homoscedastic regression. In practice, the design region will not be unlimited, and depending on θ , the optimal design may put equal weight on one boundary point and an internal point or on the two boundary points of \mathcal{X} .

In addition to the plots of the probability of success μ against x for three values of θ_1 , 0.5, 1 and 2, Figure 13.1 also shows the D-optimal design points. For the three parameter values, we obtain design points of ± 3.0863 , ± 1.5434 and ± 0.7717 . The D-efficiencies of each for the other set of parameter values are in Table 13.1. The most striking feature of the table is that the design at ± 3.0863 , that is, for $\theta_1 = 0.5$, has an efficiency of only 5.72% when $\theta_1 = 2$. The explanation of this low efficiency is clear from Figure 13.1; the design point is so extreme that the value of μ when $\theta_1 = 2$ is virtually 0 or 1, so that the experiment yields little information. Clearly we need a compromise design when the parameters are not approximately known. Note that this comparison does not include the value of θ_0 , changes in which will make the design asymmetrical about zero. A design robust to the values of θ typically involves more support points than does the locally optimal design. We give an example in Section 13.5.1.

We conclude by showing how condition 3 of the equivalence theorem in Section 13.3.1 can be used both to check the optimality of designs and to suggest improvements to non-optimal designs.

TABLE 13.1

D-Efficiencies (%) of Designs for One-Variable Logistic Regression as θ_1 Varies with $\theta_0 = 0$

θ_1	Design Point x^* (13.30)		
	3.086	1.5434	0.7717
0.5	100	74.52	41.52
1	57.56	100	74.52
2	5.72	57.56	100

Figure 13.2 plots the derivative functions $\psi(x, \xi)$ in (13.25) for the three designs, assessed for $\theta_1 = 1$. The results therefore relate to the second line of Table 13.1. For $\theta_1 = 1$, the minimum of the function is at zero, the minimum occurring at the points of support of the design. The design is therefore D-optimal.

For the other two designs, the minima are both well below zero. The most central set of design points is for $\theta_1 = 2$. Here, the minima are around ± 2 , indicating that the design points should be more extreme than they are. Likewise, for $\theta_1 = 0.5$, the design points are too extreme, and the minimum of the function at the centre of the design region indicates that the design should be shrunk. An important feature of the plot is that, for all three designs, $\psi(x, \xi) = 0$ at the points of support of the design. In order to check the optimality of designs, it is necessary to search over \mathcal{X} and determine the minimum of the function, rather than just checking at the design points.

Although we have illustrated the use of $\psi(x, \xi)$ in the assessment of designs, it can also be used in their construction. Sequential addition of design points at the minimum of $\psi(x, \xi)$ leads to the D-optimal design. An example of such sequential design construction is in Section 11.2 of Atkinson et al. (2007). A straightforward extension is to the adaptive construction of designs mentioned in Section 13.3.4 where the parameters are re-estimated as observations become available and one, or a few, trials added at the point maximizing the criterion function. For D-optimality and addition of single observations, this is achieved after N trials by addition of the point minimizing $\psi(x, \xi_N)$ (Box and Hunter 1965).

13.5 Optimal Designs for Two or More Explanatory Variables

13.5.1 Binomial Data

13.5.1.1 Induced Design Region

As the information matrix (13.18) is of a weighted form, design for the additive linear predictor,

$$\eta(x) = \theta_0 + \sum_{j=1}^k \theta_j x_j, \quad (13.31)$$

is equivalent to (unweighted) design for the linear model

$$E(y) = \theta_0\sqrt{u(x)} + \sum_{j=1}^k \theta_j\sqrt{u(x)}x_j, = \theta_0z_0 + \sum_{j=1}^k \theta_jz_j, \tag{13.32}$$

where $u(x)$ is defined in (13.3). Hence, the original design region \mathcal{X} can be transformed to the induced design region \mathcal{Z} for the induced variables z_0, \dots, z_k . Clearly, \mathcal{Z} depends on both \mathcal{X} and θ .

Ford et al. (1992) used this relationship with linear model design to provide geometric insight into the structure of designs for single-variable GLMs. With one explanatory variable, \mathcal{X} is a line segment $a \leq x \leq b$. However, because in (13.32) $z_0 = \sqrt{u(x)}$, \mathcal{Z} is of dimension 2. For the models for binomial data of Section 13.2.3 and $x \in \mathbb{R}$, \mathcal{Z} is a closed curve similar to the *design locus* in Figure 2 of Box and Lucas (1959), the exact shape of which will vary with θ . The results of Ford et al. (1992) require the careful enumeration of a large number of criteria and special cases. For D-optimality, they use results on the relationship between D-optimal designs and the ellipse of minimum volume centred at the origin that contains \mathcal{Z} (Silvey and Titterton 1973; Sibson 1974; Silvey 1980).

Ford et al. (1992) are concerned with exponential family linear models with a single explanatory variable. Wu and Stufken (2014) provide results for a single-variable model with a quadratic predictor. Mathematical results for the linear predictor (13.32) with more than one explanatory variable are not generally available. An important limitation on the class of models for which results can be expected comes from the form of (13.32) which excludes interaction and higher-order polynomial terms. We have written $z_j = \sqrt{u(x)}x_j$ and so $z_k = \sqrt{u(x)}x_k$. But $z_jz_k \neq u(x)x_jx_k$.

In Section 13.5.1.2, we compute locally optimal designs for binomial data with linear predictor (13.31) for two variables. Views of the designs in \mathcal{X} and in \mathcal{Z} are quite distinct, but both are informative about the structure of the designs. Those in \mathcal{X} relate the design points to the values of the response, whereas those in \mathcal{Z} show that the design points are at the extremes of the region and, for second-order models, near centres of edges and at the centre of the region. The relationship with response surface designs is clear, even if the mirror is distorted.

13.5.1.2 First-Order Models

The properties of designs for response surface models, that is, with two or more continuous explanatory variables, depend much more on the experimental region than those where there is only one factor.

Although it was assumed in the previous section that the experimental region \mathcal{X} was effectively unbounded, the design was constrained by the weight u to lie in a region in which μ was not too close to zero or one. But with more than one explanatory variable, constraints on the region are necessary. For example, for the two-variable first-order model

$$\log \left\{ \frac{\mu}{(1 - \mu)} \right\} = \eta = \theta_0 + \theta_1x_1 + \theta_2x_2, \tag{13.33}$$

with $\theta' = (\theta_0, \theta_1, \theta_2)$, all design points for which $x_1 + x_2 = 0$ yield a value of 0.5 for μ , however extreme the values of x .

We now explore designs for the linear predictor (13.33) with the logistic link for a variety of parameter values and $\mathcal{X} = [-1, 1]^2$. These and succeeding designs were found by a numerical search with a quasi-Newton algorithm. The constraints to ensure that $x \in \mathcal{X}$ and on the design weights were enforced using the trigonometric transformations listed in Atkinson et al. (2007, Section 9.5).

Four sets of parameter values are given in Table 13.2. D-optimal designs for the sets B1 and B2 are listed in Table 13.3. The parameter values of B1 (0, 1, 1) are closest to zero. The table shows that the design has support at the points of the 2^2 factorial, although the design weights are not quite equal. They are so for the normal theory model and become so for the logistic model as θ_1 and $\theta_2 \rightarrow 0$ with $\theta_0 > 0$. At those factorial points for which $x_1 + x_2 = 0, \mu = 0.5$ since $\theta_1 = \theta_2$. At the other design points, $\mu = 0.119$ and 0.881 , slightly more extreme values than the values of 0.176 and 0.824 for the experiment with a single variable.

An interesting feature of our example is that the number of support points of the design depends upon the values of the parameter θ . From the discussion of Carathéodory's theorem in Section 13.3.1, the maximum number of support points required by an optimal design is usually $p(p + 1)/2$ (Pukelsheim 1993, Section 8.3). Our second set of parameters, B2 in which $\theta' = (0, 2, 2)$, gives two four-point optimal designs, with weights given by $w_i^{(1)}$ and $w_i^{(2)}$ in Table 13.3 and support points where $\mu = 0.146, 0.5$ and 0.854 . Any convex combination of these two designs, $\alpha w_i^{(1)} + (1 - \alpha)w_i^{(2)}$ with $0 \leq \alpha \leq 1$, will also be optimal, and will have six support points, which is the value of the usual bound when $p = 3$. These two component designs arise from the symmetry of the design problem; not only does $\theta_1 = \theta_2$ but also the design region is symmetrical in x_1 and x_2 .

TABLE 13.2
Sets of Parameter Values for First-Order Linear Predictors in Two Variables with the Logistic Link

Set	θ_0	θ_1	θ_2
B1	0	1	1
B2	0	2	2
B3	2	2	2
B4	2.5	2	2

TABLE 13.3
D-Optimal Designs for Logistic Models with the Sets of Parameter Values B1 and B2 of Table 13.2; w_i Design Weights

i	Design for B1				Design for B2				
	x_{1i}	x_{2i}	w_i	μ_i	x_{1i}	x_{2i}	$w_i^{(1)}$	$w_i^{(2)}$	μ_i
1	-1	-1	0.204	0.119	0.1178	-1.0000		0.240	0.146
2					1.0000	-0.1178		0.240	0.854
3	1	-1	0.296	0.500	1.0000	-1.0000	0.327	0.193	0.500
4	-1	1	0.296	0.500	-1.0000	1.0000	0.193	0.327	0.500
5					-1.0000	0.1178		0.240	0.146
6	1	1	0.204	0.881	-0.1178	1.0000		0.240	0.854

TABLE 13.4

D-Optimal Designs for Logistic Models with the Parameter Values B3 and B4 of Table 13.2; w_i Design Weights

i	Parameter Set B3				Parameter Set B4			
	x_{1i}	x_{2i}	w_i	μ_i	x_{1i}	x_{2i}	w_i	μ_i
1	-1.0000	-0.7370	0.169	0.186	-1.0000	0.5309	0.333	0.827
2	-1.0000	0.7370	0.331	0.814	-1.0000	-1.0000	0.333	0.182
3	-0.7370	-1.0000	0.169	0.186	0.5309	-1.0000	0.333	0.827
4	0.7370	-1.0000	0.331	0.814				

The D-optimal designs for the two remaining sets of parameters in Table 13.2 are given in Table 13.4. These designs have, respectively, 4 and 3 points of support. When $\theta' = (2, 2, 2)$, the design points are where $\mu = 0.186$ and 0.814 . For $\theta' = (2.5, 2, 2)$, the minimum value of μ is 0.182 at $(-1, -1)$, and the experimental values of μ are 0.182 and 0.827 . For this three-point design for a three-parameter model, the design weights $w_i = 1/3$. A useful general indication is that an informative experiment should have $0.15 < \mu < 0.85$. This bound is included in the plots of designs in Figure 13.3.

The relationship between the design points and the values of μ are shown, for parameter values B2 and B3, in Figure 13.3. For $\theta' = (0, 2, 2)$, one four-point design for B2 is depicted by open circles and the other by filled circles; the symmetry of the designs is evident. For $\theta' = (2, 2, 2)$, there are again four support points of the design, which now lie somewhat away from the boundaries of the regions of high and low values of μ .

13.5.1.3 Induced Design Region for the Logistic Link

For the first-order model in $k = 2$ factors (13.33), for which $p = 3$, the induced design space \mathcal{Z} is of dimension three. Two examples, projected onto z_1 and z_2 and thus ignoring $z_0 = \sqrt{u(x)}$, are given in Figure 13.4 for \mathcal{X} the unit square. In the left-hand panel of the




FIGURE 13.3 Support points for D-optimal designs for logistic models with parameter values B2 and B3 in Table 13.2. In the lightly shaded area, $\mu \leq 0.15$, whereas in the darker region, $\mu \geq 0.85$. In the left-hand panel, one four-point design for B2 is depicted by open circles and the other by filled circles. The symmetry of the designs is evident.




FIGURE 13.4

Support points for D-optimal designs for logistic models with parameter values B2 (a) and B3 (b) of Table 13.2 in the induced design region \mathcal{Z} . For these first-order models, all design points lie on the boundary of \mathcal{Z} .

figure, $\theta' = (0, 2, 2)$ so that at the corner of \mathcal{X} for which $x_1 = x_2 = 1$, $\eta = 4$ and $\mu = 0.982$. This is well beyond the range for informative experiments, and the projection of the induced design space appears to be folded over. As a consequence, experiments at extreme positions in \mathcal{Z} are not at extreme points in \mathcal{X} . The results in the other panel for $\theta' = (2, 2, 2)$ are similar but more extreme. For both sets of parameter values, the design points lie, as they should, on the boundary of the induced design region.

These examples show the importance of both the design region and the value of μ in determining the optimal design. In order to reveal the structure of the designs as clearly as possible, the designs considered have all had $\theta_1 = \theta_2$ and so are symmetrical in x_1 and x_2 . When $\theta_1 \neq \theta_2$, both the design region and the values of μ are important in the resulting asymmetrical designs. Asymmetrical designs also arise when the log-log and complementary log-log links are used, since these links are not symmetrical functions.

13.5.1.4 Theoretical Results for a First-Order Predictor

Various authors have derived optimal designs for $k > 1$ for first-order linear predictor (13.31) for some special cases; see Sitter and Torsney (1995a,b) and Torsney and Gunduz (2001). Yang et al. (2011) considered the case when the assumption of a bounded design region \mathcal{X} is relaxed by allowing one variable, say x_k , to take any value in \mathbb{R} . As the D-optimality objective function is still bounded in this case, the authors were able to provide a methodology to obtain a locally D-optimal design analytically. We restate their corollary 1 for the special case of logistic regression and D-optimality.

Theorem 13.1 (Yang et al. 2011) *For the logistic model with linear predictor (13.31) and $\mathcal{X} = [-1, 1]^{k-1} \times \mathbb{R}$, a D-optimal design has 2^k equally weighted support points,*

$$x_l^* = \begin{cases} (x_{1l}, \dots, x_{(k-1)l}, a_l^*) & \text{for } l = 1, \dots, 2^{k-1} \\ (x_{1l}, \dots, x_{(k-1)l}, -a_l^*) & \text{for } l = 2^{k-1} + 1, \dots, 2^k \end{cases}$$

where

$$x_{jl} = \begin{cases} -1 & \text{if } \lceil \frac{l}{2^{k-1-j}} \rceil \text{ is odd} \\ 1 & \text{if } \lceil \frac{l}{2^{k-1-j}} \rceil \text{ is even} \end{cases} \quad j = 1, \dots, k - 1,$$

$\lceil a \rceil$ is the smallest integer greater than or equal to a , the numerator of the fractions is l and $a_l^* = \eta^* - \eta_l^*(k)$. Here, η^* maximizes

$$\eta^2 \left\{ \frac{dh(\eta)/d\eta}{h(\eta)[1 - h(\eta)]} \right\}^{k+1},$$

where $h = g^{-1}$, the inverse of the logistic link function, and $\eta_l^*(k) = \theta_0 + \sum_{j=1}^{k-1} \theta_j x_{jl}$.

This result has a fairly straightforward interpretation. If we fix the values of $k - 1$ variables in the support points at the level combinations of a 2^{k-1} factorial design, then the selection of an optimal design reduces to a one-variable problem (the choice of the values to be taken by x_k). Note that \mathcal{X} is such that each variable lies in an interval, rather than just taking the two values ± 1 .

To illustrate this result, Table 13.5 gives D-optimal designs on $[-1, 1] \times \mathbb{R}$ for the sets of parameter values in Table 13.2. These designs are quite different from the designs in Tables 13.3 and 13.4, where the restriction of x_2 to the interval $[-1, 1]$ results in the designs having different numbers of support points, different values for the support points and different weights.

TABLE 13.5

D-Optimal Support Points from Theorem 13.1 (Yang et al. 2011) for Parameter Values in Table 13.2

Parameters	Support Points	
	x_1	x_2
B1	-1	-2.2229
	-1	2.2229
	+1	-0.2229
	+1	0.2229
B2	-1	1.6115
	-1	-1.6115
	+1	-0.3886
	+1	0.3886
B3	-1	0.6115
	-1	-0.6115
	+1	-1.3886
	+1	1.3886
B4	-1	-0.3615
	-1	0.3615
	+1	-1.6386
	+1	1.6386

Note: For each design, the support points are equally weighted.

13.5.1.5 Second-Order Response Surface

This section extends the results of earlier sections to the second-order response surface model, again with two factors and again with the logistic link. The purpose is to show the relationship with, and differences from, designs for regression models. The D-optimal designs are found, as before, by minimizing $-\log |\mathcal{M}(\xi; \theta)|$. With $n = 9$ and 2D x , the numerical search is in 26 dimensions. However, once the structure of the designs is established, with 8 of the design points on the edges of \mathcal{X} (see Figure 13.5), the search can be reduced to 18 dimensions, with the equivalence theorem providing a check on this reduction.

To explore how the design changes with the parameters of the model, we look at a series of designs for the family of linear predictors

$$\eta = \theta_0 + \gamma(\theta_1x_1 + \theta_2x_2 + \theta_{12}x_1x_2 + \theta_{11}x_1^2 + \theta_{22}x_2^2) \quad \text{with } \gamma \geq 0, \tag{13.34}$$

and design region $\mathcal{X} = [-1, 1]^2$. The parameter γ provides a family of similarly shaped linear predictors which increasingly depart, in a proportional way, from a constant value as γ increases. When $\gamma = 0$, the result of Cox (1988) shows that the design is the D-optimal design for the second-order regression model, the unequally weighted 3^2 factorial design given in Section 13.5.2.

For numerical exploration, we take $\theta_0 = 1, \theta_1 = 2, \theta_2 = 2, \theta_{12} = -1, \theta_{11} = -1.5$ and $\theta_{22} = 1.5$. As γ varies from 0 to 2, the shape of the response surface becomes increasingly complicated.

Figure 13.5 shows the support points of the D-optimal designs as γ increases from zero in steps of 0.1. The design points are labelled, for $\gamma = 0$, in standard order for the 3^2 factorial, with x_1 changing more frequently. The figure shows how all but one of the design points stay on the boundary of the design region; the circles and black dots are the support




FIGURE 13.5 D-optimal designs for second-order logistic model as γ varies ($0 \leq \gamma \leq 2$). Support points: numbers, $\gamma = 0$; circles, $\gamma = 1$; black dots, $\gamma = 2$; grey dots, intermediate values.

points for $\gamma = 1$ and 2, respectively, with the grey dots indicating intermediate values. There is little change in the location of the centre point, point 5, over this range of values for γ . Initially, the design has nine points, but the weight on point 8 decreases to zero when $\gamma = 0.3$. Thereafter, the design has eight support points until $\gamma = 1.4$ when the weight on observation 6 becomes zero.

Figure 13.6 serves to help interpret the behaviour of the design as γ increases, showing the values of x_2 for the three design points (3, 6 and 9) in Figure 13.5 for which $x_1 = 1$. Initially, the values of x_2 are those for the 3^2 factorial, and they remain virtually so until $\gamma = 0.6$. Thereafter, they gradually converge towards three more central values.

The relationship between the support points of the design and the values of μ is highlighted in Figure 13.7 where, as in Figure 13.3, the pale areas are regions in which $\mu \leq 0.15$, with the dark regions as the complementary ones where $\mu \geq 0.85$. The left-hand panel of




FIGURE 13.6 D-optimal designs for second-order logistic model. The values of x_2 for support points 3, 6 and 9 of Figure 13.5 as γ varies between zero and two. The coding of the symbols is the same as in the earlier figure.




FIGURE 13.7 Support points for D-optimal designs for the second-order logistic model. (a) $\gamma = 1$, (b) $\gamma = 2$. In the lightly shaded area, $\mu \leq 0.15$, whereas in the darker region, $\mu \geq 0.85$.

Figure 13.7, for $\gamma = 1$, shows that the 8-point design is a distortion of a standard response surface design, with most points in the white area and the remainder on the boundary of the design region, some close to the contours of $\mu = 0.15$ or 0.85 . In the numbering of Figure 13.5, points 2 and 6 are on the edge of the design region where μ is close to 0.5 . Points 3 and 9 are at higher values of μ .

A similar pattern is clear in the seven-point design for $\gamma = 2$ in the right-hand panel of the figure; four of the seven points are on the edges of the white region, one is in the centre and only points 3 and 9 are at more extreme values of μ .

The two panels of Figure 13.7 taken together explain the trajectories of the points in Figure 13.5 as γ varies. For example, points 1 and 4 move away from $(-1, -1)$ as the value of μ at that point decreases, point 3 remains at $(1, -1)$ until γ is close to one and point 8 at $(0, 1)$ is rapidly eliminated from the design as the value of μ there increases with γ .

Further insight into the structure of the designs can be obtained from consideration of the induced design region introduced in Section 13.5.1. Although, as stated earlier, extension of the procedure based on (13.32) to second-order models such as (13.34) is not obvious, it is still informative to look at the plot of the designs in \mathcal{Z} space. The left-hand panel of Figure 22.8 of Atkinson et al. (2007) shows the eight-point design for $\gamma = 1$ plotted against z_1 and z_2 ; seven points lie on the edge of this region, well spaced and far from the centre, which is where the eighth point is. The right-hand panel for $\gamma = 2$ shows six points similarly on the edge of \mathcal{Z} ; the centre point is hidden under the seemingly folded-over region near the origin.

In the induced design region, these designs are reminiscent of response surface designs, with a support point at the centre of the region and others at remote points. However, the form of \mathcal{Z} depends on the unknown parameters of the linear predictor, so this description is not helpful in constructing designs. In the original space \mathcal{X} , we have described the designs for this second-order model as a series of progressive distortions of designs with support at the points of the 3^2 factorial. For small values of γ , the unweighted 3^2 factorial provides an efficient design, with a D-efficiency of 97.4% when $\gamma = 0$. However, the efficiency of this design declines steadily with γ , being 74.2% for $\gamma = 1$ and a low 38.0% when $\gamma = 2$. If appreciable effects of the factors are expected, the special experimental design methods of this section need to be used. Further discussion of designs for two-variable logistic models is given by Sitter and Torsney (1995a) with particular emphasis on the structure of \mathcal{Z} .

13.5.1.6 Bayesian D-Optimal Designs

We can also find Bayesian D-optimal designs, maximizing (13.24), for response surface designs and binomial data. Motivated by a food technology example, Woods et al. (2006) found designs for logistic regression with linear predictor

$$\eta(\mathbf{x}) = \theta_0 + \sum_{i=1}^3 \theta_i x_i + \sum_{i=1}^3 \sum_{j \geq i}^3 \theta_{ij} x_i x_j,$$

with $x_i \in [-1.2782, 1.2782]$. Here, we find a Bayesian D-optimal design assuming independent uniform prior densities for the parameters in θ , defined on the support

$$\theta_1, \theta_2 \in [2, 6], \quad \theta_0, \theta_3, \theta_{ij} \in [-2, 2] \quad \text{for } i, j = 1, 2, 3.$$




FIGURE 13.8 Two-dimensional projections of the Bayesian D-optimal design for a logistic regression model (lower diagonal) and a CCD (upper diagonal).

We approximate the expectation (13.24) by the sample average across a 20-run Latin hypercube sample (see, e.g., Chapter 17 and Santner et al. 2003, Chapter 6). Gotwalt et al. (2009) and Woods and van de Ven (2011) discuss and compare some alternative approaches for this approximation.

For this example, a simulated annealing algorithm (Woods 2010) was used to find an exact design, d_{16} , with $n = 16$ points. The design is given in Figure 13.8 and, in fact, has $t = n = 16$ support points and no replication. For reference, Figure 13.8 also gives the design points of a 16-run central composite design (CCD; see Chapter 5), d_{ccd} , with 8 factorial points, 6 axial points with $x_j = \pm 1.2782$ and two centre points; see, for example, Box and Draper (2007). This design is standard for normal theory response surface studies and is an obvious comparator for the Bayesian GLM design. In fact, a CCD had been employed for the food technology example in previous experimentation.

Whilst there are some familiar features to the Bayesian D-optimal design, including (near) centre points, there are also some distinct differences from the CCD. These include the presence of extreme corner points in the D-optimal design and, for x_1 and x_2 , fewer distinct levels (~ 3 for each of these two variables).

The relative performance of the designs was assessed via a simulation study. A sample of 1000 parameter vectors, $\theta^{(l)}$, was drawn from the prior distribution for θ . For each vector, we calculated the relative D-efficiency

$$Eff_D(l) = \left\{ \frac{|M(d_{16}); \theta^{(l)}|}{|M(d_{ccd}); \theta^{(l)}|} \right\}^{\frac{1}{10}}, \quad l = 1, \dots, 1000. \tag{13.35}$$

The empirical cumulative distribution of the relative efficiency, Figure 13.9, shows a dramatic difference in performance between the two designs. The Bayesian D-optimal design is more efficient than the CCD for about 85% of the sampled parameter vectors




FIGURE 13.9

Empirical cumulative distribution function of the relative D-efficiency of the Bayesian D-optimal design compared to the CCD.

and has efficiency of 1.75 or more for 50% of the sampled vectors. The maximum relative efficiency is close to 7.

13.5.2 Gamma Data

13.5.2.1 Theoretical Results for First-Order Designs with the Power Link

We use numerical examples to illustrate the great difference between designs for gamma models and those for binomial responses of Section 13.5.1. The examples are calculated using the power link. In this introductory section, we present the theoretical results of Burrige and Sebastiani (1994) for first-order models. Our numerical example of a first-order model is in Section 13.5.2 with an example for a response surface model in Section 13.5.2. We conclude with a brief comparison of designs for the power link with those for the Box–Cox link.

If the design region for uncoded variables is of the form $a_j \leq x_j^u \leq b_j$, the region can be coded so that $0 \leq x_j \leq 1$ for $j = 1, \dots, k$. The requirement that $\mu > 0$ for all non-negative x leads to a canonical form of the original problem with $\theta_j \geq 0$ for all $j = 0, \dots, k$, with at least one inequality. Since the weights (13.16) are monotonic in η , the support points of D-optimal designs as θ varies must then be some of the points of the 2^k factorial. Which points have non-zero weights depends on the values of the θ_j . For effects large relative to θ_0 , Burrige and Sebastiani (1994) provide the following theorem.

Theorem 13.2 (Burrige and Sebastiani 1994) *For the coded design variables x_j , the design which puts weights $1/(k+1)$ at each of the $k+1$ points*

$$(0, \dots, 0)', (1, 0, \dots, 0)', (0, 1, 0, \dots, 0)', \dots, (0, \dots, 0, 1)'$$

is D -optimal for gamma regression, the power link and a first-order linear predictor if, and only if, for all $i, j = 1, \dots, k$,

$$\theta_0^2 \leq \theta_i \theta_j.$$

Thus, for large effects, a *one-factor-at-a-time* approach is optimal. However, as the effects become smaller, the design approaches the 2^k factorial in line with the result of Cox (1988) discussed in Section 13.3.5. Of course, the weights w_i have to be found numerically. However, the numerical search is greatly simplified by being restricted to the support of the 2^k factorial. It is also a great simplification that, as we showed in Section 13.2.5, the designs do not depend on the value of the power κ .

13.5.2.2 Examples of First-Order Designs with the Power Link

In both examples, there are two explanatory variables. We only look at symmetrical designs generated with $\theta_1 = \theta_2$ having the three values 0.1, 0.5 and 1. In all calculations, $\theta_0 = 1$. The resulting designs are in Table 13.6.

These results nicely illustrate the theoretical results of Section 13.5.2. We have parameterized the problem with $\theta_1 = \theta_2 = \chi$. For $\chi = 0.1$, that is, with small effects, the design is virtually the 2^2 factorial, with weights close to 1/4 ranging from 0.225 to 0.271. Increasing χ to 0.5 leaves the support points unchanged, but now the weights range, symmetrically of course in x_1 and x_2 , from 0.125 to 0.3125, with the lowest weight on (1,1). When $\chi = 1$, we have $\theta_0^2 = \theta_1 \theta_2$, so that we are at the lowest value of χ for which we obtain a design with three support points. All weights are, of course, equal.

These designs were found numerically using a quasi-Newton algorithm combined with the transformations given in Section 9.5 of Atkinson et al. (2007). The general equivalence theorem was used to check the designs by evaluation of the derivative function $\psi(x, \xi)$ over a fine grid in \mathcal{X} . One point in the construction of these designs is that with $\chi = 1$, the optimization algorithm had not quite converged to the theoretical value after the default limit of 100 iterations, whereas around 10 iterations were needed for the other values of χ . The effect on the minimum value of $\psi(x, \xi)$ was negligible. A second point is that, to five significant values, the weights for $\chi = 0.5$ were exactly 5/16, 9/32, 9/32 and 1/8. Such simple weights can be an indication that theoretical results are possible. See Atkinson (2010) and Dette et al. (2012) for an example in discrimination between polynomial regression models.

TABLE 13.6
 D-Optimal Designs for Two-Variable First-Order Model for Gamma Responses with the Power Link; $\theta = (1, \chi, \chi)'$

x_1	x_2	Design Weights w_i		
		$\chi = 0.1$	$\chi = 0.5$	$\chi = 1$
0	0	0.271	0.313	1/3
0	1	0.252	0.281	1/3
1	0	0.252	0.281	1/3
1	1	0.225	0.125	0




FIGURE 13.10

Support points in the induced design region \mathcal{Z} of D-optimal designs for two-variable first-order model for gamma responses with the power link; • design points.

To conclude our discussion of this example, we look at the plots of the design in the induced design region \mathcal{Z} defined in (13.32). As the results of Burrige and Sebastiani (1994) show, the boundary of \mathcal{Z} is formed, for $k = 2$, by straight line segments, as is shown in Figure 13.10. There is none of the curving over of the space of \mathcal{X} that is caused by the nonlinear nature of the GLM weights $u(x_i)$ for the binomial distribution that is evident in Figure 13.4.

For small χ , the induced design region is virtually square, becoming less so as χ increases. The left-hand panel of Figure 13.10 is for $\chi = 0.5$, for which the weight at $x = (1, 1)'$ is $1/8$. As χ increases, the weight on this value decreases. Insight about the case $\chi = 1$ comes from the results of Silvey and Titterington (1973) relating D-optimality to minimum volume ellipsoids enclosing design regions. It follows that when $\chi = 1$, the values of z_1 and z_2 at $x = (1, 1)'$, say $z_1(1, 1)$ and $z_2(1, 1)$, must be the same distance from the origin as $z_1(1, 0)$ and $z_2(1, 0)$ (or $z_1(0, 1)$ and $z_2(0, 1)$). Hence, $z_1(1, 1) = z_2(1, 1) = (\sqrt{2})/4$. For larger values of χ , the values of z_1 and z_2 lie inside the circle, and \mathcal{Z} becomes increasingly triangular. The design does not change as χ increases above 1.

In these calculations, we have taken $\kappa = 1$. From the form of $u(x)$ in (13.16), other values of κ lead to identical figures, but with different numerical values on the axes.

13.5.2.3 Second-Order Response Surface with the Power Link

Atkinson and Riani (2000, Section 6.9) use a gamma model to analyse data from Nelson (1981) on the degradation of insulation due to elevated temperature at a series of times. A second-order model is required in the two continuous variables, and a gamma model fits well with a power link with $\kappa = 0.5$. We scale the variables to obtain design region $\mathcal{X} = [-1, 1]^2$. The linear predictor is the quadratic

$$\eta = \theta_0 + \theta_1x_1 + \theta_2x_2 + \theta_{11}x_1^2 + \theta_{22}x_2^2 + \theta_{12}x_1x_2, \tag{13.36}$$

that is, (13.34) with $\gamma = 1$. Then the standard D-optimal design for the normal theory regression model has unequally weighted support at the points of the 3^2 factorial: weight

TABLE 13.7

Parameter Values for Linear Predictor (13.36) with $\kappa = 0.5$

Parameter Set	θ_0	θ_1	θ_2	θ_{11}	θ_{22}	θ_{12}
G1	3.7	-0.46	-0.65	-0.19	-0.45	-0.57
G2	3.7	-0.23	-0.325	-0.095	-0.225	-0.285

0.1458 at the four corners of the design region, 0.0802 at the centre points of the sides and 0.0960 at the centre of the region. This design is optimal for the gamma model with log link and for the model with the power link as the parameters in (13.36), apart from θ_0 , become small. We take θ to have the values given in Table 13.7, G1 being rounded from an analysis of Nelson’s data.

The exact optimal 9-trial design for G1, found by searching over a grid of candidates with steps of 0.01 in x_1 and x_2 , is in Table 13.8. This shows that, at the points of the design, the minimum value of μ is 1.90 and the maximum 14.59. The parameter values are thus such that we satisfy the requirement $\mu > 0$.

As the left-hand half of Table 13.8 shows, the design has seven support points. The points of the 2^2 factorial are in the upper part of the table. All are included in the design, two being replicated. The other three support points are slight distortions of some remaining points of the support of the 3^2 factorial. Figure 13.11 makes clear the virtually symmetrical nature of the design, although the parameters are not quite symmetrical in value for x_1 and x_2 .

To illustrate the approach of the design to the 3^2 factorial as the parameter values decrease, we also found the D-optimal 9-point design for the set of parameter values G2 in Table 13.7 in which all parameters, other than θ_0 , have half the values they have for design G1. As Table 13.8 shows, the range of means at the design points is now 6.45 – 14.71, an appreciable reduction in the ratio of largest to smallest response. The support points of the design for G2 are shown in Figure 13.11 by the symbol X. There are now nine distinct support points close to those of the 3^2 factorial. For G2, the three design points in the lower half of the table for G1 are moved in the direction of the full factorial design. For linear regression, the unweighted 3^2 factorial is D-optimal.

TABLE 13.8

Exact D-Optimal Designs for the Parameter Sets G1 and G2 of Table 13.7

i	Design for G1				Design for G2			
	x_{1i}	x_{2i}	n_i	μ_i	x_{1i}	x_{2i}	n_i	μ_i
1	-1.00	-1.00	1	12.96	-1.00	-1.00	1	13.32
2	-1.00	1.00	2	11.83	-1.00	1.00	1	12.74
3	1.00	-1.00	2	14.59	1.00	-1.00	1	14.14
4	1.00	1.00	1	1.90	1.00	1.00	1	6.45
5	0.11	0.15	1	12.46	-1.00	0.00	1	14.71
6	0.26	1.00	1	5.38	-0.01	-1.00	1	14.44
7	1.00	0.29	1	7.07	0.07	0.09	1	13.33
8					0.08	1.00	1	9.66
9					1.00	0.09	1	11.01




FIGURE 13.11

Points for D-optimal 9-point designs for gamma models in Table 13.8: +, the points of the 3² factorial; o, G1 and x, G2. Points for G1 which are replicated twice are darker.

13.5.2.4 Efficient Standard Designs for Gamma Models

We conclude our analysis of designs for the gamma model with second-order predictor and a power link by briefly exploring how well the unweighted 3² factorial performs when there are nine trials by comparing both it and the design for G2 with that for G1.

The D-optimal design for the less extreme parameter set G2 of Table 13.7 has efficiency 97.32%, whilst the equi-replicated 3² factorial has efficiency 96.35%. The main feature of these designs is how efficient they are for the gamma model, both efficiencies being greater than 95%. The design for parameters G2 is for a model with smaller effects than G1, so that the design and its efficiency are between those for G1 and the factorial design.

An indication of this example with a gamma response is that standard designs may be satisfactory for second-order response surfaces. However, Burrige and Sebastiani (1994) show that, for first-order models, full 2² factorial designs, or their regular fractions (see Chapter 7), can be very inefficient when the effects are strong and the optimal designs have only k + 1 points of support.

13.5.2.5 Designs for the Power Link and for the Box-Cox Link

We now briefly explore the relationship between designs for the power link with weights given by (13.16) and those for the Box-Cox link (13.14). We relate the two through their dependence on the linear predictor η.

From (13.16), the weights for the power link can be written as

{u(x)}^{-0.5} = η, (13.37)

since the constant value of κ is ignored. For the Box-Cox link, on the other hand,

{u(x)}^{-0.5} = μ^λ. (13.38)

However, from (13.13),

$$\mu = (1 + \lambda\eta)^{1/\lambda},$$

so that for the Box–Cox link,

$$\{u(\mathbf{x})\}^{-0.5} = (1 + \lambda\eta) = 1 + \lambda(\theta_0 + \theta_1x_1 + \cdots + \theta_kx_k).$$

The condition in Theorem 13.2 of Section 13.5.2.1 that the one-factor-at-a-time design is optimal therefore becomes

$$(1 + \lambda\theta_0)^2 \leq \lambda^2\theta_i\theta_j. \tag{13.39}$$

An advantage of the Box–Cox link in data analysis is that it is continuous at $\lambda = 0$, becoming the log link. The search over suitable links to describe the data therefore does not contain any discontinuity. In designing experiments, on the other hand, a locally optimal design will be selected for a particular λ . The results of Section 13.2.5 show that, if the power link is used, a value of κ does not have to be stated *a priori*. However, prior values will be required for θ . These will typically be ascertained from guessed responses as the factors vary. Despite the absence of explicit reference to κ in the design criterion, the value will enter implicitly through the relationship between μ and η (13.15). Finally, (13.39) shows that as $\lambda \rightarrow 0$, the one-factor-at-a-time design will not be optimal. Further, from (13.38), it follows that under these conditions, the weights $u(\mathbf{x}) \rightarrow 1$ and the design will tend to the 2^k factorial, even for large values of the θ_j .

13.5.3 Poisson Data

D-optimal designs for Poisson regression with the log link, $\log \mu = \eta$, share some similarities with the gamma designs in Section 13.5.2. In particular, for log linear models with a first-order linear predictor,

$$\log \mu = \eta = \theta_0 + \sum_{i=1}^k \theta_i x_i, \tag{13.40}$$

the optimal design has a similar structure to those from Theorem 13.2.

There are only a moderate number of results on designs for Poisson regression in the literature. For (13.40) and $k = 1$, Minkin (1993) found locally optimal designs for estimating θ_1 ; see also Chapter 14 and the references therein for more general results on models with one variable. For $k = 1, 2$, Wang et al. (2006) investigated the dependence of locally optimal designs on functions of the parameter values, and Wang et al. (2006) developed sequential designs. For a single variable, Ford et al. (1992) used a transformation of the design space to a canonical form, together with geometrical arguments along the lines of Section 13.5.1.1, to find locally optimal designs for a class of nonlinear models that included Poisson regression.

Russell et al. (2009) addressed the problem of D-optimal design for (13.40) with $k \geq 1$ and provided the following theorem.

Theorem 13.3 (Russell et al. 2009) *A D-optimal design for Poisson regression (13.40) with $l_i \leq x_{ij} \leq u_i$ and $|\theta_i(u_i - l_i)| \geq 2$ ($i = 1, \dots, k; j = 1, \dots, t$) has the $t = k + 1$ equally weighted support points,*

$$\begin{aligned}
 x_i &= c - \frac{2}{\theta_i} e_i, \quad i = 1, \dots, k \\
 x_{k+1} &= c,
 \end{aligned}
 \tag{13.41}$$

for e_i the i th column vector of the $k \times k$ identity matrix, $i = 1, \dots, k$, and $c = (c_1, \dots, c_k)'$, where $c_i = u_i$ if $\theta_i > 0$ and $c_i = l_i$ if $\theta_i < 0$.

The proof of Theorem 13.3 is via a canonical transformation and an application of the general equivalence theorem. Note that the D-optimal design does not depend on the value of the intercept, θ_0 , and is invariant to permutation of the factor labels. The requirement $|\theta_i(u_i - l_i)| \geq 2$ is not overly restrictive in practice; $\mathcal{X} = [-1, 1]^k$ requires $|\theta_i| \geq 1, i = 1, \dots, k$. In Figure 13.12, we give the support points for $k = 2$ and a number of example parameter vectors, $\theta = (0, \chi, \chi)'$. Notice the one-factor-at-a-time structure of the design and how the support points tend towards $(1, 1)$ as χ increases. Figure 13.13 gives the support points in the induced design space \mathcal{Z} , projected into z_1, z_2 and defined from (13.32). The optimal support points lie on the boundary of the induced space. Not only do the values of the




FIGURE 13.12

Support points for locally D-optimal designs for Poisson regression with (13.40) and $\theta = (0, \chi, \chi)$. Key: $\circ \chi = 1$; $\triangle \chi = 2$; $+ \chi = 3$; $\times \chi = 4$; $\diamond \chi = 5$. All designs include the point $(1, 1)$.




FIGURE 13.13

Support points in the induced design space \mathcal{Z} for D-optimal designs for Poisson regression with (13.40) and $\theta = (0, \chi, \chi)$. (a) $\chi = 2$; (b) $\chi = 3$; (c) $\chi = 4$; (d) $\chi = 5$.

z_i increase with χ , but also the induced design region itself becomes more elongated as χ increases.

For $|\theta_i(u_i - l_i)| < 2$, numerical investigation has found that both the optimal support points and weights depend on θ_0 , in addition to the other parameters. As expected, as $|\theta_i/\theta_0|$ tends to zero, for $i = 1, \dots, k$, the D-optimal design tends to the equally weighted factorial design.

For more general linear predictors, for example, containing interactions or quadratic terms, numerical search is required to find optimal designs. This is also the case for Bayesian D-optimal designs except for the special case of minimally-supported designs for (13.40), that is, designs with $t = k + 1$ support points. McGree and Eccleston (2012) provided theoretical results for minimally supported designs robust to a set of models of the form (13.40) defined through a discrete set of parameter vectors. We extend their result to Bayesian minimally supported D-optimal designs.

Theorem 13.4 *Assume a Poisson regression model with linear predictor (13.40). The Bayesian D-optimal design amongst the class of minimally supported designs, minimizing (13.24), is the locally D-optimal design (13.41) for the parameter vector $\theta^* = E(\theta)$ provided $|E(\theta_i)(u_i - l_i)| \geq 2$.*

Proof: For a minimally-supported design, the model matrix

$$X = [f(x_1), \dots, f(x_p)]'$$

is $p \times p$. Now the objective function (13.24) can be written as

$$\Psi(\xi) = - \int_{\Theta} \log |\mathbf{M}(\xi; \theta)| p(\theta) d\theta \tag{13.42}$$

$$= - \int_{\Theta} 2 \log |\mathbf{X}| p(\theta) d\theta - \int_{\Theta} \log \prod_{j=1}^t w_j \exp(\eta_j) p(\theta) d\theta$$

$$= -2 \log |\mathbf{X}| - \sum_{j=1}^t \left[\log w_j + \int_{\Theta} \eta_j p(\theta) d\theta \right]$$

$$= -2 \log |\mathbf{X}| - \sum_{j=1}^t \left[\log w_j - \eta_j^* \right]$$

$$= - \log |\mathbf{M}(\xi; \theta^*)|, \tag{13.43}$$

where $\theta^* = E(\theta)$ and $\eta_j^* = \theta_0^* + \sum_{i=1}^k \theta_i^* x_{ij}$. The fourth line above follows as η_j is a linear function of θ . The equality of (13.42) and (13.43) establishes that, provided $|E(\theta_i)(u_i - l_i)| \geq 2$, design (13.41) is Bayesian D-optimal amongst the class of minimally supported designs.

To illustrate Theorem 13.4, we find a Bayesian minimally supported D-optimal design for (13.40) and $k = 5$ factors with $x_i \in [-1, 1]$, $\theta_0 = 0$ and each $\theta_i \sim U(a, b)$ ($i = 1, \dots, 5$). The values of a and b are given in Table 13.9 in terms of a common parameter α . Increasing α leads to more diffuse prior densities. However, for any $\alpha \geq 2$, the Bayesian minimally supported D-optimal design is given by the locally D-optimal design for $\theta_0 = 0$ and $\theta_i = (a + b)/2 = (-1)^{(i+1)}(1 + \alpha/2)$; see Table 13.10.

We assess the performance of these designs through simulation of 10,000 parameter vectors from the uniform distributions defined by Table 13.9 for $\alpha = 2, 5, 10$ and 20. For each parameter vector, we derive the locally D-optimal design from Theorem 13.3 and then calculate the D-efficiency (13.26) for the design in Table 13.10. The induced empirical cumulative distribution functions are given in Figure 13.14.

For relatively precise prior information ($\alpha = 2$), high efficiency is maintained for all the samples from the prior distribution, with minimum efficiency of 79% and median of 93%. As the prior distribution becomes more diffuse (with increasing α), the induced efficiency distribution also becomes more diffuse. For $\alpha = 5$, the minimum and median efficiencies

TABLE 13.9
Bayesian Minimally Supported D-Optimal Design: Ranges for the Uniform, $U(a, b)$, Prior Densities for $\theta_1, \dots, \theta_5$

Parameter	Limits	
	a	b
θ_1	1	$1 + \alpha$
θ_2	$-1 - \alpha$	-1
θ_3	1	$1 + \alpha$
θ_4	$-1 - \alpha$	-1
θ_5	1	$1 + \alpha$

TABLE 13.10

Bayesian Minimally Supported D-Optimal Design: Equally Weighted Support Points; $\beta = [(\alpha - 2)/(\alpha + 2)]$ for $\alpha = 2, 5, 10, 20$

	x_1	x_2	x_3	x_4	x_5
1	β	-1	1	-1	1
2	1	$-\beta$	1	-1	1
3	1	-1	β	-1	1
4	1	-1	1	$-\beta$	1
5	1	-1	1	-1	β
6	1	-1	1	-1	1




FIGURE 13.14

Empirical cumulative distributions for the D-efficiency of the Bayesian minimally supported D-optimal design for four different prior distributions; see Tables 13.9 and 13.10.

are 53% and 85%, respectively; the corresponding summaries for $\alpha = 10$ are 34% and 80% and for $\alpha = 20$ are 21% and 75%. The high median efficiencies maintained for more diffuse distributions are not typical of minimally supported designs in general (e.g., see van de Ven and Woods 2014, for binary data designs). However, it seems the structure of the Poisson designs, with all the support points on the boundary, makes Bayesian minimally supported designs an efficient choice (see also McGree and Eccleston 2012).

13.6 Designs with Dependent Data

There is a variety of practical experiments in which the observed responses may be dependent; see also Chapter 6. Most obviously, and perhaps most importantly, are experiments where there is a natural grouping in the experimental units such that,

given the same treatment, two observations from within the same group are expected to be more similar than two observations in different groups. Common examples for non-normal data include longitudinal studies in clinical and pharmaceutical trials (Diggle et al. 2002) and blocked experiments in industry (Robinson et al. 2004). Our exposition in this section focuses on block designs but is equally appropriate for other grouped experiments.

For experiments involving the observation of a continuous response that is appropriately described by a normal distribution, there is, of course, a wealth of literature on the design and analysis of blocked experiments for linear models (see Section II). There is also a considerable literature on design for nonlinear models and dependent data, much of it stemming from the world of pharmacokinetics/pharmacodynamics (PK/PD) and the seminal work of Mentré et al. (1997). For experiments with discrete data, with a binary or count response, there are rather fewer results available, although the results we describe below naturally share some similarities with the PK/PD literature, particularly the so-called first-order approximation (see Retout and Mentré 2003; Bazzoli et al. 2010).

As with linear models, the first decision is whether to model the block effects as fixed or random. We will choose the latter option and take a mixed model approach as: (1) in industrial experiments, the blocks are usually a nuisance factor and not of interest in themselves; (2) random block effects allow prediction of the response for unobserved blocks and (3) pragmatically, when finding optimal designs for nonlinear models, the use of random block effects reduces the number of unknown parameters for which prior information is required. For example, with b blocks, a fixed-effect design would require the specification of $b - 1$ block effects. See Stufken and Yang (2012a) for locally optimal designs with fixed group effects and a single variable.

13.6.1 Random Intercept Model

To model the responses from a blocked experiment, we adopt the framework of GLMMs (Breslow and Clayton 1993) and, in particular, apply random intercept models. We develop our framework for b blocks, each of equal size m . For the j th unit in the i th block,

$$y_{ij}|\gamma_i \sim \pi(\mu_{ij}), \quad \text{for } i = 1, \dots, b; j = 1, \dots, m,$$

where $\pi(\cdot)$ is a distribution from the exponential family with mean μ_{ij} and

$$g(\mu_{ij}) = f'(x_{ij})\boldsymbol{\theta} + \gamma_i.$$

Here, $g(\cdot)$ is the link function, and x_{ij} is the i, j th combination of variable values. As earlier, the vector $f(x_{ij})$ holds known functions of the k variables, and $\boldsymbol{\theta}$ holds the p unknown regression parameters. The unobservable random block effects γ_i are assumed to follow independent $N(0, \sigma_\gamma^2)$ distributions, with σ_γ^2 known. Under this model, observations in different blocks are independent. More general models including additional random effects, such as random slopes, may also be defined.

13.6.2 Continuous Block Designs

We choose to generalize (13.17) to include blocks of fixed size m through

$$\xi = \left\{ \begin{matrix} \zeta_1 & \dots & \zeta_t \\ w_1 & \dots & w_t \end{matrix} \right\},$$

where $\zeta_l \in \mathcal{X}^m$ is the set of design points that form the l th block (or *support point*) and $0 < w_l \leq 1$ is the corresponding weight; $\sum_{l=1}^t w_l = 1$. See Cheng (1995) and Woods and van de Ven (2011). For example, if $k = m = 2$, a possible continuous design is

$$\xi = \left\{ \begin{matrix} \zeta_1 = \{(-1, -1), (1, 1)\} & \zeta_2 = \{(1, -1), (-1, 1)\} \\ 0.5 & 0.5 \end{matrix} \right\},$$

that is, one-half of the b blocks in a realized exact design would contain design points $x_1 = x_2 = -1$ and $x_1 = x_2 = 1$, and the other half would contain design points $x_1 = 1, x_2 = -1$ and $x_1 = -1, x_2 = 1$.

13.6.3 Information Matrix for a Random Intercept Model

To apply the usual design selection criteria, for example, D-optimality, we need to derive and evaluate the information matrix for θ . As observations in different blocks are independent,

$$M(\xi; \tau) = \sum_{l=1}^t w_l M(\zeta_l; \tau),$$

where $\tau = (\theta', \sigma_\gamma^2)$, and the information matrix for the l th block is, by definition,

$$M(\zeta_l, \tau) = E_{y_l} \left\{ -\frac{\partial^2 \log p(\mathbf{y}_l | \tau, \zeta_l)}{\partial \theta \partial \theta'} \right\}, \tag{13.44}$$

where the m -vector $\mathbf{y}_l = (y_{l1}, \dots, y_{lm})'$ holds the responses from block ζ_l .

Direct calculation of the expectation in (13.44) is possible for small experiments with binary data through

$$M(\zeta_l, \tau) = \sum_{\mathbf{y}_l \in \{0,1\}^m} -\frac{\partial^2 \log p(\mathbf{y}_l | \tau, \zeta_l)}{\partial \theta \partial \theta^T} p(\mathbf{y}_l | \tau, \zeta_l),$$

although both the marginal likelihood and its derivative will require numerical approximation (see Waite, 2013). For more practically sized experiments, this direct calculation will be computationally infeasible.

13.6.4 Approximating the Information Matrix Using Estimating Equations

For model estimation, the marginal likelihood can be approximated using methods such as quadrature, Markov chain Monte Carlo or the EM algorithm (McCulloch et al. 2008). However, for the purposes of finding an optimal design using numerical search, repeated evaluation of the information matrix, or some alternative, is required. Hence, a fast approximation to (13.44) is needed.

An approximate variance–covariance matrix for θ is available from the theory of estimating equations (see, e.g., Godambe 1991). For a GLMM, standard unbiased estimating equations are an extension of the score equations for a GLM and have the form

$$\sum_{l=1}^t w_l X_l' \Delta_l V_l^{-1} (y_l - \mu_l) = 0,$$

where $\Delta_l = \text{diag} [d\mu_{ij}/d\eta_{ij}]$, $\eta_{ij} = f'(x_{ij})\theta$, X_l and μ_l are the $m \times p$ model matrix and $m \times 1$ mean vector defined for the l th block and V_l is a $m \times m$ weight matrix for the observations from the l th block. Depending on the approximation, μ_l may be either the *marginal* mean response or the *conditional* mean response given $\gamma_l = 0$.

The approximate variance–covariance matrix of the resulting estimators is given by

$$\text{Var}(\hat{\theta}) \approx \left(\sum_{l=1}^t w_l X_l' \Delta_l V_l^{-1} \Delta_l X_l \right)^{-1}. \tag{13.45}$$

The inverse of this variance–covariance matrix can be used as an approximation to the information matrix $M(\xi, \tau)$.

Various different choices of V_l and μ_l have been proposed, of which we will discuss the following three:

1. *Quasi-likelihood (QL)*: $V_l = \text{Var}(y_l)$ and $\mu_l = E(y_l)$, the marginal variance and mean; see Wedderburn (1974). The marginal variance–covariance matrix for y_l is generally not available in closed form for non-normal data. One notable exception is for the Poisson distribution, for which Niaparast (2009) and Niaparast and Schwabe (2013) used QL to find D-optimal designs.
2. *Marginal quasi-likelihood (MQL)*: $V_l = \text{diag} [\text{Var}(y_{lj})] + \Delta_l J \Delta_l \sigma_\gamma^2$, with $J = \mathbf{11}'$, and $\mu_l = E(y_l | \gamma_l = 0)$; see Breslow and Clayton (1993). Here, a linear mixed model approximation is used for the marginal variance–covariance matrix of y_l . This approximation has been used to find designs for binary data by authors such as Moerbeek and Maas (2005), Tekle et al. (2008) and Ogungbenro and Aarons (2011).
3. *Generalized estimating equations (GEEs)*:

$$V_l = \{ \text{diag} [\text{Var}(y_{lj})] \}^{1/2} R \{ \text{diag} [\text{Var}(y_{lj})] \}^{1/2},$$

with R as an intra-block marginal *working correlation* matrix, and $\mu_l = E(y_l | \gamma_l = 0)$; see Liang and Zeger (1986). The matrix R is assumed to be independent of x and usually chosen to have a standard form, for example, exchangeable or known

up to a small number of correlation parameters. For discrete data, it is in fact often impossible for either of these two assumptions to hold, but the resulting estimators can be relatively efficient compared to a full maximum likelihood approach (Chaganty and Joe 2004). Methodology for D-optimal designs using this approach was developed by Woods and van de Ven (2011).

Note that for each of these approximations, if $\sigma_\gamma^2 \rightarrow 0$ (or equivalently, $\mathbf{R} \rightarrow \mathbf{I}$), the variance-covariance matrix for \mathbf{y}_l reverts to that for a simple GLM. Waite and Woods (2015) developed, assessed and compared a variety of methods of approximating \mathbf{M} to find D-optimal designs for a GLMM.

13.6.5 Comparison of Approximations

We use a small example to perform a simple comparison of designs from the three approximations in the previous section. Consider an experiment in blocks of size $m = 2$ with a single variable x to collect count data. Conditional on the random block effect, we assume $y_{ij}|\gamma_i \sim \text{Poisson}(\mu_{ij})$ ($i = 1, \dots, b; j = 1, \dots, m$), and we choose a second-order predictor and the log link

$$\log(\mu_{ij}) = \eta_{ij} = \gamma_i + \theta_0 + \theta_1 x_{ij} + \theta_2 x_{ij}^2,$$

for $x_{ij} \in [-1, 1]$. For the purposes of finding designs, we assume point prior information and set $\theta_0 = 0, \theta_1 = 5$ and $\theta_2 = 1$. The random block effect has distribution $\gamma_i \sim N(0, \sigma_\gamma^2)$ for $i = 1, \dots, b$ and $\sigma_\gamma^2 = 0.5$.

For the log link, $\Delta_l = \text{diag} \{ \mu_{lj} \}$. We consider each of the three approximations, with μ_{lj} representing either the marginal or conditional mean response for the j th point in the l th block of support ($l = 1, \dots, t; j = 1, \dots, m$).

Quasi-likelihood (QL): Here, $\mu_{lj} = E(y_{lj}) = \exp(\eta_{lj} + \sigma_\gamma^2/2)$, the marginal mean response, and

$$\mathbf{V}_l = \text{diag} \left\{ \exp(\eta_{lj} + \sigma_\gamma^2/2) \right\} + \left\{ \exp(\sigma_\gamma^2) - 1 \right\} \bar{\boldsymbol{\mu}}_l \bar{\boldsymbol{\mu}}_l'$$

where $\bar{\boldsymbol{\mu}}_l' = \left\{ \exp(\eta_{lj} + \sigma_\gamma^2/2) \right\}_{j=1}^m$.

Marginal quasi-likelihood (MQL): For this approximation, $\mu_{lj} = \exp(\eta_{lj})$, the conditional mean response given $\gamma_l = 0$, and

$$\mathbf{V}_l = \text{diag} \left\{ \exp(\eta_{lj}) \right\} + \sigma_\gamma^2 \boldsymbol{\mu}_l \boldsymbol{\mu}_l'$$

where $\boldsymbol{\mu}_l' = \left\{ \exp(\eta_{lj}) \right\}_{j=1}^m$.

Generalized estimating equations (GEEs): Now, $\mu_{lj} = \exp(\eta_{lj})$, the conditional mean response given $\gamma_l = 0$, and

$$\mathbf{V}_l = \text{diag} \left\{ \exp(\eta_{lj}) \right\}^{\frac{1}{2}} \mathbf{R} \text{diag} \left\{ \exp(\eta_{lj}) \right\}^{\frac{1}{2}},$$

TABLE 13.11

D-Optimal Continuous Block Designs with Blocks of Size $m = 2$ for a Poisson Example and QL, MQL and GEE Approaches (to 2 Decimal Places)

	Support Blocks		
	Block 1	Block 2	Block 3
QL/MQL	(0.88, 0.10)	(1, 0.75)	—
Weights	0.5	0.5	—
GEE	(0.84, 0.02)	(0.72, 1)	(1, 0.26)
Weights	0.38	0.35	0.27

where the working correlation matrix for this example is

$$R = \begin{pmatrix} 1 & \alpha \\ \alpha & 1 \end{pmatrix}.$$

For the GEE design, we redefine $\tau = (\theta', \alpha)$.

Locally D-optimal designs under the three approximations are given in Table 13.11. Note that the same design was found under the QL and MQL approximations and that, for the GEE design, $\alpha = 0.5$ was chosen so that the working correlation closely matched the average intra-block correlation (≈ 0.49) from the other design. Table 13.12 gives the D-efficiencies (13.26) of each design under each approximation; the GEE design is 87% efficient under the QL and MQL approximations, whilst the QL/MQL design is 90% efficient under the GEE approximation.

Optimality of these designs can be confirmed, as in Section 13.3.1, via the application of a multivariate equivalence theorem; see Atkinson (2008b) and Woods and van de Ven (2011). A necessary and sufficient condition for a design ξ^* to be locally D-optimal is

$$\psi(x, \xi^*; \tau) = p - \text{trace} \left\{ M(\zeta; \tau) M^{-1}(\xi^*; \tau) \right\} \geq 0, \quad (13.46)$$

for all $\zeta \in \mathcal{X}^m$. This condition can be verified numerically; Figure 13.15 plots the derivative function for each of the three approximations, with the support points of the optimal designs marked. Notice that (1) the support points occur at minima of the derivative surface, with $\psi(x, \xi^*; \tau) = 0$; (2) with blocks of size $m = 2$, the derivative function must be symmetric about the line $x_{11} = x_{12}$; and (3) the derivative surfaces for QL and MQL are very similar.

TABLE 13.12

Efficiencies of Three Optimal Designs under Three Approximations to the Information Matrix

Design	Approximation	
	QL/MQL	GEE
QL/MQL	1	0.90
GEE	0.87	1

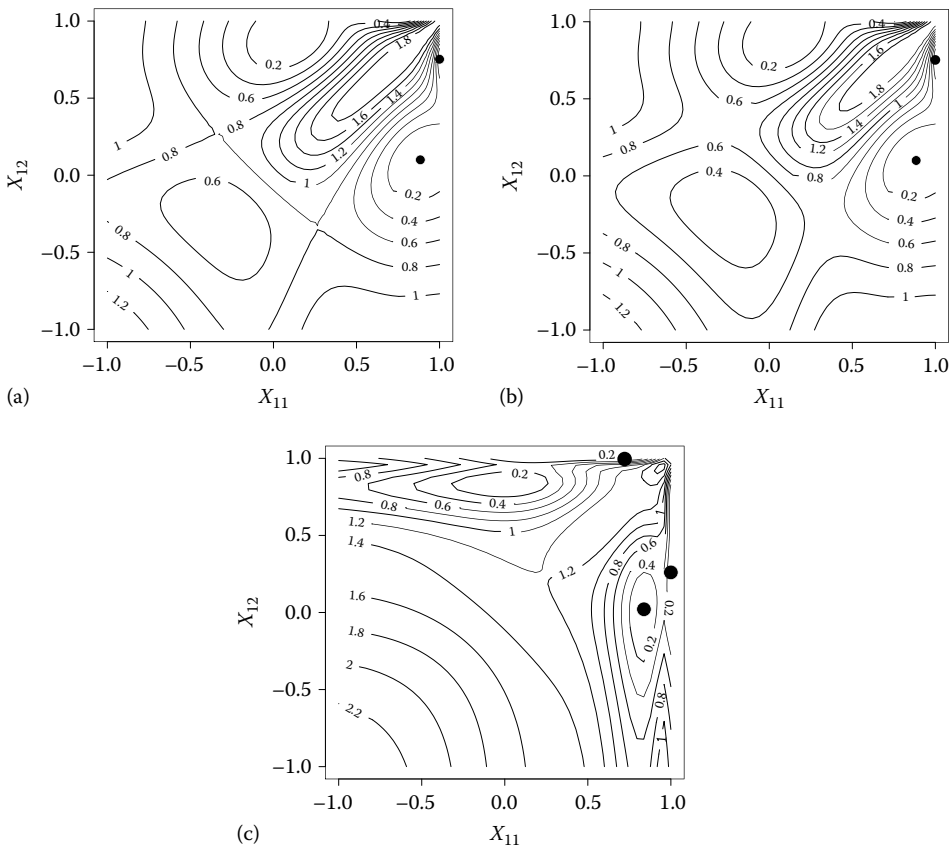


FIGURE 13.15 Derivative surfaces, $\psi(x, \xi^*; \tau)$, and D-optimal designs for QL (a), MQL (b) and GEE (c) approximations.

13.7 Extensions and Further Reading

Although we have described designs for three types of response and several link functions, interest in designs for GLMs continues to focus on binary responses and the logistic link. Much of this reflects the rapid growth of applications of the discrete choice models described in detail in Chapter 22.

There is also appreciable interest in design for logistic regression in medical statistics, particularly drug development. A typical problem is to find the dose giving a specified probability of toxicity. The natural way to proceed is the use of a sequential design as described in Section 13.3.4.2. The appropriate design criterion is c-optimality, which designs to estimate this dose with minimum variance. Such designs for nonlinear regression models are exemplified by Atkinson et al. (2007, Section 17.5) and by Ford et al. (1992) for GLMs. Often, however, the locally optimal design has a single support point at the estimated dose. The sequential design may then fail to provide sufficient information to guarantee identification of the required dose (Pronzato 2000; O’Quigley and Conaway 2010).

The designs need to provide sufficient perturbation in the experimental conditions to ensure convergence.

The designs we have exemplified, particularly for first-order models, often have the number of support points equal to the number of parameters. They therefore fail to provide any information for model checking and choice. Designs for discriminating between two regression models were introduced by Fedorov and Maljutov (1972) and by Atkinson and Fedorov (1975) who called them T-optimal. Ponce de Leon and Atkinson (1992) and Waterhouse et al. (2008) extend T-optimality to GLMs. A general design criterion for discrimination between models using Kullback–Leibler distances is that of López-Fidalgo et al. (2007). A potential disadvantage of these designs is that they focus exclusively on model discrimination. Compound designs for the joint problem of parameter estimation and model discrimination, called DT-optimal, are given, for regression models, by Atkinson (2008a). Waterhouse et al. (2008) also attend to the quality of parameter estimates, but not through use of a compound criterion. D-optimal designs robust to the form of the linear predictor were developed by Woods et al. (2006) and Woods and Lewis (2011).

In some applications, it is not necessary to establish a model that holds globally. In the context of dose finding, O’Quigley and Conaway (2010) recommend the use of a model that has adequate local behaviour. A more formal approach to model uncertainty is in Chapter 20. In particular, Li and Wiens (2011) consider approximate models in dose–response experiments, whereas Wiens (2009) provides robust designs for discrimination between regression models.

In Section 13.2.5, we mentioned that response transformation is sometimes an alternative to the use of a gamma model. Atkinson and Cook (1997) find optimal designs for estimation of the transformation in linear models, whereas Atkinson (2005) studies designs for transforming both sides of a nonlinear model.

Finally, we note that some interesting models for non-normal responses have a structure as in Section 13.2.1 but with a predictor which is nonlinear in the parameters. Optimal designs for one particular generalized nonlinear model are given by Biedermann and Woods (2011).

Acknowledgments

The authors thank Dr. T.W. Waite (University of Southampton, United Kingdom) for assistance with the computation in Section 13.6. D.C. Woods was supported by a fellowship from the U.K. Engineering and Physical Sciences Research Council (EP/J018317/1), and the work was partly undertaken whilst the authors were visiting fellows at the Isaac Newton Institute for Mathematical Sciences (Cambridge, United Kingdom).

References

- Agresti, A. (2002). *Categorical Data Analysis*, 2nd edn. New York: Wiley.
- Atkinson, A. C. (2005). Robust optimum designs for transformation of the responses in a multivariate kinetic model. *Technometrics* 47, 478–487.

- Atkinson, A. C. (2006). Generalized linear models and response transformation. In A. I. Khuri (Ed.), *Response Surface Methodology and Related Topics*, Chapter 8, pp. 173–202. Singapore: World Scientific.
- Atkinson, A. C. (2008a). D-optimum designs for model discrimination and parameter estimation. *Journal of Statistical Planning and Inference* 138, 56–64.
- Atkinson, A. C. (2008b). Examples of the use of an equivalence theorem in constructing optimum experimental designs for random-effects nonlinear regression models. *Journal of Statistical Planning and Inference* 138, 2595–2606.
- Atkinson, A. C. (2010). The non-uniqueness of some designs for discriminating between two polynomial models in one variable. In A. Giovagnoli, A. C. Atkinson, and B. Torsney (Eds.), *mODA9—Advances in Model-Oriented Design and Analysis*, pp. 9–16. Heidelberg, Germany: Springer-Verlag.
- Atkinson, A. C. and R. D. Cook (1995). D-optimum designs for heteroscedastic linear models. *Journal of the American Statistical Association* 90, 204–212.
- Atkinson, A. C. and R. D. Cook (1997). Designing for a response transformation parameter. *Journal of the Royal Statistical Society B* 59, 111–124.
- Atkinson, A. C., A. N. Donev, and R. D. Tobias. (2007). *Optimum Experimental Designs, with SAS*. Oxford, U.K.: Oxford University Press.
- Atkinson, A. C. and V. V. Fedorov (1975). The design of experiments for discriminating between two rival models. *Biometrika* 62, 57–70.
- Atkinson, A. C., V. V. Fedorov, A. M. Herzberg, and R. Zhang (2014). Elemental information matrices and optimal design for generalized regression models. *Journal of Statistical Planning and Inference* 144, 81–91.
- Atkinson, A. C. and M. Riani (2000). *Robust Diagnostic Regression Analysis*. New York: Springer-Verlag
- Bazzoli, C., S. Retout, and F. Mentré (2010). Design evaluation and optimisation in multiple response nonlinear mixed effect models: Pfm 3.0. *Computer Methods and Programs in Biomedicine* 98, 55–65.
- Berger, M. and W. K. Wong (2009). *An Introduction to Optimal Designs for Social and Biomedical Research*. New York: Wiley.
- Biedermann, S. and D. C. Woods (2011). Optimal designs for generalized non-linear models with application to second-harmonic generation experiments. *Journal of the Royal Statistical Society C* 60, 281–299.
- Bliss, C. I. (1935). The calculation of the dosage-mortality curve. *Annals of Applied Biology* 22, 134–167.
- Box, G. E. P. and D. R. Cox (1964). An analysis of transformations (with discussion). *Journal of the Royal Statistical Society B* 26, 211–246.
- Box, G. E. P. and N. R. Draper (1963). The choice of a second order rotatable design. *Biometrika* 50, 335–352.
- Box, G. E. P. and N. R. Draper (2007). *Response Surfaces, Mixtures, and Ridge Analysis*, 2nd edn. New York: Wiley.
- Box, G. E. P. and W. G. Hunter (1965). Sequential design of experiments for nonlinear models. In *Proceedings IBM Scientific Computing Symposium: Statistics*, pp. 113–137. New York: IBM.
- Box, G. E. P. and H. L. Lucas (1959). Design of experiments in nonlinear situations. *Biometrika* 46, 77–90.
- Breslow, N. E. and D. G. Clayton (1993). Approximate inference in generalized linear mixed models. *Journal of the American Statistical Association* 88, 9–25.
- Burridge, J. and P. Sebastiani (1994). D-optimal designs for generalised linear models with variance proportional to the square of the mean. *Biometrika* 81, 295–304.
- Chaganty, N. R. and H. Joe (2004). Efficiency of generalized estimating equations for binary responses. *Journal of the Royal Statistical Society B* 66, 851–860.
- Chaloner, K. and K. Larntz (1989). Optimal Bayesian design applied to logistic regression experiments. *Journal of Statistical Planning and Inference* 21, 191–208.
- Chaloner, K. and I. Verdini (1995). Bayesian experimental design: A review. *Statistical Science* 10, 273–304.

- Cheng, C.-S. (1995). Optimal regression designs under random block-effects models. *Statistica Sinica* 5(2), 485–497.
- Collett, D. (2002). *Modelling Binary Data*, 2nd edn. London, U.K.: Chapman & Hall/CRC.
- Cox, D. R. (1988). A note on design when response has an exponential family distribution. *Biometrika* 75, 161–164.
- Dette, H., V. B. Melas, and P. Shpilev (2012). T-optimal designs for discrimination between two polynomial models. *Annals of Statistics* 40, 188–205.
- Diggle, P. J., P. Heagerty, K.-Y. Liang, and S. Zeger (2002). *Analysis of Longitudinal Data*, 2nd edn. Oxford, U.K.: Oxford University Press.
- Dobson, A. (2001). *An Introduction to Generalized Linear Models*, 2nd edn. London, U.K.: Chapman & Hall.
- Dror, H. A. and D. M. Steinberg (2006). Robust experimental design for multivariate generalized linear models. *Technometrics* 48, 520–529.
- Dror, H. A. and D. M. Steinberg (2008). Sequential experimental designs for generalized linear models. *Journal of the American Statistical Association* 103, 288–298.
- Fedorov, V. and M. Malyutov (1972). Optimal designs in regression problems. *Mathematical Operationsforsch Statist* 3, 281–308.
- Fedorov, V. V. (1972). *Theory of Optimal Experiments*. New York: Academic Press.
- Fedorov, V. V. and P. Hackl (1997). *Model-Oriented Design of Experiments*. Lecture Notes in Statistics 125. New York: Springer Verlag.
- Fedorov, V. V. and S. L. Leonov (2014). *Optimal Design for Nonlinear Response Models*. Boca Raton, FL: Chapman & Hall/CRC Press.
- Firth, D. (1993). Bias reduction of maximum likelihood estimates. *Biometrika* 80, 27–38.
- Ford, I., B. Torsney, and C. F. J. Wu (1992). The use of a canonical form in the construction of locally optimal designs for non-linear problems. *Journal of the Royal Statistical Society B* 54, 569–583.
- Godambe, V. P. (Ed.) (1991). *Estimating Functions*. Oxford, U.K.: Oxford Science Publications.
- Goos, P. and B. Jones (2011). *Optimal Design of Experiments: A Case Study Approach*. New York: Wiley.
- Gotwalt, C. M., B. A. Jones, and D. M. Steinberg (2009). Fast computation of designs robust to parameter uncertainty for nonlinear settings. *Technometrics* 51, 88–95.
- Khuri, A. I., B. Mukherjee, B. K. Sinha, and M. Ghosh (2006). Design issues for generalized linear models: A review. *Statistical Science* 21, 376–399.
- Kiefer, J. (1959). Optimum experimental designs (with discussion). *Journal of the Royal Statistical Society B* 21, 272–319.
- Kiefer, J. and J. Wolfowitz (1960). The equivalence of two extremum problems. *Canadian Journal of Mathematics* 12, 363–366.
- King, J. and W. K. Wong (2000). Minimax D-optimal designs for the logistic model. *Biometrics* 56, 1263–1267.
- Li, P. and D. P. Wiens (2011). Robustness of design in dose–response studies. *Journal of the Royal Statistical Society B* 73, 215–238.
- Liang, K. and S. Zeger (1986). Longitudinal data analysis using generalized linear models. *Biometrika* 73, 13–22.
- López-Fidalgo, J., C. Trandafir, and C. Tommasi (2007). An optimal experimental design criterion for discriminating between non-normal models. *Journal of the Royal Statistical Society B* 69, 231–242.
- Magnus, J. R. and H. Neudecker (1988). *Matrix Differential Calculus with Applications in Statistics and Econometrics*. Chichester, U.K.: Wiley.
- McCullagh, P. and J. A. Nelder (1989). *Generalized Linear Models*, 2nd edn. London, U.K.: Chapman & Hall.
- McCulloch, C. E., S. R. Searle, and J. M. Neuhaus (2008). *Generalized, Linear and Mixed Models*, 2nd ed. Hoboken, NJ: Wiley.
- McGree, J. M. and J. A. Eccleston (2012). Robust designs for Poisson regression models. *Technometrics* 54, 64–72.

- Mentré, F., A. Mallet, and D. Baccar (1997). Optimal design in random-effects regression models. *Biometrika* 84, 429–442.
- Minkin, S. (1993). Experimental design for clonogenic assays in chemotherapy. *Journal of the American Statistical Association* 88, 410–420.
- Moerbeek, M. and C. J. M. Maas (2005). Optimal experimental designs for multilevel logistic models with two binary predictors. *Communications in Statistics: Theory and Methods* 34, 1151–1167.
- Morgan, B. J. T. (1992). *Analysis of Quantal Response Data*. London, U.K.: Chapman & Hall.
- Muirhead, R. J. (1982). *Aspects of Multivariate Statistical Theory*. New York: Wiley.
- Myers, R. H., D. C. Montgomery, G. G. Vining, and T. J. Robinson (2010). *Generalized Linear Models: With Applications in Engineering and the Sciences*, 2nd edn. New York: Wiley.
- Nelson, W. (1981). The analysis of performance-degradation data. *IEEE Transactions on Reliability* R-30, 149–155.
- Niaparast, M. (2009). On optimal design for a Poisson regression model with random intercept. *Statistics and Probability Letters* 79, 741–747.
- Niaparast, M. and R. Schwabe (2013). Optimal design for quasi-likelihood estimation in Poisson regression with random coefficients. *Journal of Statistical Planning and Inference* 143, 296–306.
- Nyquist, H. (2013). Convergence of an algorithm for constructing minimax designs. In D. Uciński, A. C. Atkinson, and M. Patan (Eds.), *mODA 10—Advances in Model-Oriented Design and Analysis*, pp. 187–194. Heidelberg, Germany: Springer.
- Ogungbenro, K. and L. Aarons (2011). Population Fisher information matrix and optimal design of discrete data responses in population pharmacodynamic experiments. *Journal of Pharmacokinetics and Pharmacodynamics* 38, 449–469.
- O’Quigley, J. and M. Conaway (2010). Continual reassessment and related dose-finding designs. *Statistical Science* 25, 202–216.
- Pázman, A. (1986). *Foundations of Optimum Experimental Design*. Dordrecht, the Netherlands: Reidel.
- Ponce de Leon, A. M. and A. C. Atkinson (1992). The design of experiments to discriminate between two rival generalized linear models. In L. Fahrmeir, B. Francis, R. Gilchrist, and G. Tutz (Eds.), *Advances in GLIM and Statistical Modelling: Proceedings of the GLIM92 Conference, Munich*, pp. 159–164. New York: Springer.
- Pronzato, L. (2000). Adaptive optimization and D-optimum experimental design. *Annals of Statistics* 28, 1743–1761.
- Pronzato, L. and A. Pázman (2013). *Design of Experiments in Nonlinear Models*. New York: Springer.
- Pukelsheim, F. (1993). *Optimal Design of Experiments*. New York: Wiley.
- Retout, S. and F. Mentré (2003). Further developments of the Fisher information matrix in nonlinear mixed effects models with evaluation in population pharmacokinetics. *Biopharmaceutical Statistics* 13, 209–227.
- Robbins, H. and S. Monro (1951). A stochastic approximation method. *Annals of Mathematical Statistics* 22, 400–407.
- Robinson, T., R. Myers, and D. Montgomery (2004). Analysis considerations in industrial split-plot experiments with non-normal responses. *Journal of Quality Technology* 36(2), 180–191.
- Russell, K. G., D. C. Woods, S. M. Lewis, and J. A. Eccleston (2009). D-optimal designs for Poisson regression models. *Statistica Sinica* 19, 721–730.
- Santner, T. J., B. Williams, and W. Notz (2003). *The Design and Analysis of Computer Experiments*. New York: Springer-Verlag.
- Sibson, R. (1974). D_A -optimality and duality. In J. Gani, K. Sarkadi, and I. Vincze (Eds.), *Progress in Statistics, Proceedings of the Ninth European Meeting of Statisticians, Budapest*, Vol. 2. Amsterdam, the Netherlands: North-Holland.
- Silvey, S. D. (1980). *Optimum Design*. London, U.K.: Chapman & Hall.
- Silvey, S. D. and D. M. Titterton (1973). A geometric approach to optimal design theory. *Biometrika* 60, 15–19.
- Sitter, R. R. (1992). Robust designs for binary data. *Biometrics* 48, 1145–1155.
- Sitter, R. R. and B. Torsney (1995a). Optimal designs for binary response experiments with two variables. *Statistica Sinica* 5, 405–419.

- Sitter, R. S. and B. Torsney (1995b). D-optimal designs for generalized linear models. In C. P. Kitsos and W. G. Müller (Eds.), *MODA 4—Advances in Model-Oriented Data Analysis*, pp. 87–102. Heidelberg: Physica-Verlag.
- Smith, K. (1918). On the standard deviations of adjusted and interpolated values of an observed polynomial function and its constants and the guidance they give towards a proper choice of the distribution of observations. *Biometrika* 12, 1–85.
- Stufken, J. and M. Yang (2012a). On locally optimal designs for generalized linear models with group effects. *Statistica Sinica* 22, 1765–1786.
- Stufken, J. and M. Yang (2012b). Optimal designs for generalized linear models. In K. Hinkelmann (Ed.), *Design and Analysis of Experiments, Special Designs and Applications*, Chapter 4, pp. 137–165. New York: Wiley.
- Tekle, F. B., F. E. S. Tan, and M. P. F. Berger (2008). Maximin D-optimal designs for binary longitudinal responses. *Computational Statistics and Data Analysis* 52, 5253–5262.
- Torsney, B. and N. Gunduz (2001). On optimal designs for high dimensional binary regression models. In A. C. Atkinson, B. Bogacka, and A. Zhigljavsky (Eds.), *Optimal Design 2000*, pp. 275–285. Dordrecht, the Netherlands: Kluwer.
- van de Ven, P. M. and D. C. Woods (2014). Optimal blocked minimum-support designs for non-linear models. *Journal of Statistical Planning and Inference* 144, 152–159.
- von Eye, A., E. Mun, and P. Mair (2011). Log-linear modeling. *Wiley Interdisciplinary Reviews: Computational Statistics* 4, 218–223.
- Waite, T. W. (2013). Designs for experiments with mixed effects and discrete responses plus related topics. PhD thesis, University of Southampton, Southampton, U.K.
- Waite, T. W. and D. C. Woods (2015). Designs for generalized linear models with random block effects via information matrix approximations. *Biometrika*, in press.
- Wang, Y., R. H. Myers, E. P. Smith, and K. Ye (2006). D-optimal designs for Poisson regression models. *Journal of Statistical Planning and Inference* 136, 2831–2845.
- Wang, Y., E. P. Smith, and K. Ye (2006). Sequential designs for a Poisson regression model. *Journal of Statistical Planning and Inference* 136, 3187–3202.
- Waterhouse, T. H., D. C. Woods, J. A. Eccleston, and S. M. Lewis (2008). Design selection criteria for discrimination/estimation for nested models and a binomial response. *Journal of Statistical Planning and Inference* 138, 132–144.
- Wedderburn, R. W. M. (1974). Quasi-likelihood functions, generalized linear models, and the Gauss-Newton method. *Biometrika* 61, 439–447.
- Whittle, P. (1973). Some general points in the theory of optimal experimental design. *Journal of the Royal Statistical Society, Series B* 35, 123–130.
- Wiens, D. P. (2009). Robust discrimination designs. *Journal of the Royal Statistical Society, Series B* 71, 805–829.
- Winkelmann, R. (2008). *Econometric Analysis of Count Data*. New York: Springer.
- Woods, D. C. (2010). Robust designs for binary data: Applications of simulated annealing. *Journal of Statistical Computation and Simulation* 80, 29–41.
- Woods, D. C. and S. M. Lewis (2011). Continuous optimal designs for generalized linear models under model uncertainty. *Journal of Statistical Theory and Practice* 5, 137–145.
- Woods, D. C., S. M. Lewis, J. A. Eccleston, and K. G. Russell (2006). Designs for generalized linear models with several variables and model uncertainty. *Technometrics* 48, 284–292.
- Woods, D. C. and P. van de Ven (2011). Blocked designs for experiments with non-normal response. *Technometrics* 53, 173–182.
- Wu, H.-P. and J. Stufken (2014). Locally ϕ_p -optimal designs for generalized linear models with a single-variable quadratic polynomial predictor. *Biometrika* 101, 365–375.
- Yang, M., B. Zhang, and S. Huang (2011). Optimal designs for generalized linear models with multiple design variables. *Statistica Sinica* 21, 1415–1430.

14

Designs for Selected Nonlinear Models

Stefanie Biedermann and Min Yang

CONTENTS

14.1	Introduction.....	515
14.2	Classical Methods.....	518
14.2.1	Methods Based on the Equivalence Theorem.....	519
14.2.2	Geometric Approach.....	522
14.2.2.1	Elfving's Theorem and Its Implications for c -Optimality.....	523
14.2.2.2	Characterization of ϕ_p -Optimal Designs via Covering Ellipses.....	525
14.2.3	Functional Approach.....	527
14.3	General Solutions.....	528
14.3.1	Algebraic Method.....	529
14.3.2	Method Based on Chebyshev Systems.....	531
14.4	Further Examples.....	533
14.4.1	Two-Parameter Exponential Model.....	533
14.4.2	Emax Model.....	534
14.4.3	Heteroscedastic Linear Model.....	536
14.5	Model Discrimination for Nonlinear Models.....	538
14.6	Parameter Robust Approaches.....	540
14.6.1	Response-Adaptive Sequential/Batch Sequential Experimentation.....	542
14.6.2	Bayesian/Pseudo-Bayesian Designs.....	542
14.6.3	(Standardized) Maximin Optimal Designs.....	543
14.6.4	Cluster Designs.....	544
14.7	Summary.....	544
	References.....	544

14.1 Introduction

This chapter is an example-based guide to optimal design for nonlinear regression models. For clarity, we restrict ourselves to models with only one continuous explanatory variable. The classical theory presented in Section 14.2 also holds for multivariable models, whereas the results shown in Section 14.3 have been developed for single-variable models. In practice, designs for multivariable models are usually found numerically due to the increased complexity (see, e.g., Yu 2011, or Yang et al. 2013 for some recent developments on algorithms). Further information on algorithms for design search is given in Chapters 13 and 21.

Some analytical results on finding optimal designs for multivariable models can be found in Biedermann et al. (2011), Yang et al. (2011) and references therein.

Throughout this chapter, we assume that we can make n observations y_1, \dots, y_n , from a nonlinear model, at experimental conditions $x_1, \dots, x_n \in \mathcal{X}$, respectively. More specifically,

$$y_i = \eta(x_i, \theta) + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \sigma^2), \quad \sigma^2 > 0, \quad i = 1, \dots, n, \quad (14.1)$$

where $\eta(x_i, \theta)$ is the nonlinear regression function, known up to a vector of m unknown parameters, θ , and the errors $\epsilon_i, i = 1, \dots, n$, are independent and identically distributed. The design space \mathcal{X} is usually an interval on the real axis, and the variance parameter σ^2 is assumed to be a nuisance parameter.

Suppose without loss of generality that $x_1, \dots, x_t, t \leq n$, are the distinct points among x_1, \dots, x_n . We consider approximate designs of the form $\xi = \{(x_1, w_1), \dots, (x_t, w_t)\}$, where the weight w_i gives the proportion of observations to be made at the corresponding support point $x_i, i = 1, \dots, t$. We thus require $0 < w_i \leq 1$, for $i = 1, \dots, t$, and $\sum_{i=1}^t w_i = 1$. Note that nw_i is not restricted to be an integer to avoid cumbersome discrete optimization problems. In order to run an approximate design in practice, a rounding procedure (see, e.g., Pukelsheim and Rieder 1992), is used. Approximate designs are also known as continuous designs in the literature (see Chapter 13).

We are concerned with the *optimal* choice of a design. A decision rule stating what is deemed optimal is provided by an optimality criterion, which is selected to reflect the purpose of the experiment. In what follows, we assume that we want to estimate the mean model parameters θ as accurately as possible and that the estimation is either through maximum likelihood or nonlinear least squares. It is therefore natural to consider optimality criteria that are concerned with minimizing some function of the (asymptotic) covariance matrix of the estimator $\hat{\theta}$ or, equivalently, maximizing some function of the Fisher information matrix, M_ξ , which for model (14.1) is given by

$$M_\xi = \sum_{i=1}^t w_i f(x_i, \theta) f'(x_i, \theta) = \mathbf{X}' \mathbf{W} \mathbf{X}, \quad (14.2)$$

where

$$f(x, \theta) = \left(\frac{\partial \eta(x, \theta)}{\partial \theta_1}, \dots, \frac{\partial \eta(x, \theta)}{\partial \theta_m} \right)' \quad (14.3)$$

is the vector of partial derivatives of $\eta(x, \theta)$ with respect to θ , $\mathbf{X} = [f(x_1, \theta), \dots, f(x_t, \theta)]'$, and \mathbf{W} is the diagonal matrix holding the weights.

Example 14.1

The Michaelis–Menten model is typically used to describe the rate of an enzymatic reaction as a function of the concentration of a substrate. This model has expected response

$$\eta(x, \theta) = \frac{\theta_1 x}{\theta_2 + x}, \quad \theta_1, \theta_2 > 0, \quad x \geq 0,$$




FIGURE 14.1 Plot of the expected response $\eta(x, \theta)$ for the Michaelis–Menten model with parameter vector $\theta' = (1, 0.6)$. For $x \rightarrow \infty$, $\eta(x, \theta)$ is asymptoting at $\theta_1 = 1$, and half of its supremum is attained at $x = \theta_2 = 0.6$.

which is depicted in Figure 14.1 for parameter values $\theta' = (1, 0.6)$. The parameter θ_1 gives the supremum of the curve, whereas θ_2 can be interpreted as the value of x at which half the supremum is attained.

For the Michaelis–Menten model, the Fisher information of a design ξ is

$$M_{\xi} = \sum_{i=1}^t w_i \frac{x_i^2}{(\theta_2 + x_i)^2} \begin{bmatrix} 1 & \frac{-\theta_1}{(\theta_2 + x_i)} \\ \frac{-\theta_1}{(\theta_2 + x_i)} & \frac{\theta_1^2}{(\theta_2 + x_i)^2} \end{bmatrix}. \tag{14.4}$$

There is no total ordering on the nonnegative definite matrices of size $(m \times m)$. Hence, they are mapped to the real axis through an objective function to make them comparable. A popular class of optimality criteria is the ϕ_p -criteria (see, e.g., Pukelsheim 1993), which maximize the corresponding matrix means. The p th matrix mean, $p \in [-\infty, 1]$, is defined as

$$\phi_p(\xi) = \begin{cases} \left(\frac{1}{m} \text{trace} M_{\xi}^p \right)^{1/p} & p \neq -\infty, 0 \\ |M_{\xi}|^{1/m} & p = 0 \\ \lambda_{\min}(M_{\xi}) & p = -\infty, \end{cases}$$

where $\lambda_{\min}(M_{\xi})$ is the minimal eigenvalue of M_{ξ} . Well-known special cases are the D -, A -, and E -criteria, where $p = 0, -1, -\infty$, respectively. A D -optimal design minimizes the volume of an asymptotic confidence ellipsoid for $\hat{\theta}$, and an A -optimal design minimizes the average of the asymptotic variances for the estimators of the individual parameters. For an interpretation of the ϕ_p -criteria in terms of the eigenvalues of the information matrix, see

Chapter 3. Note that the values for p therein correspond to $-p$ in the preceding definition, since the inverse of the information matrix is considered.

If interest is in estimating a linear combination of the parameters, $c'\theta$, for a given vector c , we use the c -optimality criterion, which minimizes the objective function $\phi_c(\xi) = c'M_{\xi}^{-}c$, where M_{ξ}^{-} is a generalized inverse of the information matrix M_{ξ} . This corresponds to minimizing the asymptotic variance of $c'\hat{\theta}$.

We note that in nonlinear models at least some of the partial derivatives, and thus the Fisher information, depend on the unknown parameter vector θ . An optimal design with respect to some optimality criterion will therefore only be optimal for a specific value of θ and is called a locally optimal design (see, e.g., Chernoff 1953 or Chapter 13). An important subclass of nonlinear models, the partially nonlinear models, are defined by Hill (1980) and Khuri (1984) as models where some of the parameters appear linearly. For D -optimality, these linear parameters do not affect the maximization problem, and thus D -optimal designs depend only on the nonlinear parameters. The Michaelis–Menten model, for example, is a partially nonlinear model where the parameter θ_1 appears linearly.

This chapter is organized as follows. In Section 14.2, we review three classical methods for finding optimal designs. Section 14.3 is devoted to a recent approach, shedding light on optimal design problems from a more general perspective. All these methods are illustrated through a running example, the Michaelis–Menten model. Further models are discussed in Section 14.4. For each of these, only the most suitable method is applied, including a discussion of the drawbacks of the other methods for this particular situation. While the main focus of this chapter is optimal design for parameter estimation, Section 14.5 gives a brief overview of optimal design when the purpose of the experiment is discrimination between two or more models. All designs provided in this chapter are locally optimal in the sense of Chernoff (1953), that is, they depend on a best guess of the unknown model parameters. In Section 14.6, we briefly discuss approaches to overcome this problem. In each section, we point the interested reader to further relevant articles from the recent literature on optimal design for nonlinear models.

14.2 Classical Methods

In this section, we distinguish between three approaches to facilitate the computation of optimal designs.

The standard method in many situations is the use of an appropriate equivalence theorem in order to find certain properties, usually the number of support points and possibly the inclusion of end points of \mathcal{X} in the support of the optimal design. Equivalence theorems are available for all commonly applied optimality criteria based on the Fisher information, for example, the ϕ_p -criteria or c -optimality (see Pukelsheim 1993). The practical applications of an equivalence theorem are not restricted to design construction, but it also allows a check of the optimality of a given candidate design.

Similarly powerful methods, summarized as the geometric approach, use the visualization of what is called the induced design space, a combination of the model and the design space \mathcal{X} . Again, this often leads to finding the number of support points of an optimal design and to results concerning the inclusion of boundary points of \mathcal{X} . Results are available for ϕ_p - and c -optimality; see Biedermann et al. (2006) and Elfving (1952), respectively. Since the plots used for visualization have as many axes as the underlying

model has parameters, this approach is most useful for models with two or at most three parameters.

A further method is the functional approach (see, e.g., Melas 2006). The main idea of this approach is to express the support points (and sometimes also the weights) of optimal designs as implicit functions of some auxiliary parameters. In many cases these functions, being real and analytic, can be expanded into Taylor series, and recursive formulae are available for the coefficients of these. Results in this area cover the D -, E -, and c -criteria and some parameter robust criteria.

Unlike some more recent methods, see Section 14.3, which aim at finding complete classes of optimal designs that are dominating with respect to the Loewner ordering, the classical methods usually solve one design problem at a time. Some of these approaches, however, allow conclusions for a particular class of optimality criteria, the ϕ_p -criteria.

In some situations, the first two methods provide (some of) the support points of an optimal design, but usually no characterization of the optimal weights. For the situation where the optimal design has the minimal number of support points, m , to estimate θ , Pukelsheim and Torsney (1991) have developed a method to find optimal weights given the support points, applicable to many optimality criteria including the ϕ_p -criteria. A similar result is available for c -optimality.

Recall the definition of $\mathbf{X} = [f(x_1, \theta), \dots, f(x_m, \theta)]'$, and let $\mathbf{V} = (\mathbf{X}\mathbf{X}')^{-1}\mathbf{X}$. Then, for $p \in (-\infty, 1]$, the ϕ_p -optimal weights w_i of a design with support points x_1, \dots, x_m can be obtained by solving the system of equations

$$w_i = \frac{\sqrt{u_{ii}}}{\sum_{j=1}^m \sqrt{u_{jj}}}, \quad i = 1, \dots, m, \tag{14.5}$$

where u_{ii} is the i th diagonal element of the matrix $\mathbf{U} = \mathbf{V}\mathbf{M}_\xi^{p+1}\mathbf{V}'$. Using (14.2), for $p = -1$ (i.e., A -optimality), we get an explicit solution since \mathbf{U} reduces to $(\mathbf{X}\mathbf{X}')^{-1}$, and thus the right-hand side of (14.5) does not depend on the weights. Similarly, for $p = 0$ (i.e., D -optimality), \mathbf{U} simplifies to \mathbf{W} ; hence, all weights are equal to $1/m$. For c -optimality with respect to a vector c , we also obtain an explicit solution $w_i = |v_i| / \sum_{j=1}^m |v_j|$, $i = 1, \dots, m$, where the vector v is defined as $v = \mathbf{V}c$.

14.2.1 Methods Based on the Equivalence Theorem

The idea behind an equivalence theorem is the following. From real variable calculus, we know that if a function, $h(x)$ say, has a local maximum at $x = x_0$, then its derivative $h'(x)$ equals zero at x_0 . A similar reasoning can be applied to the objective functions of optimality criteria, which take information matrices as their argument. In this case, the Fréchet directional derivative is used, and at the optimal information matrix, it has to be nonpositive in all directions. (From the top of the mountain, you cannot go further up.)

In what follows, we consider the situation where interest is in the whole parameter vector θ . Moreover, we restrict attention to the ϕ_p -optimality criteria where $p > -\infty$. Specific equivalence results for subsets of θ and further criteria that are information functions in the sense of Pukelsheim (1993) can be found, for example, in Chapter 7 of that book. A general version of the equivalence theorem is also presented in Chapters 2 and 13.

The equivalence theorem for ϕ_p -optimality is as follows.

Theorem 14.1 *The design ξ^* is ϕ_p -optimal on \mathcal{X} if and only if*

$$f'(x, \theta)M_{\xi^*}^{p-1}f(x, \theta) - \text{trace}(M_{\xi^*}^p) \leq 0 \quad \forall x \in \mathcal{X}. \quad (14.6)$$

Moreover, equality applies in (14.6) for the support points of ξ^* .

Note that while the ϕ_p -optimal information matrix is unique for $p > -\infty$, there may be more than one design to achieve this matrix. In this situation, equality in (14.6) is attained at the support points of any of these designs.

Example 14.2 (Example 14.1 continued)

Suppose we seek a D -optimal design for the Michaelis–Menten model. This problem is tackled in three standard steps:

- Step 1: Use the equivalence theorem to find the number of support points of the D -optimal design on $\mathcal{X} = [0, B]$ (for any value of θ).
- Step 2: Show that the larger support point of the D -optimal design is given by B , the upper boundary of the design region \mathcal{X} .
- Step 3: Find the smaller support point of the D -optimal design.

Step 1: We require that ξ has at least two support points to have nonsingular Fisher information M_ξ . For D -optimality, Theorem 14.1 simplifies to the following.

Corollary 14.1 *The design ξ^* is D -optimal for θ if and only if the inequality*

$$d(\xi^*, x, \theta) = f'(x, \theta)M_{\xi^*}^{-1}f(x, \theta) - m \leq 0 \quad (14.7)$$

holds for all $x \in \mathcal{X}$, with equality in the support points of ξ^* .

Consider inequality (14.7) for the Michaelis–Menten model with arbitrary parameter value θ and a D -optimal design ξ^* . In this case,

$$d(\xi^*, x, \theta) = \frac{x^2}{(\theta_2 + x)^2} m_{1,1} - \frac{2\theta_1 x^2}{(\theta_2 + x)^3} m_{1,2} + \frac{\theta_1^2 x^2}{(\theta_2 + x)^4} m_{2,2} - 2,$$

where m_{ij} , $i, j = 1, 2$, denotes the (i, j) -element of $M_{\xi^*}^{-1}$ and $m_{2,1} = m_{1,2}$. Multiplying (14.7) through with $(\theta_2 + x)^4$, we obtain a polynomial of degree four, $p_4(x)$ say, on the left-hand side. We now count its possible number of roots, the support points. A polynomial of degree four can have at most four roots. However, if $p_4(x)$ had four roots, at least the two middle ones would have to be turning points, since $p_4(x)$ must not become positive on \mathcal{X} .




FIGURE 14.2 Plot of a polynomial of degree four, which is nonpositive on $[0, 1]$ and attains its maximum, zero, at three points, including the end points 0 and 1.

Hence, the derivative of $p_4(x)$, a polynomial of degree three, would have at least five roots, which is a contradiction. Now suppose $p_4(x)$ has three roots on \mathcal{X} . By the same argument as before, only the middle root may be a turning point, so the other two roots have to be the end points of \mathcal{X} . A schematic of such a polynomial is depicted in Figure 14.2. Now substitute the lower end point into $d(\xi^*, x, \theta)$. Since $f(0, \theta) = 0$, we find that $d(\xi^*, 0, \theta) = -2 \neq 0$, so 0 cannot be a support point of the D -optimal design, which contradicts the assumption of a three-point design. Hence, the D -optimal design is supported on exactly two points.

Step 2: By (14.5), a D -optimal design with the number of support points t equal to m , the number of model parameters, (often called a saturated design) must have equal weights $w_1 = \dots = w_m = 1/m$. Hence, using (14.4), the objective function becomes

$$\phi_D(\xi, \theta) = |M_\xi| = \frac{1}{4} \frac{\theta_1^2 x_1^2 x_2^2 (x_2 - x_1)^2}{(\theta_2 + x_1)^4 (\theta_2 + x_2)^4}. \tag{14.8}$$

We note that the linear parameter θ_1 comes out as a factor and does therefore not affect the maximization of (14.8) with respect to the design. Without loss of generality, let x_2 be the larger support point, that is, $x_2 > x_1$. For the derivative of ϕ_D with respect to x_2 , we obtain

$$\frac{\partial \phi_D(\xi, \theta)}{\partial x_2} = \frac{\theta_1^2 x_1^2 x_2 (x_2 - x_1) [\theta_2 (x_2 - 0.5x_1) + 0.5x_1 x_2]}{(\theta_2 + x_1)^4 (\theta_2 + x_2)^5} > 0,$$

so ϕ_D is increasing in x_2 and is thus maximized at the upper boundary, B , of \mathcal{X} .




FIGURE 14.3

Plot of $d(\xi^*, x, \theta)$ for the Michaelis–Menten model with parameter vector $\theta' = (1, 0.6)$ and the D -optimal design ξ^* on the design space $\mathcal{X} = [0, 1]$.

Step 3: Substitute $x_2 = B$ into (14.8) and solve $\partial\phi_D(\xi, \theta)/\partial x_1 = 0$ for x_1 . There are three solutions: 0 , $B\theta_2/(B + 2\theta_2)$, and B . For both $x_1 = 0$ and $x_1 = B$, the objective function becomes zero, so these points correspond to local (and global) minima. Hence, the point $x_1 = B\theta_2/(B + 2\theta_2)$, situated between the two, has to correspond to the only local maximum, which is also global on \mathcal{X} since the values attained at the end points are minima.

We finally check for one example that the design obtained is indeed D -optimal. Figure 14.3 shows $d(\xi^*, x, \theta)$ for the D -optimal design ξ^* with parameter vector $\theta' = (1, 0.6)$ and design region $\mathcal{X} = [0, 1]$. The conditions of the equivalence theorem are clearly satisfied.

Note that the same strategy can be applied to search for other ϕ_p -optimal designs. One major difference is that, unlike the D -optimal design, other ϕ_p -optimal designs may depend on the value of the linear parameter θ_1 . The other difference is that the optimal weights are not readily available in closed form, which requires either the use of formula (14.5) or a further variable over which to optimize in Step 3.

14.2.2 Geometric Approach

Elfving's theorem (see Elfving 1952) is a powerful tool for characterizing c -optimal designs. Namely, the optimal design can be found at the intersection of a straight line representing the vector c and the boundary of a convex set referred to as the Elfving's set. Elfving's set is determined by the design space \mathcal{X} and the regression model. Similar results were established for the D -optimality criterion (Dette 1993) and for Bayesian optimality criteria (Dette 1996). A related approach based on covering ellipses was introduced by Sibson (1972), Silvey (1972), and Silvey and Titterton (1973) for D -optimality and subsequently used by, for example, Ford et al. (1992) and Haines (1993). This method was extended to E -optimality for linear and nonlinear models by Dette and Studden (1993a,b) and




FIGURE 14.4 Parametric plot of the induced design space \mathcal{G} (solid line) for the Michaelis–Menten model with parameter vector $\theta' = (1, 0.6)$ and design space $\mathcal{X} = [0, 1]$ and its reflection $-\mathcal{G}$ (dotted line). Horizontal axis, $\partial\eta(x, \theta)/\partial\theta_1$; vertical axis, $\partial\eta(x, \theta)/\partial\theta_2$.

Dette and Haines (1994). In an integrated approach, Biedermann et al. (2006) generalized this method, for two-parameter models, to the class of ϕ_p -optimality criteria. We will briefly review the results by Biedermann et al. (2006) and Elfving (1952), and illustrate them through an example.

Both approaches use the concept of an induced design space, \mathcal{G} , where $\mathcal{G} = \{f(x, \theta), x \in \mathcal{X}\}$ with $f(x, \theta)$ defined in (14.3). In what follows, we require \mathcal{G} to be compact and, for every admissible value of θ , $f(x, \theta)$ to be continuous in x . The former assumption is trivially satisfied if the latter assumption holds and if \mathcal{X} is compact.

Figure 14.4 shows a parametric plot of the induced design space for the Michaelis–Menten model with parameter vector $\theta' = (1, 0.6)$ and design space $\mathcal{X} = [0, 1]$ as a solid line. Its reflection $-\mathcal{G}$ has been added as a dotted line. The axes are given by the entries of the vector $f(x, \theta)$, that is, the horizontal axis is $\partial\eta(x, \theta)/\partial\theta_1$, and the vertical axis is $\partial\eta(x, \theta)/\partial\theta_2$.

14.2.2.1 Elfving’s Theorem and Its Implications for c -Optimality

Consider a vector c and designs ξ such that $c \in \text{range}(M_\xi)$ to ensure estimability of $c'\theta$ (see, e.g., Pukelsheim 1980). Define Elfving’s set, \mathcal{E} , by

$$\mathcal{E} = \text{co}(\mathcal{G} \cup -\mathcal{G}),$$

where $\text{co}(A)$ means the convex hull of a set $A \subset \mathbb{R}^m$. Elfving’s theorem characterizes a c -optimal design in terms of the intersection of $\{\gamma c \mid \gamma > 0\}$ with the boundary of \mathcal{E} .

Theorem 14.2 A design $\xi^* = \{(x_1, w_1), \dots, (x_t, w_t)\}$ is c -optimal for estimating $c'\theta$ if and only if there exists a positive number $\gamma^* > 0$ and real numbers $\varepsilon_1, \dots, \varepsilon_t \in \{-1, 1\}$ such that the point $\gamma^*c = \sum_{i=1}^t w_i \varepsilon_i f(x_i, \theta)$ is a boundary point of Elfving's set \mathcal{E} .

To see how this result can be used to find c -optimal designs, consider the following example.

Example 14.3 (Example 14.1 continued)

Suppose an experimenter is interested in estimating percentiles x_r of the Michaelis–Menten curve, that is, values of x , for which one expects a proportion r of the supremum of the reaction rate, θ_1 , to be attained. For $r \in (0, 1)$, we solve the equation

$$\frac{\theta_1 x}{\theta_2 + x} = r\theta_1$$

for x to obtain $x_r = r\theta_2/(1 - r)$. The problem of estimating x_r for fixed but arbitrary $r \in (0, 1)$ is therefore equivalent to estimating θ_2 , and we seek a c -optimal design where $c = (0, 1)'$.

Now consider the shape of Elfving's set \mathcal{E} depicted in Figure 14.5 (for $\theta = (1, 0.6)'$ and $\mathcal{X} = [0, 1]$). Note that $\{\gamma(0, 1)' \mid \gamma > 0\}$ intersects the boundary of \mathcal{E} for some positive γ^* at the point $(0, \gamma^*)'$ on the vertical axis. This point is a convex combination of two points (shown as circles): one at the right end point of \mathcal{G} and one on $-\mathcal{G}$, since they are all on the same straight line.




FIGURE 14.5

Parametric plot of Elfving's set, \mathcal{E} , for the Michaelis–Menten model with parameter vector $\theta' = (1, 0.6)$ and design space $\mathcal{X} = [0, 1]$. Solid line, induced design space \mathcal{G} ; dotted line, its reflection $-\mathcal{G}$; dashed line, convex hull of $\mathcal{G} \cup -\mathcal{G}$; vertical arrow, vector c ; circles, c -optimal support points (or their reflections). Horizontal axis, $\partial\eta(x, \theta)/\partial\theta_1$; vertical axis, $\partial\eta(x, \theta)/\partial\theta_2$.

This tells us that the c -optimal design must have two support points, x_1 and x_2 (without loss of generality, let $x_1 < x_2$), satisfying

$$\begin{aligned} \begin{pmatrix} 0 \\ \gamma^* \end{pmatrix} &= w_1 \varepsilon_1 \begin{pmatrix} \frac{x_1}{\theta_2+x_1} \\ -\theta_1 x_1 \\ (\theta_2+x_1)^2 \end{pmatrix} + (1-w_1) \varepsilon_2 \begin{pmatrix} \frac{x_2}{\theta_2+x_2} \\ -\theta_1 x_2 \\ (\theta_2+x_2)^2 \end{pmatrix} \\ &= -w_1 \begin{pmatrix} \frac{x_1}{0.6+x_1} \\ -x_1 \\ (0.6+x_1)^2 \end{pmatrix} + (1-w_1) \begin{pmatrix} \frac{1}{0.6+1} \\ -1 \\ (0.6+1)^2 \end{pmatrix}. \end{aligned}$$

We have used that $x_2 = 1$ since the right end point of \mathcal{G} is attained at the upper bound of \mathcal{X} and $\varepsilon_1 = -1, \varepsilon_2 = 1$ since they correspond to points on $-\mathcal{G}$ and \mathcal{G} , respectively. As we do not know the value of γ^* , we cannot use this system of equations to find x_1 and w_1 , but we can substitute $x_2 = 1$ into the objective function, $\phi_c(\xi) = c' M_\xi^- c$ (see Section 14.1), and minimize either analytically or numerically with respect to x_1 and w_1 . Alternatively, we can use the weight formula by Pukelsheim and Torsney (1991) to find the optimal weight w_1 in terms of x_1 to reduce the number of variables in the optimization problem. For $\theta = (1, 0.6)'$ and $\mathcal{X} = [0, 1]$, we obtain $x_1 = 0.210$ and $w_1 = 0.707$ (to three decimal places).

14.2.2.2 Characterization of ϕ_p -Optimal Designs via Covering Ellipses

Biedermann et al. (2006) express the ϕ_p -optimal design problem for two-parameter models in terms of a dual problem. They define a weighted volume, called v_{2q} -content, for ellipses where the weight depends on the choice of p . They then show that finding a ϕ_p -optimal design is equivalent to finding a centered ellipse, which covers the induced design space and which has minimal v_{2q} -content.

Theorem 14.3 *Let N be a nonnegative definite matrix with eigenvalues λ_1 and λ_2 , and let q be determined by the equation $p + q = pq$. Define the v_{2q} -content of the ellipse $E_N = \{u \in \mathbb{R}^2 | u'Nu \leq 1\}$ as*

$$v_{2q}(E_N) = \frac{\text{Vol}(E_N)}{l_{2q}(E_N)} = \frac{\pi/\sqrt{\lambda_1\lambda_2}}{\{[(2/\sqrt{\lambda_1})^{2q} + (2/\sqrt{\lambda_2})^{2q}]/2\}^{1/(2q)'}}$$

where $\text{Vol}(E_N)$ denotes the volume of the ellipse E_N and $l_{2q}(E_N)$ is the l_{2q} -mean of the lengths of its major and minor diameters.

Then the ϕ_p -optimal design problem is the dual of finding a centered ellipse E_N , which covers the induced design space \mathcal{G} and has minimal v_{2q} -content. Moreover, if x_i^* is a support point of a ϕ_p -optimal design, the ellipse E_N touches \mathcal{G} at $f(x_i^*, \theta)$.

For $p = 0$, we have $q = 0$ where $l_{2 \times 0}(E_N)$ is defined as $\lim_{q \downarrow 0} l_{2q}(E_N)$. Using L'Hôpital's rule, we can show that $\lim_{q \downarrow 0} l_{2q}(E_N) = 2(\lambda_1\lambda_2)^{-1/4} = 2\sqrt{\text{Vol}(E_N)/\pi}$. In this case, the dual problem reduces to the well-known geometric interpretation of the D -optimal design problem.

We consider the following example in order to illustrate how the calculation of optimal designs can be facilitated by this method.

Example 14.4 (Example 14.1 continued)

For arbitrary $p \in [-\infty, 1]$, suppose the aim is to find a ϕ_p -optimal design for the Michaelis–Menten model. From Figure 14.4, we can see that a centered ellipse that covers the induced design space \mathcal{G} must touch \mathcal{G} in exactly two points to have minimal v_{2q} -content. One of these points is the right end point of \mathcal{G} corresponding to the upper boundary of the design space \mathcal{X} . This general form of design does not depend on the value of p . An example showing the D -optimal ($p = 0$) covering ellipse is depicted in Figure 14.6.

This approach provides a geometric characterization of the optimal support points, but not of the optimal weights. Biedermann et al. (2006) present an example where all ϕ_p -optimal designs have the same covering ellipse with minimal v_{2q} -content, but different weights. The standard strategy for finding a ϕ_p -optimal design would therefore follow similar steps as we have seen in the section on the equivalence theorem:

1. Establish the number of support points by visual inspection (two, in this example).
2. Identify as many support points as possible (the upper boundary of \mathcal{X} , in this example).
3. Substitute this information into the objective function, and optimize with respect to the remaining support points and the weights.

Again, if the optimal design is saturated, the method by Pukelsheim and Torsney (1991) can be used to find the optimal weights.




FIGURE 14.6

Parametric plot of the induced design space \mathcal{G} (solid line) for the Michaelis–Menten model with parameter vector $\theta' = (1, 0.6)$ and design space $\mathcal{X} = [0, 1]$, its reflection $-\mathcal{G}$ (dotted line), with D -smallest covering ellipse (dashed line) and D -optimal support points. Horizontal axis, $\partial\eta(x, \theta)/\partial\theta_1$; vertical axis, $\partial\eta(x, \theta)/\partial\theta_2$.

14.2.3 Functional Approach

The idea behind the functional approach is the following. Suppose the design problem has been reduced to an optimization problem, for which the optimum is attained in the interior of its domain. Hence, the solution can be found by setting the gradient (with respect to the design variables) of the objective function to zero and solving for the design variables. By *design variables* we mean all support points and weights of the optimal design that require calculation. Now this gradient, $g(\tau, \theta)$ say, depends on two sets of *variables*, the design variables, τ say, and the model parameters held in the vector θ . Under some regularity assumptions, the implicit function theorem implies that in the neighborhood U of an arbitrary vector θ_0 for which we have a vector τ_0 with $g(\tau_0, \theta_0) = 0$, there exists a function $\tau = \tau(\theta)$ such that for all $\theta \in U$, we obtain $g(\tau(\theta), \theta) = 0$, and $\tau(\theta)$ can be expanded into a Taylor series. The coefficients for this series can be obtained by recursive formulae provided in Melas (2006).

This means that once we have found an optimal design τ_0 with respect to one parameter vector θ_0 , we can approximate optimal designs $\tau(\theta)$ for different values of θ by the first few terms of their Taylor polynomials about θ_0 . An essential assumption for this approach is that there are as many design variables as there are model parameters to make the Jacobian of $\tau(\theta)$ invertible. A natural application is therefore to find saturated D -optimal designs. In many other situations, the approach can still be applied after using some properties of the optimal design and/or restricting/transforming the parameter space. More details can be found in Melas (2006) and the references therein. The method is easy to implement in software allowing symbolic calculations, such as Mathematica (Wolfram 2012) or Maple (Monagan et al. 2005), and usually only a few coefficients are required for a good approximation to the true function $\tau(\theta)$, provided the interval for each component of θ is not too wide.

Example 14.5 (Examples 14.1, 14.2, 14.4 continued)

Suppose we seek D -optimal designs for the Michaelis–Menten model on the design space $\mathcal{X} = [0, 1]$, for different values of the parameter vector θ . Further assume we have already established that 1, the upper bound of \mathcal{X} , is a support point and that the optimal designs depend only on θ_2 . Hence, we aim to approximate the smaller support point $x_1 = \tau(\theta_2)$ as a function of θ_2 . We note that for this particular example, $\tau(\theta_2) = \theta_2/(1 + 2\theta_2)$ can be found explicitly by a simple calculation (see, e.g., Step 3 in Example 14.2), so an approximation would not be necessary in practice. It is still useful for illustration of the method.

From (14.8), with $x_2 = 1$, we have to maximize $x_1^2(1 - x_1)^2/(\theta_2 + x_1)^4$, so we set its derivative with respect to x_1 equal to zero, which, after some algebra and observing that $x_1 \neq 0, 1 - x_1 \neq 0$, and $\theta_2 + x_1 \neq 0$, is equivalent to

$$g(x_1, \theta_2) = x_1(1 + 2\theta_2) - \theta_2 = 0.$$

Assume we expand $g(x_1, \theta_2)$ about $\theta_{2,0} = 0.6$ (a different value could be chosen here); then the corresponding $\tau_0 = \tau(\theta_{2,0})$ is $3/11$. The recursive formulae from Melas (2006, p. 34), yield for the $(s + 1)$ th coefficient of the Taylor expansion

$$\tilde{\tau}_{s+1} = -J_0^{-1} \frac{1}{(s + 1)!} \frac{\partial^{s+1} g(\hat{\tau}_s(\theta_2), \theta_2)}{\partial \theta_2^{s+1}}, \quad s = 0, 1, \dots,$$

where J_0 is the Jacobian of $g(\tau, \theta_2)$ with respect to τ , evaluated at $(\tau_0, \theta_{2,0})$, and

$$\hat{\tau}_s(\theta_2) = \tilde{\tau}_0 + \sum_{j=1}^s \tilde{\tau}_j(\theta_2 - \theta_{2,0})^j$$




FIGURE 14.7

Taylor series approximations to the smaller support point, x_1 , of the D -optimal design for the Michaelis–Menten model with design space $\mathcal{X} = [0, 1]$ in terms of the parameter θ_2 . Solid line, true curve; dotted line, Taylor polynomial of degree four; dashed line, Taylor polynomial of degree two. The Taylor series are centered about $\theta_2 = 0.6$.

is the Taylor expansion from the s th step. Therefore, the function $g(\hat{\tau}_s(\theta_2), \theta_2)$ depends on θ_2 in both components (the first component is a polynomial in θ_2), which must be taken into account when working out its partial derivatives with respect to θ_2 .

We obtain $J_0 = 1 + 2\theta_{2,0} = 2.2$, so $-J_0^{-1} = -1/2.2 = -0.45$ and $\hat{\tau}_1 = -0.45(2\tau_0 - 1) = 0.2066$ to 4 decimal places. Hence, $\hat{\tau}_1(\theta_2) = 3/11 + 0.2066(\theta_2 - 0.6)$. For $\hat{\tau}_2$ we require the second derivative of $g(3/11 + 0.2066(\theta_2 - 0.6), \theta_2) = 3/11 + 0.2066(\theta_2 - 0.6) + 2\theta_2[3/11 + 0.2066(\theta_2 - 0.6)] + \theta_2$ with respect to θ_2 . This is given by $4 \times 0.2066 = 0.8264$, and we obtain for the second coefficient $\hat{\tau}_2 = -J_0^{-1} \times 0.8264/2! = -0.1878$ (4 decimal places). Continuing in this manner yields the next coefficients $\hat{\tau}_3 = 0.1708$, $\hat{\tau}_4 = -0.1552$, $\hat{\tau}_5 = 0.1411$, and $\hat{\tau}_6 = -0.1283$.

Figure 14.7 shows two Taylor approximations about $\theta_2 = 0.6$ to the true function $x_1 = \tau(\theta_2)$ on the domain $\theta_2 \in [0.1, 1.1]$. The Taylor polynomial of degree four is virtually identical to $\tau(\theta_2)$ across this interval. On a smaller interval, for example, $\theta_2 \in [0.4, 0.8]$, the Taylor polynomial of degree two is already a good approximation.

14.3 General Solutions

The *classical* methods have in common that design problems are solved on a case-by-case basis. Each combination of model and optimality criterion requires its own proof. There is a recent development toward more general solutions of design problems. Yang and Stufken (2009) considered nonlinear models with two parameters, and algebraically derived conditions, under which for each design ξ there is a design $\tilde{\xi}$ from a simple class, which is

at least as good as ξ in terms of the Loewner ordering of the corresponding information matrices, that is, $M_{\tilde{\xi}} - M_{\xi}$ is positive semidefinite. For any reasonable optimality criterion maximizing an objective function $\phi(\xi)$, $M_{\tilde{\xi}} \geq M_{\xi}$ implies that $\phi(\tilde{\xi}) \geq \phi(\xi)$, that is, optimality criteria should be isotonic relative to the Loewner ordering. The objective functions in the sense of Pukelsheim (1993) satisfy this condition.

These results were subsequently generalized to models with an arbitrary number of parameters by Yang (2010), using algebraic methods, and Dette and Melas (2011), using properties of Chebyshev systems (see, e.g., Karlin and Studden 1966). In particular, for large classes of models and arbitrary optimality criteria based on the information matrix, these papers provide considerably tighter upper bounds on the maximal number of support points than Carathéodory’s bound, which is $m(m + 1)/2 + 1$ (see, e.g., Silvey 1980, or Pukelsheim 1993). These results greatly reduce the computational effort required to find optimal designs. In many situations, the optimal designs are saturated, that is, they have as many support points as the model has parameters to be estimated. For a p th degree polynomial model, de la Garza (1954) showed that for any n -point design where $n > p + 1$, there exists an alternative design with exactly $p + 1$ support points such that the two designs have identical information matrices. This result has subsequently been named the de la Garza phenomenon. A further extension of the algebraic method, which can result in finding even smaller complete classes for optimal designs, can be found in Yang and Stufken (2012). An alternative proof for these results, based on Chebyshev systems, is presented in Dette and Schorning (2013).

14.3.1 Algebraic Method

The method proposed in Yang and Stufken (2009) and Yang (2010) uses a transformation of the information matrix for a design ξ , of the form

$$M_{\xi} = P(\theta) \left[\sum_{i=1}^t w_i C(\theta, z_i) \right] P(\theta)', \tag{14.9}$$

where

$$C(\theta, z_i) = \begin{bmatrix} \Psi_{1,1}(z_i) & \cdots & \Psi_{1,m}(z_i) \\ \vdots & \ddots & \vdots \\ \Psi_{1,m}(z_i) & \cdots & \Psi_{m,m}(z_i) \end{bmatrix},$$

say, and $P(\theta)$ is a $m \times m$ nonsingular matrix that may depend on the value of θ , but on no other parameters. For fixed θ and an interval design space \mathcal{X} , the map from $x \in \mathcal{X}$ to $z \in [Z_l, Z_u]$ is one to one and onto, and a design ξ can be expressed in terms of its transformed support points z_i , $i = 1, \dots, t$, and its weights. The functions $\Psi_{l,s}$, $l, s = 1, \dots, m$, depend on θ , except, perhaps, for models where optimal design questions do not depend on θ . This dependence could occur just through z . The matrices $P(\theta)$ and $C(\theta, z)$ are not unique, and are chosen in such a way that the functions in $C(\theta, z)$ are as simple as possible.

We now reduce the number of functions $\Psi_{l,s}$ we need to consider. To this end, keep only a maximal set of nonconstant linearly independent functions $\Psi_{l,s}$, $1 \leq l \leq s \leq m$. Rename these to Ψ_1, \dots, Ψ_k , such that the last one, Ψ_k , is one of the diagonal elements of $C(\theta, z)$ and that there is no other function $\Psi_{l,s} = \Psi_k$ in the matrix $C(\theta, z)$. The idea behind this

approach is to show that for each design $\xi = \{(z_1, w_1), \dots, (z_t, w_t)\}$, there exists a design $\tilde{\xi} = \{(\tilde{z}_1, \tilde{w}_1), \dots, (\tilde{z}_t, \tilde{w}_t)\}$ from a simple class, for which $\sum_{i=1}^t \tilde{w}_i \Psi_j(\tilde{z}_i) = \sum_{i=1}^t w_i \Psi_j(z_i)$ for $j = 1, \dots, k-1$ and $\sum_{i=1}^t \tilde{w}_i \Psi_k(\tilde{z}_i) \geq \sum_{i=1}^t w_i \Psi_k(z_i)$, which makes $\tilde{\xi}$ at least as good as ξ in the Loewner ordering. A *simple class* of designs here means a class of designs that have a considerably tighter upper bound on the number of support points than given by Carathéodory's bound and in some cases also include one or both of the boundary points of the design space; see also Theorem 14.4.

If Ψ_1, \dots, Ψ_k are differentiable arbitrarily often, define the functions $\psi_{l,s}$ recursively as

$$\psi_{l,s}(z) = \begin{cases} \Psi'_l(z), & s = 1, l = 1, \dots, k, \\ \left(\frac{\psi_{l,s-1}(z)}{\psi_{s-1,s-1}(z)} \right)', & 2 \leq s \leq k, s \leq l \leq k. \end{cases} \quad (14.10)$$

A more detailed chart describing how to obtain the functions $\psi_{l,s}(z)$ in (14.10) is provided in Yang (2010). If the functions $\psi_{l,l}$, $l = 1, \dots, k$, have no root in the transformed design space $[Z_l, Z_u]$, the following result holds (see Yang 2010).

Theorem 14.4 Let $\Phi(z) = \prod_{l=1}^k \psi_{l,l}(z)$, $z \in [Z_l, Z_u]$. For any given design ξ , there exists a design $\tilde{\xi}$, such that $M_\xi \leq M_{\tilde{\xi}}$ in the Loewner ordering.

- When k is odd and $\Phi(z) < 0$, $\tilde{\xi}$ can be taken to have at most $(k+1)/2$ support points including point Z_l .
- When k is odd and $\Phi(z) > 0$, $\tilde{\xi}$ can be taken to have at most $(k+1)/2$ support points including point Z_u .
- When k is even and $\Phi(z) > 0$, $\tilde{\xi}$ can be taken to have at most $k/2 + 1$ support points including points Z_l and Z_u .
- When k is even and $\Phi(z) < 0$, $\tilde{\xi}$ can be taken to have at most $k/2$ support points.

Note that the general formula for computing $\Phi(z)$ can easily be implemented in software that is capable of symbolic calculations. Furthermore, even if $\Phi(z)$ has a complicated structure, making it impossible to check directly if this function is positive/negative on $[Z_l, Z_u]$, we can easily obtain this information from visual inspection of its graph.

The extension by Yang and Stufken (2012) uses a similar idea where, instead of just one entry, $\Psi_k(z)$, a lower principal submatrix (of arbitrary size $m_1 \times m_1$ where $1 \leq m_1 < m$) of the matrix $C(\theta, z)$ is considered. Thus, the results presented earlier are included as a special case where $m_1 = 1$. The more general method can be applied to many models, which do not satisfy the conditions in Yang (2010). Moreover, for some models for which answers could be obtained by earlier work, the extension finds smaller simple classes of designs.

To fix ideas on how to apply Theorem 14.4, consider the following example.

Example 14.6 (Example 14.1 continued)

Consider the Michaelis–Menten model. After some algebra, we find that the information matrix (14.4) can be written in form (14.9) with

$$P(\theta) = \begin{bmatrix} 1/\theta_1 & 0 \\ -1/\theta_2 & 1/(\theta_1\theta_2) \end{bmatrix}, \quad C(\theta, z) = \begin{bmatrix} z^2 & z^3 \\ z^3 & z^4 \end{bmatrix},$$

where $z = \theta_1 x / (\theta_2 + x)$. Let $\Psi_1(z) = \Psi_{1,1}(z) = z^2$, $\Psi_2(z) = \Psi_{1,2}(z) = z^3$, and $\Psi_3(z) = \Psi_{2,2}(z) = z^4$. Then $k = 3$, $\Psi_k = \Psi_{l,l}$ for $l = 2$, and there is no $l < s$ with $\Psi_{l,s} = \Psi_k$. From (14.10), we find that

$$\begin{aligned} \psi_{1,1} &= \Psi'_1(z) = 2z \\ \psi_{2,2} &= \left(\frac{\psi_{2,1}(z)}{\psi_{1,1}(z)}\right)' = \left(\frac{\Psi'_2(z)}{\Psi'_1(z)}\right)' = \left(\frac{3z^2}{2z}\right)' = \frac{3}{2} \\ \psi_{3,3} &= \left(\frac{\psi_{3,2}(z)}{\psi_{2,2}(z)}\right)' = \left(\frac{[\Psi_{3,1}(z)/\psi_{1,1}(z)]'}{3/2}\right)' = \left(\frac{[\Psi'_3(z)/\Psi'_1(z)]'}{3/2}\right)' = \left(\frac{[4z^3/2z]'}{3/2}\right)' = \frac{8}{3}. \end{aligned}$$

Hence, from Theorem 14.4, $\Phi(z) = 2z \times 3/2 \times 8/3 = 8z > 0$ if $z > 0$. Since an observation in $x = 0$ (so $z = 0$) does not give any increase in information, we can choose $\mathcal{X} = [A, B]$ where A is small but positive instead of the interval $[0, B]$ we used before. Therefore $x > 0$, and we obtain that $z > 0$ since $\theta_1 > 0$ and $\theta_2 > 0$. That means we are in the situation of case (b) in Theorem 14.4. Hence, for each optimality criterion based on the information matrix, there exists an optimal design with no more than $(k + 1)/2 = 2$ support points, one of which is Z_u , which translates into B by the inverse map. This confirms our results from Examples 14.2 and 14.4. It remains to select an appropriate optimality criterion, and to use analytical or numerical search to obtain an optimal design from the class found.

Note that for larger values of k , we do not necessarily obtain the existence of a saturated optimal design. For example, if $m = 6$, k can be as large as 21, provided all $\Psi_{l,s}$, $1 \leq l \leq s \leq m$ are linearly independent. If $\Phi(z) \neq 0$ on $[Z_l, Z_u]$, there exists an optimal design with at most $(k + 1)/2 = 11$ support points, one of which is already determined. This leaves an optimization problem in 20 dimensions, 10 for the remaining support points and 10 for the weights. This is still a considerable improvement on Carathéodory's bound, where in the preceding situation we would have to solve an optimization problem in 41 dimensions (21 support points and 20 weights). However, for many models, some $\Psi_{l,s}$ occur repeatedly in the information matrix, and replicates can be excluded due to linear dependence. For example, for polynomials of degree five, we have $m = 6$, and the information matrix holds only 11 different functions of $z = x$, the monomials up to degree ten. The monomial of degree zero, that is, the function constant to 1, does not depend on the design, and hence, $k = 10$, and there exists an optimal design with six support points, two of which are the end points of the design interval \mathcal{X} . This result extends the result given by the de la Garza phenomenon. In the next subsection, we will further investigate for which classes of models the de la Garza phenomenon holds.

14.3.2 Method Based on Chebyshev Systems

Roughly speaking, Dette and Melas (2011) use properties of Chebyshev systems to establish the de la Garza phenomenon (or substantial improvements over Carathéodory's bound on the number of support points). Following Karlin and Studden (1966), a set of $k + 1$ continuous functions $\{u_0, \dots, u_k\}$ is called a Chebyshev system on the interval $[Z_l, Z_u]$ if

$$\begin{vmatrix} u_0(z_0) & \cdots & u_0(z_k) \\ \vdots & \ddots & \vdots \\ u_k(z_0) & \cdots & u_k(z_k) \end{vmatrix} > 0 \tag{14.11}$$

for all $Z_l \leq z_0 < \dots < z_k \leq Z_u$. Note that if the determinant in (14.11) is negative, then the set $\{u_0, \dots, -u_k\}$ is a Chebyshev system on $[Z_l, Z_u]$.

As in Section 14.3.1, denote the linearly independent elements of the transformed information matrix $C(\theta, z)$ defined in (14.9) as Ψ_1, \dots, Ψ_k , such that Ψ_k is one of the diagonal elements, and that there is no $\Psi_{l,s} = \Psi_k$ for $l < s$. Assume that Ψ_1, \dots, Ψ_k are all continuous, so infinite differentiability of the Ψ_i as in Yang (2010) is not required. Furthermore, let $\Psi_0(z) = 1$, and define the index, $I(\xi)$, of a design ξ , on the interval $[Z_l, Z_u]$ as the number of support points, where the boundary points, Z_l and Z_u , are only counted by 1/2.

Suppose that the sets $\{\Psi_0, \Psi_1, \dots, \Psi_{k-1}\}$ and $\{\Psi_0, \Psi_1, \dots, \Psi_k\}$ are Chebyshev systems. Then the following result holds; see Theorem 3.1 in Dette and Melas (2011).

Theorem 14.5 *For any design ξ , there exists a design $\tilde{\xi}$ with at most $(k+2)/2$ support points such that $M_{\tilde{\xi}} \geq M_{\xi}$. If $I(\xi) < k/2$, then $\tilde{\xi} = \xi$. Otherwise:*

- (a) *If k is odd, so that $\tilde{\xi}$ has at most $(k+1)/2$ support points, one point can be taken as Z_u .*
- (b) *If k is even, so that $\tilde{\xi}$ has at most $k/2 + 1$ support points, two points can be taken as Z_l and Z_u .*

If the sets $\{\Psi_0, \Psi_1, \dots, \Psi_{k-1}\}$ and $\{\Psi_0, \Psi_1, \dots, -\Psi_k\}$ are Chebyshev systems, a similar result holds, with the point Z_u in (a) replaced by Z_l , $k/2 + 1$ in (b) replaced by $k/2$, and assurance about end points being part of the support removed in (b).

Several sets of functions, for example, the monomials up to degree m for any integer m , are known to be Chebyshev systems. If this information is not available, using the definition given in (14.11) can be unwieldy, in particular if k is large. In this situation, it is usually easier to check the condition on $\Phi(z)$ from the algebraic method described in Section 14.3.1.

Example 14.7 (Examples 14.1 and 14.6 continued)

To apply this result to the Michaelis–Menten model, we need to check if the sets of functions $S_2 = \{\Psi_0(z), \Psi_1(z), \Psi_2(z)\}$ and $S_3 = S_2 \cup \{\Psi_3(z)\}$ are Chebyshev systems on $[Z_l, Z_u]$ where $Z_l \geq 0$, $\Psi_0(z) = 1$, $\Psi_1(z) = z^2$, $\Psi_2(z) = z^3$, and $\Psi_3(z) = z^4$. For S_2 we obtain

$$\begin{vmatrix} 1 & 1 & 1 \\ z_0^2 & z_1^2 & z_2^2 \\ z_0^3 & z_1^3 & z_2^3 \end{vmatrix} = (z_2 - z_1)(z_2 - z_0)(z_1 - z_0)(z_1 z_2 + z_1 z_0 + z_2 z_0) > 0$$

for $Z_l \leq z_0 < z_1 < z_2 \leq Z_u$. Similarly, the determinant for S_3 is $(z_3 - z_1)(z_3 - z_0)(z_3 - z_2)(z_2 - z_1)(z_2 - z_0)(z_1 - z_0)(z_0 z_1 z_2 + z_1 z_2 z_3 + z_0 z_2 z_3 + z_0 z_1 z_3)$, which is also positive for $Z_l \leq z_0 < z_1 < z_2 < z_3 \leq Z_u$. Hence, S_2 and S_3 are both Chebyshev systems on $[Z_l, Z_u]$.

Here, $k = 3$, so $(k+2)/2 = 2.5$, and for any design ξ , the dominating design $\tilde{\xi}$ has at most 2 support points. Only a design ξ with support points Z_l and Z_u or a one-point design can achieve an index $I(\xi) < k/2 = 1.5$. Hence, such a design is dominated by itself. Any other design will be dominated by a design $\tilde{\xi}$ with at most $(k+1)/2 = 2$ support points, one of which is the upper bound of the transformed design interval, Z_u , which translates into the upper bound, B , of the original design space \mathcal{X} . Note that if interest is in estimating both model parameters, any dominating design must have exactly two support points to ensure estimability. It remains to select an appropriate optimality criterion and to use analytical evaluation or numerical search to obtain an optimal design from the class found.

14.4 Further Examples

In this section, we will apply the methods described earlier to further examples. Unlike the previous sections, we will not apply every method to every model but present only a combination of the most suitable methods for each situation.

14.4.1 Two-Parameter Exponential Model

Exponential growth models with expected response $\eta(x, \theta)$ of the form

$$\eta(x, \theta) = \sum_{l=1}^L a_l e^{-b_l x}, \quad a_l > 0, \quad l = 1, \dots, L, \tag{14.12}$$

occur in chemical kinetics (see, e.g., Gibaldi and Perrier 1982), with particular emphasis on toxicology (see Becka et al. 1992, 1993) and microbiology (see Alvarez et al. 2003). Locally D -, c -, and E -optimal designs for this class of models have been found in Dette et al. (2006a), Dette et al. (2006b), and Ermakov and Melas (1995).

Example 14.8

For the purpose of this example, we assume that $L = 1$ in (14.12), and for consistency rename the parameters to obtain $\eta(x, \theta) = \theta_1 e^{-\theta_2 x}$. We further let $\theta_2 > 0$ and $\mathcal{X} = [0, B]$ for some $B > 0$. Note that knowledge of the sign of θ_2 is not a restrictive assumption, since the experimenter will usually know whether to expect growth or decline. By (14.2), the information matrix for this model for a design ξ is given by

$$M_\xi = \sum_{i=1}^t w_i \begin{bmatrix} e^{-2\theta_2 x_i} & -\theta_1 x_i e^{-2\theta_2 x_i} \\ -\theta_1 x_i e^{-2\theta_2 x_i} & \theta_1^2 x_i^2 e^{-2\theta_2 x_i} \end{bmatrix}. \tag{14.13}$$

Before selecting an optimality criterion, we seek the complete class of dominating designs in the Loewner ordering sense. Following the approaches presented in Section 14.3, we simplify the functions in the information matrix (14.13) using the transformation $z = \theta_2 x$, $z \in [0, \theta_2 B]$, and defining

$$P(\theta) = \begin{bmatrix} 1 & 0 \\ 0 & -\theta_1/\theta_2 \end{bmatrix}.$$

This yields the functions $\Psi_1(z) = e^{-2z}$, $\Psi_2(z) = ze^{-2z}$, and $\Psi_3(z) = z^2 e^{-2z}$. The algebraic method from Section 14.3.1 involves checking if $\Phi(z)$ as defined in Theorem 14.4, a function consisting of ratios of Ψ_1, Ψ_2, Ψ_3 and their derivatives, is positive/negative on $[0, \theta_2 B]$. We can see that the exponential term, e^{-2z} , will cancel in these ratios and therefore expect $\Phi(z)$ to have a simple form. Alternatively, we could use the method described in Section 14.3.2 and show that $\{1, \Psi_1, \Psi_2\}$ and $\{1, \Psi_1, \Psi_2, \pm\Psi_3\}$ are Chebyshev systems, which appears to be harder. It turns out that $\Phi(z) = -4e^{-2z}$, which is negative for all z . Hence, we stick to the algebraic method for this example.

Here, $k = 3$ is odd and $\Phi(z) < 0$ for all $z \in [Z_l, Z_u]$. Hence, from Theorem 14.4, we obtain that the dominating designs have at most $(k + 1)/2 = 2$ support points, one of which is Z_l or zero in the original design space. It thus remains to find the other support point (if necessary) and one weight. Note that if interest is in estimating both parameters, any optimal design will have two support points. If, however, we seek a c -optimal

design, this may have just one support point, zero, depending on where the vector c intersects the boundary of Elfving's set \mathcal{E} . Observing that $\eta(0, \theta) = \theta_1$, we find that only the parameter θ_1 , that is, the single linear combination $c'\theta$ where $c = (1, 0)'$, will be estimable by the one-point design supported at $x = 0$.

Assume the experimenter is only interested in estimating the rate, θ_2 , of exponential decay. In this case, the c -optimality criterion with $c = (0, 1)'$ will be appropriate, and the optimal design will have two support points. Recall the definition of X and $V = (XX')^{-1}X$ in Section 14.2. We find that the vector $v = Vc$ is given by $(1/(\theta_1 x_2), e^{\theta_2 x_2}/(\theta_1 x_2))'$. Therefore, from (14.5), the weight at $x = 0$ is $w = |v_1|/(|v_1| + |v_2|) = 1/(1 + e^{\theta_2 x_2})$. We note that w does not depend on the value of θ_1 and on θ_2 only through the product $\theta_2 x_2$.

We substitute the expression for w into the objective function $\phi_c(\xi, \theta) = c'M_{\xi}^{-1}c$ and obtain

$$\phi_c(\xi, \theta) = \frac{(1-w)e^{-2\theta_2 x_2} + w}{w(1-w)e^{-2\theta_2 x_2} \theta_1^2 x_2^2} = \frac{(1 + e^{\theta_2 x_2})^2}{\theta_1^2 x_2^2}.$$

Setting the derivative with respect to x_2 equal to zero is equivalent to solving

$$e^{\theta_2 x_2} (\theta_2 x_2 - 1) = 1,$$

which yields $x_2 = 1.278/\theta_2$ (to 3 decimal places). Inspection of the second derivative reveals that this is indeed a minimum. Hence, $\theta_2 x_2 = 1.278$ is constant, and the weight $w = 0.2178$ is constant, too, for any combination of θ_2 and the corresponding optimal value of x_2 .

If θ_2 is relatively small, the optimal x_2 becomes large and may not be included in the design interval $\mathcal{X} = [0, B]$. In this case, inspection of the first derivative of the objective function with respect to x_2 reveals that ϕ_c is strictly decreasing on \mathcal{X} , and therefore, the second support point has to be B . The corresponding optimal weight is then given by $w = 1/(1 + e^{\theta_2 B})$ and depends on the value of θ_2 .

In this example, substituting the expression for the optimal weight into the objective function resulted in a considerable simplification of the optimization problem. This is not necessarily always the case. For example, the objective function for A -optimality for the exponential model appears to become rather more complicated, possibly because of the square root terms involved in the weight formula (14.5).

14.4.2 Emax Model

Example 14.9

The Emax model is a generalization of the Michaelis–Menten model and is widely used in dose–response studies. The expected response is given by

$$\eta(x, \theta) = \theta_0 + \frac{\theta_1 x}{\theta_2 + x}, \quad \theta_1, \theta_2 > 0, \quad x \in [0, B],$$

where θ_0 represents the placebo response, θ_1 (often called E_{\max}) is the maximum achievable increase above the placebo response, and θ_2 is the dose that produces 50% of the E_{\max} effect.

Dette et al. (2010) show that D - and ED_r -optimal designs for this model have the same support points but different weights, where ED_r is the smallest effective dose that achieves a proportion of r , $0 < r < 1$, of the maximum effect in the observed dose range, and an ED_r -optimal design minimizes the variance of the ML-estimator for ED_r . Optimal

designs for estimating the minimum effective dose, that is, the smallest dose producing a practically relevant response, are given in Dette et al. (2008).

Again, we start with writing the information matrix in the form (14.9). Following Yang (2010), we use the transformation $z = 1/(\theta_2 + x)$ and

$$P(\theta) = \begin{bmatrix} 1 & 0 & 0 \\ 1 & -\theta_2 & 0 \\ 0 & -\theta_1 & \theta_1\theta_2 \end{bmatrix}$$

to obtain $\Psi_1(z) = z$, $\Psi_2(z) = z^2$, $\Psi_3(z) = z^3$, and $\Psi_4(z) = z^4$. We know that the monomials of degree $0, \dots, l$ form a Chebyshev system for every integer l ; hence, we can apply Theorem 14.5 immediately.

Here, $k = 4$, so for each design ξ , there exists a dominating design $\tilde{\xi}$ with at most $(k + 2)/2 = 3$ support points. In particular, provided the index $I(\xi) \geq k/2 = 2$, from part (b), we obtain that Z_l and Z_u can be chosen as support points, which translate back into the end points of the design interval \mathcal{X} . A design with index strictly less than 2 has strictly less than three support points, and thus produces a singular information matrix. If interest is in estimating all parameters, we can thus restrict design search to designs with three support points, including zero and B .

It may be of interest to know how well an optimal design for the Emax model performs if the true model is the Michaelis–Menten model. This corresponds to a situation where a placebo effect was anticipated and therefore taken into account when designing the experiment, but then it turned out that this parameter was unnecessary in the model, so a Michaelis–Menten model would be used in the analysis. For comparison with the designs found in Example 14.2, we seek the D -optimal design for the Emax model, on a design space $\mathcal{X} = [0, B]$. Since this model is partially nonlinear, the D -optimal design will not depend on the linear parameters θ_0 and θ_1 .

We know that the weights of this saturated D -optimal design will be equal, that is, $w_1 = w_2 = w_3 = 1/3$ (see, e.g., Silvey 1980 or (14.5)). Substituting these weights, together with the known support points, $x_1 = 0$ and $x_3 = B$, into the objective function, we obtain

$$\phi_D(\xi, \theta) = |M_\xi| = \frac{1}{3^3} \frac{\theta_1^2 B^2 x_2^2 (B - x_2)^2}{(\theta_2 + x_2)^4 (\theta_2 + B)^4}.$$

This is proportional to the objective function of the Michaelis–Menten model (14.8) and thus is also maximized by $x_2 = \theta_2 B / (B + 2\theta_2)$.

We find that two of the support points of the D -optimal design, ξ^* say, for the Emax model coincide with those of the D -optimal design for the Michaelis–Menten model. The third support point of ξ^* , $x_1 = 0$, however, does not provide any information for estimating the Michaelis–Menten model, that is, the information matrix in this point is the zero matrix. Another way to see this is by observing that $\eta(0, \theta) = 0$ and thus does not depend on the model parameters. To assess the performance of ξ^* in the Michaelis–Menten model, we compute its D -efficiency, where the D -efficiency of a design ξ is defined as

$$\text{eff}_D(\xi) = \left(\frac{|M_\xi|}{|M_{\xi_D}|} \right)^{1/m}, \tag{14.14}$$

with ξ_D the D -optimal design for the true scenario.

For this example, ξ_D is the D -optimal design for the Michaelis–Menten model and $m = 2$. We straightforwardly obtain that $M_{\xi^*} = 2/3 \times M_{\xi_D}$ for all eligible values of θ_2 and B , provided these are the same for both designs. Hence, regardless of the parameter values or the upper end point of the design interval, we have that $\text{eff}_D(\xi^*) = 2/3$. This is not surprising, since one-third of the observations, that is, those at $x_1 = 0$, are not used for inference of θ under the Michaelis–Menten model.

If the D -efficiency of a design ξ is equal to $\text{eff}_D(\xi)$, then this means that the number of observations in a D -optimal design can be limited to $100\text{eff}_D(\xi)\%$ of that in ξ to obtain the same accuracy under the D -optimality criterion. For example, if we conducted a clinical trial with 300 patients, using the D -optimal design for the Emax model, but then it turned out that there is no placebo effect and the Michaelis–Menten model is appropriate, we could get estimates for θ_1 and θ_2 with the same precision from a trial with 200 patients using the D -optimal design for the Michaelis–Menten model. In practice, however, it would not be known before analyzing the data from the trial that the placebo effect is not significant. At the planning stage, there are therefore two possible scenarios (placebo effect/no placebo effect) and two possible decisions (design for Emax/Michaelis–Menten model) for the experimenter to make. Taking into account that the D -optimal design for the Michaelis–Menten model is not capable of estimating/testing the presence/absence of the placebo effect, using design ξ^* seems to be the safer bet, even if some efficiency is lost if the smaller model is correct. In practice, a compromise design could be employed, for example, putting only weight 0.2 at point zero and weight 0.4 at the other two support points. This design has 80% D -efficiency in the smaller model and is capable of estimating/testing all parameters. There will, however, be some loss in efficiency if the Emax model is correct.

14.4.3 Heteroscedastic Linear Model

In some situations, it is not realistic to assume that the variability of observations is constant throughout the design region. The variance function may be of the form $\text{Var}(\epsilon_i) = \sigma^2(x_i, \alpha)$ for some parameter vector α , where the functional form of $\sigma^2(x_i, \alpha)$ is known. If this is not taken into account when planning the experiment, inefficient inference may result. Consider the class of heteroscedastic models where observations are described by

$$y_i = \eta(x_i, \theta) + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \sigma^2(x_i, \alpha)), \quad i = 1, \dots, n. \quad (14.15)$$

The function $\lambda(x_i, \alpha) = 1/\sigma^2(x_i, \alpha)$ is called the efficiency or intensity function. We consider the situation where there is no overlap between the vector of mean parameters, θ , and the vector of variance parameters, α . Moreover, we assume that α is a vector of nuisance parameters, and we are thus not interested in precise estimation of these; see also Chapter 2, Section 2.3.2. The more general case where α may overlap θ , or may be of interest, is beyond the scope of this text; see Fedorov and Hackl (1997) for an example.

For our situation, the information matrix for model (14.15) is block diagonal, where the first block is equal to the information matrix for θ only, that is,

$$M_{\xi} = \sum_{i=1}^t w_i \lambda(x_i, \alpha) f(x_i, \theta) f'(x_i, \theta),$$

and the second block is the information matrix for α only. The covariance matrix of $\hat{\theta}$ is thus the inverse of M_{ξ} , so in order to estimate θ with high precision, we consider ϕ_p -optimality with respect to M_{ξ} and not the complete information matrix for all model parameters.

We note that even in this simple case, where the parameters in the efficiency function do not overlap with the parameter vector θ in the expected response and are not of interest, optimal designs will be affected by heteroscedasticity. In particular, design problems for linear models with nonconstant variance resemble those for nonlinear models in that they depend on the nuisance parameters in α . In what follows, we will consider a simple example.

Example 14.10

Let $\eta(x, \theta) = \theta_0 + \theta_1x + \theta_2x^2$, and $\lambda(x, \alpha) = e^{-\alpha x}$ for some $\alpha > 0$ and $x \in \mathcal{X} = [0, \infty)$. This means that we assume that the model is a linear model and that the variance is increasing exponentially as x increases. Here the information matrix for estimating $\theta = (\theta_0, \theta_1, \theta_2)'$ is given by

$$M_{\xi} = \sum_{i=1}^t w_i \lambda(x_i, \alpha) \begin{bmatrix} 1 & x_i & x_i^2 \\ x_i & x_i^2 & x_i^3 \\ x_i^2 & x_i^3 & x_i^4 \end{bmatrix} = \sum_{i=1}^t w_i e^{-\alpha x_i} \begin{bmatrix} 1 & x_i & x_i^2 \\ x_i & x_i^2 & x_i^3 \\ x_i^2 & x_i^3 & x_i^4 \end{bmatrix},$$

which clearly depends on α . Replacing $\alpha x = z, z \in [Z_l, Z_u) = [0, \infty)$, and letting

$$P(\theta) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1/\alpha & 0 \\ 0 & 0 & 1/\alpha^2 \end{bmatrix},$$

we obtain $\Psi_1(z) = e^{-z}, \Psi_2(z) = ze^{-z}, \Psi_3(z) = z^2e^{-z}, \Psi_4(z) = z^3e^{-z}$, and $\Psi_5(z) = z^4e^{-z}$. We observe that the design problem closely resembles the problem for the exponential model. In fact, if we dropped the squared term, θ_2x^2 , from the model equation, we would get a problem identical to that in Example 14.8 in terms of finding the simple complete class of designs under the Loewner ordering.

Applying Theorem 14.4, we find that $\Phi(z) = -24e^{-z} < 0$, which implies that the optimal design will have at most 3 support points, one of which is Z_l , translating into $x_1 = 0$ in the original design interval. Suppose interest is in estimating all three parameters in θ , and we select the D -criterion. In this case, the optimal design will have exactly three support points, one of which is zero, and equal weights. The objective function to maximize is

$$|M_{\xi}| = \frac{1}{27} e^{-\alpha(x_2+x_3)} x_2^2 x_3^2 (x_3 - x_2)^2.$$

A straightforward maximization yields the remaining D -optimal support points $x_2 = (3 - \sqrt{3})/\alpha$ and $x_3 = (3 + \sqrt{3})/\alpha$.

Lau and Studden (1988) show a more general result, finding D -optimal designs for polynomials of any degree with efficiency function $\lambda(x) = e^{-x}$ and several further functions commonly used to model heteroscedasticity. In particular, they find that D -optimal designs on $\mathcal{X} = [0, \infty)$ for model (14.15) where $\eta(x, \theta) = \theta_0 + \theta_1x + \dots + \theta_kx^k, k \geq 1$, and $\lambda(x) = e^{-x}$ have $k + 1$ equally weighted support points. These are given by the roots of $xL_k^{(1)}(x)$ where $L_k^{(1)}(x)$ is the k th generalized Laguerre polynomial. For more information on generalized Laguerre and other classical orthogonal polynomials, see, for example, Szegö (1975).

By a simple transformation, we see that for efficiency function $\lambda(x, \alpha) = e^{-\alpha x}$, the support points of the D -optimal design are the roots of $xL_k^{(1)}(\alpha x)$. We apply this to our example for comparison. The generalized Laguerre polynomial $L_2^{(1)}(x)$ is given by $L_2^{(1)}(x) = 0.5x^2 - 3x + 3$. Solving

$$xL_2^{(1)}(\alpha x) = x(0.5\alpha^2 x^2 - 3\alpha x + 3) = 0$$

confirms the results we found before.

Note that Lau and Studden (1988) used a method not yet described in this chapter, the method of expressing the objective function in terms of canonical moments. We will only give a brief description of this approach. For further reading, the interested reader is referred to Dette and Studden (1997) and the references therein.

The idea behind this approach is as follows. The entries in the information matrix M_ξ can be viewed as the moments of the design ξ . This is particularly evident for polynomial models with constant variance. If we tried to maximize the determinant of M_ξ with respect to the moments, we would get into trouble due to the complicated structure of the moment spaces. For example, the possible range for the second moment will depend on the value of the first moment in a nontrivial way.

Canonical moments are transformations of the ordinary moments of a probability measure. Roughly speaking, a canonical moment determines the relative position of the corresponding ordinary moment in its moment space, given the lower-order moments. The big advantage of canonical moments is the simple structure of their moment spaces. In particular, they do not depend on the values of the lower-order canonical moments. If it is possible to express the objective function in terms of canonical moments, it can be optimized over each of these quantities independently, which usually results in a considerable simplification of the problem. Once the optimal canonical moments have been found, the corresponding optimal design can be determined by applying results on continued fractions, the Stieltjes transform, and orthogonal polynomials. The major limitation of canonical moments is that the objective function can only be expressed as a function of canonical moments in a few special cases, for example, for D - or D_s -optimality for polynomial models with certain efficiency functions or trigonometric models with constant variance.

14.5 Model Discrimination for Nonlinear Models

To discriminate between two nested models, a popular optimality criterion is D_s -optimality; see also Chapter 2. Intuitively, this is related to D -optimality for those s entries in the parameter vector θ by which the models differ. In particular, a D_s -optimal design minimizes the volume of the confidence ellipsoid for these s parameters. Without loss of generality, let $\theta' = (\theta'_{(1)}, \theta'_{(2)})$ where the s additional parameters for the larger model are held in $\theta_{(1)}$. Then a D_s -optimal design maximizes

$$\phi_{D_s}(\xi) = |(K'M_\xi^-K)^{-1}|, \quad (14.16)$$

where $K' = [I_s \mathbf{0}_{s \times (m-s)}]$ and M_ξ^- denotes a generalized inverse of the information matrix for θ .

The blocks in K' are the identity matrix of size $(s \times s)$ and the zero matrix of size $(s \times (m-s))$, respectively.

To discriminate between more than two nested models, compound or constrained criteria can be used. A compound criterion (see, e.g., Läuter 1974) typically optimizes a combination of l objective functions of the form $\phi(\xi) = \phi_1(\xi)^{\beta_1} \times \dots \times \phi_l(\xi)^{\beta_l}$, where $l > 1$ is an integer and the weights β_1, \dots, β_l sum to one. The weights are chosen to reflect the importance of each criterion, and the objective functions should be appropriately standardized to avoid some of them dominating the others just because they take values on different scales. In the context of model discrimination, the individual criteria could be D_s -criteria for different pairs of models. Constrained criteria optimize one objective function subject to the constraints that the resulting design achieves at least given efficiencies for all the other criteria. A constrained optimal design may not exist for certain combinations of lower bounds for the efficiencies. An application of this method to discriminate within a class of linear models can be found in Biedermann et al. (2009).

Another popular optimality criterion for model discrimination, which does not require the models to be nested and can be applied directly for discriminating between more than two models, is T -optimality (see Atkinson and Fedorov 1975a,b). Suppose the aim of the experiment is to discriminate between models $\eta_1(x, \theta_1)$ and $\eta_2(x, \theta_2)$ for any $\theta_1 \in \Theta_1$ and $\theta_2 \in \Theta_2$. Atkinson and Fedorov (1975a) suggest to fix a model, $\eta_1(x, \theta_1)$ say. A T -optimal design then maximizes the minimal deviation between the model η_1 and the class of models defined by η_2 , that is,

$$\phi_T(\xi, \theta_1) = \inf_{\theta_2 \in \Theta_2} \int_{\mathcal{X}} (\eta_1(x, \theta_1) - \eta_2(x, \theta_2))^2 d\xi(x).$$

If the models are not nested, it may be difficult to decide which of them should be fixed and thus assumed to be the *true* model. In this situation, a compound design for the two T -criteria with each model fixed in turn could be applied. If both models are linear, that is, they can be expressed as $\eta_1(x, \theta_1) = \theta_1' f_1(x)$ and $\eta_2(x, \theta_2) = \theta_2' f_2(x)$, where $f_1(x)$ and $f_2(x)$ are known vectors of regression functions, and if in addition $f_1(x) = (f_0(x), f_2'(x))'$, where $f_0(x)$ is a scalar function, then the T -optimal design coincides with the D_1 -optimal design for the coefficient of $f_0(x)$ (see Atkinson and Fedorov 1975a, and Dette and Titoff 2009 for a derivation and further discussion). Further properties of T -optimal designs in the context of approximation theory are derived in Dette and Titoff (2009). Generally, analytical results for this criterion are hard to obtain, and usually optimal designs have to be found numerically.

A feature of both D_s - and T -optimal designs is that in some situations, these designs have fewer support points than there are parameters in the larger model, so this model cannot be estimated if found preferable by the likelihood ratio test. Sometimes not even the smaller model can be estimated, as illustrated in Example 14.11. In such a situation, again compound or constrained optimal designs, where the additional criteria are D -efficiencies for estimating each model, can be useful. If these turn out to be difficult to find, hybrid designs (see, e.g., Waterhouse et al. 2009) can be a good compromise. Hybrid designs include design points from both the T -optimal design and the D -optimal designs in such a way that the weight from each point is adjusted to account for the priority given to the criterion that point represents. Compound optimal designs combining T -optimality for model discrimination and D -optimality for estimation (also called DT -optimal designs) are described in Atkinson (2008), which also gives an overview of similar criteria used in the literature.

The optimal designs depend on the values of the unknown model parameters. For examples of discrimination designs for the Michaelis–Menten model and exponential models, respectively, which have been made robust to parameter misspecifications, see Biedermann et al. (2007) and Dette et al. (2005).

For models with nonnormal errors, López-Fidalgo et al. (2007) suggest an optimality criterion based on the Kullback–Leibler distance and show that this is consistent with T -optimality. To discriminate between different link functions for GLMs, see, for example, Waterhouse et al. (2008), who consider the difference in deviances for the rival models.

Example 14.11 (Examples 14.1 and 14.9 continued)

We briefly discuss the discrimination problem between the Michaelis–Menten and the Emax models with normally distributed errors. These are nested models, so either the T - or the D_s -criterion can be used. For the T -criterion, an optimal design has to be found numerically. We will focus on D_s -optimality, since this criterion is based on the information matrix, so this will allow us to use results from previous sections.

For this example, $K' = (1, 0, 0)$. Hence, the D_s -criterion (14.16) corresponds to the c -criterion for estimating θ_0 in the Emax model. From Example 14.9, any optimal design ξ^* will have at most three support points, including the end points of $\mathcal{X} = [0, B]$, that is, $\xi^* = \{(0, w_1), (x_2, w_2), (B, 1 - w_1 - w_2)\}$ with x_2, w_1 , and w_2 to be determined. Substituting this design into the objective function (14.16) yields that $\phi_c(\xi) = 1/w_1$, which is minimized for $w_1 = 1$. The optimal design is thus a one-point design in $x_1 = 0$.

We observe that this design does not allow estimation of either of the two models. We consider hybrid designs, that is, weighted averages of the D_s -optimal design for discrimination and the D -optimal designs in either model. In practice, the weighting is often selected to achieve certain values for the individual efficiencies. For example, if the resulting design has equal weights on the support points $0, \theta_2 B / (2\theta_2 + B)$, and B , its D_s -efficiency is $1/3$, its D -efficiency for the Michaelis–Menten model is $2/3$, and its D -efficiency for the Emax model is 1 . Giving more weight to the point 0 , for example, $1/2$ and $1/4$ to each of the other two support points, improves the efficiency for model discrimination to $1/2$, at the expense of reducing the D -efficiencies for estimating the Michaelis–Menten and the Emax model, respectively, to $1/2$ and 0.945 .

14.6 Parameter Robust Approaches

All design problems discussed so far have in common that the optimal designs found depend on at least some of the unknown model parameters. This leads to a chicken and egg situation: To get a good design and thus precise estimates, we need to know the very quantities we actually want to estimate from the data before these are collected. In this section, we will illustrate the effects of parameter misspecification when designing experiments and then briefly discuss strategies to make designs robust. A more detailed investigation of this issue can be found in Chapter 20; see also the discussion in Section 13.3.4 of Chapter 13.

Example 14.12 (Examples 14.1, 14.2, 14.8 continued)

Consider the Michaelis–Menten model, and suppose an experiment was designed to be D -optimal for a specific value of θ_2 . However, the *true* value of this parameter is θ_2^* . We can then find the D -efficiency of the misspecified design ξ , relative to the *correct* D -optimal design ξ^* . Substituting

$$\xi = \left\{ \left(\frac{\theta_2 B}{2\theta_2 + B}, 0.5 \right), (B, 0.5) \right\}, \quad \xi^* = \left\{ \left(\frac{\theta_2^* B}{2\theta_2^* + B}, 0.5 \right), (B, 0.5) \right\},$$

into the expression (14.14) for D -efficiency, we obtain that

$$\text{eff}_D(\xi) = \frac{\theta_2 \theta_2^* (\theta_2 + B)(\theta_2^* + B)}{[\theta_2 \theta_2^* + B(\theta_2 + \theta_2^*)/2]^2}.$$

Similarly, the D -efficiency of a misspecified D -optimal design ξ for the exponential model is given by

$$\text{eff}_D(\xi) = \frac{\theta_2^*}{\theta_2} e^{1 - \theta_2^*/\theta_2}.$$

Figure 14.8 shows D -efficiencies of the locally optimal designs for the Michaelis–Menten model and the exponential model, respectively, when the value of θ_2 has been misspecified across a range of $\theta_2 \in [0.1, 2]$. In the upper panel, we see two scenarios for the Michaelis–Menten model, where the true value, θ_2^* , is 0.3 and 0.6, respectively. We see that the efficiencies appear to be reasonable even on this relatively wide range, with 0.764 and 0.628 the respective minimal efficiencies.




FIGURE 14.8

Upper panel, efficiencies of the locally D -optimal designs for the Michaelis–Menten model with parameter $\theta_2 \in [0.1, 2]$. Left, true parameter $\theta_2^* = 0.3$. Right, true parameter $\theta_2^* = 0.6$. Lower panel, efficiencies of the locally D -optimal designs for the exponential model with parameter $\theta_2 \in [0.1, 2]$. Left, true parameter $\theta_2^* = 0.6$. Right, true parameter $\theta_2^* = 1.2$.

The lower panel shows two scenarios for the exponential model, with θ_2^* given by 0.6 and 1.2, respectively. Here, the drop in efficiency is dramatic. For example, if $\theta_2^* = 1.2$, but the experimenter designed the experiment for $\theta_2 = 0.1$, the efficiency of the D -optimal design is only 0.0002. This is intuitive, since the observations are taken at points 0 and 10, which is appropriate for a relatively slow decay with rate 0.1. However, if the true rate of decay is 1.2, the mean response decreases much faster than expected and is almost zero at $x = 10$. Hence, the design *misses* the *interesting* part of the experimental region. This can also be seen from the information matrix, whose entries are almost zero for $x = 10$ since $e^{-1.2 \times 10} = 6.144 \times 10^{-6}$, and thus $x = 10$ provides almost zero information for the estimation of the model parameters.

These results show that parameter misspecification can be a serious issue and robust designs are sought for experiments in practice. There are several different approaches, of which we will briefly introduce the four most common ones.

14.6.1 Response-Adaptive Sequential/Batch Sequential Experimentation

If the nature of the experiment permits observations to be taken sequentially, future experimental conditions can be optimized based on the observations already made. Starting with an initial design ξ_0 , which can, for example, be a locally optimal design, a robust design as described in the following, or an equidistant uniform design, we take some observations, from which the parameter vector θ is estimated. This estimate, $\hat{\theta}_1$ say, is then substituted into the objective function, to find the design ξ_1 for the next stage, such that the combined design $\xi_0 + \xi_1$ optimizes $\phi(\xi, \hat{\theta}_1)$. One or a whole batch of observations will be made according to the design ξ_1 , from which an updated estimate for θ is obtained. This procedure is repeated until the total sample size has been reached.

It is expected that by adopting this strategy, the quality of the design, and thus the estimate, can be improved successively. However, there are a few drawbacks. For example, observations at experimental conditions from ξ_1 depend on the estimate $\hat{\theta}_1$ obtained from previous observations, that is, the data from a sequential experiment are dependent, making inference more difficult. For many situations, it has been shown that under mild conditions, the estimators obtained from such experiments are asymptotically efficient and that the sequential designs converge to the locally optimal design for the *true* parameter value (see, e.g., Roy et al. 2008). Dette et al. (2013) show analytically that one can expect a benefit from two-stage adaptive designs over nonadaptive designs for sufficiently large sample sizes. However, for small sample sizes, the adaptive design may still be outperformed by nonadaptive designs, in particular if the initial design has been chosen poorly. Another open question in this context is how to choose the number of batches and observations per batch in the sequential procedure. Usually, extensive simulations are required prior to experimentation to determine a good strategy.

14.6.2 Bayesian/Pseudo-Bayesian Designs

Bayesian (also called fully Bayesian) designs are optimized for Bayesian inference and are beyond the scope of this chapter. In the frequentist literature, a Bayesian (often called pseudo-Bayesian) ϕ -optimal design optimizes an objective function of the form

$$\int \phi(\xi, \theta) \pi(\theta) d\theta, \quad (14.17)$$

where $\phi(\xi, \theta)$ is the objective function of a local optimality criterion and $\pi(\theta)$ is a prior distribution summarizing the available knowledge for θ .

This means that the local objective function $\phi(\xi, \theta)$ is averaged over the plausible values for θ . The prior $\pi(\theta)$ is specified solely for the purpose of finding a design that performs reasonably well across its domain and is not used for data analysis. For further reading on relationships and differences of fully Bayesian and pseudo-Bayesian designs, we refer to the review paper Chaloner and Verdinelli (1995). Chapter 13 provides further discussion of Bayesian design methods in the context of generalized linear models.

A potential problem with pseudo-Bayesian designs is the choice of prior distribution. If the domain is far from the *true* parameter value, the same problems as for locally optimal designs arise. Moreover, even if the *true* value of θ is contained in the domain of $\pi(\theta)$, it is not guaranteed that it can be estimated efficiently, since other values of the parameter vector may dominate the weighted average in (14.17).

Pseudo-Bayesian D -optimal designs for the Michaelis–Menten model are found in Matthews and Allcock (2004). Note that for numerical computation of a pseudo-Bayesian optimal design, the integral in (14.17) is usually replaced by a finite sum that approximates the integral. For guidance on the choice of values for θ to be used in the summation, see, for example, Gotwalt et al. (2009).

14.6.3 (Standardized) Maximin Optimal Designs

This approach is more cautious than the pseudo-Bayesian and addresses the problem of possibly low design efficiency in some regions within the domain of the prior $\pi(\theta)$ by optimizing the design for the worst-case scenario. Moreover, it is not necessary to specify a prior distribution on θ , but only a plausible range Θ .

Let $\phi(\xi, \theta)$ be the objective function of a local optimality criterion, which without loss of generality must be maximized. Then a maximin ϕ -optimal design maximizes $\min_{\theta \in \Theta} \phi(\xi, \theta)$. These designs have been criticized by Dette (1997) because the scale of $\phi(\xi, \theta)$ can be considerably affected by the value of θ . To address this issue, Dette (1997) suggests a standardization approach. A standardized maximin ϕ -optimal design maximizes

$$\min_{\theta \in \Theta} \frac{\phi(\xi, \theta)}{\phi(\xi_{\theta}^*, \theta)}, \tag{14.18}$$

where ξ_{θ}^* is the locally ϕ -optimal design for the value θ . Note that if $\phi(\xi, \theta)$ were to be minimized, the numerator and denominator in (14.18) would change places. For many local optimality criteria, the ratio in (14.18) is also the definition of ϕ -efficiency of the design ξ , or a one-to-one mapping to this quantity. For example, for D -optimality, the ratio would be raised to the power $1/m$ to obtain the D -efficiency as defined in (14.14). Notice that there is an alternative equivalent definition of maximin design called minimax design (see Chapter 13, Section 13.3.4.3).

One drawback of standardized maximin optimal designs is that they do not necessarily perform well if the *true* parameter value for θ is outside the range Θ specified by the experimenter. If it is inside this range, we have a lower bound for the ϕ -efficiency of the design for each value of $\theta \in \Theta$. However, the efficiencies of standardized maximin optimal designs, in particular for large ranges Θ , tend to be flat, that is, they may be close to the lower bound across the whole range. Another issue, preventing more widespread use of

maximin and standardized maximin optimal designs, is that they are usually very hard to compute (see Nyquist 2013 for a recent algorithm). Examples where standardized maximin D -optimal designs with minimum support are found analytically for the Michaelis–Menten model, the exponential model and polynomials with several different efficiency functions are given in Dette and Biedermann (2003), Imhof (2001), and Biedermann and Dette (2003).

14.6.4 Cluster Designs

These designs are used as an alternative if pseudo-Bayesian and standardized maximin optimal designs are difficult to compute. A sample of J (J large, e.g., 1000) values of the parameter vector θ is drawn according to a prior distribution $\pi(\theta)$. A clustering algorithm is then applied to the support points of the corresponding J locally optimal designs. The cluster design is formed by taking the centroids of the resulting clusters as equally weighted support points.

This basic method was first introduced in Dror and Steinberg (2006) and later modified by Biedermann and Woods (2011) to take non-equal weights of the locally optimal designs into account. Cluster designs are easy to compute but suffer potentially from the same drawbacks as pseudo-Bayesian designs. There is no general rule on how to select the number of support points for these designs. This is currently done through summary statistics for efficiencies relative to a large number of locally optimal designs and simulations on a case-by-case basis.

14.7 Summary

We have outlined the most popular methods for finding optimal designs for nonlinear models and illustrated them through examples. Some methods are particularly useful in specific situations. The general strategy, however, is as follows: First, apply one of the *new* methods described in Section 14.3 to identify a simple class for the optimal design. Second, select an appropriate optimality criterion and use one of the *classical* methods to further narrow down this class (if possible). Finally, use an optimization procedure, either analytically or numerically, to find the optimal design. Some powerful new algorithms for numerical calculation of optimal designs have been suggested in the recent literature; see, for example, Yu (2011) for D -optimal designs and Yang et al. (2013) for ϕ_p -optimal designs, the latter of which is also capable of finding response-adaptive optimal designs and optimal designs for subsets or functions of the unknown model parameters.

References

- Alvarez, I., R. Virto, J. Raso, and S. Condon, Comparing predicting models for the *Escherichia coli* inactivation by pulsed electric fields, *Innovat. Food Sci. Emerg. Technol.*, 4, 195–202, 2003.
- Atkinson, A.C., DT -optimum designs for model discrimination and parameter estimation, *J. Stat. Plann. Infer.*, 138, 56–64, 2008.

- Atkinson, A.C. and V.V. Fedorov, The design of experiments for discriminating between two rival models, *Biometrika*, 62, 57–70, 1975a.
- Atkinson, A.C. and V.V. Fedorov, Optimal design: Experiments for discriminating between several models, *Biometrika*, 62, 289–303, 1975b.
- Becka, M., H.M. Bolt, and W. Urfer, Statistical analysis of toxicokinetic data by nonlinear regression (example: inhalation pharmacokinetics of propylene), *Arch. Toxicol.*, 66, 450–453, 1992.
- Becka, M., H.M. Bolt, and W. Urfer, Statistical evaluation of toxicokinetic data, *Environmetrics*, 4, 311–322, 1993.
- Biedermann, S. and H. Dette, A note on maximin and Bayesian D -optimal designs in weighted polynomial regression, *Math. Methods Stat.*, 12, 358–370, 2003.
- Biedermann, S., H. Dette, and P. Hoffmann, Constrained optimal discriminating designs for Fourier regression models, *Ann. Inst. Stat. Math.*, 61, 143–157, 2009.
- Biedermann, S., H. Dette, and A. Pepelyshev, Optimal discrimination designs for exponential regression models, *J. Stat. Plann. Infer.*, 137, 2579–2592, 2007.
- Biedermann, S., H. Dette, and D.C. Woods, Optimal design for additive partially nonlinear models, *Biometrika*, 98, 449–458, 2011.
- Biedermann, S., H. Dette, and W. Zhu, Optimal designs for dose-response models with restricted design spaces, *J. Am. Stat. Assoc.*, 101, 747–759, 2006.
- Biedermann, S. and D.C. Woods, Optimal designs for generalised nonlinear models with application to second harmonic generation experiments, *J. R. Stat. Soc., Ser. C*, 60, 281–299, 2011.
- Chaloner, K. and I. Verdinelli, Bayesian experimental design: A review, *Stat. Sci.*, 10, 273–304, 1995.
- Chernoff, H. Locally optimal designs for estimating parameters, *Ann. Math. Stat.*, 24, 586–602, 1953.
- de la Garza, A. Spacing of information in polynomial regression, *Ann. Math. Stat.*, 25, 123–130, 1954.
- Dette, H. Elfving's theorem for D -optimality, *Ann. Stat.*, 21, 753–766, 1993.
- Dette, H. A note on Bayesian c - and D -optimal designs in nonlinear regression models, *Ann. Stat.*, 24, 1225–1234, 1996.
- Dette, H. Designing experiments with respect to standardized optimality criteria, *J. R. Stat. Soc., Ser. B*, 59, 97–110, 1997.
- Dette, H. and S. Biedermann, Robust and efficient designs for the Michaelis–Menten model, *J. Am. Stat. Assoc.*, 98, 679–686, 2003.
- Dette, H., B. Bornkamp, and F. Bretz, On the efficiency of two-stage response-adaptive designs, *Stat. Med.*, 32, 1646–1660, 2013.
- Dette, H., F. Bretz, A. Pepelyshev, and J. Pinheiro, Optimal designs for dose-finding studies, *J. Am. Stat. Assoc.*, 103, 1225–1237, 2008.
- Dette, H. and L.M. Haines, E -optimal designs for linear and nonlinear models with two parameters, *Biometrika*, 81, 739–754, 1994.
- Dette, H., C. Kiss, M. Bevanda, and F. Bretz, Optimal designs for the e_{max} , log-linear and exponential models, *Biometrika*, 97, 513–518, 2010.
- Dette, H. and V.B. Melas, A note on the de la Garza phenomenon for locally optimal designs, *Ann. Stat.*, 39, 1266–1281, 2011.
- Dette, H., V.B. Melas, and A. Pepelyshev, Local c - and E -optimal designs for exponential regression, *Ann. Inst. Stat. Math.*, 58, 407–426, 2006a.
- Dette, H., V.B. Melas, and W.K. Wong, Locally D -optimal designs for exponential regression models, *Stat. Sinica*, 16, 789–803, 2006b.
- Dette, H., V.B. Melas, and W.K. Wong, Optimal design for goodness-of-fit of the Michaelis–Menten enzyme kinetic function, *J. Am. Stat. Assoc.*, 100, 1370–1381, 2005.
- Dette, H. and K. Schorning, Complete classes of designs for nonlinear regression models and principal representations of moment spaces, *Ann. Stat.*, 41, 1260–1267, 2013.
- Dette, H. and W.J. Studden, Geometry of E -optimality, *Ann. Stat.*, 21, 416–433, 1993a.
- Dette, H. and W.J. Studden, A geometric solution of the Bayesian E -optimal design problem," in *Statistical Decision Theory and Related Topics V*, eds. S.S. Gupta and J.O. Berger. Springer-Verlag, New York, pp. 157–170, 1993b.

- Dette, H. and W.J. Studden, *The Theory of Canonical Moments with Applications in Statistics, Probability and Analysis*, John Wiley & Sons, New York, 1997.
- Dette, H. and S. Titoff, Optimal discrimination designs, *Ann. Stat.*, 37, 2056–2082, 2009.
- Dror, H.A. and D.M. Steinberg, Robust experimental design for multivariate generalized linear models, *Technometrics*, 48, 520–529, 2006.
- Elfving, G. Optimum allocation in linear regression theory, *Ann. Math. Stat.*, 23, 255–262, 1952.
- Ermakov, S.M. and V.B. Melas, *Design and Analysis of Simulation Experiments*, Kluwer Academic Publishers, Dordrecht, London, U.K., 1995.
- Fedorov, V.V. and P. Hackl, *Model-Oriented Design of Experiments*, Lecture notes in statistics, Springer, New York, 1997.
- Ford, I., B. Torsney and C. F. J. Wu, “The use of a canonical form in the construction of locally optimal designs for nonlinear problems,” *J. R. Stat. Soc., Ser. B*, 54, 569–583, 1992.
- Gibaldi, M. and D. Perrier, *Pharmacokinetics*, Marcel Dekker, New York, 1982.
- Gotwalt, C.M., B.A. Jones, and D.M. Steinberg, Fast computation of designs robust to parameter uncertainty for nonlinear settings, *Technometrics*, 51, 88–95, 2009.
- Haines, L.M. Optimal designs for nonlinear regression models, *Commun. Stat., Theory Methods*, 22, 1613–1627, 1993.
- Hill, P.D.H. D -optimal designs for partially nonlinear regression models, *Technometrics*, 22, 275–276, 1980.
- Imhof, L. Maximin designs for exponential growth models and heteroscedastic polynomial models, *Ann. Stat.*, 29, 561–576, 2001.
- Karlin, S. and W.J. Studden, *Tchebycheff systems: With applications in analysis and statistics*, Pure and applied mathematics series, vol. XV, John Wiley & Sons, New York, 1966.
- Khuri, A.I. A note on D -optimal designs for partially nonlinear regression models, *Technometrics*, 26, 59–61, 1984.
- Läuter, E. Experimental design in a class of models, *Math. Operationsforsch. Stat.*, 5, 379–398, 1974.
- Lau, T.S. and W.J. Studden, On an extremal problem of Fejer, *J. Approx. Theory*, 53, 184–194, 1988.
- López-Fidalgo, J., C. Tommasi, and P. C. Trandafir, An optimal experimental design criterion for discriminating between nonnormal models, *J. R. Stat. Soc., Ser. B*, 69, 231–242, 2007.
- Matthews, J.N. and G.C. Allcock, Optimal designs for Michaelis-Menten kinetic studies, *Stat. Med.*, 23, 477–491, 2004.
- Melas, V.B. *Functional Approach to Optimal Experimental Design*, Lecture Notes in Statistics, Springer, New York, 2006.
- Monagan, M.B., K. O. Geddes, K. M. Heal, G. Labahn, S. M. Vorkoetter, J. McCarron, and P. DeMarco, *Maple Advanced Programming Guide*, Maplesoft, Waterloo, Ontario, Canada, 2005.
- Nyquist, H. Convergence of an algorithm for constructing minimax designs, *Proceedings of the 10th International Workshop in Model-Oriented Design and Analysis (moda)*, eds. D. Ucinski, A.C. Atkinson, and M. Patan, Springer, New York, pp. 187–194, 2013.
- Pukelsheim, F. On linear regression designs which maximize information, *J. Stat. Plann. Infer.*, 4, 339–364, 1980.
- Pukelsheim, F. *Optimal design of experiments*, John Wiley & Sons, Inc., New York, 1993.
- Pukelsheim, F. and S. Rieder, Efficient rounding of approximate designs, *Biometrika*, 79.4, 763–770, 1992.
- Pukelsheim, F. and B. Torsney, Optimal weights for experimental designs on linearly independent support points, *Ann. Stat.*, 19, 1614–1625, 1991.
- Roy, A., S. Ghosal, and W. Rosenberger, Convergence properties of sequential Bayesian D -optimal designs, *J. Stat. Plann. Infer.*, 139, 425–440, 2008.
- Sibson, R. Contribution to discussion of papers by H. P. Wynn and P. J. Laycock, *J. R. Stat. Soc., Ser. B*, 34, 181–183, 1972.
- Silvey, S.D. Contribution to discussion of papers by H. P. Wynn and P. J. Laycock, *J. R. Stat. Soc., Ser. B*, 34, 174–175, 1972.
- Silvey, S.D. *Optimal Design*, Chapman and Hall, New York, 1980.

- Silvey, S.D. and D.M. Titterton, A geometric approach to optimum design theory, *Biometrika*, 60, 21–32, 1973.
- Szegő, G. *Orthogonal Polynomials*, 4th edition, American Mathematical Society Colloquium Publications, 23, Providence, RI, 1975.
- Waterhouse, T.H., D.C. Woods, J.A. Eccleston, and S.M. Lewis, Design selection criteria for discrimination/estimation for nested models and a binomial response, *J. Stat. Plann. Infer.*, 138, 132–144, 2008.
- Waterhouse, T.H., J.A. Eccleston, and S.B. Duffull, Optimal design criteria for discrimination and estimation in nonlinear models, *J. Biopharm. Stat.*, 19, 386–402, 2009.
- Wolfram Research, Inc., *Mathematica, Version 9.0*, Champaign, IL, 2012.
- Yang, M. On the de la Garza phenomenon, *Ann. Stat.*, 38, 2499–2524, 2010.
- Yang, M., S. Biedermann and E. Tang, On optimal designs for nonlinear models: A general and efficient algorithm, *J. Am. Stat. Assoc.*, 108, 1411–1420, 2013.
- Yang, M. and J. Stufken, Support points of locally optimal designs for nonlinear models with two parameters, *Ann. Stat.*, 37, 518–541, 2009.
- Yang, M. and J. Stufken, Identifying locally optimal designs for nonlinear models: A simple extension with profound consequences, *Ann. Stat.*, 40, 1665–1685, 2012.
- Yang, M., B. Zhang, and S. Huang, Optimal designs for generalized linear models with multiple design variables, *Stat. Sinica*, 21, 1415–1430, 2011.
- Yu, Y. *D*-optimal designs via a cocktail algorithm, *Stat. Comput.*, 21, 475–481, 2011.

15

Optimal Design for Spatial Models

Zhengyuan Zhu and Evangelos Evangelou

CONTENTS

15.1 Introduction.....	549
15.2 Model-Free Spatial Sampling Design.....	550
15.2.1 Space-Filling Designs.....	550
15.2.2 Spatial Probability-Based Sampling Design.....	552
15.3 Model-Based Spatial Sampling Designs.....	556
15.3.1 Model-Based Geostatistics.....	556
15.3.1.1 Estimation of Variance Parameters.....	557
15.3.1.2 Spatial Prediction.....	558
15.3.2 Designs for Estimating the Covariance Structure.....	559
15.3.2.1 Designs Based on the Generalized Least Squares Method.....	559
15.3.2.2 Designs Based on the Maximum Likelihood Method.....	560
15.3.3 Designs for Spatial Prediction.....	561
15.3.3.1 Designs Based on the Prediction Variance.....	561
15.3.3.2 Entropy-Based Designs.....	564
15.3.4 Bayesian Design.....	568
15.3.5 Dynamic Designs.....	569
15.4 Further Reading.....	570
References.....	571

15.1 Introduction

Spatial sampling design plays an important role in many applications such as environmental monitoring, natural resource survey, ecological studies, and water resource management. For example, in the National Resources Inventory survey conducted by the Natural Resources Conservation Service at USDA to monitor the status and change of the soil, water, and related resources on nonfederal land in the United States, sample segments were selected in space in the first stage of a two-stage area sample, and spatial sampling design techniques were used to achieve geographic spread, which greatly increased the design efficiency (Nusser et al. 1998). In another example described by Zidek et al. (2000), the authors considered the problem of extending an existing pollutant monitoring network in Southern Ontario by adding monitoring stations at new spatial locations. To reduce the uncertainty in predicting multiple pollutants at unobserved locations, a model-based spatial sampling design method was used that maximized an entropy criterion based on a

Gaussian model. In both applications, spatial information is an important part of the data and has to be taken into account in design.

Let \mathbb{X} be the spatial domain of interest and let $\mathcal{Z} = \{z_x, x \in \mathbb{X}\}$ be the quantity of interest. In most applications, \mathbb{X} is a subset of two-dimensional Euclidian space \mathbb{R}^2 , though other types of spatial domains do arise from different applications. For example, for those who are interested in sampling stream networks, the spatial domains are one-dimensional tree-structured curves. For those who are interested in environmental sampling on a global scale, the spatial domain is a two-dimensional manifold (sphere). For meteorologists who are interested in sampling the wind fields for modeling, the domain is \mathbb{R}^3 . In this chapter, we will focus on the case when $\mathbb{X} \subset \mathbb{R}^2$, and the design concepts can be adapted to other spatial domains as well.

In most applications, we can only observe z_x on $x \in d$, where d is a discrete subset of \mathbb{X} with finite size. The optimal design problem is to select d under certain constraints such that one can make good inference about \mathcal{Z} based on the finite sample on d . There are two major differences between spatial sampling and classical design problems: (1) an important part of the modeling takes place in the correlation structure based on spatial locations, and (2) it is often not possible to take repeated samples from the same location. For example, in a typical agriculture survey, a segment of land is the sampling unit, and variables such as acres of cultivated land on that segment are collected. There is no need to sample at the same location more than once. Thus, many of the design strategies that are optimal for independent observations may not be applicable for spatial sampling. Similar to classical design problems, the objective for spatial sampling design is often optimal estimation and prediction. Both of them take more complicated forms under spatial models due to the spatial dependence in data and are more difficult to analyze. For example, the uncertainty in estimation and spatial prediction often depends on parameters in a highly nonlinear way, making both theoretical study and numerical optimization more challenging. In this chapter, we review the general methodology and commonly used numerical methods for spatial sampling design.

Broadly speaking, spatial sampling design strategies can be classified into two large groups. One is the model-free approach, for which no specific spatial model is assumed for the data except the assumption that they are spatially correlated. It includes geometrically based space-filling designs and probability-based sampling for spatial data developed from survey sampling methods. The latter includes approaches such as systematic sampling (SYS), stratified sampling with random tessellation, and Markov sampling. The former approach requires that a spatial model is assumed and the optimal sampling strategy is designed under that model. In this chapter, we will introduce both approaches, with a focus on the model-based approach. For the model-based approach, we restrict ourselves largely to Gaussian random field (RF) models, though a generalization to spatial generalized linear mixed model is briefly mentioned toward the end.

15.2 Model-Free Spatial Sampling Design

15.2.1 Space-Filling Designs

In many practical applications, there is little prior knowledge about z_x on \mathbb{X} , and it is desirable to have a design that has good overall coverage of the study area with no big

unsampled “holes.” Let $dist(x, x')$ be a distance function defined on $\mathbb{X} \times \mathbb{X}$, and for design $d \subset \mathbb{X}$ with fixed size n , define $dist(x, d) = \min_{x' \in d} dist(x, x')$. A design d^* is defined as a minimax distance design (Johnson et al. 1990) if

$$\max_{x \in \mathbb{X}} dist(x, d^*) = \min_d \max_{x \in \mathbb{X}} dist(x, d) = dist^*.$$

Intuitively, a minimax distance design guarantees that all points in \mathbb{X} are within $dist^*$ of at least one point in the design, that is, there are no “holes” with diameters bigger than $dist^*$.

Another viewpoint of spatial sampling is that because of the spatial dependence, points sampled in close proximity provide redundant information. To increase the efficiency of the design, it may be desirable to have the points in the design as far away from each other as possible. One can define the maximin distance design (Johnson et al. 1990) as the design d° such that

$$\min_{x, x' \in d^\circ} dist(x, x') = \max_d \min_{x, x' \in d} dist(x, x') = dist^\circ.$$

Both minimax and maximin distance designs are examples of space-filling design for which the design points cover the whole study area with no big holes. The difference is that maximin distance design pushes points to the boundary compared to minimax distance design. Note that there are usually multiple designs that are minimax/maximin distance designs. Johnson et al. (1990) introduced the notion of index for maximin/minimax designs based on the number of point pairs separated by the maximum/minimum distance and show that the designs with the highest/lowest index are asymptotically optimal under G - and D -criteria. In practice, minimax distance design is typically better for spatial prediction, and maximin distance design is better for estimating regression parameters.

There are often multiple structurally different designs that have the minimax or maximin property. The numerical problem of constructing the minimax/maximin distance design is in general very difficult for arbitrary sample size and shape of the design region. Numerical algorithms such as exchange (Royle and Nychka 1998), simulated annealing (Morris and Mitchell 1995), and coffee-house algorithm (Müller 2001) can be used to find approximate solutions. Royle and Nychka implemented the exchange algorithm for minimax design in the R package *fields* (i.e., function *cover.design*). When the sample size is large, the regular exchange algorithm becomes computationally infeasible, and short cuts such as the nearest neighbor option can be used to speed up the computation. Figure 15.1 shows an example of a minimax Euclidean distance design on $[0, 1] \times [0, 1]$ with $n = 50$.

In some applications of spatial sampling design, such as the design of computer experiments, the design space \mathbb{X} is typically a hypercube of higher dimension, corresponding to multiple factors for the computer experiment. In this context, many of the factors may have little effect on the outcome, and for the few factors that influence the outcome, the relationship is typically nonlinear, and it is often useful to treat the factors the same as spatial locations in the statistical model. Thus, it is desirable to have a design that has good space-filling properties not only in the original hypercube but also on its lower dimensional projections. The original minimax distance designs do not have such properties, and minimax distance Latin hypercube designs (Morris and Mitchell 1995) are preferred, in which a minimax distance design is selected from all Latin hypercube designs to guarantee that its one-dimensional projections have good space-filling properties (see also Chapter 17). A similar approach can be applied to orthogonal array-based Latin hypercubes (Tang 1993) to control




FIGURE 15.1
Example of minimax distance design.

higher dimensional projections. See Joseph and Hung (2008) for examples of orthogonal Latin hypercube designs with good space-filling properties. See Chapter 17 for more discussion on designs for computer experiments.

15.2.2 Spatial Probability-Based Sampling Design

In agriculture and natural resource surveys, there is a long history of spatial sampling design in the form of area sampling. See, for example, King (1945), Das (1950), Sukhatme (1957), and Monroe and Finkner (1959). The predominant method used in this field is probability-based sampling, in which the underlying spatial process is assumed to be fixed and the randomness is introduced by using a probability sampling design. A probability sampling design consists of two parts: a well-defined frame, which is a set of all elements in the population of interest, and a probability measure defined on the frame, which gives the probability of selecting any subset. In this section, we will assume that the population has finite number of elements. The statistical properties of the estimators for functions of the population are solely based on the probability measure that defines the design and are commonly referred to as design-based inference.

Typical probability-based sampling designs include simple random sampling (SRS), SYS, stratified sampling, and cluster sampling. The fundamental building block of probability-based sampling design is SRS, in which each element in the frame is selected with equal probability and the selection is independent of the selection of any other elements. The primary objective of probability-based sampling in spatial sampling design is to make inference about a population quantity such as a mean, total, variance, or proportion, within a certain spatial domain. For example, one may be interested in estimating the total area of wetland in a certain hydrologic unit or the average yield of corn in a county. One main advantage of a probability-based design is that it does not require any model assumptions about the underlying process, thus avoiding the possibility of selection bias

based on a potentially incorrect model. See Cochran (2007) or Fuller (2011), for example, for a more comprehensive introduction to the theory and practice of probability-based sampling design.

In the context of spatial sampling design, it has long been recognized that in the presence of spatial dependence, SRS is often an inefficient design because of its poor coverage in the space, which leads to large estimation variance. One alternative is SYS. In one-dimensional space, a SYS can be defined as follows. Let the length of the one-dimensional sample space be l and the sample size be n . Let $a = l/n$ be the interval length. The first element of the SYS is selected from a uniform distribution on interval $[0, a)$, and the rest of the sample is drawn by taking every sample to have distance a away from the previous sample. In d -dimensional space, the samples are taken from a regular grid to ensure they are equally spaced, with the starting point of the SYS chosen from a uniform distribution on the first cube $[0, a)^d$, where $a = l/n^{1/d}$. Figures 15.2 and 15.3 show examples of SRS and SYS on $[0, 1] \times [0, 1]$. It is apparent that SYS is more spatially balanced and has better space-filling properties compared to SRS. There is also theoretical evidence (Bellhouse 1977) that SYS is optimal under certain superpopulation models in the sense of minimum average variance when the sample mean is used to estimate the population mean.

One potential drawback of the SYS is that it may introduce significant bias if the population being sampled has a periodic pattern that is aligned with the SYS. For example, if one takes land use samples at 1 mile interval in an area in which the road grid is also designed to be 1 mile apart, one would have no sample of the road if the starting point is not on a road. One alternative is to compromise between SRS and SYS by using a one-per-stratum sampling design, a form of stratified random sampling, in which the region of interest is divided into n strata of equal size and an SRS of size one is selected from each stratum. A special case of such design is the random-tessellation stratified (RTS) design (Overton and Stehman 1993), which requires that each stratum has the same shape. These designs are spatially balanced and do not have the potential alignment issue of SYS. Figure 15.4 gives an example of an RTS design on $[0, 1] \times [0, 1]$.




FIGURE 15.2
Simple random sample on $[0, 1] \times [0, 1]$.




FIGURE 15.3
Systematic sampling with random starting point.




FIGURE 15.4
Random-tessellation stratified design.

Note that RTS design does not necessarily have good space-filling properties. Because the selection of the sample within each stratum is done independently, one may by chance have points in the design that are arbitrarily close to each other, potentially reducing the efficiency of the design. Markov chain designs for one-per-stratum sampling (Breidt 1995) can be used to address this problem. In a Markov chain design on one-dimensional space, each stratum is further divided into several smaller substrata, and the samples are selected sequentially, starting with an SRS on the first stratum. A transition probability matrix is specified to determine the selection probability of the next sample conditional on the



FIGURE 15.5
Markov chain design.

location of the previous sample. This design includes both SYS and RTS as special cases, and one can choose an appropriate transition matrix to have good space-filling properties while avoiding the alignment problem of SYS. See Figure 15.5 for an example of Markov chain design on one-dimensional space.

In practice, one often needs to allow the inclusion probability of elements in a frame to vary based on relevant covariate information, sampling costs, and other considerations. One drawback of the spatially balanced designs we have discussed so far (SYS, RTS, and Markov chain designs) is that it is difficult to modify them to accommodate variable inclusion probability. The generalized random-tessellation stratified (GRTS) design is a popular spatially balanced design that allows for unequal inclusion probability (Stevens and Olsen 2004). The basic idea of GRTS is to find a function f that maps points in \mathbb{R}^2 to an interval in \mathbb{R} and preserves the proximity relationships between points in \mathbb{R}^2 as much as possible and to use a SYS with random start on the interval to select samples. The resulting design would be spatially balanced due to the construction of the mapping. Without loss of generality, let's consider the problem of mapping points on $[0, 1) \times (0, 1]$ onto the unit interval $(0, 1]$. A recursive approach can be defined as follows: divide the unit square into 2×2 smaller squares of equal size, randomly permute them, and assign them to intervals $(0, 1/4]$, $(1/4, 1/2]$, $(1/2, 3/4]$, and $(3/4, 1]$. Keep dividing the squares into smaller 2×2 squares and randomly assign them to the corresponding subintervals until no square contains more than one population element, and a proximity-preserving mapping is constructed by placing the points to the center of the assigned subintervals. Variable inclusion probability can be accommodated by associating each point in the frame with an interval that has length proportional to its inclusion probability. The GRTS design is implemented in R package *spsurvey* (Kincaid and Olsen 2013). Figure 15.6 shows an example of GRTS with




FIGURE 15.6
GRTS design with unequal inclusion probability.

unequal inclusion probability. The number of points on the upper half of the unit square is one-third of those on the lower half.

15.3 Model-Based Spatial Sampling Designs

In this section, we discuss optimal spatial designs for geostatistical models. We distinguish two types of designs: (1) designs for estimation of the variance parameters and (2) designs for prediction of the spatial RF. We first give an overview of the geostatistical model before we describe these designs in detail.

15.3.1 Model-Based Geostatistics

We consider a stochastic process $\mathcal{Z} = \{z_x, x \in \mathbb{X}\}$ defined over \mathbb{X} . In most geostatistics applications, \mathbb{X} is a continuous spatial domain, \mathcal{Z} is called the spatial RF, and z_x denotes the value of the RF at location x . Each component z_x is decomposed into a linear mean $f'_x \boldsymbol{\beta}$, a zero-mean spatial term s_x , and a zero-mean microscale (or nugget) term ϵ_x , that is,

$$z_x = f'_x \boldsymbol{\beta} + s_x + \epsilon_x, \quad x \in \mathbb{X}, \quad (15.1)$$

where f'_x is a vector of k regressor functions at x and $\boldsymbol{\beta}$ is a k -dimensional vector of regressor coefficients.

The term s_x determines the spatial nature of the process. This process is characterized by its correlation function $\varrho(x, x')$ between two components s_x and $s_{x'}$ and is parametrized by ψ . We denote by σ^2 the variance of this process, which does not depend on x . We also assume that the process s_x is isotropic, namely, that the correlation function $\varrho(x, x')$ depends only on the distance h between the locations x and x' . Typically h would correspond to the Euclidean distance between x and x' although other choices are also possible, for example, space deformation is used for modeling nonstationary fields. In this case, we write the correlation function as $\varrho(h)$. A popular choice is the so-called Matérn correlation function:

$$\varrho(h) = \{2^{\kappa-1} \Gamma(\kappa)\}^{-1} (h/\phi)^{\kappa} K_{\kappa}(h/\phi),$$

where Γ denotes the gamma function and K_{κ} is the modified Bessel function of order κ (Abramowitz and Stegun 1965). In this case, $\psi = (\phi, \kappa)$ are the range and smoothness parameters, respectively.

The microscale term ϵ_x will be used here to denote an uncorrelated error term with variance τ^2 termed the nugget. We will refer to the parameters associated with the spatial and microscale components as the variance parameters and denote these by $\boldsymbol{\theta}$, that is, $\boldsymbol{\theta} = (\sigma^2, \tau^2, \psi)$, and we denote the number of elements in $\boldsymbol{\theta}$ by p .

Another relevant function that frequently arises in estimation is the semivariogram,

$$\gamma(h; \boldsymbol{\theta}) = \frac{1}{2} \text{Var}\{z_{x+h} - z_x\},$$

and the variogram $2\gamma(h)$. It can be readily verified that if $\lim_{h \rightarrow \infty} \varrho(h) = 0$,

$$\lim_{h \rightarrow 0^+} \gamma(h; \theta) = \tau^2 \quad \text{and} \quad \lim_{h \rightarrow \infty} \gamma(h; \theta) = \sigma^2 + \tau^2,$$

while the correlation function $\varrho(h)$ affects the shape of the semivariogram. The variogram is useful for understanding the dependence structure of the RF and for inference (see next section). In particular, the uncertainty in the estimation and prediction can be expressed in terms of the variogram.

In practice, we can only make inference about \mathcal{Z} over a finite subset $\mathcal{X} \subset \mathbb{X}$. To that end, suppose \mathcal{Z} is observed at locations $d = \{x_1, \dots, x_n\} \subset \mathcal{X}$. The set d denotes the sampling design or sampled sites. The data will be used to estimate the model parameters and predict \mathcal{Z} at the remaining locations $d^c = \mathcal{X} \setminus d$, the unsampled sites. Next, we discuss approaches to these tasks.

15.3.1.1 Estimation of Variance Parameters

There are two main approaches in the estimation of the covariance parameters: the method of moments estimator and the maximum likelihood (ML) estimator.

Consider estimation of the variogram at a given lag h by the method of moments. To that end, let $r_h \subset d \times d$ consist of the unique pairs of locations separated by distance h (in practice, we use pairs whose distance falls within some neighborhood of h) and let $|r_h|$ denote the number of elements in r_h . Then the method of moments estimator for the variogram $2\gamma(h; \theta)$ under the constant mean model is

$$2\hat{\gamma}(h) = \frac{1}{|r_h|} \sum_{r_h} (z_x - z_{x'})^2, \tag{15.2}$$

where the summation is over $(x, x') \in r_h$. Realistically, we can compute (15.2) for those lags h for which r_h is of significant size. Suppose that there are m lags h_1, \dots, h_m for which the variogram can be estimated and let $\hat{\gamma}$ the m -dimensional vector with i th element $\hat{\gamma}(h_i)$. In practice, $\hat{\gamma}(h_i)$ is plotted against h_i , and the shape of the scatterplot determines the correlation function to be used.

The vector $\hat{\gamma}$ can be regarded as a random variable with mean $\gamma(\theta)$, the semivariogram at the same lags when the true parameter value is θ , and variance $V(\theta)$, the $m \times m$ covariance matrix of the estimator (15.2) corresponding to the m lags. In particular, the (i, j) element of $V(\theta)$ is given by (Cressie 1993)

$$V(\theta)_{ij} = \text{tr}(A_{h_i} \Sigma A_{h_j} \Sigma),$$

where Σ denotes the $n \times n$ variance–covariance matrix of \mathcal{Z} at d and A_h is a positive semidefinite $n \times n$ matrix whose diagonal elements are the number of times the corresponding location appears in r_h and the off-diagonal is -1 or 0 depending on whether the corresponding pair belongs to r_h or not. Then we define the following:

Definition 15.1 The generalized least squares estimator for the parameter θ is defined by

$$\hat{\theta}^{\text{GLS}} = \underset{\theta}{\operatorname{argmin}} (\hat{\gamma} - \gamma(\theta))' V(\theta)^{-1} (\hat{\gamma} - \gamma(\theta)). \quad (15.3)$$

An alternative method for estimating the covariance parameters is by the restricted ML method. For this method, we assume that the processes s_x and ϵ_x in (15.1) are normally distributed. Let z_d denote the value of Z over the design d , F the $n \times k$ matrix of covariates associated with d , and let $\Sigma = \Sigma(\theta)$ be the corresponding variance–covariance matrix as earlier. Then the restricted log-likelihood function for θ becomes, up to a constant (Harville 1974) (see Chapter 2),

$$\ell(\theta) = -\frac{1}{2} \log |\Sigma| - \frac{1}{2} \log |F' \Sigma^{-1} F| - \frac{1}{2} W(\theta)^2,$$

where

$$W(\theta)^2 = z_d' \{ \Sigma^{-1} - \Sigma^{-1} F (F' \Sigma^{-1} F)^{-1} F' \Sigma^{-1} \} z_d$$

is the generalized residual sum of squares. Then the restricted ML estimator for θ is obtained by maximizing

$$\hat{\theta}^{\text{ML}} = \underset{\theta}{\operatorname{argmax}} \ell(\theta). \quad (15.4)$$

15.3.1.2 Spatial Prediction

We now shift focus to the question of predicting the value of the RF at the unsampled locations d^c , having observed it at the sampled locations d . The term kriging is used to describe the methodology outlined in the following, pioneered by Krige (1951).

Fix $x \in d^c$ and consider the minimum mean square predictor \hat{z}_x of z_x in the sense of minimizing $E(\hat{z}_x - z_x)^2$ subject to the unbiasedness constraint $E(\hat{z}_x - z_x) = 0$. Denote by f_x the vector of covariates at x , by c_x the variance $\operatorname{Var} z_x$, and by c the covariance between z_x and z_d . Then the following result holds (Cressie 1993; Section 3.4.5).

Theorem 15.1 The minimum mean square unbiased predictor \hat{z}_x of z_x for given data z_d is $\hat{z}_x = \lambda' z_d$, where

$$\lambda = \Sigma^{-1} [c + F(F' \Sigma^{-1} F)^{-1} (f_x - F' \Sigma^{-1} c)], \quad (15.5)$$

with prediction variance

$$\operatorname{Var} \hat{z}_x = c_x - c' \Sigma^{-1} c + (f_x' - c' \Sigma^{-1} F) (F' \Sigma^{-1} F)^{-1} (f_x - F' \Sigma^{-1} c). \quad (15.6)$$

Remark: Note that the predictor \hat{z}_x as well as its variance (15.6) depends on the covariance parameters θ . On the other hand, (15.6) does not depend on the possibly unobserved data z_d , which can be advantageous later when we consider optimal designs for prediction.

15.3.2 Designs for Estimating the Covariance Structure

Consider the spatial domain of interest \mathbb{X} over which we define an RF \mathcal{Z} . We wish to choose n locations in \mathbb{X} from which we will observe \mathcal{Z} and use these observations to estimate the variance parameters of the model θ .

Designs for the purpose of estimating the covariance parameters aim in optimizing with respect to some information functional associated with the estimator for θ in the sense described in Chapter 2.

15.3.2.1 Designs Based on the Generalized Least Squares Method

Consider the generalized least squares estimator (15.3). Then the following result is given in Bogaert and Russo (1999) and Müller and Zimmerman (1999).

Theorem 15.2 *Let $G(\theta)$ be the $m \times p$ matrix with (i, j) element $\partial\gamma(h_i; \theta)/\partial\theta_j$. The information matrix for the generalized least squares estimator (15.3) is*

$$M_d^{\text{GLS}}(\theta) = G(\theta)'V(\theta)^{-1}G(\theta). \tag{15.7}$$

On the other hand, any optimality criterion would depend on the unknown parameter to be estimated θ . One way to overcome this is to replace it by some preliminary estimate, say $\hat{\theta}_0$. Thus, in practice, we derive a “locally” optimal design corresponding to $\hat{\theta}_0$. Another idea is to adopt a Bayesian model and compute the design criterion by integrating with respect to the distribution of θ although this tends to be computationally intensive. In the following, we will discuss mainly the first approach; we briefly describe the Bayesian design at the end of Section 15.3.3 and in Section 15.3.4.

Consider, for example, the D -optimality criterion

$$\phi_D = -\log |M_d^{\text{GLS}}|.$$

As noted by Müller and Zimmerman (1999), the assumed correlation function has little impact in the optimal design, though the range parameter does have a considerable impact. In particular, we note that smaller values of the range parameter tend to create clustered designs. A typical D -optimal design is shown in Figure 15.7. There is an apparent clustering of sampling locations in the design that accounts for the estimation of the variogram at small lags.

On the other hand, there is a drawback in computing (15.7) directly for every design. It requires the inversion of a large $m \times m$ matrix where m could potentially be as large as $\binom{n}{2}$. Consider, for example, the situation where a new site x is to be selected to augment an existing design d_n of size n to $d_{n+1} = \{x\} \cup d_n$. For every point that might be added in the design, there are potentially n new pairs created, and the variance of the variogram estimator changes from V_n of dimension $m \times m$ to V_{n+1} with dimension at most $(m + n) \times (m + n)$. This can easily become computationally demanding. The trick here is to use a well-known formula for computing the determinant of a partitioned matrix. Writing

$$V_{n+1} = \begin{pmatrix} V_x & V'_{n,x} \\ V_{n,x} & V_n \end{pmatrix} \quad G_{n+1} = \begin{pmatrix} G_x \\ G_n \end{pmatrix}$$




FIGURE 15.7
D-optimal design for covariance estimation for the Matérn correlation function with $\kappa = 0.5$ and $\tau^2/\sigma^2 = 0.5$. Each design consists of 40 points. Optimal design with range being 20% of the domain length (a) and 150% of the domain length (b).

for the matrices $V(\theta)$ and $G(\theta)$ corresponding to the design d_{n+1} , and

$$G_{x|n} = G_x - V'_{n,x} V_n^{-1} G_n \quad V_{x|n} = V_x - V'_{n,x} V_n^{-1} V_{n,x},$$

then it is easy to see that

$$|M_{d_{n+1}}| = \frac{|V_{x|n} + G'_{x|n} M_{d_n}^{-1} G_{x|n}|}{|V_{x|n}|} |M_{d_n}|,$$

which makes calculations significantly faster. Note that this result can be easily generalized when augmenting by more than one point.

15.3.2.2 Designs Based on the Maximum Likelihood Method

We now consider the ML estimator (15.4) and let $\Sigma_j(\theta) = \partial \Sigma(\theta) / \partial \theta_j$. Then the (i, j) element of the Fisher information matrix $M_d^{ML}(\theta)$ is (see Chapter 2)

$$M_d^{ML}(\theta)_{ij} = \frac{1}{2} \text{trace} \left\{ \Sigma(\theta)^{-1} \Sigma_i(\theta) \Sigma(\theta)^{-1} \Sigma_j(\theta) \right\}. \tag{15.8}$$

As in the GSL case, the information matrix depends on the unknown parameter to be estimated. The solution to this problem is again to replace it by a preliminary estimate $\hat{\theta}_0$. Consequently, the *D*-optimality criterion becomes

$$\underset{d \in \mathcal{X}}{\text{Maximize}} \log |M_d^{ML}|. \tag{15.9}$$

However, there is a more significant problem that needs to be resolved. The use of the Fisher information matrix is justified because asymptotically the variance of the ML estimator is its inverse. In spatial statistics, it is far from clear whether this is true asymptotically because, in theory, the notion of an infinite sample can occur in two different ways. In the first, the sampling sites increase, but their minimum distance remains bounded away from zero, so the domain over which the observations take place becomes unbounded. This is known as the increasing domain asymptotic regime, and asymptotic results hold in analogy to time series (Mardia and Marshall 1984). On the other hand, one may consider the domain fixed and the sampling becoming denser as the sample size increases. This is known as the infill asymptotic regime where eventually the minimum distance between sites converges to zero resulting in a strong type of dependence between observations. It has been shown (Zhang 2004) that under the infill regime, the established likelihood theory does not hold, and the Fisher information matrix does not approach the variance of the ML estimator asymptotically. Even so, it can be shown (Zhu and Stein 2005) that (15.9) ranks candidate designs in roughly the correct order, making it an appropriate criterion to use.

Designs constructed using the ML approach have similar features to the ones derived from GLS (Figure 15.7). It is also possible to consider estimation of the regressor coefficients jointly. As discussed elsewhere in this handbook (e.g., Chapter 2), it is the case that the Fisher information matrix for the joint estimators (β, θ) has a block-diagonal structure that does not depend on β .

15.3.3 Designs for Spatial Prediction

In this case, we can choose to observe the RF \mathcal{Z} at some locations d and predict it at different locations d^c . The measures used to assess discrepancy of prediction are not the same measures used to assess estimation, and designs for this purpose can be very different from the ones derived for the purpose of estimation. In this section, we will discuss two criteria for deriving optimal designs for prediction.

15.3.3.1 Designs Based on the Prediction Variance

Recall from Theorem 15.1 that the prediction variance, or mean square prediction error (MSPE), for predicting the RF at a given location given its value at other locations is given by (15.6). Equation (15.6) can be used as a measure of the uncertainty in prediction, and the goal is to make it as small as possible. In most circumstances, this can be achieved by sampling very close to x . On the other hand, it is rarely the case that we want to predict at only one location.

Often a finite set $\mathcal{X} \subset \mathbb{X}$ of size N is prescribed out of which n sites d are to be sampled and the remaining $N - n$ unsampled sites are to be predicted. An overall criterion that encompasses uncertainty over all unsampled sites is needed. Subsequently, we denote by $\phi\{v(x) : x \in \mathcal{X}\}$ the criterion chosen to minimize where $v(x)$ is some measure of uncertainty at location x , or simply by $\phi\{\mathcal{X}\}$ when $v(x)$ is understood.

One such criterion is the so-called integrated MSPE:

$$\phi\{\mathcal{X}\} = \frac{1}{N - n} \sum_{x \in d^c} \text{Var } \hat{z}_x,$$

which can be interpreted as minimizing the average uncertainty in the prediction. Another criterion discussed in the literature is the maximum prediction variance,

$$\phi\{\mathcal{X}\} = \max_{x \in d^c} \text{Var } \hat{z}_x,$$

while a third criterion is the average standard deviation,

$$\phi\{\mathcal{X}\} = \frac{1}{N - n} \sum_{x \in d^c} \sqrt{\text{Var } \hat{z}_x},$$

which is interpreted as minimizing the average length of the prediction interval at locations d^c . For all three criteria mentioned earlier, the optimal design $d \subset \mathcal{X}$ will be the one that gives the smallest value.

Of course, each criterion might suggest a different optimal design; however, the designs that arise from each are very similar. Two typical designs for the minimization of the integrated MSPE are shown in Figure 15.8. In the case where the mean is constant over the field, these designs tend to be evenly spread in space as shown in the left part of Figure 15.8. In the right, the optimal design when the mean increases linearly with the location from southwest to northeast is shown. In this case, the criterion tends to push designs at the boundary of the domain and mainly at the opposite corners where the mean differs the most.

In the discussion so far in this section, we implicitly assumed knowledge of the variance parameters of the RF. In practice, this is rarely the case, and the usual practice is to replace it with some consistent estimator. In other words, let $V_0(x; \theta)$ denote the prediction variance (15.6) when the variance parameter is θ and let $V_0(x; \hat{\theta})$ be the so-called plug-in estimate of the variance. Then the following theorem is given in Zimmerman and Cressie (1992).




FIGURE 15.8 Optimal 18-point designs for prediction that minimize the integrated MSPE for constant mean (a), mean increasing linearly from southwest to northeast (b).

Theorem 15.3 Suppose that $E\Sigma(\hat{\theta}) = \Sigma(\theta)$. Then $EV_0(x; \hat{\theta}) < V_0(x; \theta)$.

The implication of Theorem 15.3 is that the plug-in method underestimates the true MSPE of the RF.

To derive a correction, write $\hat{z}_x(\theta)$ for the prediction under θ and similarly $\hat{z}_x(\hat{\theta})$ for the plug-in prediction. Then

$$\begin{aligned} E(\hat{z}_x(\hat{\theta}) - z_x)^2 &= E(\hat{z}_x(\theta) - z_x)^2 + E(\hat{z}_x(\hat{\theta}) - \hat{z}_x(\theta))^2 \\ &= V_0(x; \theta) + E(\hat{z}_x(\hat{\theta}) - \hat{z}_x(\theta))^2. \end{aligned}$$

Using an approximation to the second term derived by Harville and Jeske (1992), we arrive at the criterion

$$V_1(x; \theta) = V_0(x; \theta) + \text{trace} \left\{ \mathbf{M}^{-1}(\theta) \left(\frac{\partial \lambda}{\partial \theta} \right)' \Sigma(\theta) \left(\frac{\partial \lambda}{\partial \theta} \right) \right\},$$

where $\mathbf{M}(\theta) = \mathbf{M}_d^{ML}(\theta)$ is given in (15.8) and $\partial \lambda / \partial \theta$ denotes the matrix of derivatives of the vector λ from (15.5) with respect to θ .

As it stands, $V_1(x; \theta)$ is not directly useful because it still depends on the unknown parameter θ . In practice, we could replace θ by some preliminary estimate $\hat{\theta}$, so in practice we obtain a “locally” optimal design corresponding to $\hat{\theta}$. However, by doing so, we ignore the uncertainty in the estimation of θ . By an application of Taylor expansion, it can be shown that (Zhu and Stein 2006) $\text{Var } V_0(x; \hat{\theta}) \approx V_2(x; \hat{\theta})$, where

$$V_2(x; \theta) = \left(\frac{\partial V_0(x; \theta)}{\partial \theta} \right)' \mathbf{M}(\theta)^{-1} \left(\frac{\partial V_0(x; \theta)}{\partial \theta} \right).$$

This gives rise to a family of criteria of the form

$$\phi \left\{ V_1(x; \hat{\theta}) + \kappa_x V_2(x; \hat{\theta}) : x \in \mathcal{X} \right\}, \tag{15.10}$$

where $\kappa_x, x \in \mathcal{X}$ is a parameter of the criterion. A recommendation made by Zhu and Stein (2006) is $\kappa_x = 1/(2V_0(x; \hat{\theta}))$.

One might argue that κ_x should depend on the desired coverage probability α of the prediction interval. For example, if we wish to predict with very high confidence (i.e., α is very close to 1), we should give more weight to the uncertainty in the covariance matrix; hence, the weight κ_x should be large. In accordance to this, we may choose as a penalty factor $\kappa_x = \zeta_\alpha^2 / (4V_0(x; \hat{\theta}))$ where ζ_α denotes the α -quantile of the standard normal distribution. In this case, the choice $\kappa_x = 1/(2V_0(x; \hat{\theta}))$ would give optimal designs that minimize the length of the 84% prediction interval for the spatial RF. Designs produced by this criterion tend to be evenly spread with small clusters (Figure 15.9).

Another possibility is to use $V_2(x; \hat{\theta})$ to introduce constraints in the optimization. Note that minimizing $V_1(x; \hat{\theta})$ tends to spread the sampling locations uniformly, while




FIGURE 15.9

Optimal 29-point predictive design with estimated parameters for the Matérn model with $\kappa = 0.5$ and $\tau^2/\sigma^2 = 0.1$ using the criterion of Zhu and Stein (2006).

minimizing $V_2(x; \hat{\theta})$ tends to create clusters. One may specify an upper bound L^2 on $V_2(x; \hat{\theta})$ and then try to make $V_1(x; \hat{\theta})$ as small as possible, that is,

$$\text{Minimize } \phi\{V_1(x; \hat{\theta}) : x \in \mathcal{X}\} \text{ subject to } \max_{x \in d^c} V_2(x; \hat{\theta}) < L^2.$$

15.3.3.2 Entropy-Based Designs

In this section, we describe the entropy-based approach to spatial design theory. This approach is closely related to the information provided by an experiment proposed by Lindley (1956) that we now describe.

Definition 15.2 The information pertaining to a random variable y with probability density/mass function (pdf) $f(y)$ is defined by

$$\text{Inf } y = E \log f(y). \tag{15.11}$$

In particular, for a normally distributed n -dimensional random variable z with variance $\text{Var } z = \Sigma$,

$$\text{Inf } z = -\frac{n}{2} \log(2\pi e) - \frac{1}{2} \log |\Sigma|. \tag{15.12}$$

This definition can be extended to the conditional information of y given z with pdf $f(y|z)$:

$$\text{Inf } (y|z) = E \log f(y|z),$$

where the expectation is taken with respect to the joint distribution of (y, z) . Thus, we have the following decomposition of the joint information about (y, z) .

Lemma 15.1

$$\text{Inf}(y, z) = \text{Inf}(y|z) + \text{Inf} z. \tag{15.13}$$

Remark: In fact, for a continuous random variable, (15.11) does not correspond to an information measure since it does not attain some of the desired properties of information, namely, it can become positive for certain random variables, and it is not invariant under monotone transformations. To that end, Jaynes (1968) proposed using

$$\text{Inf}_m y = E \log \frac{f(y)}{m(y)},$$

where m is an appropriate reference measure representing complete ignorance. However, m can be chosen somewhat arbitrarily in such a way so that it satisfies (15.13). In the following, we will ignore the role of the measure m .

The connection between information theory and experimental design was made by Lindley (1956). Suppose we are interested on a random quantity y with pdf $f(y)$. An experiment \mathcal{E} will be conducted to gain information about the value of y , which will result data z . This leads to the following definition.

Definition 15.3 *The amount of information gained about y from the data z is defined to be*

$$\text{Inf}(y|z) - \text{Inf} y. \tag{15.14}$$

Bernardo (1979) showed that (15.14) can be interpreted, under some assumptions, as the expected utility about the experiment. It is therefore sensible to rank designs according to (15.14) and select the experiment that corresponds to the highest value.

Recall that in the context of spatial prediction, we are interested in predicting the RF over \mathcal{X} , $z_{\mathcal{X}}$, by sampling at n sites $d = \{x_1, \dots, x_n\} \subset \mathcal{X}$ where d corresponds to the experimental design.

In the following, we partition $z_{\mathcal{X}} = (z_d, z_{d^c})$ with variance partitioned similarly as

$$\Sigma_{\mathcal{X}} = \begin{pmatrix} \Sigma_d & \Sigma_{dd^c} \\ \Sigma_{d^c d} & \Sigma_{d^c} \end{pmatrix}$$

and let

$$\Sigma_{d^c|d} = \Sigma_{d^c} - \Sigma_{d^c d} \Sigma_d^{-1} \Sigma_{d d^c}$$

be the Schur complement of the block corresponding to d^c . In the special case where $z_{\mathcal{X}}$ is normally distributed, $\Sigma_{d^c|d}$ is the conditional variance of $z_{d^c}|z_d$ and

$$\text{Inf } z_{d^c} = -\frac{N-n}{2} \log(2\pi e) - \frac{1}{2} \log |\Sigma_{d^c}|$$

and

$$\text{Inf}(z_{d^c}|z_d) = -\frac{N-n}{2} \log(2\pi e) - \frac{1}{2} \log |\Sigma_{d^c|d}|.$$

By (15.14), the optimal design for predicting z_{d^c} from z_d is obtained by

$$\text{Maximize}_{d \subset \mathcal{X}} \text{Inf}(z_{d^c}|z_d) - \text{Inf } z_{d^c}, \quad (15.15)$$

which is interpreted as the gain in information about z_{d^c} from observing z_d . Another way to see (15.15) is as the mutual information between z_d and z_{d^c} , which is interpreted as the reduction in the uncertainty in z_{d^c} due to the knowledge of z_d . In the case of the normal distribution, this amounts to

$$\text{Minimize}_{d \subset \mathcal{X}} \log |\Sigma_{d^c|d}| - \log |\Sigma_{d^c}|. \quad (15.16)$$

On the other hand, the criterion in (15.15) has the disadvantage that it ignores the information at the sampled sites and only maximizes the information at the unsampled locations. However, since we are interested in the RF in both the sampled and unsampled locations, we should consider the gain in information compared to the uncertainty over all sites of interest. To that end, the criterion (15.14) becomes

$$\text{Inf}(z_{d^c}|z_d) - \text{Inf } z_{\mathcal{X}}. \quad (15.17)$$

The optimal choice for d will then be the one that maximizes (15.17). Noting the decomposition (15.13), we have

$$\text{Inf}(z_{d^c}|z_d) - \text{Inf } z_{\mathcal{X}} = -\text{Inf } z_d;$$

hence, according to criterion (15.17), the optimal choice for d is derived by

$$\text{Minimize}_{d \subset \mathcal{X}} \text{Inf } z_d. \quad (15.18)$$

A different interpretation of (15.18), derived from (15.13) by noting that $\text{Inf } z_{\mathcal{X}}$ is constant over different choices of d , is the design that maximizes $\text{Inf}(z_{d^c}|z_d)$, that is, the information about z_{d^c} contained in z_d , or the design that minimizes the entropy of $z_{d^c}|z_d$, that is, the uncertainty in z_{d^c} by observing z_d . In the case where $z_{\mathcal{X}}$ is normally distributed, by (15.12), the criterion becomes simply

$$\text{Maximize}_{d \subset \mathcal{X}} |\Sigma_d|.$$

We can naturally extend this criterion by taking a Bayesian approach. Consider, for example, the case where at location x

$$z_x = \mathbf{f}'_x \boldsymbol{\beta} + s_x, \quad x \in \mathbb{X},$$

for k -dimensional \mathbf{f}_x and $\boldsymbol{\beta}$, and s_x being a zero-mean, isotropic Gaussian RF with $\text{Var } s_x = \sigma^2$ for all $x \in \mathbb{X}$. Denote by \mathbf{F} the $N \times k$ matrix with rows \mathbf{f}'_x and \mathbf{s} the N -dimensional vector with elements $s_x, x \in \mathcal{X}$, and let $\text{Var } \mathbf{s} = \sigma^2 \mathbf{R}$ where \mathbf{R} is the correlation matrix of \mathbf{s} . If an error term is added to the model, then \mathbf{R} is replaced by $\mathbf{R}^* = \mathbf{R} + (\tau^2/\sigma^2)\mathbf{I}$ where τ^2 is the error variance. Assume further that $\boldsymbol{\beta}$ is normally distributed with covariance matrix $\text{Var } \boldsymbol{\beta} = \nu^2 \mathbf{I}_k$. Then $\boldsymbol{\Sigma} = \text{Var } \mathbf{z} = \nu^2 \mathbf{F}\mathbf{F}' + \sigma^2 \mathbf{R}$. Let $\boldsymbol{\Sigma}_d, \mathbf{R}_d, \mathbf{F}_d$ be the blocks of the corresponding matrices associated with the design d so that

$$|\boldsymbol{\Sigma}_d| = (\sigma^2)^n |\mathbf{R}_d| |\mathbf{I}_k + \alpha \mathbf{F}'_d \mathbf{R}_d^{-1} \mathbf{F}_d|,$$

where $\alpha = \nu^2/\sigma^2$ is the signal-to-noise ratio. When $\alpha = 0$, we have the simple Gaussian model. However, as α increases, the criterion behaves more like $|\mathbf{F}'_d \mathbf{R}_d^{-1} \mathbf{F}_d|$, which is the D -optimality criterion for $\boldsymbol{\beta}$, and as such, it tends to select sites toward the boundary of the region of interest.

In the derivations previously, the knowledge of the covariance matrix of the RF is assumed, but we may consider a fully Bayesian approach where the covariance matrix is also unknown. We may also extend the model to allow for replicated measurements at each location, possibly at different times. Such models could be used for spatiotemporal fields, for example, air-pollution fields (Zidek et al. 2000). For example, consider the following model:

$$z_{tx} = \mathbf{f}'_t \boldsymbol{\beta}_x + s_{tx}, \quad x \in \mathbb{X}, \quad t = 1, \dots, T,$$

where t denotes the time point, \mathbf{f}_t is a vector of k regressors that does not vary in \mathbb{X} , $\boldsymbol{\beta}_x$ is a k -dimensional vector of regressor coefficients, and s_{tx} is a zero-mean Gaussian RF for all t such that s_{t_1x} is independent of s_{t_2x} for $t_1 \neq t_2$. Let $\mathbf{z}_t, \mathbf{s}_t$ be the N -dimensional vectors containing the values of z_{tx} and s_{tx} over \mathcal{X} , respectively, and similarly define the $N \times k$ matrix \mathbf{B} with rows $\boldsymbol{\beta}'_x$ for $x \in \mathcal{X}$. Furthermore, let $\text{Var } \mathbf{s}_t = \boldsymbol{\Sigma}$. Assume further the following conjugate prior distributions for the parameters \mathbf{B} and $\boldsymbol{\Sigma}$ (Anderson 2003; Section 7.7):

$$\begin{aligned} \mathbf{B} | \boldsymbol{\Sigma} &\sim N_{N,k}(\mathbf{B}_0, \boldsymbol{\Sigma} \otimes \mathbf{Q}^{-1}), \\ \boldsymbol{\Sigma} &\sim W_N^{-1}(\boldsymbol{\Psi}, \delta), \end{aligned}$$

with \mathbf{Q} a $k \times k$ known matrix and where $N_{m,n}(\mathbf{M}, \mathbf{U} \otimes \mathbf{V})$ denotes the $m \times n$ matrix normal distribution with mean \mathbf{M} and variance $\mathbf{U} \otimes \mathbf{V}$ and $W_m^{-1}(\mathbf{U}, \delta)$ denotes the m -dimensional inverse Wishart distribution with scale matrix \mathbf{U} and degrees of freedom δ (Dawid 1981).

Suppose data are obtained at sites $d \subset \mathcal{X}$ that provide data $\mathbf{z}_{td}, t = 1, \dots, T$. Then the predictive distribution of \mathbf{z}_{t^*} for some $t^* > T$ is the matrix t distribution (Dawid 1981). The unconditional distribution of $\{\mathbf{z}_{td} : t = 1, \dots, T\}$ is also a matrix t with information

$$\text{Inf}(\{\mathbf{z}_{td} : t = 1, \dots, T\}) = -\frac{1}{2} \log |\boldsymbol{\Psi}| + \text{constant}.$$

The details of the derivations are outlined in Le and Zidek (1992) and in Section 9.1 of Le and Zidek (2006). Therefore, the optimal design is obtained by

$$\underset{d \subset \mathcal{X}}{\text{Minimize}} |\Psi|.$$

Remark: It is worth pointing out that the design criterion does not depend on the time t^* that we wish to predict that could potentially change in the future. This is a significant advantage of entropy-based designs as they are optimal for every objective for which we intend to use on the response on \mathcal{X} .

15.3.4 Bayesian Design

Choosing a good design can be put into a decision-theoretic framework. Let $d \subset \mathcal{X}$ be a sampling design and let z be the data on \mathcal{X} according to the model $f(z|\theta)$ where θ denotes the parameters of the model. Then one can define a utility function $u(d, z, \theta)$ for the given triplet (d, z, θ) , and the choice of the optimal d can be seen as a decision problem. With that perspective, the optimal d would be the one that maximizes the expected utility $E u(d, z, \theta)$ where the expectation is taken with respect to the distribution of z .

In the Bayesian context, the parameter θ is a random variable with prior density $f(\theta)$. Then the best decision is the Bayes, rule, that is, the one that minimizes the Bayes risk, defined by

$$U(d) = - \int u(d, z, \theta) f(z|\theta) f(\theta) dz d\theta.$$

The view of optimal spatial sampling as a decision problem has been developed by Sansó and Müller (1999) and Müller et al. (2004), among others. Suppose that we are interested in the value of the RF at sites \mathcal{X} and we observe z_d at sites $d \subset \mathcal{X}$. The sample will be used for predicting z_{d^c} at the remaining sites $d^c = \mathcal{X} \setminus d$. Denote the prediction at site $x \in d^c$ by \hat{z}_x and define the following utility function:

$$u(d, z, \theta) = C \sum_{x \in d^c} \mathbb{I}(|z_x - \hat{z}_x| < \delta) - \sum_{x \in d} c(x) + C_0, \tag{15.19}$$

although other choices are possible depending on the objective of the network. In (15.19), the expression $\mathbb{I}(A)$ denotes the indicator function, taking the value 1 if event A occurs and the value 0 otherwise. For this particular choice, C is a parameter denoting the payout for prediction accuracy within δ , $c(x)$ for $x \in \mathcal{X}$ denotes the cost of sampling at site x , and C_0 is a fixed parameter chosen to make (15.19) positive. For example, suppose we wish to design a network for monitoring air pollution. One of the aims is to be able to detect sites where the pollution level exceeds a given standard z^* . We also wish to minimize the overall prediction error, and finally, we wish to minimize the running cost of the network. Following Müller et al. (2004), we deduce the following utility function:

$$u(d, z, \theta) = C_1 \sum_{x \in d^c} \mathbb{I}(z_x > z^*) \mathbb{I}(\hat{z}_x > z^*) + C_2 / \sum_{x \in d^c} (z_x - \hat{z}_x)^2 - \sum_{x \in d} c(x) + C_0,$$

where the coefficients C_1 and C_2 denote the importance of each aim relative to the costs.

We proceed by augmenting the original probability model for (z, θ) to include the design d as a random variable. To that end, consider the joint density

$$h(d, z, \theta) \propto u(d, z, \theta)f(z|\theta)f(\theta);$$

thus, the distribution of d is defined from its joint density $h(d, z, \theta)$ in terms of the utility function. The idea is to use $h(d, z, \theta)$ to draw random samples of (d, z, θ) using, for example, a Markov chain Monte Carlo (MCMC) algorithm. Even so, it is a nontrivial task to choose the optimal design, that is, the design that minimizes the Bayes, risk. One approach is to perform clustering on the random samples for d and choose a design that is a member of the largest cluster. A better approach is to consider the proposal density

$$h_J(d, z_1, \theta_1, \dots, z_J, \theta_J) \propto \prod_{j=1}^J u(d, z_j, \theta_j)f(z_j|\theta_j)f(\theta_j),$$

where J increases regularly as the MCMC algorithm progresses, an idea inspired by the simulated annealing algorithm with the cooling corresponding to $T = 1/J$. In this case for the marginal density of d , $h_J(d) \propto U(d)^J$, which, for large J , is concentrated around its maximum. Typically, an algorithm would start with $J = 1$ and gradually increment J to a sufficiently large value. The random samples at the final J would correspond to the optimal design. This has the advantage that there is no need to perform a separate maximization step as this is incorporated into the sampling.

The advantages of this method are that it can be used for very complicated decision problems, which combine parameter estimation, prediction, and cost constraints. Such situations arise in the context of designing large monitoring networks for air pollution. Although it is a very general method, it is also highly computationally intensive.

15.3.5 Dynamic Designs

The designs described in the previous sections were *static*, in the sense that the sampling sites are decided before the experiment takes place and cannot be changed during the experiment. This approach would be reasonable if the RF is constant over time; however, there are examples, such as monitoring the weather or air pollution, where the RF is not constant.

Here, we extend our definition of an RF and suppose that $\mathcal{Z} = \{z_{tx}, x \in \mathcal{X}, t = 1, 2, \dots\}$ is a spatiotemporal process and let $z_t = \{z_{tx}, x \in \mathcal{X}\}$ denote the value of the RF at time t , having dimension N . The process is observed with an error at times $t = 1, 2, \dots$ and at sites $d_t = \{x_1, \dots, x_{n_t}\} \subset \mathcal{X}$ giving data $y_t = \{y(x_1, t), \dots, y(x_{n_t}, t)\}$, according to the following hidden Markov model:

$$y_t = K_t z_t + \epsilon_t, \tag{15.20}$$

$$z_t = H_t z_{t-1} + \eta_t, \tag{15.21}$$

where ϵ_t is an n_t -dimensional zero-mean measurement error process with variance-covariance matrix Σ_ϵ , η_t is an N -dimensional zero-mean spatial process with variance-covariance matrix Σ_η , the matrix K_t is an $n_t \times N$ matrix of 0s and 1s indicating which

components of the RF z_t are observed, and H_t is a general $N \times N$ matrix related to changes in the dynamic nature of the process with time.

In the context of air pollution, z_t could denote the actual concentration of the pollutant at all sites of interest at time t and y_t the noisy measurement at the monitoring sites. Knowledge of z_t could help assess adherence to regulatory standards. However, for reasons of economy, we may not operate the whole network simultaneously, especially if the process exhibits high autocorrelation in time, so we may turn on sensors selectively. This motivates the use of a dynamic design approach.

Consider the purpose of predicting the value of the RF z_t given data y_1, \dots, y_t . Under this model, the conditional prediction variance becomes (Wikle and Royle 1999)

$$\begin{aligned} A_t &= \text{Var}(z_t | y_1, \dots, y_t) \\ &= B_t - B_t K_t' (K_t B_t K_t' + \Sigma_\epsilon)^{-1} K_t B_t \end{aligned} \quad (15.22)$$

$$\begin{aligned} B_t &= \text{Var}(z_t | y_1, \dots, y_{t-1}) \\ &= H_t A_{t-1} H_t' + \Sigma_\eta. \end{aligned} \quad (15.23)$$

These equations are evaluated recursively for $t = 1, 2, \dots$. To start the recursion, A_0 is set to Σ_η , the unconditional variance of z_t . For simplicity, assume that $H_t = H$ for all t , which is sensible if there is very little information about the underlying process.

Assuming that H , Σ_ϵ , and Σ_η are available up to time $t - 1$, B_t is obtained iteratively from (15.22) and (15.23), which may be used in a design criterion either from a prediction variance or from an information perspective.

15.4 Further Reading

Designs for estimating the covariance structure using generalized least squares were discussed in Bogaert and Russo (1999) and Müller and Zimmerman (1999), while the ML approach is developed by Zhu and Stein (2005), Zimmerman (2006), and Xia et al. (2006), among others.

Predictive designs by minimizing the prediction variance were considered by McBratney et al. (1981), Cressie et al. (1990), and Heuvelink et al. (2010), among others, while Zhu and Stein (2006) and Smith and Zhu (2004) make further extensions to account for parameter uncertainty. Information-based designs were originally proposed by Caselton and Zidek (1984) and Shewry and Wynn (1987) with contributions by Caselton et al. (1992) and Zidek et al. (2000). Bueso et al. (1998) provide an extension of this approach to spatiotemporal fields.

Bayesian designs in the context of decision theory were developed by Sansó and Müller (1999) and Müller et al. (2004), while dynamic designs were proposed by Wikle and Royle (1999). The idea of integrating over the model parameters in computing the design criterion was elaborated by Diggle and Lophaven (2006).

Texts covering spatial sampling design include Le and Zidek (2006), Müller (2007), and Mateu and Müller (2012).

References

- Abramowitz, M. and Stegun, I. (1965). *Handbook of Mathematical Functions*. Dover Publications, New York.
- Anderson, T. W. (2003). *An Introduction to Multivariate Statistical Analysis*, 3rd edn. Wiley series in probability and statistics. John Wiley & Sons, New York.
- Bellhouse, D. (1977). Some optimal designs for sampling in two dimensions. *Biometrika*, 64(3): 605–611.
- Bernardo, J. M. (1979). Expected information as expected utility. *The Annals of Statistics*, 7(3):686–690.
- Bogaert, P. and Russo, D. (1999). Optimal spatial sampling design for the estimation of the variogram based on a least squares approach. *Water Resources Research*, 35(4):1275–1289.
- Breidt, F. J. (1995). Markov chain designs for one-per-stratum spatial sampling. In *Proceedings of the Section on Survey Research Methods, American Statistical Association*, Washington, DC, pp. 356–361.
- Bueso, M., Angulo, J., and Alonso, F. (1998). A state-space model approach to optimum spatial sampling design based on entropy. *Environmental and Ecological Statistics*, 5(1):29–44.
- Caselton, W. F., Kan, L., and Zidek, J. V. (1992). Quality data networks that minimize entropy. In Walden, A. T. and Guttorp, P., eds. *Statistics in the Environmental and Earth Sciences*, pp. 10–38. Edward Arnold, London, U.K.
- Caselton, W. F. and Zidek, J. V. (1984). Optimal monitoring network designs. *Statistics & Probability Letters*, 2(4):223–227.
- Cochran, W. G. (2007). *Sampling Techniques*. Wiley series in probability and statistics. John Wiley & Sons, New York.
- Cressie, N., Gotway, C. A., and Grondona, M. O. (1990). Spatial prediction from networks. *Chemometrics and Intelligent Laboratory Systems*, 7(3):251–271.
- Cressie, N. A. C. (1993). *Statistics for Spatial Data*. John Wiley & Sons, New York.
- Das, A. (1950). Two dimensional systematic sampling and the associated stratified and random sampling. *Sankhyā: The Indian Journal of Statistics*, 10(1/2):95–108.
- Dawid, A. P. (1981). Some matrix-variate distribution theory: Notational considerations and a Bayesian application. *Biometrika*, 68(1):265–274.
- Diggle, P. and Lophaven, S. (2006). Bayesian geostatistical design. *Scandinavian Journal of Statistics*, 33(1):53–64.
- Fuller, W. A. (2011). *Sampling Statistics*. Wiley series in survey methodology. John Wiley & Sons, Hoboken, NJ.
- Harville, D. A. (1974). Bayesian inference for variance components using only error contrasts. *Biometrika*, 61(2):383.
- Harville, D. A. and Jeske, D. R. (1992). Mean squared error of estimation or prediction under a general linear model. *Journal of the American Statistical Association*, 87(419):724–731.
- Heuvelink, G. B. M., Jiang, Z., De Bruin, S., and Twenhöfel, C. J. W. (2010). Optimization of mobile radioactivity monitoring networks. *International Journal of Geographical Information Science*, 24(3):365–382.
- Jaynes, E. T. (1968). Prior probabilities. *IEEE Transactions on Systems Science and Cybernetics*, 4(3): 227–241.
- Johnson, M. E., Moore, L. M., and Ylvisaker, D. (1990). Minimax and maximin distance designs. *Journal of Statistical Planning and Inference*, 26(2):131–148.
- Joseph, V. R. and Hung, Y. (2008). Orthogonal-maximin latin hypercube designs. *Statistica Sinica*, 18(1):171.
- Kincaid, T. M. and Olsen, A. R. (2013). spsurvey: Spatial survey design and analysis. R package version 2.6. URL: <http://www.epa.gov/nheerl.arm/>.
- King, A. J. (1945). The master sample of agriculture: I. Development and use. *Journal of the American Statistical Association*, 40(229):38–45.

- Krige, D. G. (1951). A statistical approach to some mine valuation and allied problems on the Witwatersrand. Master's thesis, University of the Witwatersrand, Johannesburg, South Africa.
- Le, N. D. and Zidek, J. V. (1992). Interpolation with uncertain spatial covariances: A Bayesian alternative to kriging. *Journal of Multivariate Analysis*, 43(2):351–374.
- Le, N. D. and Zidek, J. V. (2006). *Statistical Analysis of Environmental Space-Time Processes*. Springer series in statistics. Springer Verlag, New York.
- Lindley, D. V. (1956). On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, 27(4):986–1005.
- Mardia, K. V. and Marshall, R. J. (1984). Maximum likelihood estimation of models for residual covariance in spatial regression. *Biometrika*, 71(1):135.
- Mateu, J. and Müller, W. G., eds. (2012). *Spatio-Temporal Design: Advances in Efficient Data Acquisition*. John Wiley & Sons, Chichester, U.K.
- McBratney, A. B., Webster, R., and Burgess, T. M. (1981). The design of optimal sampling schemes for local estimation and mapping of regionalized variables—I. Theory and method. *Computers and Geosciences*, 7(4):331–334.
- Monroe, J. and Finkner, A. (1959). *Handbook of Area Sampling*. Chiton Co. Book Division, Philadelphia, PA.
- Morris, M. D. and Mitchell, T. J. (1995). Exploratory designs for computational experiments. *Journal of Statistical Planning and Inference*, 43(3):381–402.
- Müller, P., Sansó, B., and De Iorio, M. (2004). Optimal Bayesian design by inhomogeneous Markov chain simulation. *Journal of the American Statistical Association*, 99(467):788–798.
- Müller, W. G. (2001). Coffee-house designs. In Atkinson, A., Bogacka, B., and Zhigljavsky, A. eds. *Optimum Design 2000, Nonconvex Optimization and Its Applications*, Vol. 51, pp. 241–248. Springer, New York.
- Müller, W. G. (2007). *Collecting Spatial Data: Optimum Design of Experiments for Random Fields*. Springer Verlag, Heidelberg, Germany.
- Müller, W. G. and Zimmerman, D. L. (1999). Optimal designs for variogram estimation. *Environmetrics*, 10(1):23–37.
- Nusser, S., Breidt, F., and Fuller, W. (1998). Design and estimation for investigating the dynamics of natural resources. *Ecological Applications*, 8(2):234–245.
- Overton, W. S. and Stehman, S. V. (1993). Properties of designs for sampling continuous spatial resources from a triangular grid. *Communications in Statistics—Theory and Methods*, 22(9):251–264.
- Royle, J. A. and Nychka, D. (1998). An algorithm for the construction of spatial coverage designs with implementation in splus. *Computers and Geosciences*, 24(5):479–488.
- Sansó, B. and Müller, P. (1999). Redesigning a network of rainfall stations. In Gatsonis, C., Carlin, B., Kass, R. E., West, M., Carriquiry, A., Gelman, A., and Verdinelli, I., eds., *Case Studies in Bayesian Statistics, Volume IV*, volume 140 of Lecture notes in statistics, pp. 383–394. Springer, New York.
- Shewry, M. C. and Wynn, H. P. (1987). Maximum entropy sampling. *Journal of Applied Statistics*, 14(2):165–170.
- Smith, R. L. and Zhu, Z. (2004). Asymptotic theory for kriging with estimated parameters and its application to network design. Preliminary version, available from <http://www.stat.unc.edu/postscript/rs/supp5.pdf>. Accessed: June 18, 2014.
- Stevens, D. L. and Olsen, A. R. (2004). Spatially balanced sampling of natural resources. *Journal of the American Statistical Association*, 99(465):262–278.
- Sukhatme, P. V. (1957). *Sampling Theory of Surveys with Applications*. ISAS, Delhi, India.
- Tang, B. (1993). Orthogonal array-based latin hypercubes. *Journal of the American Statistical Association*, 88(424):1392–1397.
- Wikle, C. K. and Royle, J. A. (1999). Space-time dynamic design of environmental monitoring networks. *Journal of Agricultural, Biological, and Environmental Statistics*, 4:489–507.
- Xia, G., Miranda, M. L., and Gelfand, A. E. (2006). Approximately optimal spatial design approaches for environmental health data. *Environmetrics*, 17(4):363–385.

- Zhang, H. (2004). Inconsistent estimation and asymptotically equal interpolations in model-based geostatistics. *Journal of the American Statistical Association*, 99(465):250–261.
- Zhu, Z. and Stein, M. L. (2005). Spatial sampling design for parameter estimation of the covariance function. *Journal of Statistical Planning and Inference*, 134(2):583–603.
- Zhu, Z. and Stein, M. L. (2006). Spatial sampling design for prediction with estimated parameters. *Journal of Agricultural, Biological, and Environmental Statistics*, 11(1):24–44.
- Zidek, J. V., Sun, W., and Le, N. D. (2000). Designing and integrating composite networks for monitoring multivariate Gaussian pollution fields. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 49(1):63–79.
- Zimmerman, D. L. (2006). Optimal network design for spatial prediction, covariance parameter estimation, and empirical prediction. *Environmetrics*, 17(6):635–652.
- Zimmerman, D. L. and Cressie, N. A. (1992). Mean squared prediction error in the spatial linear model with estimated covariance parameters. *Annals of the Institute of Statistical Mathematics*, 44(1):27–43.

Section V

Computer Experiments

16

Design of Computer Experiments: Introduction and Background

Max D. Morris and Leslie M. Moore

CONTENTS

16.1	Introduction.....	577
16.1.1	Basic Notation.....	578
16.2	Numerical Integration.....	579
16.2.1	Stratified Sampling.....	579
16.2.2	Low-Discrepancy Sequences.....	581
16.3	Sensitivity Analysis.....	581
16.3.1	Sampling-Based SA.....	582
16.3.1.1	OA-Based SA.....	583
16.3.2	Rank Regression.....	583
16.4	Metamodels.....	583
16.4.1	GaSP Introduction.....	584
16.4.2	LHS as a Default Design.....	585
16.4.3	Design Refinements.....	585
16.4.3.1	Additional Structure.....	585
16.4.3.2	Maximin and Minimax Distance.....	586
16.4.3.3	Uniform Design Theory.....	587
16.5	Sequential Experiments.....	587
16.5.1	Inverse Problems and Sequential Improvement.....	588
16.5.2	Models and Reality.....	588
16.6	Conclusion.....	589
	References.....	589

16.1 Introduction

The traditional idea of an *experiment* involves observation of a system of interest under controlled conditions, with the intent of learning something about that system. The system of interest varies by discipline: engineers and physicists may be interested in systems involving physical material, biologists may focus on living organisms (or collections or components of them), while social scientists may be interested in experiments involving

the behavior of human beings. In contrast, the system of interest in a computer experiment is often a computer model, usually a mathematical description of a *real* system of interest, represented as a computer program. While not critical to anything written here, the computer representation is usually necessary due to the complexity of the model. Experimental goals are often similar to those in traditional experiments. While the computer model must, in principle, be fully known, it is generally so complex that a useful understanding of its behavior requires the empirical approach of an experiment.

Computer experiments have been performed, in one form or another, throughout the entire history of computers, with physicists being some of the earliest experimenters. One early example, undertaken by Enrico Fermi and colleagues at Los Alamos National Laboratory (LANL), was a simulation of the nonlinear interaction of atoms in a crystal (Fermi et al. 1965). Some of the earliest active involvement of statisticians in analyzing the output of computer models was also related to physics models. For example, in the mid-1970s, the Nuclear Regulatory Commission funded work in LANL's Statistical Sciences Group to collaborate with developers of nuclear reactor safety codes (models), to understand various accident scenarios and potential consequences in nuclear power plants. Growing concern with risk of accidents such as the one at the Browns Ferry Nuclear Plant in Alabama on March 22, 1975, motivated interest in risk and safety assessment. At Browns Ferry, an accidental fire was started by a candle that workers were using to search for air leaks in a repair area of the cable spreading room. The fire subsequently spread to the Unit 1 reactor building causing damage there as well as to control cables for the Unit 2 reactor. Statistical analysis based on nuclear reactor safety codes represented a shift from using sampling as an *internal* component of the calculations implemented in the computer model (e.g., Monte Carlo numerics) to using sampling of the input space of a computer model to gain understanding of consequences of different accident scenarios with different associated risks. Sampling to achieve good numerical calculation required substantial computing resources for each execution of the code, and this in turn resulted in practical limits on the number of executions that could be included in a study. Hence, efficient experiment design became a focus for sensitivity and uncertainty studies using a computer model to evaluate accident consequences for safety studies and risk assessment. From the early years of research supporting nuclear safety studies, a rich statistical research area emerged for developing statistical sampling, experimental design, and analysis methods for complex computer experiments. Motivation for this fundamental work came from applications in the areas of nuclear safety, environmental impact, weapon reliability assessment in the test ban era, and homeland security assurance, all studied via computer models. Various related publications and technical reports co-authored by Michael McKay are available at <http://cybermesa.com/~michaelm/>.

Our discussion is organized around four types of computer experiments. Numerical integration (Section 16.2) and sensitivity analysis (SA) (Section 16.3) are most often set as investigations of the properties of random outputs resulting from random inputs. In contrast, our introduction to metamodels (Section 16.4) and sequential experiments (Section 16.5) focuses on settings in which inputs are regarded as fixed.

16.1.1 Basic Notation

In this chapter, we shall denote model outputs as y (quantities to be regarded as *observations*), inputs as elements of a k -dimensional vector x (quantities defining the controlled conditions to be simulated), and the model itself as the function M . The relationship between these objects is

$$y = M(x, \boldsymbol{\phi}), \quad x \in \mathcal{X}, \quad (16.1)$$

where $\boldsymbol{\phi}$ is a vector of model parameters that are part of the definition of the simulated system and are regarded as constants in many applications and \mathcal{X} denotes the space within which x must lie for valid and/or meaningful use of M . For most of this chapter, we will regard $\boldsymbol{\phi}$ as simply being part of the definition of the model and shorten the notation appropriately:

$$y = M(x), \quad x \in \mathcal{X}. \quad (16.2)$$

We will focus attention on models that are deterministic—multiple executions of M for a given x always yield the same y .

16.2 Numerical Integration

Numerical integration was the first, and perhaps the simplest, kind of computer experiment in which statisticians were heavily involved. Suppose that the input vector x is regarded as a random variable with a specified continuous probability distribution function $F(x)$. The goal of the experiment is to evaluate, at least approximately, the expectation of the model output with respect to the distribution induced by F :

$$\int_{\mathcal{X}} M(x) dF(x). \quad (16.3)$$

In the next section, this is generalized to the expectation of any transformation of the model output:

$$\int_{\mathcal{X}} g(M(x)) dF(x). \quad (16.4)$$

16.2.1 Stratified Sampling

One of the first topics encountered by many statisticians in learning about computer experiments is the *Latin hypercube sample* (LHS), see Chapter 17 for an in-depth treatment. The LHS was introduced by McKay et al. (1979) in the context of computer experiments in which inputs are chosen randomly from some specified distribution, and analysis focuses on estimating properties, such as the mean or specified quantiles, of the resulting distribution of the outputs. In this kind of experiment, the values of inputs actually selected are generally not used in the estimation exercise, that is, n input vectors are randomly selected, the computer model is executed for each of them, and the analysis is based only on the resulting random sample of output values. McKay et al. (1979) focused in particular on averages of functions of the output:

$$T = \frac{1}{n} \sum_{i=1}^n g(y_i), \quad (16.5)$$

where y_i , $i = 1, 2, 3, \dots, n$ is the value of the output of interest resulting from execution of the model with the i th selected set of inputs. In this setting, g is an arbitrary function that accommodates a useful variety of output statistics. For example, $g(y) = y$ leads to the sample mean, $g(y) = y^m$ for integer m yields the m th noncentral sample moment, and $g(y) = 1$ for $y < y^*$ and 0 otherwise results in the empirical distribution function evaluated at y^* .

Latin hypercube sampling is based on the idea that a joint probability distribution has been specified for the input vector, $F(x)$, and that the elements of x are independent so that the joint distribution can be written as the product of the marginals, $F(x) = \prod_{i=1}^k F_i(x_i)$. Values of the inputs are selected individually. For the i th input, the range of x_i is partitioned into n non-overlapping intervals, each of probability $1/n$ under F_i , and one value of x_i is drawn conditionally from each of these intervals. After n values have been thus chosen for each input, they are combined randomly (with equal probability for each possible arrangement) to form the n input vectors each of order k . When n is large, the conditional sampling from equal-probability intervals is often ignored, and values are simply taken from a grid. Figure 16.1a displays an LHS in $k = 2$ inputs and $n = 48$ points.

The basic result presented by McKay et al. (1979) compares the efficiency of estimation for LHS to that for simple random sampling (SRS) and can be easily stated. For a fixed sample size n , let T_{SRS} be the quantity of interest calculated from outputs resulting from a simple random sample of inputs and let T_{LHS} be the same quantity resulting from an LHS. Then if the computer model is such that y is a monotonic function of each of the inputs, and g is a monotonic function of y , then $\text{Var}(T_{LHS}) \leq \text{Var}(T_{SRS})$. Stein (1987) showed that so long as $E(g(y)^2)$ is finite, the asymptotic (large n) variance of T_{LHS} is no larger than that of T_{SRS} without the monotonicity requirements and that the asymptotic efficiency of T_{LHS} relative to T_{SRS} is governed by how well the computer model can be approximated by a linear function in x .




FIGURE 16.1

Designs in $k = 2$ inputs and $n = 48$ runs. (a) Latin hypercube sample. (b) Sobol sequence. (c) 4×4 Orthogonal array-based Latin hypercube with approximate maximin distance. (d) Approximate maximin distance Latin hypercube.

16.2.2 Low-Discrepancy Sequences

There are, of course, nonstatistical arguments for how a finite number of function evaluations should be selected so as to best approximate the integral of that function. One approach is through specification of a so-called minimum discrepancy sequence, motivated in the following way. Suppose the intent is to select a sampling design $d = \{x_1, x_2, \dots, x_n\}$ in such a way as to spread the n design points throughout \mathcal{X} as *uniformly* as possible (in some useful sense). Thinking of the n points in k -dimensional space as a random sample (even though, for these purposes, it actually is not), the Kolmogorov–Smirnov statistic for characterizing the departure (or discrepancy) of this collection of points from any specified probability distribution function F is

$$\max_{x \in \mathcal{X}} |\hat{F}(x) - F(x)|, \quad (16.6)$$

where \hat{F} is the empirical probability distribution function for its argument.

$$\hat{F} = \frac{1}{n} \sum_{j=1}^n I(x_j < x). \quad (16.7)$$

A more general measure of discrepancy indexed by p is

$$\phi_p(d) = \left[\int_{\mathcal{X}} |\hat{F}(x) - F(x)|^p \right]^{1/p}, \quad (16.8)$$

for which the Kolmogorov–Smirnov statistic, or *star discrepancy*, is recovered when $p \rightarrow \infty$. Motivation for this approach relies on a deterministic Koksma–Hlawka bound:

$$\left| T - \int_{\mathcal{X}} g(y(x)) dx \right| \leq \Phi_{\infty}(d) \times \psi(g(y)), \quad (16.9)$$

where $\psi(g(y))$, sometimes called the total variation of g , is a sum of integrated absolute derivatives of the function and is not dependent on the design. Selecting a design that minimizes ϕ thus minimizes the bound on this error (if not the error itself). Figure 16.1b displays a $k = 2$ dimensional Sobol sequence of $n = 48$ points.

Quasi-Monte Carlo is a phrase often used to describe the approach presented here to select design points for numerical integration. Niederreiter (1992) and Lemieux (2009), for example, offer substantial background on quasi-Monte Carlo and a number of low-discrepancy sequences for this purpose.

16.3 Sensitivity Analysis

Informally, the phrase “sensitivity analysis” generally refers to an investigation that has as its goal the determination of which elements of x are most influential in determining the value of y (or changes in y). This kind of activity can be carried out in a

number of different ways, with differences largely determined by how the word *influential* is interpreted and the assumptions about M that are acceptable. In the following sections, we briefly describe two approaches to sensitivity analysis (SA) that have been widely used across a number of application areas, both of which are based on the formal assumption that the input vector can be regarded as a random variable. In this setting, $y = M(x)$ is also clearly random, and standard statistical measures of variability and correlation may be used as the basis for SA. (And since variance and correlation are defined as integrals when y is continuous, the goals of SA can be viewed as a specialized form of those described in the previous section.) In practice, the distribution assigned to x may reflect knowledge of the physical system; in other cases, $F(x)$ is regarded as a formalism that represents neither physical *noise* nor any real expression of knowledgeable uncertainty, but is arbitrarily specified simply to provide a basis for SA.

16.3.1 Sampling-Based SA

Suppose, as in Section 16.2, that $F(x) = \prod_{i=1}^k F_i(x_i)$ is the joint distribution assigned to x . This implies that the variation observed among realizations of y that follow from independent draws of x is actually *propagated* through M from the variation in the inputs, individually and in combination. Unconditionally, suppose the induced distribution of y has variance $\sigma^2 = \text{Var}(y)$. One way to assess the *importance* of each input is through a measure of how much variation in y is associated with that input. Two sets of unit less indices are commonly used for this purpose:

1. Main effect sensitivity of x_i : $S_i = \text{Var}_i[E_{(i)}(y)]/\sigma^2$
2. Total sensitivity of x_i : $T_i = E_{(i)}[\text{Var}_i(y)]/\sigma^2$

where by Var_i , we mean the variance taken with respect to x_i , and by $E_{(i)}$, we mean the expectation taken with respect to all inputs *except* x_i . For any i , $T_i \geq S_i$. This relationship is made clear by considering the *ANOVA decomposition* of y as a function of x (e.g., Sobol 2001). In this expansion, S_i reflects variability in y attributable to x_i alone, while T_i reflects variability transmitted to y from x_i both alone and in conjunction with other inputs (i.e., through interactions). Related to this, $\sum_{i=1}^k S_i \leq 1$ since variation associated with interactions between inputs is not included. On the other hand, $\sum_{i=1}^k T_i \geq 1$, because variation associated with interactions is included in the total sensitivity index of each involved inputs. Hence, T_i is sometimes described as the more appropriate index when interest centers on the effect of removing uncertainty associated with one input, while S_i is more appropriate where there is interest in knowing the variability attributable to only one input. Chapter 18 describes the development of these indices and their use in more detail, and Saltelli et al. (2000) offer a broader treatment of this approach to SA.

Simple sampling plans that support estimation of S_i and/or T_i can be constructed based on conditional sampling that follows the form of each index (e.g., Saltelli et al. 2010). For example, T_i can be estimated using a hierarchical sampling scheme in which a value for $x_{(i)}$ (including all inputs except x_i) is randomly drawn, and with this fixed, several values of x_i are drawn. The sample variance of outputs resulting from these runs is an estimate of $\text{Var}_i(y|x_{(i)})$. If the process is repeated, and several such estimates are computed, their average is an estimate of T_i . A *replicated* LHS (McKay 1995), in which n draws of each input value are collected, but randomly combined into input vectors a times (for a total of $a \times n$

input vectors), supports estimation of all main-effect sensitivity indices, $S_i, i = 1, 2, 3, \dots, k$. Saltelli (2002) shows how both sets of indices can be estimated using a combined sampling design.

Morris (1991) described a different restricted random sampling procedure for screening inputs to computer models. This approach relies on clusters of input points that take the structure of simple one-factor-at-a-time designs. Such plans are notoriously inefficient in contexts involving random error in y . But in this context, they lead to unambiguous information about the local effect of each input on y and are used repeatedly at randomly selected locations in \mathcal{X} to identify both linear and nonlinear effects.

16.3.1.1 OA-Based SA

As noted, replicated LHSs are constructed by repeatedly using a relatively small number of unique values for each input, combined in different ways to specify different input vectors. When this repeated combination of input values is accomplished through a simple random process, some values of, say, two inputs can appear together in multiple input vectors. This results in bias in the estimates of main-effect sensitivity coefficients, which were characterized by Morris et al. (2006). The same authors (2008) examined how deterministic assignment of input values in *permuted column samples* (which contain replicated LHSs as a special case) can eliminate this bias. Specifically, they showed that permuted column samples of a given size that lead to unbiased estimates of sensitivity coefficients and that employ the smallest number of unique values for each input can be constructed as orthogonal arrays (OAs) of strength 2.

16.3.2 Rank Regression

Another route to characterizing the sensitivity of a y to model inputs, based on an assumed probability structure for the latter, is through correlations between each x_i and y . Iman and Conover (1980) used Latin hypercube sampling to generate input vectors and assessed sensitivity by correlations or regression based on the ranks of the variables (to make the process more robust to nonlinear relationships). Note that while this approach shares the random-input structure with the other sampling-based SA methods, it is based on the tacit assumption that $y = M(x)$ is a continuous function and in fact can only be expected to work well when the relationships between y and each x_i are monotonic. The variance-based approach described in Section 16.3.1 does not require this assumption but typically requires substantially larger sample sizes to produce estimates of useful precision.

16.4 Metamodels

As described in the previous section, the goal of SA is to determine the degree to which changes in individual inputs or groups of inputs influence model outputs and to summarize these relationships with informative indices. In contrast, a *metamodel*, sometimes also called a *surrogate* or an *emulator*, is a complete approximation of the function M over \mathcal{X} . Metamodels can be used for many purposes, but a driving motivation in most cases is the need to compute an approximate y more quickly (sometimes much more quickly) than is

possible through direct evaluation of M . Metamodels are constructed using a large variety of modeling techniques, including neural networks, spline representations, and least squares fits of polynomials (e.g., Barton 1998), although when M is deterministic, the interpretation of metamodels that do not interpolate known data is obviously problematic. In this section, we will briefly outline how Gaussian stochastic processes (GaSP) are used for this purpose. GaSP metamodels yield point predictions of y that do interpolate known data and also provide a principled basis for quantifying the uncertainty of those predictions.

16.4.1 GaSP Introduction

For our purposes, a *stochastic process* is a functional generalization of a *random variable*. In our case, the random function is defined with argument $x \in \mathcal{X}$. A GaSP is fully defined by a mean function $\mu(x)$ and a covariance function $C(x_1, x_2)$; the representation of the random function at x is a Gaussian random variable with mean $\mu(x)$ and variance $C(x, x)$, and at every finite collection of x 's in \mathcal{X} , the representation of the random function is multivariate Gaussian with covariances determined by C . In most applications to computer experiments, the mean structure is simplified (as in linear regression) to a linear combination of known basis functions,

$$\mu(x) = \sum_i \beta_i f_i(x), \tag{16.10}$$

or further simplified to a single value (referred to as a stationary mean),

$$\mu(x) = \mu. \tag{16.11}$$

Further common simplifications are the adoption of a single variance,

$$C(x, x) = \sigma^2, \tag{16.12}$$

and a *stationary* correlation structure

$$C(x_1, x_2) = \sigma^2 R(x_1 - x_2; \theta), \tag{16.13}$$

where elements of the vector θ are correlation parameters associated with the *roughness* of the process realizations and the rate at which positive correlation *dies off* with distance in each direction within \mathcal{X} .

An experimental design for a computer experiment is a list of values of the input set, $d = \{x_1, x_2, \dots, x_n\}$; if each execution of the computer model results in one output, we may summarize these as the n -element vector y . For known parameters (μ , σ , and θ), the *trained* metamodel is the *conditional* Gaussian process *fitted to* the collected data. For an arbitrary input x , the full prediction of the model output is a Gaussian random variable with mean and variance:

$$E[y(x)|y] = \mu + r'_x R^{-1} (y - \mu \mathbf{1}) \quad \text{Var}[y(x)|y] = \sigma^2 (1 - r'_x R^{-1} r_x). \tag{16.14}$$

where elements of the vector r_x and matrix R are correlations defined by C .

More realistically, since the parameter set is generally not known, estimates are often substituted for the parameters (sometimes called the *empirical Bayes* or *plug-in* approach,

e.g., Currin et al. 1991), or a more comprehensive Bayesian inference is performed (e.g., Bayarri et al. 2007). Regardless of how the estimation-and-prediction inference is carried out, the challenge of experimental design is the specification of d so that the metamodel is effective for its intended application.

16.4.2 LHS as a Default Design

As described in Section 16.2, the original justification for the Latin hypercube was as a stratified sampling plan, requiring a specified probability distribution for x . However, the *structure* of the LHS is frequently used in metamodeling applications where x is not regarded as random, resulting in the *Latin hypercube design* (Chapter 17). Intuitive appeal for this approach to design for metamodeling includes the following:

1. *One-dimensional stratification*: In an LHS (or design), each input takes on n distinct values spread across its experimental range. This is not particularly appealing in physical experimentation since it implies that the design cannot have factorial structure or include replication. However, in experiments with deterministic models, there is no uncontrolled *noise* as such; replication is not needed to estimate uncontrolled variation, and the benefits of factorial structure that come from maximizing *signal to noise* in the analysis are not relevant. The n -value one-dimensional projections of a Latin hypercube provide (at least in some cases) the information needed to map out more complex functional y -to- x behavior than can be supported with designs that rely on a small number of unique values for each input.
2. *Space-filling potential*: The modeling techniques that are most appropriate in this context are *data interpolators* rather than *data smoothers*; they generally perform best when the n points of an experimental design *fill the space*, as opposed to being arranged so that there are relatively large subregions of \mathcal{X} containing no points (as is the case, for example, with factorial designs with relatively few levels attached to each input). While Latin hypercube designs do not necessarily have good space-filling properties, they *can* be made to fill \mathcal{X} effectively through judicious (nonrandom) arrangement of the combinations of input values used. Section 16.4.3 describes a few more specific ways in which the space-filling character of experimental designs can be defined and quantified.

16.4.3 Design Refinements

16.4.3.1 Additional Structure

Tang (1993) and Owen (1994) are among the authors who have proposed restrictions on LHSs to make them more effective as experimental designs. The Latin hypercube structure was originally proposed as a *sampling plan*, but its use in computer experiments designed to create metamodels does not require the randomization necessary for sampling inference. While every Latin hypercube offers good one-dimensional projections, both Tang and Owen restrict the construction of LHSs to improve coverage in two-dimensional (or higher-dimensional) projections. Figure 16.1c displays an OA-based Latin Hypercube arrangement of 48 points, constrained in such a way that 3 points lie in each of $4 \times 4 = 16$ grid squares (forming an *overall* 4^2 pattern). Additionally, the design displayed in Figure 16.1c is maximin (Euclidean) distance as described in the next section, from among 10,000 randomly drawn samples with this 4×4 OA-based grid feature.

16.4.3.2 Maximin and Minimax Distance

The use of the GaSP provides a mathematical structure under which optimal design arguments can be formulated. Specifically, we might want to consider designs that are optimal in one of two ways:

1. *D-optimal*: Such that predictions of y made at any finite set of points in \mathcal{X} have the smallest possible generalized variance (determinant of the joint variance matrix)
2. *G-optimal*: Such that the largest prediction variance over all points in \mathcal{X} is as small as possible

While numerically challenging, such designs can be constructed for GaSP models in which parameters are regarded as known, for example, Shewry and Wynn (1987) and Currin et al. (1991) who constructed *D-optimal* designs (which are called *entropy optimal*) for such processes.

A practical difficulty with constructing these optimal designs is the fact that they are dependent on the parameter values, in particular, the parameters of the correlation function. Johnson et al. (1990) proposed an elegant extension to these ideas by developing asymptotic forms of these criteria. In particular, write the covariance function C for a stationary process as

$$C(\mathbf{x}_1, \mathbf{x}_2) = \sigma^2 e^{-\theta d(\mathbf{x}_1, \mathbf{x}_2)}, \quad (16.15)$$

where $d(\mathbf{x}_1, \mathbf{x}_2)$ is a *correlation distance* whose form determines the general character (e.g., differentiability) of the GaSP and the value of the parameter θ determines the rate at which correlation *dies off* with increasing distance. They show that in the limit as θ approaches ∞ (i.e., all correlations become *locally weak*),

1. *D-optimal* designs approach *maximin distance designs* (those for which the closest two points are as far apart as possible) of *minimum index* (with the smallest number of point pairs separated by this minimum distance).
2. *G-optimal* designs approach *minimax distance designs* (those for which the greatest distance between the design and any point not included in the design is minimized) of *minimum index* (with the smallest number of unsampled points separated from the design by this maximum distance).

These distance-based criteria are intuitive in that they lead to designs that *fill* \mathcal{X} in their respective senses and provide a convenient, although asymptotic, link between this intuition and the performance of the GaSP metamodel. Figure 16.1d displays the maximin (Euclidean) distance Latin hypercube arrangement of $n = 48$ points, from among 10,000 randomly drawn samples.

To summarize, the four designs displayed in Figure 16.1 have the following minimum (Euclidean) distance values of 0.030 for design A (LHS), 0.046 for design B (Sobol sequence), 0.077 for design C (4×4 grid OA-based LHS, with maximin distance), and 0.067 for design D (maximin distance). The 4×4 grid structure for the four designs are described in Table 16.1, which lists the number of $n = 48$ design points in each of the 16 grid cell.

TABLE 16.1

Numbers of Points in Each 4×4 Cell of the Designs Displayed in Figure 16.1

Design A				Design B				Design C				Design D			
2	4	3	3	4	2	3	3	3	3	3	3	4	2	2	4
5	1	3	3	3	4	2	3	3	3	3	3	4	3	1	4
2	3	2	5	3	3	3	2	3	3	3	3	2	3	6	1
3	4	4	1	3	3	3	4	3	3	3	3	2	4	3	3

(A) Latin hypercube sample; (B) Sobol sequence; (C) 4×4 OA-based Latin Hypercube with approximate maximin distance; (D) Approximate maximin distance Latin Hypercube.

16.4.3.3 Uniform Design Theory

A more direct approach to constructing a design that *fills* \mathcal{X} is that taken with uniform design theory. Here, the design is constructed as a k -dimensional minimum discrepancy sequence of n points, as described in Section 16.2.2, often minimizing the discrepancy measure of the design with respect to the k -dimensional uniform distribution. As is the case with LHSs used as designs, the original justification for uniform designs is less directly relevant for the purpose of fitting metamodels than it is for numerical integration, but the result often leads to an arrangement of points that is effective for fitting Gaussian process metamodels. The uniform design criterion (as well as the maximin distance criterion) is often applied within a class of attractive designs, such as Latin hypercubes. One less-than-desirable characteristic of uniform designs not shared by distance-based designs is that uniform designs are not generally rotation or reflection invariant. For example, if $\mathcal{X} = [0, 1]^k$, $d_1 = \{x_1, x_2, \dots, x_n\}$, and $d_2 = \{\mathbf{1} - x_1, x_2, \dots, x_n\}$, designs d_1 and d_2 are not necessarily equivalent under uniform design criteria. In contrast, most criteria based on the quality of inference that might be expected from a stationary GaSP model would regard d_1 and d_2 as equivalent. Nonetheless, the uniform design approach typically results in point arrangements that, for practical purposes, fill the design space effectively.

16.5 Sequential Experiments

Sequential experimentation has distinct advantages over one-stage experimentation in a number of different contexts. To the extent that timely analysis is possible and in the interest of conserving computational resources, it is desirable to collect as few runs as required, possibly one at a time until sufficient information is gained to address problems of interest. Of course defining a stopping rule for sufficient data collection is often difficult, and it is generally agreed that some minimal ensemble of runs is initially required, although the minimal required number is not obvious. The goal of SA inherently suggests the value of sequential experimentation and analysis as a way to reduce the dimensionality of the input space and thereby improve the quality of a fitted metamodel. Further, assessed lack of predictive capability of a metamodel demands improvement that requires additional samples, and sequential experimentation in this context can be focused on areas in which the metamodel is initially least accurate. Another possible circumstance of computer experiments is the availability of new information such as field data and physical experimental data that are a basis for model validation or calibration or for reduction of the input space

to a region of greater relevance. The circumstance of available field data poses questions of resource allocation in that the aggregate experimental budget must be divided based on a limited, fixed budget and differential costs and benefits of physical and computer experiments. Although the resource allocation problem based on availability of physical data is beyond our scope, a similar resource allocation problem for sequential computer experimentation arises when dividing an experimental budget between an initial experiment, possibly for conducting SA, and additional sampling, possibly in sequential batches of runs, for metamodel construction and evaluation. The batch size (number of runs in a batch) of a sequential strategy is more commonly dictated by the complexity of setup of an experiment or availability of computing resources than by experiment design optimality.

16.5.1 Inverse Problems and Sequential Improvement

As described earlier, computer models are written to compute an output value y that follows from an input vector x —what is sometimes called a *forward problem*. In many applied settings, however, the question of real interest poses an associated *inverse problem*; given conditions on y , what values of x satisfy these conditions? In many cases, *inverse models* cannot be written to solve this problem directly, and sequential strategies based on evaluations of the forward model are employed.

Function optimization is a widely encountered inverse problem; what value or values of x within \mathcal{X} lead to the largest (or smallest) value of y ? Schonlau et al. (1998) and Jones et al. (1998) each described sequential strategies for function optimization based on the concept of *expected improvement*. In the setting of function maximization, if a new run of the computer model yields an output value that is greater than the largest previously observed output, the difference between these values is the *improvement*; if the output at the new run is no greater than what has been observed, the improvement is zero. Before the new run is executed, the improvement obviously cannot be evaluated, but using a GaSP metamodel trained to the data already acquired, the *expected improvement* can be computed for any new run that might be contemplated. Both Schonlau et al. and Jones et al. use the expected improvement as a criterion to sequentially select runs that, in this sense, are most likely to yield the greatest improvement (increase) in y .

The concept of expected improvement has been extended and adapted to a number of different inverse problems; see Chapter 19 for an in-depth treatment of this topic. One interesting example was given by Ranjan et al. (2008), who developed a sequential design strategy for estimating a *level set* for model output, essentially a contour map in inputs x associated with a selected value of the output y . The motivating application for this work was based on a model of a computer network, where the output of interest was the average delay for user jobs in the system. In this setting, a maximum acceptable value of the average delay may be specified, and the level set used to separate regions of the input space that correspond to acceptable and unacceptable performance.

16.5.2 Models and Reality

All that has been discussed to this point has focused on the computer model and experiments that are designed to understand its behavior, approximate its outputs, etc. However, the greatest interest in most computer models stems from their value in explaining or predicting physical processes or systems. The joint consideration of data from physical processes or systems *and* the computer models written to mimic them leads to model validation and calibration.

A widely used framework for dealing with computer models and the reality they represent was described by Kennedy and O'Hagan (2001). Briefly, they develop a Bayesian analysis around a model of the following form:

$$R(x) = M(x, \Phi) + \delta(x), \quad (16.16)$$

where R represents the *reality* that would occur under conditions specified by x , Φ is a set of model parameters that are not known with certainty, and *discrepancy* δ represents the error of model M to accurately represent reality (even with the correct value of Φ). By observing physical data, generally modeled as R plus measurement error, and executing runs of M for various values of x and Φ , joint inferences for model calibration and correction are made, leading to improved prediction of R at untried values of x .

Williams et al. (2011) and Ranjan et al. (2011) developed a number of sequential design strategies for the Kennedy and O'Hagan framework, in which values of both x and Φ are controlled via criteria based on expected improvement and integrated mean square error of prediction, respectively. In each case, the aim is that of reaching *model maturity*, which Williams et al. define as "stability [that] is reached when the collection of additional field data results in minimal changes to the predictive density as measured by an appropriate metric."

16.6 Conclusion

The field of design for computer experiments is rapidly growing and offers numerous challenges for researchers. For readers interested in more detailed descriptions of the topics described here, the books by Saltelli et al. (2000) and Santner et al. (2003) offer good background in the areas of sampling-based SA and metamodels, respectively.

The following three chapters of this volume offer current and in-depth treatments of three topics briefly introduced here. Chapter 17 discusses Latin hypercube designs, especially the extensions and generalizations that have been made since introduction of Latin hypercube sampling by McKay et al. (1979). Chapter 18 focuses on SA, which is especially important in examining models with relatively large numbers of inputs. The fundamental idea of expected improvement and the numerous ways in which it has been exploited in computer experiments is the topic of Chapter 19.

References

- Barton, R.R. (1998). Simulation metamodels, *Proceedings of the 1998 Winter Simulation Conference*, Medeiros, D.J., E.F. Watson, J.S. Carlson, and M.S. Manivannan, eds., Institute of Electrical and Electronics Engineers, Piscataway, NJ.
- Bayarri, M.J., J.O. Berger, R. Paulo, J. Sacks, J.A. Cafeo, J. Cavendish, C.H. Lin, and J. Tu (2007). A framework for validation of computer models, *Technometrics* **49**, 138–154.
- Currin, C., T.J. Mitchell, M.D. Morris, and D. Ylvisaker. (1991). Bayesian prediction of deterministic functions, with applications to the design and analysis of computer experiments, *Journal of the American Statistical Association* **86**, 953–963.

- Fermi, E., J. Pasta, and S. Ulam. (1955). Studies of nonlinear problems, Report LA-1940. Los Alamos National Laboratory, Los Alamos, NM. Republished in E. Serge', ed., *Collected Papers of Enrico Fermi*, Vol. 2, University of Chicago Press, Chicago, IL, (1965).
- Iman, R.L. and W.J. Conover. (1980). Small sample sensitivity analysis techniques for computer models with an application to risk assessment, *Communications in Statistics A* **9**, 1749–1842.
- Johnson, M.E., L.M. Moore, and D. Ylvisaker. (1990). Minimax and maximin distance designs, *Journal of Statistical Planning and Inference* **26**, 131–148.
- Jones, D., Schonlau, M., and Welch, W. (1998). Efficient global optimization of expensive black-box functions. *Journal of Global Optimization* **13**, 455–492.
- Kennedy, M. and A. O'Hagan. (2001). Bayesian calibration of computer models, (with discussion), *Journal of the Royal Statistical Society B* **68**, 425–464.
- Lemieux, C. (2009). *Monte Carlo and Quasi-Monte Carlo Sampling*, Springer, New York.
- McKay, M.D. (1995). Evaluating prediction uncertainty, Report NUREG/CR-6311, LA-12915-MS. Los Alamos National Laboratory, Los Alamos, NM.
- McKay, M.D., R.J. Beckman, and W.J. Conover. (1979). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code, *Technometrics* **21**, 239–245.
- Morris, M. D. (1991). Factorial sampling plans for preliminary computational experiments, *Technometrics* **33**, 161–174.
- Morris, M. D., L. M. Moore, and M. D. McKay. (2006). Sampling plans based on balanced incomplete block designs for evaluating the importance of computer model inputs, *Journal of Statistical Planning and Inference* **136**, 3203–3220.
- Morris, M. D., L. M. Moore, and M. D. McKay. (2008). Using orthogonal arrays in the sensitivity analysis of computer models, *Technometrics* **50**, 205–215.
- Niederreiter, H. (1992). *Random Number Generation and Quasi-Monte Carlo Methods*, Society for Industrial and Applied Mathematics, Philadelphia, PA.
- Owen, A.B. (1994). Controlling correlations in latin hypercube samples, *Journal of the American Statistical Association* **89**, 1517–1522.
- Ranjan, P., D. Bingham, and G. Michalidis. (2008). Sequential experiment design for contour estimation from complex computer codes, *Technometrics* **50**, 527–541.
- Ranjan, P., W. Lu, D. Bingham, S. Reese, B. J. Williams, C.-C. Chou, F. Doss, M. Grosskopf, and J.P. Holloway. (2011). Follow-up experimental designs for computer models and physical processes, *Journal of Statistical Theory and Practice* **5**, 119–136.
- Saltelli, A. (2002). Making best use of model evaluations to compute sensitivity indices, *Computer Physics Communications* **145**, 280–297.
- Saltelli, A., K. Chan, and E.M. Scott. (eds.) (2000). *Sensitivity Analysis*, John Wiley & Sons, New York.
- Saltelli, A., P. Annoni, I. Azzini, F. Campolongo, M. Ratto, and S. Tarantola. (2010). Variance based sensitivity analysis of model output: Design and estimator for the total sensitivity index, *Computer Physics Communications* **181**, 259–270.
- Santner, T.J., B.J. Williams, and W.I. Notz (2003). *The Design and Analysis of Computer Experiments*, Springer-Verlag, New York.
- Schonlau, M., Welch, W.J., and Jones, D.R. (1998). Global versus local search in constrained optimization of computer models, in *New Developments and Applications in Experimental Design*, eds. Flournoy, N., Rosenberger, W.F., and Wong, W.K., Institute of Mathematical Statistics, Vol. 34, pp. 11–25.
- Shewry, M.C. and H.P. Wynn. (1987). Maximum entropy sampling, *Journal of Applied Statistics* **14**, 165–170.
- Sobol, I.M. (2001). Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates, *Mathematics and Computers in Simulation* **55**, 271–280.

- Stein, M. (1987). Large sample properties of simulations using latin hypercube sampling, *Technometrics* **29**, 143–151.
- Tang, B. (1993). Orthogonal array-based latin hypercubes, *Journal of the American Statistical Association* **88**, 1392–1397.
- Williams, B.J., J.L. Loeppky, L.M. Moore, and M.S. Macklem (2011). Batch sequential design to achieve predictive maturity with calibrated computer model, *Reliability Engineering and System Safety* **96**, 1208–1219.

17

Latin Hypercubes and Space-Filling Designs

C. Devon Lin and Boxin Tang

CONTENTS

17.1 Introduction.....	593
17.2 Latin Hypercube Designs.....	595
17.2.1 Introduction and Examples.....	595
17.2.2 Latin Hypercube Designs Based on Measures of Distance.....	597
17.2.2.1 Maximin Latin Hypercube Designs.....	598
17.2.3 Orthogonal Array-Based Latin Hypercube Designs.....	601
17.2.4 Orthogonal and Nearly Orthogonal Latin Hypercube Designs.....	604
17.2.4.1 Construction Method Based on an Orthogonal Array and a Small Orthogonal Latin Hypercube.....	605
17.2.4.2 Recursive Construction Method.....	608
17.2.4.3 Construction Method Based on Small Orthogonal Designs and Small Orthogonal Latin Hypercubes.....	610
17.2.4.4 Existence of Orthogonal Latin Hypercubes	613
17.3 Other Space-Filling Designs.....	615
17.3.1 Orthogonal Designs with Many Levels.....	615
17.3.2 Low-Discrepancy Sequences and Uniform Designs.....	616
17.3.2.1 (t, m, s) -Nets and (t, s) -Sequences.....	617
17.3.2.2 Uniform Designs.....	619
17.4 Concluding Remarks.....	621
References.....	622

17.1 Introduction

This chapter discusses a general design approach to planning computer experiments, which seeks design points that fill a bounded design region as uniformly as possible. Such designs are broadly referred to as *space-filling designs*.

The literature on the design for computer experiments has focused mainly on deterministic computer models; that is, running computer code with the same inputs always produces the same outputs (see Chapter 16). Because of this feature, the three fundamental design principles, *randomization*, *replication*, and *blocking*, are irrelevant in computer experiments. The true relationship between the inputs and the responses is unknown and often very complicated. To explore the relationship, one could use traditional regression models. But the most popular is Gaussian process models; see Chapter 16 for details. However, before data are collected, quite often little *a priori* or background knowledge is available

about which model would be appropriate, and designs for computer experiments should facilitate diverse modeling methods. For this purpose, a space-filling design is the best choice. The design region in which to make prediction may be unspecified at the data collection stage. Therefore, it is appropriate to use designs that represent all portions of the design region. When the primary goal of experiments is to make prediction at unsampled points, space-filling designs allow us to build a predictor with better average accuracy.

One most commonly used class of space-filling designs for computer experiments is that of Latin hypercube designs. Such designs, introduced by McKay et al. (1979), do not have repeated runs. Latin hypercube designs have one-dimensional uniformity in that, for each input variable, if its range is divided into the same number of equally spaced intervals as the number of observations, there is exactly one observation in each interval. However, a random Latin hypercube design may not be a good choice with respect to some optimality criteria such as maximin distance and orthogonality (discussed later). The maximin distance criterion, introduced by Johnson et al. (1990), maximizes the smallest distance between any two design points so that no two design points are too close. Therefore, a maximin distance design spreads out its points evenly over the entire design region. To further enhance the space-filling property for each individual input of a maximin distance design, Morris and Mitchell (1995) proposed the use of maximin Latin hypercube designs.

Many applications involve a large number of input variables. Finding space-filling designs with a limited number of design points that provide a good coverage of the entire high-dimensional input space is a very ambitious, if not hopeless, undertaking. A more reasonable approach is to construct designs that are space filling in the low-dimensional projections. Moon et al. (2011) constructed designs that are space filling in the two-dimensional projections and demonstrated empirically that such designs also perform well in terms of the maximin distance criterion in higher dimensions. Other designs that are space filling in the low-dimensional projections are randomized orthogonal arrays (OAs) (Owen 1992) and OA-based Latin hypercubes (Tang 1993). Another important approach is to construct orthogonal Latin hypercube designs. The basic idea of this approach is that orthogonality can be viewed as a stepping stone to constructing designs that are space filling in low-dimensional projections (Bingham et al. 2009).

Originating as popular tools in numerical analysis, low-discrepancy nets, low-discrepancy sequences, and uniform designs have also been well recognized as space-filling designs for computer experiments. These designs are chosen to achieve uniformity in the design space based on the discrepancy criteria such as the L_p discrepancy (see Section 17.3.2).

As an alternative to the use of space-filling designs, one could choose designs that perform well with respect to some model-dependent criteria such as the minimum integrated mean square error and the maximum entropy (Sacks et al. 1989; Shewry and Wynn 1987). One drawback of this approach is that such designs require the prior knowledge of the model. For instance, to be able to construct maximum entropy designs and integrated mean square error optimal designs, one would need the values of the parameters in the correlation function when a Gaussian process is used to model responses. One could also consider a Bayesian approach (Leatherman et al. 2014). A detailed account of model-dependent designs can be found in Santner et al. (2003), Fang et al. (2006), and the references therein.

This chapter is organized as follows. Section 17.2 gives a detailed review of Latin hypercube designs and discusses three important types of Latin hypercube designs (Latin hypercube designs based on measures of distance OA-based Latin hypercube designs; orthogonal and nearly orthogonal Latin hypercube designs). Section 17.3 describes other space-filling designs that are not Latin hypercube designs. Concluding remarks are provided in Section 17.4.

17.2 Latin Hypercube Designs

17.2.1 Introduction and Examples

A Latin hypercube of n runs for k factors is represented by an $n \times k$ matrix, each column of which is a permutation of n equally spaced levels. For convenience, the n levels are taken to be $-(n - 1)/2, -(n - 3)/2, \dots, (n - 3)/2, (n - 1)/2$. For example, design L in Table 17.1 is a Latin hypercube of 5 runs for 3 factors. Given an $n \times k$ Latin hypercube $L = (l_{ij})$, a Latin hypercube design D in the design space $[0, 1)^k$ can be generated, and the design matrix of D is an $n \times k$ matrix with the (i, j) th entry being

$$d_{ij} = \frac{l_{ij} + (n - 1)/2 + u_{ij}}{n}, \quad i = 1, \dots, n, j = 1, \dots, k, \tag{17.1}$$

where u_{ij} 's are independent random numbers from $[0, 1)$. If each u_{ij} in (17.1) is taken to be 0.5, the resulting design D is termed *lattice sample* due to Patterson (1954). For each factor, Latin hypercube designs have exactly one point in each of the n intervals $[0, 1/n), [1/n, 2/n), \dots, [(n - 1)/n, 1)$. This property is referred to as *one-dimensional uniformity*. For instance, design D in Table 17.1 is a Latin hypercube design based on the L in the table, and its pairwise plot in Figure 17.1 illustrates the one-dimensional uniformity. When the five points are projected onto each axis, there is exactly one point in each of the five equally spaced intervals.

The popularity of Latin hypercube designs was largely attributed to their theoretical justification for the variance reduction in numerical integration. Consider a function $y = f(\mathbf{x})$ where f is known, $\mathbf{x} = (x_1, \dots, x_k)$ has a uniform distribution in the unit hypercube $[0, 1)^k$, and $y \in R$. (More generally, when x_j follows a continuous distribution with a cumulative distribution function F_j , then the inputs of x_j can be selected via the quantile transformation $F_j^{-1}(u_j)$ where u_j follows a uniform distribution in $[0, 1)$.) The expectation of y ,

$$\mu = E(y), \tag{17.2}$$

is of interest. When the expectation μ cannot be computed explicitly or its derivation is unwieldy, one can resort to approximate methods. Let x_1, \dots, x_n be a sample of size n . One estimate of μ in (17.2) is

$$\hat{\mu} = \frac{1}{n} \sum_{i=1}^n f(x_i). \tag{17.3}$$

TABLE 17.1
5 × 3 Latin Hypercube L and a Latin Hypercube Design D Based on L

L			D		
2	0	-2	0.9253	0.5117	0.1610
1	-2	-1	0.7621	0.1117	0.3081
-2	2	0	0.1241	0.9878	0.4473
0	-1	2	0.5744	0.3719	0.8270
-1	1	1	0.3181	0.7514	0.6916




FIGURE 17.1
The pairwise plot of the Latin hypercube design D in Table 17.1 for the three factors $x_1, x_2,$ and x_3 .

The approach that takes x_1, \dots, x_n independently from the uniform distribution in $[0, 1)^k$ is *simple random sampling*. McKay et al. (1979) suggested an approach based on a Latin hypercube sample x_1, \dots, x_n . Denote the estimator $\hat{\mu}$ in (17.3) of μ under simple random sampling and Latin hypercube sampling by $\hat{\mu}_{srs}$ and $\hat{\mu}_{lhs}$, respectively. Note that $\hat{\mu}_{srs}$ and $\hat{\mu}_{lhs}$ have the same form, but $\hat{\mu}_{srs}$ uses a simple random sample and $\hat{\mu}_{lhs}$ a Latin hypercube sample. Both samples are denoted by x_1, \dots, x_n , for convenience. McKay et al. (1979) established the following theorem.

Theorem 17.1 *If $y = f(x)$ is monotonic in each of its input variables, then $Var(\hat{\mu}_{lhs}) \leq Var(\hat{\mu}_{srs})$.*

Theorem 17.1 says that when the monotonicity condition holds, Latin hypercube sampling yields a smaller variance of the sample mean than simple random sampling. Theorem 17.2 (Stein 1987) provides some insights into the two methods of sampling.

Theorem 17.2 We have that for $\mathbf{x} \in [0, 1]^k$,

$$\text{Var}(\hat{\mu}_{srs}) = \frac{1}{n} \text{Var}[f(\mathbf{x})]$$

and

$$\text{Var}(\hat{\mu}_{lhs}) = \frac{1}{n} \text{Var}[f(\mathbf{x})] - \frac{1}{n} \sum_{j=1}^k \text{Var}[f_j(x_j)] + o\left(\frac{1}{n}\right),$$

where x_j is the j th input of \mathbf{x} , $f_j(x_j) = E[f(\mathbf{x})|x_j] - \mu$, and $o(\cdot)$ is little o notation.

The term $f_j(x_j)$ in Theorem 17.2 is the main effect of the j th input variable. Theorem 17.2 tells us that the variance of the sample mean under Latin hypercube sampling is smaller than the counterpart under simple random sampling by an amount contributed by main effects. The extent of the variance reduction depends on the extent to which the function f is additive in the inputs. Asymptotic normality and a central limit theorem of Latin hypercube sampling were established in Stein (1987) and Owen (1992), respectively. A related approach is that of quasi-Monte Carlo methods, which select design points in a deterministic fashion (see Niederreiter 1992 and Chapter 19, Section 19.3.2).

A randomly generated Latin hypercube design does not necessarily perform well with respect to criteria such as those of *space filling* or *orthogonality* alluded to in Section 17.1. For example, when projected onto two factors, design points in a random Latin hypercube design may roughly lie on the diagonal as in the plot of x_1 versus x_2 in Figure 17.1, leaving a large area in the design space unexplored. In this case, the corresponding two columns in the design matrix are highly correlated. Examples of Latin hypercube designs with desirable properties are maximin Latin hypercube designs, OA-based Latin hypercube designs, and orthogonal or nearly orthogonal Latin hypercube designs; these will be discussed throughout the chapter.

17.2.2 Latin Hypercube Designs Based on Measures of Distance

To construct space-filling Latin hypercube designs, one natural approach is to make use of distance criteria. In what follows, we review several measures of distance.

Let $\mathbf{u} = (u_1, \dots, u_k)$ and $\mathbf{v} = (v_1, \dots, v_k)$ be two design points in the design space $\chi = [0, 1]^k$. For $t > 0$, define the interpoint distance between \mathbf{u} and \mathbf{v} to be

$$d(\mathbf{u}, \mathbf{v}) = \left(\sum_{j=1}^k |u_j - v_j|^t \right)^{1/t}. \tag{17.4}$$

When $t = 1$ and $t = 2$, the measure in (17.4) becomes the rectangular and Euclidean distances, respectively. The *maximin distance* criterion seeks a design \mathbf{D} of n points in the design space χ that maximizes the smallest interpoint distance; that is, it maximizes

$$\min_{\substack{\mathbf{u}, \mathbf{v} \in \mathbf{D} \\ \mathbf{u} \neq \mathbf{v}}} d(\mathbf{u}, \mathbf{v}), \tag{17.5}$$

where $d(\mathbf{u}, \mathbf{v})$ is defined as in (17.4) for any given t . This criterion attempts to place the design points such that no two points are too close to each other.

A slightly different idea is to spread out the design points of a design D in such a way that every point in the design space χ is close to some point in D . This is the *minimax distance* criterion that seeks a design D of n points in χ that minimizes the maximum distance between an arbitrary point $x \in \chi$ and the design D ; that is, it minimizes

$$\max_{x \in \chi} d(x, D),$$

where $d(x, D)$, representing the distance between x and the closest point in D , is defined as $d(x, D) = \min_{x_i \in D} d(x, x_i)$ and $d(x, x_i)$ is given in (17.4) for any given t .

Audze and Eglais (1977) introduced a distance criterion similar in spirit to the maximin distance criterion by using

$$\sum_{1 \leq i < j \leq n} d(x_i, x_j)^{-2}, \quad (17.6)$$

where x_1, \dots, x_n are the design points. This criterion of minimizing (17.6) was used by Liefvendahl and Stocki (2006).

Moon et al. (2011) defined a two-dimensional maximin distance criterion. Let the interpoint distance between two design points $\mathbf{u} = (u_1, \dots, u_k)$ and $\mathbf{v} = (v_1, \dots, v_k)$ projected onto dimensions h and l be

$$d_{h,l}^{(2)}(\mathbf{u}, \mathbf{v}) = (|u_h - v_h|^t + |u_l - v_l|^t)^{1/t}, \quad t > 0.$$

Then the minimum interpoint distance of a design D over all two-dimensional subspaces is

$$d_{\min}^{(2)} = \min_{\substack{\mathbf{u}, \mathbf{v} \in D \\ \mathbf{u} \neq \mathbf{v}, h \neq l}} d_{h,l}^{(2)}(\mathbf{u}, \mathbf{v}). \quad (17.7)$$

The two-dimensional maximin distance criterion selects a design that maximizes $d_{\min}^{(2)}$ in (17.7). Moon et al. (2011) showed by examples that optimal Latin hypercube designs based on this criterion also perform well under the maximin distance criterion (17.5).

17.2.2.1 Maximin Latin Hypercube Designs

We now focus on maximin distance criterion. The Gaussian process model was introduced in Chapter 16 and can be written as

$$y(x) = \mu + Z(x), \quad (17.8)$$

where μ is the unknown but constant mean function and $Z(x)$ is a stationary Gaussian process with mean 0, variance σ^2 , and correlation function $R(\cdot|\theta)$.

A popular choice for the correlation function is the power exponential correlation:

$$R(\mathbf{h}|\theta) = \exp\left(-\theta \sum_{j=1}^k |h_j|^p\right), \quad 0 < p \leq 2,$$

where h_j is the j th element of \mathbf{h} . Johnson et al. (1990) showed that as the correlation parameter θ goes to infinity, a maximin design maximizes the determinant of the correlation matrix, where the correlation matrix refers to that of the outputs from running the computer model at the design points. That is, a maximin design is asymptotically D -optimal under the model in (17.8) as the correlations become weak. Thus, a maximin design is also asymptotically optimal with respect to the maximum entropy criterion (Shewry and Wynn 1987).

The problem of finding maximin designs is referred to as the maximum facility dispersion problem (Erkut 1990) in location theory. It is closely related to the sphere packing problem in the field of discrete and computational geometry (Melissen 1997; Conway et al. 1999). The two problems are, however, different as explained in Johnson et al. (1990).

An extended definition of a maximin design was given by Morris and Mitchell (1995). Define a distance list (d_1, \dots, d_m) and an index list (J_1, \dots, J_m) respectively in the following way. The distance list contains the distinct values of interpoint distances, sorted from the smallest to the largest, and J_i in the index list is the number of pairs of design points in the design separated by the distance d_i , $i = 1, \dots, m$. Note that $1 \leq m \leq \binom{n}{2}$. In Morris and Mitchell (1995), a design is called a maximin design if it sequentially maximizes d_i s and minimizes J_i s in the order $d_1, J_1, d_2, J_2, \dots, d_m, J_m$. They further introduced a computationally efficient scalar-value criterion

$$\Phi_q = \left(\sum_{i=1}^m \frac{J_i}{d_i^q} \right)^{1/q}, \quad (17.9)$$

where q is a positive integer. Minimizing Φ_q with a large q results in a maximin design. Values of q are chosen depending on the size of the design searched for, ranging from 5 for small designs to 20 for moderate-sized designs to 50 for large designs.

Maximin designs tend to place design points toward or on the boundary. For example, Figure 17.2 exhibits a maximin Euclidean distance design and a maximin rectangular distance design, both of seven points. Maximin designs are likely to have clumped projections onto one dimension. Thus, such designs may not possess desirable one-dimensional uniformity, which is guaranteed by Latin hypercube designs. To strike the balance, Morris and Mitchell (1995) examined maximin designs within Latin hypercube designs. Although this idea sounds simple, generating maximin Latin hypercube designs is a challenging task particularly for large designs. The primary approach for obtaining maximin Latin hypercube designs is using the algorithms summarized in Table 17.2. These algorithms search for maximin Latin hypercube designs that have u_{ij} in (17.1) being a constant. For example, Moon et al. (2011) used $u_{ij} = 0.5$, which corresponds to the midpoint of the interval $[(i-1)/n, i/n]$ for $i = 1, \dots, n$. For detailed descriptions of these algorithms, see the respective references.

Some available implementations of the algorithms in Table 17.2 include the Matlab code provided in Viana et al. (2010), the function *maximinLHS* in the R package *lhs* (Carnell 2009), and the function *lhsdesign* in the Matlab statistics toolbox. The function in R uses random




FIGURE 17.2 Maximin designs with $n = 7$ points on $[0, 1]^2$: (a) Euclidean distance and (b) rectangular distance.

TABLE 17.2
Algorithms for Generating Maximin Latin Hypercube Designs

Article	Algorithm	Criterion
Morris and Mitchell (1995)	Simulated annealing	ϕ_q^a
Ye et al. (2000)	Columnwise-pairwise	ϕ_q and entropy
Jin et al. (2005)	Enhanced stochastic evolutionary algorithm	ϕ_q , entropy, and L_2 discrepancy
Liefvendahl and Stocki (2006)	Columnwise-pairwise and genetic algorithm	Maximin and the Audze-Eglais function ^b
van Dam et al. (2007)	Branch-and-bound	Maximin with Euclidean distance
Forrester et al. (2008)	Evolutionary operation	ϕ_q
Grosso et al. (2009)	Iterated local search heuristics	ϕ_q
Viana et al. (2010)	Translational propagation	ϕ_q
Zhu et al. (2011)	Successive local enumeration	Maximin
Moon et al. (2011)	Smart swap algorithm	$d_{\min}^{(2)c}$, ϕ_q
Chen et al. (2013)	Particle swarm algorithm	ϕ_q

^a ϕ_q is given as in (17.9).
^b The Audze-Eglais function in (17.6).
^c $d_{\min}^{(2)}$ is given in (17.7).

u_{ij} s in (17.1), while the function in Matlab allows both random u_{ij} s and $u_{ij} = 0.5$. It should be noted, however, that these designs are approximate maximin Latin hypercube designs. No theoretical method is available to construct exact maximin Latin hypercube designs of flexible run sizes except Tang (1994) and van Dam et al. (2007). These two articles provided methods for constructing exact maximin Latin hypercube designs of certain run sizes and numbers of input variables. Tang (1994) constructed a Latin hypercube based on a single replicate full factorial design (see Chapter 1 and also Wu and Hamada 2011, Chapter 4)

and showed that the constructed Latin hypercube has the same rectangular distance as the single replicate full factorial design, where the rectangular distance of a design is given by (17.5) with $t = 1$. Van Dam et al. (2007) constructed two-dimensional maximin Latin hypercubes with the distance measures with $t = 1$ and $t = \infty$ in (17.4). For the Euclidean distance measure with $t = 2$, van Dam et al. (2007) used the branch-and-bound algorithm to find maximin Latin hypercube designs with $n \leq 70$.

17.2.3 Orthogonal Array–Based Latin Hypercube Designs

Tang (1993) proposed OA-based Latin hypercube designs, also known as U designs, which guarantee multidimensional space filling. Recall the definition of an s -level OA of n runs, k factors, and strength r , denoted by $OA(n, s^k, r)$ in Chapter 9. The s levels are taken to be $1, 2, \dots, s$ in this chapter. By the definition of OAs, a Latin hypercube of n runs for k factors is an $OA(n, n^k, 1)$.

The construction of OA-based Latin hypercubes in Tang (1993) works as follows. Let A be an $OA(n, s^k, r)$. For each column of A and $m = 1, \dots, s$, replace the n/s positions with entry m by a random permutation of $(m - 1)n/s + 1, (m - 1)n/s + 2, \dots, mn/s$. Denote the design after the aforementioned replacement procedure by A' . In our notation, an OA-based Latin hypercube is given by $L = A' - (n + 1)J/2$, where J is an $n \times k$ matrix of all 1s. An OA-based Latin hypercube design in the design space $[0, 1)^k$ can be generated via (17.1). In addition to the one-dimensional uniformity, an $OA(n, s^k, r)$ -based Latin hypercube has the r -dimensional projection property that when projected onto any r columns, it has exactly $\lambda = n/s^r$ points in each of the s^r cells \mathcal{P}^r where $\mathcal{P} = \{[0, 1/s], [1/s, 2/s), \dots, [1 - 1/s, 1)\}$. Example 17.1 illustrates this feature of an $OA(9, 3^4, 2)$ -based Latin hypercube.

Example 17.1

Table 17.3 displays an OA-based Latin hypercube L based on the $OA(9, 3^4, 2)$ in the table. Figure 17.3 depicts the pairwise plot of this Latin hypercube. In each subplot, there is exactly one point in each of the nine dot-dash line boxes.

A generalization of OA-based Latin hypercubes using asymmetrical OA (see Chapter 9) can be readily made. For example, Figure 17.4a displays a Latin hypercube design based on an asymmetrical OA of six runs for two factors with three levels in the

TABLE 17.3
 $OA(9, 3^4, 2)$ and a Corresponding OA-Based Latin Hypercube

$OA(9, 3^4, 2)$				L			
1	1	1	1	-2	-2	-4	-2
1	2	2	3	-4	0	1	2
1	3	3	2	-3	4	2	1
2	1	2	2	-1	-4	-1	-1
2	2	3	1	1	-1	4	-3
2	3	1	3	0	2	-3	4
3	1	3	3	3	-3	3	3
3	2	1	2	2	1	-2	0
3	3	2	1	4	3	0	-4




FIGURE 17.3

The pairwise plot of an OA-based Latin hypercube design based on the Latin hypercube in Table 17.3 for the four factors x_1, \dots, x_4 .

first factor and two levels in the second factor. Note that each of the six cells separated by dot-dash lines contains exactly one point. By contrast, in the six-point randomly generated Latin hypercube design displayed in Figure 17.4b, two out of six such cells do not contain any point.

OA have been used directly to carry out computer experiments (see, e.g., Joseph et al. 2008). Compared with OA, OA-based Latin hypercubes are more favorable for computer experiments. When projected onto lower dimensions, the design points in OAs often have replicates. This is undesirable at the early stages of experimentation when relatively few factors are likely to be important.

The construction of OA-based Latin hypercubes depends on the existence of OAs. For the existence results of OAs, see, for example, Hedayat et al. (1999) and Mukerjee and




FIGURE 17.4 Plots of six-point Latin hypercube designs: (a) a Latin hypercube design based on an asymmetrical OA and (b) a random Latin hypercube design.

Wu (2006). A library of OAs is freely available on the N.J.A. Sloane website and the MktEx macro using the software SAS (Kuhfeld 2009). It should be noted that for certain given run sizes and numbers of factors, OAs of different numbers of levels and different strengths may be available. For instance, an $OA(16, 4^5, 2)$, an $OA(16, 2^5, 4)$, and an $OA(16, 2^5, 2)$ all produce OA-based Latin hypercubes of 16 runs for 5 factors. However, OAs with more levels and/or higher strength are preferred because they lead to designs with better projection space-filling properties.

Given an OA, the construction of Tang (1993) can produce many OA-based Latin hypercubes. There arises the problem of choosing a preferable OA-based Latin hypercube. Leary et al. (2003) presented an algorithm for finding optimal OA-based Latin hypercubes that minimize (17.6) using the Euclidean distance between design points. The optimization was performed via the simulated annealing algorithm (Morris and Mitchell 1995) and the columnwise–pairwise algorithm (Li and Wu 1997).

Recall the problem of estimating the mean μ in (17.2) of a known function $y = f(\mathbf{x})$ using a design with n points $\mathbf{x}_1, \dots, \mathbf{x}_n$ in Section 17.2.1. Latin hypercube sampling stratifies all univariate margins simultaneously and thus achieves a variance reduction compared with simple random sampling, as quantified in Theorem 17.2. Theorem 17.3 from Tang (1993) shows that a further variance reduction is achieved by OA-based Latin hypercube sampling.

Theorem 17.3 Suppose that f is continuous on $[0, 1]^k$. Let $\hat{\mu}_{oahls}$ denote the $\hat{\mu}$ in (17.3) with a randomly selected OA-based Latin hypercube design of n points. Then we have that

$$\text{Var}(\hat{\mu}_{oahls}) = \frac{1}{n} \text{Var}[f(\mathbf{x})] - \frac{1}{n} \sum_{j=1}^k \text{Var}[f_j(x_j)] - \frac{1}{n} \sum_{i < j}^k \text{Var}[f_{ij}(x_i, x_j)] + o\left(\frac{1}{n}\right),$$

where x_j is the j th input of \mathbf{x} , $f_j(x_j) = E[f(\mathbf{x})|x_j] - \mu$, and $f_{ij}(x_i, x_j) = E[f(\mathbf{x})|x_i, x_j] - \mu - f_i(x_i) - f_j(x_j)$.

To better understand this result, we write

$$f(\mathbf{x}) = \mu + \sum_{j=1}^k f_j(x_j) + \sum_{i<j}^k f_{ij}(x_i, x_j) + r(\mathbf{x}),$$

where the terms on the right side of the equation are uncorrelated with each other. Thus, the variance decomposition of the function f is

$$\text{Var}[f(\mathbf{x})] = \sum_{j=1}^k \text{Var}[f_j(x_j)] + \sum_{i<j}^k \text{Var}[f_{ij}(x_i, x_j)] + \text{Var}[r(\mathbf{x})].$$

We see that Latin hypercube sampling achieves a variance reduction by removing the variances of the main effects $f_j(x_j)$ from $\text{Var}[f(\mathbf{x})]$, and OA-based Latin hypercube sampling further removes the variances of the interactions $f_{ij}(x_i, x_j)$.

We conclude this section by mentioning that randomized OAs proposed by Owen (1992) also enjoy good space-filling properties in the low-dimensional projections. Results similar to Theorem 17.3 are given in Owen (1992).

17.2.4 Orthogonal and Nearly Orthogonal Latin Hypercube Designs

This section discusses the properties and constructions of Latin hypercube designs that have zero or small columnwise correlations in all two-dimensional projections. Such designs are called orthogonal and nearly orthogonal Latin hypercube designs. Orthogonal Latin hypercube designs are directly useful in fitting data using main effect models because they allow uncorrelated estimates of linear main effects. Another rationale of obtaining orthogonal or nearly orthogonal Latin hypercube designs is that they may not be space filling, but space-filling designs should be orthogonal or nearly orthogonal. Thus, we can search for space-filling designs within the class of orthogonal or nearly orthogonal Latin hypercube designs. Other justifications are given in Iman and Conover (1982), Owen (1994), Tang (1998), Joseph and Hung (2008), among others.

Extensive research has been carried out on the construction of orthogonal or nearly orthogonal Latin hypercube designs. Ye (1998) initiated this line of research and constructed orthogonal Latin hypercubes with $n = 2^m$ or $2^m + 1$ runs and $k = 2m - 2$ factors where $m \geq 2$. Ye's construction was extended by Cioppa and Lucas (2007) to obtain more columns for given run sizes. Steinberg and Lin (2006) constructed orthogonal Latin hypercubes of the run sizes $n = 2^{2^m}$ by rotating groups of factors in two-level 2^{2^m} -run regular fractional factorial designs. This idea was generalized by Pang et al. (2009) who constructed orthogonal Latin hypercubes of p^{2^m} runs and up to $(p^{2^m} - 1)/(p - 1)$ factors by rotating groups of factors in p -level p^{2^m} -run regular factorial designs, where p is a prime. Lin (2008) obtained orthogonal Latin hypercube designs of small run sizes ($n \leq 20$) through an algorithm that adds columns sequentially to an existing design. More recently, various methods have been proposed to construct orthogonal Latin hypercubes of more flexible run sizes and with large factor-to-run-size ratios. Here we review constructions in Lin et al. (2009), Sun et al. (2009), and Lin et al. (2010). These methods are general and simple to implement and cover the results in Table 17.10. For other methods, see Georgiou (2009), Sun et al. (2010), and Yang and Liu (2012).

We first give some notation and background. A vector $\mathbf{a} = (a_1, \dots, a_n)$ is said to be balanced if its distinct values have equal frequency. For an $n_1 \times k_1$ matrix \mathbf{A} and an $n_2 \times k_2$ matrix \mathbf{B} , their Kronecker product $\mathbf{A} \otimes \mathbf{B}$ is the $(n_1 n_2) \times (k_1 k_2)$ matrix:

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} a_{11}\mathbf{B} & a_{12}\mathbf{B} & \dots & a_{1k_1}\mathbf{B} \\ a_{21}\mathbf{B} & a_{22}\mathbf{B} & \dots & a_{2k_1}\mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n_1 1}\mathbf{B} & a_{n_1 2}\mathbf{B} & \dots & a_{n_1 k_1}\mathbf{B} \end{bmatrix}$$

with $a_{ij}\mathbf{B}$ itself being an $n_2 \times k_2$ matrix. For an $n \times k$ design or matrix $\mathbf{D} = (d_{ij})$, define its correlation matrix to be a $k \times k$ matrix $\mathbf{R}(\mathbf{D}) = (r_{ij})$ with

$$r_{ij} = \frac{\sum_{m=1}^n (d_{mi} - \bar{d}_i)(d_{mj} - \bar{d}_j)}{\sqrt{\sum_m (d_{mi} - \bar{d}_i)^2 \sum_m (d_{mj} - \bar{d}_j)^2}} \tag{17.10}$$

representing the correlation between the i th and j th columns of \mathbf{D} , where $\bar{d}_i = n^{-1} \sum_{m=1}^n d_{mi}$ and $\bar{d}_j = n^{-1} \sum_{m=1}^n d_{mj}$. A design or matrix \mathbf{D} is column orthogonal if $\mathbf{R}(\mathbf{D})$ is an identity matrix. A design or matrix $\mathbf{D} = (d_{ij})$ is orthogonal if it is balanced and column orthogonal.

To assess near orthogonality of design \mathbf{D} , Bingham et al. (2009) introduced two measures, the maximum correlation $\rho_M(\mathbf{D}) = \max_{i < j} |r_{ij}|$ and the average squared correlation $\rho_{ave}^2(\mathbf{D}) = \sum_{i < j} r_{ij}^2 / [(k(k-1)/2)]$, where r_{ij} is defined as in (17.10). Smaller values of $\rho_M(\mathbf{D})$ and $\rho_{ave}^2(\mathbf{D})$ imply near orthogonality. Obviously, if $\rho_M(\mathbf{D}) = 0$ or $\rho_{ave}^2(\mathbf{D}) = 0$, then an orthogonal design is obtained. For a concise presentation, we use $\text{OLH}(n, k)$ to denote an orthogonal Latin hypercube of n runs for k factors. Lin et al. (2010) established the following theorem on the existence of orthogonal Latin hypercubes.

Theorem 17.4 *There exists an orthogonal Latin hypercube of $n \geq 4$ runs with more than one factor if and only if $n \neq 4m + 2$ where m is an integer.*

Theorem 17.4 says that a Latin hypercube of run size 2, 3, 6, 10, 14, ... cannot be orthogonal. For smaller run sizes, this can be readily verified by exhaustive computer search. When orthogonal Latin hypercubes of certain run sizes exist, we want to construct such designs with as many columns as possible. We review three general construction methods. To generate design points in the region $[0, 1]^k$ from a Latin hypercube, one can use (17.1) with $u_{ij} = 0.5$, which corresponds to the midpoints of the cells.

17.2.4.1 Construction Method Based on an Orthogonal Array and a Small Orthogonal Latin Hypercube

Lin et al. (2009) constructed a large orthogonal, or nearly orthogonal, Latin hypercube by coupling an OA of index unity with a small orthogonal, or nearly orthogonal, Latin hypercube. Let \mathbf{B} be an $n \times q$ Latin hypercube, where as in Section 17.2.1, the levels are $-(n-1)/2, -(n-3)/2, \dots, (n-3)/2, (n-1)/2$. Then the elements in every column

of \mathbf{B} add up to zero, while the sum of squares of these elements is $n(n^2 - 1)/12$. Thus, the correlation matrix whose elements are defined as in (17.10) is

$$\mathbf{R}(\mathbf{B}) = \left[\frac{1}{12}n(n^2 - 1) \right]^{-1} \mathbf{B}'\mathbf{B}. \quad (17.11)$$

Let \mathbf{A} be an $\text{OA}(n^2, n^{2f}, 2)$. The construction proposed by Lin et al. (2009) proceeds as follows.

Step I: Let b_{ij} be the (i, j) th element of \mathbf{B} introduced earlier. For $1 \leq j \leq q$, obtain an $n^2 \times (2f)$ matrix A_j from \mathbf{A} by replacing the symbols $1, 2, \dots, n$ in the latter by $b_{1j}, b_{2j}, \dots, b_{nj}$, respectively, and then partition A_j as $A_j = [A_{j1}, \dots, A_{jf}]$, where each of A_{j1}, \dots, A_{jf} has two columns.

Step II: For $1 \leq j \leq q$, obtain the $n^2 \times (2f)$ matrix $L_j = [A_{j1}\mathbf{V}, \dots, A_{jf}\mathbf{V}]$, where

$$\mathbf{V} = \begin{bmatrix} 1 & -n \\ n & 1 \end{bmatrix}.$$

Step III: Obtain the matrix $\mathbf{L} = [L_1, \dots, L_q]$ of order $N \times k$, where $N = n^2$ and $k = 2qf$.

For $q = 1$, this construction is equivalent to that in Pang et al. (2009). However, by Theorem 17.4, we have $q \geq 2$ when n is not equal to 3 or $4m + 2$ for any non-negative integer m . Thus, the earlier method provides orthogonal or nearly orthogonal Latin hypercubes with an appreciably larger number of factors as compared to the method in Pang et al. (2009). For example, Pang et al. (2009) obtained $\text{OLH}(25, 6)$, $\text{OLH}(49, 8)$, $\text{OLH}(81, 40)$, $\text{OLH}(121, 12)$, and $\text{OLH}(169, 14)$, while the aforementioned construction produces $\text{OLH}(25, 12)$, $\text{OLH}(49, 24)$, $\text{OLH}(81, 50)$, $\text{OLH}(121, 84)$, and $\text{OLH}(169, 84)$.

Theorem 17.5 shows how the correlation matrix of the large Latin hypercube \mathbf{L} depends on that of the small Latin hypercube \mathbf{B} .

Theorem 17.5 *For the matrix \mathbf{L} constructed from \mathbf{B} in the previous steps I, II, and III, we have the following:*

- (i) *The matrix \mathbf{L} is a Latin hypercube.*
- (ii) *The correlation matrix of \mathbf{L} is given by*

$$\mathbf{R}(\mathbf{L}) = \mathbf{R}(\mathbf{B}) \otimes \mathbf{I}_{2f},$$

where $\mathbf{R}(\mathbf{B})$, given in (17.11), is the correlation matrix of a Latin hypercube \mathbf{B} , \mathbf{I}_{2f} is the identity matrix of order $2f$, and \otimes denotes Kronecker product.

Corollary 17.1 *If \mathbf{B} is an orthogonal Latin hypercube, then so is \mathbf{L} . In general, the maximum correlation and average squared correlation of \mathbf{L} are given by*

$$\rho_M(\mathbf{L}) = \rho_M(\mathbf{B}) \quad \text{and} \quad \rho_{ave}^2(\mathbf{L}) = \frac{q-1}{2qf-1} \rho_{ave}^2(\mathbf{B}).$$

Corollary 17.1 reveals that the large Latin hypercube \mathbf{L} inherits the exact or near orthogonality of the small Latin hypercube \mathbf{B} . As a result, the effort for searching for large orthogonal or nearly orthogonal Latin hypercubes can be focused on finding small orthogonal or nearly orthogonal Latin hypercubes that are easier to obtain via some general efficient robust optimization algorithms such as simulated annealing and genetic algorithms, by minimizing ρ_{ave}^2 or ρ_M .

Example 17.2 illustrates the actual construction of some orthogonal Latin hypercubes using the method of Lin et al. (2009). Example 17.3 is devoted to the construction of a nearly orthogonal Latin hypercube.

Example 17.2

Let n be a prime or prime power for which an $OA(n^2, n^{n+1}, 2)$ exists (Hedayat et al. 1999). For instance, consider $n = 5, 7, 8, 9, 11$. Now if we take \mathbf{B} to be an $OLH(5, 2)$, an $OLH(7, 3)$, an $OLH(8, 4)$, an $OLH(9, 5)$, or an $OLH(11, 7)$, as displayed in Table 17.4 and take \mathbf{A} , respectively, to be an $OA(25, 5^6, 2)$, an $OA(49, 7^8, 2)$, an $OA(64, 8^8, 2)$, an $OA(81, 9^{10}, 2)$, or an $OA(121, 11^{12}, 2)$, then the construction described in this section provides an

TABLE 17.4

Orthogonal Latin Hypercubes $OLH(5, 2)$, $OLH(7, 3)$, $OLH(8, 4)$, $OLH(9, 5)$, and $OLH(11, 7)$

OLH(5, 2)		OLH(7, 3)			OLH(8, 4)				OLH(9, 5)				
1	-2	-3	3	2	0.5	-1.5	3.5	2.5	-4	-2	0	-3	3
2	1	-2	0	-3	1.5	0.5	2.5	-3.5	-3	4	2	1	-2
0	0	-1	-2	-1	2.5	-3.5	-1.5	-0.5	-2	-3	-4	-1	-3
-1	2	0	-3	1	3.5	2.5	-0.5	1.5	-1	3	-2	3	4
-2	-1	1	-1	3	-3.5	-2.5	0.5	-1.5	0	-4	4	4	0
		2	1	-2	-2.5	3.5	1.5	0.5	1	2	-1	0	-4
		3	2	0	-1.5	-0.5	-2.5	3.5	2	0	3	-2	-1
					-0.5	1.5	-3.5	-2.5	3	1	1	-4	2
									4	-1	-3	2	1

OLH(11, 7)						
-5	-4	-5	-5	-3	0	0
-4	2	-1	3	4	5	4
-3	-2	4	5	-4	-2	-1
-2	3	-3	4	1	-4	-2
-1	4	2	-4	3	2	-4
0	-5	5	-2	5	-3	2
1	5	3	-3	-5	-1	5
2	-1	1	1	-2	3	-5
3	0	0	-1	0	1	-3
4	1	-4	0	2	-5	1
5	-3	-2	2	-1	4	3

TABLE 17.5

Nearly Orthogonal Latin Hypercube with 13 Rows and 12 Columns

-6	-6	-5	-4	-5	-2	2	1	-3	-2	-1	-2
-5	5	3	-5	3	4	-6	0	-4	1	-3	-1
-4	2	-4	1	2	6	5	-5	6	0	1	1
-3	1	2	4	-6	1	-2	6	2	3	2	6
-2	-2	6	-3	6	-5	3	4	4	-3	3	0
-1	-5	4	6	1	-1	0	-4	0	6	-5	-3
0	6	0	3	-4	-6	-3	-3	3	-5	0	-4
1	0	-3	5	5	0	1	2	-5	-6	-4	5
2	-1	-6	0	4	-4	-5	-2	-1	5	6	2
3	4	1	2	-1	2	6	3	-6	2	5	-6
4	-4	5	-2	-3	3	-1	-6	-2	-4	4	3
5	3	-1	-6	-2	-3	4	-1	1	4	-6	4
6	-3	-2	-1	0	5	-4	5	5	-1	-2	-5

OLH(25, 12), an OLH(49, 24), an OLH(64, 32), an OLH(81, 50), or an OLH(121, 84), respectively.

Example 17.3

Through a computer search, Lin et al. (2009) found a nearly orthogonal Latin hypercube with 13 rows and 12 columns, as given in Table 17.5. This Latin hypercube has $\rho_{ave} = 0.0222$ and $\rho_M = 0.0495$. Together with an OA(13², 13¹⁴, 2), the aforementioned procedure provides a nearly orthogonal Latin hypercube of 169 runs and 168 factors with $\rho_{ave} = 0.0057$ and $\rho_M = 0.0495$, by Corollary 17.1.

Before ending this section, we comment on the projection space-filling property of Latin hypercubes built earlier using a Latin hypercube **B** and an OA **A**. Any pair of columns obtained using different columns of **A** retains the two-dimensional projection property of **A**. When projected to those pairs of columns associated with the same column of **A**, the design points form clusters and these clusters are spread out as the corresponding two columns of **B**.

17.2.4.2 Recursive Construction Method

Orthogonal Latin hypercubes allow uncorrelated estimates of main effects in a main effect regression model. Sun et al. (2009) extended the concept of orthogonal Latin hypercubes for second-order polynomial models.

For a design **D** with columns d_1, \dots, d_k , let \tilde{D} be the $n \times [k(k+1)/2]$ matrix whose columns consist of all possible products $d_i \odot d_j$, where \odot denotes the element-wise product of vectors d_i and d_j , $i = 1, \dots, k, j = 1, \dots, k$, and $i \leq j$. Define the correlation matrix between **D** and \tilde{D} to be

$$R(D, \tilde{D}) = \begin{pmatrix} r_{11} & r_{12} & \dots & r_{1q} \\ r_{21} & r_{22} & \dots & r_{2q} \\ \vdots & \vdots & \ddots & \vdots \\ r_{k1} & r_{k2} & \dots & r_{kq} \end{pmatrix}, \tag{17.12}$$

where $q = k(k + 1)/2$ and r_{ij} is the correlation between the i th column of D and the j th column of \tilde{D} . A second-order orthogonal Latin hypercube D has the following properties: (1) the correlation matrix $R(D)$ is an identity matrix, and (2) $R(D, \tilde{D})$ in (17.12) is a zero matrix.

Sun et al. (2009) proposed the following procedure for constructing second-order orthogonal Latin hypercubes of $2^{c+1} + 1$ runs in 2^c factors for any integer $c \geq 1$. Throughout this section, let X^* represent the matrix obtained by switching the signs in the top half of the matrix X with an even number of rows.

Step I: For $c = 1$, let

$$S_1 = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad \text{and} \quad T_1 = \begin{pmatrix} 1 & 2 \\ 2 & -1 \end{pmatrix}.$$

Step II: For an integer $c \geq 2$, define

$$S_c = \begin{pmatrix} S_{c-1} & -S_{c-1}^* \\ S_{c-1} & S_{c-1}^* \end{pmatrix} \quad \text{and} \quad T_c = \begin{pmatrix} T_{c-1} & -(T_{c-1}^* + 2^{c-1}S_{c-1}^*) \\ T_{c-1} + 2^{c-1}S_{c-1} & T_{c-1}^* \end{pmatrix}.$$

Step III: Obtain a $(2^{c+1} + 1) \times 2^c$ Latin hypercube L_c as

$$L_c = \begin{pmatrix} T_c \\ \mathbf{0}_{2^c} \\ -T_c \end{pmatrix},$$

where $\mathbf{0}_{2^c}$ denotes a zero row vector of length 2^c .

Example 17.4

A second-order orthogonal Latin hypercube of 17 runs for 8 factors constructed using the earlier procedure is given by

$$\begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 \\ 2 & -1 & -4 & 3 & 6 & -5 & -8 & 7 \\ 3 & 4 & -1 & -2 & -7 & -8 & 5 & 6 \\ 4 & -3 & 2 & -1 & -8 & 7 & -6 & 5 \\ 5 & 6 & 7 & 8 & -1 & -2 & -3 & -4 \\ 6 & -5 & -8 & 7 & -2 & 1 & 4 & -3 \\ 7 & 8 & -5 & -6 & 3 & 4 & -1 & -2 \\ 8 & -7 & 6 & -5 & 4 & -3 & 2 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & -2 & -3 & -4 & -5 & -6 & -7 & -8 \\ -2 & 1 & 4 & -3 & -6 & 5 & 8 & -7 \\ -3 & -4 & 1 & 2 & 7 & 8 & -5 & -6 \\ -4 & 3 & -2 & 1 & 8 & -7 & 6 & -5 \\ -5 & -6 & -7 & -8 & 1 & 2 & 3 & 4 \\ -6 & 5 & 8 & -7 & 2 & -1 & -4 & 3 \\ -7 & -8 & 5 & 6 & -3 & -4 & 1 & 2 \\ -8 & 7 & -6 & 5 & -4 & 3 & -2 & 1 \end{pmatrix}.$$

Sun et al. (2010) further constructed second-order orthogonal Latin hypercubes of 2^{c+1} runs in 2^c factors by modifying Step III in the earlier procedure. In Step III, let $H_c = T_c - S_c/2$ and obtain L_c as

$$L_c = \begin{pmatrix} H_c \\ -H_c \end{pmatrix}.$$

Then L_c is a second-order orthogonal Latin hypercube of 2^{c+1} runs in 2^c factors.

17.2.4.3 Construction Method Based on Small Orthogonal Designs and Small Orthogonal Latin Hypercubes

This section reviews the construction from Lin et al. (2010) for constructing orthogonal and nearly orthogonal Latin hypercubes. All the proofs can be found in Lin et al. (2010). Let $A = (a_{ij})$ be an $n_1 \times k_1$ matrix with entries $a_{ij} = \pm 1$, $B = (b_{ij})$ be an $n_2 \times k_2$ Latin hypercube, $E = (e_{ij})$ be an $n_1 \times k_1$ Latin hypercube, and $F = (f_{ij})$ be an $n_2 \times k_2$ matrix with entries $d_{ij} = \pm 1$. Lin et al. (2010) construct designs via

$$L = A \otimes B + n_2 E \otimes F. \tag{17.13}$$

The resulting design L in (17.13) has $n = n_1 n_2$ runs and $k = k_1 k_2$ factors and becomes an orthogonal Latin hypercube, if certain conditions on A, B, E , and F are met.

Theorem 17.6 *Design L in (17.13) is an orthogonal Latin hypercube if*

- (i) A and F are column-orthogonal matrices of ± 1 ,
- (ii) B and E are orthogonal Latin hypercubes,
- (iii) At least one of the two, $A'E = \mathbf{0}$ and $B'F = \mathbf{0}$, is true,
- (iv) At least one of the following two conditions is true:
 - (a) A and E satisfy that for any i , if p_1 and p_2 are such that $e_{p_1 i} = -e_{p_2 i}$, then $a_{p_1 i} = a_{p_2 i}$.
 - (b) B and F satisfy that for any j , if q_1 and q_2 are such that $b_{q_1 j} = -b_{q_2 j}$, then $f_{q_1 j} = f_{q_2 j}$.

Condition (iv) in Theorem 17.6 is needed for L to be a Latin hypercube. To make L orthogonal, conditions (i), (ii), and (iii) are necessary. Choices for A and F include Hadamard matrices and OAs with levels ± 1 (see Chapter 9). Because of the orthogonality of A and F , n_1 and n_2 must be equal to two or multiples of four (Dey and Mukerjee 1999, p. 33). Theorem 17.6 requires designs B and E to be orthogonal Latin hypercubes. All known orthogonal Latin hypercubes of run sizes that are two or multiples of four can be used. As a result, Theorem 17.6 can be used to construct a vast number of orthogonal Latin hypercubes of $n = 8k$ runs. Example 17.5 illustrates the use of Theorem 17.6.

Example 17.5

Consider constructing an orthogonal Latin hypercube of 32 runs. Let $A = (1, 1)'$, B be the 16×12 orthogonal Latin hypercube in Table 17.6, $E = (1/2, -1/2)'$, and F be a

TABLE 17.6

16 × 12 Orthogonal Latin Hypercube from Steinberg and Lin (2006)

$$B = \frac{1}{2} \begin{pmatrix} -15 & 5 & 9 & -3 & 7 & 11 & -11 & 7 & -9 & 3 & -15 & 5 \\ -13 & 1 & 1 & 13 & -7 & -11 & 11 & -7 & -1 & -13 & -13 & 1 \\ -11 & 7 & -7 & -11 & 13 & -1 & -1 & -13 & 9 & -3 & 15 & -5 \\ -9 & 3 & -15 & 5 & -13 & 1 & 1 & 13 & 1 & 13 & 13 & -1 \\ -7 & -11 & 11 & -7 & 11 & -7 & 7 & 11 & 5 & 15 & -3 & -9 \\ -5 & -15 & 3 & 9 & -11 & 7 & -7 & -11 & 13 & -1 & -1 & -13 \\ -3 & -9 & -5 & -15 & 1 & 13 & 13 & -1 & -5 & -15 & 3 & 9 \\ -1 & -13 & -13 & 1 & -1 & -13 & -13 & 1 & -13 & 1 & 1 & 13 \\ 1 & 13 & 13 & -1 & -9 & 3 & -15 & 5 & 11 & -7 & 7 & 11 \\ 3 & 9 & 5 & 15 & 9 & -3 & 15 & -5 & 3 & 9 & 5 & 15 \\ 5 & 15 & -3 & -9 & -3 & -9 & -5 & -15 & -11 & 7 & -7 & -11 \\ 7 & 11 & -11 & 7 & 3 & 9 & 5 & 15 & -3 & -9 & -5 & -15 \\ 9 & -3 & 15 & -5 & -5 & -15 & 3 & 9 & -7 & -11 & 11 & -7 \\ 11 & -7 & 7 & 11 & 5 & 15 & -3 & -9 & -15 & 5 & 9 & -3 \\ 13 & -1 & -1 & -13 & -15 & 5 & 9 & -3 & 7 & 11 & -11 & 7 \\ 15 & -5 & -9 & 3 & 15 & -5 & -9 & 3 & 15 & -5 & -9 & 3 \end{pmatrix}$$

matrix obtained by taking any 12 columns from a Hadamard matrix of order 16. By Theorem 17.6, L in (17.13) with the chosen A, B, E , and F constitutes a 32×12 orthogonal Latin hypercube.

When $n_1 = n_2$, a stronger result than Theorem 17.6, as given in Theorem 17.7, can be established. It provides orthogonal Latin hypercubes with more columns than those in Theorem 17.6.

Theorem 17.7 *If $n_1 = n_2$ and A, B, E , and F are chosen according to Theorem 17.6, then design (L, U) is an orthogonal Latin hypercube with $2k_1k_2$ factors, where L is as in Theorem 17.6 and $U = -n_1A \otimes B + E \otimes F$.*

Example 17.6

To construct orthogonal Latin hypercubes of 64 runs, let $n_1 = n_2 = 8$ and take

$$A = F = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix} \quad \text{and} \quad B = E = \frac{1}{2} \begin{pmatrix} 1 & -3 & 7 & 5 \\ 3 & 1 & 5 & -7 \\ 5 & -7 & -3 & -1 \\ 7 & 5 & -1 & 3 \\ -1 & 3 & -7 & -5 \\ -3 & -1 & -5 & 7 \\ -5 & 7 & 3 & 1 \\ -7 & -5 & 1 & -3 \end{pmatrix}$$

Then design (L, U) in Theorem 17.7 is a 64×32 orthogonal Latin hypercube.

Theorem 17.8 Suppose that an OLH(n, k) is available, where n is a multiple of 4 such that a Hadamard matrix of order n exists. Then the following orthogonal Latin hypercubes, an OLH($2n, k$), an OLH($4n, 2k$), an OLH($8n, 4k$), and an OLH($16n, 8k$), can all be constructed.

We give a sketch of the proof for Theorem 17.8. The proof provides the actual construction of these orthogonal Latin hypercubes. The theorem results from an application of Theorem 17.6 and the use of orthogonal designs in Table 17.7. Note that each of the four matrices in Table 17.7 can be written as $(X', -X)'$. From such an X , define S to be the matrix obtained by choosing $x_i = 1$ for all i s. Let $A = (S', S)'$. Further let E be an orthogonal Latin hypercube derived from a matrix in Table 17.7 by letting $x_i = (2i - 1)/2$ for $i = 1, \dots, n/2$. Now we choose B to be a given OLH(n, k) and F be the matrix obtained by taking any k columns from a Hadamard matrix order n . Such matrices A, B, E , and F meet conditions (i), (ii), (iii), and (iv) in Theorem 17.6, from which Theorem 17.8 follows.

Theorem 17.8 is a very powerful result. By repeatedly applying Theorem 17.8, one can obtain many infinite series of orthogonal Latin hypercubes. For example, starting with an OLH(12, 6), we can obtain an OLH(192, 48), which can be used in turn to construct an OLH(768, 96) and so on. For another example, an OLH(256, 248) in Steinberg and Lin (2006) can be used to construct an OLH(1024, 496), an OLH(4096, 1984), and so on.

Another result from Lin et al. (2010) shows how the method in (17.13) can be used to construct nearly orthogonal Latin hypercubes.

Theorem 17.9 Suppose that condition (iv) in Theorem 17.6 is satisfied so that design L in (17.13) is a Latin hypercube. If conditions (i) and (iii) in Theorem 17.6 hold, we then have that

TABLE 17.7
Four Orthogonal Designs

		n															
		2				4				8				16			
x_1	x_1	x_2	x_1	$-x_2$	x_4	x_3	x_1	$-x_2$	$-x_4$	$-x_3$	$-x_8$	x_7	x_5	x_6			
$-x_1$	x_2	$-x_1$	x_2	x_1	x_3	$-x_4$	x_2	x_1	$-x_3$	x_4	$-x_7$	$-x_8$	$-x_6$	x_5			
	$-x_1$	$-x_2$	x_3	$-x_4$	$-x_2$	$-x_1$	x_3	$-x_4$	x_2	x_1	$-x_6$	$-x_5$	x_7	$-x_8$			
	$-x_2$	x_1	x_4	x_3	$-x_1$	x_2	x_4	x_3	x_1	$-x_2$	$-x_5$	x_6	$-x_8$	$-x_7$			
			$-x_1$	x_2	$-x_4$	$-x_3$	x_5	$-x_6$	$-x_8$	x_7	x_4	x_3	$-x_1$	$-x_2$			
			$-x_2$	$-x_1$	$-x_3$	x_4	x_6	x_5	$-x_7$	$-x_8$	x_3	$-x_4$	x_2	$-x_1$			
			$-x_3$	x_4	x_2	x_1	x_7	$-x_8$	x_6	$-x_5$	x_2	$-x_1$	$-x_3$	x_4			
			$-x_4$	$-x_3$	x_1	$-x_2$	x_8	x_7	x_5	x_6	x_1	x_2	x_4	x_3			
							$-x_1$	x_2	x_4	x_3	x_8	$-x_7$	$-x_5$	$-x_6$			
							$-x_2$	$-x_1$	x_3	$-x_4$	x_7	x_8	x_6	$-x_5$			
							$-x_3$	x_4	$-x_2$	$-x_1$	x_6	x_5	$-x_7$	x_8			
							$-x_4$	$-x_3$	$-x_1$	x_2	x_5	$-x_6$	x_8	x_7			
							$-x_5$	x_6	x_8	$-x_7$	$-x_4$	$-x_3$	x_1	x_2			
							$-x_6$	$-x_5$	x_7	x_8	$-x_3$	x_4	$-x_2$	x_1			
							$-x_7$	x_8	$-x_6$	x_5	$-x_2$	x_1	x_3	$-x_4$			
							$-x_8$	$-x_7$	$-x_5$	$-x_6$	$-x_1$	$-x_2$	$-x_4$	$-x_3$			

- (i) $\rho_{ave}^2(\mathbf{L}) = w_1\rho_{ave}^2(\mathbf{B}) + w_2\rho_{ave}^2(\mathbf{E}),$
- (ii) $\rho_M(\mathbf{L}) = \max\{w_3\rho_M(\mathbf{B}), w_4\rho_M(\mathbf{E})\},$

where $w_1, w_2, w_3,$ and w_4 are given by $w_1 = (k_2 - 1)(n_2^2 - 1)^2 / [(k_1k_2 - 1)(n^2 - 1)^2], w_2 = n_2^4(k_1 - 1)(n_1^2 - 1)^2 / [(k_1k_2 - 1)(n^2 - 1)^2], w_3 = (n_2^2 - 1) / (n^2 - 1),$ and $w_4 = n_2^2(n_1^2 - 1) / (n^2 - 1).$

Theorem 17.9 says that if \mathbf{B} and \mathbf{E} are nearly orthogonal Latin hypercubes, the resulting Latin hypercube \mathbf{L} is also nearly orthogonal. An example, illustrating the use of this result, is given in the following.

Example 17.7

Let $\mathbf{A} = (1, 1)'$ and $\mathbf{E} = (1/2, -1/2)'$. Choose a 16×15 nearly orthogonal Latin hypercube $\mathbf{B} = \mathbf{B}_0/2$ where \mathbf{B}_0 is displayed in Table 17.8, and \mathbf{B} has $\rho_{ave}^2(\mathbf{B}) = 0.0003$ and $\rho_M(\mathbf{B}) = 0.0765$. Taking any 15 columns of a Hadamard matrix of order 16 to be \mathbf{F} and then applying (17.13), we obtain a Latin hypercube \mathbf{L} of 32 runs and 15 factors. As $\rho_{ave}^2(\mathbf{E}) = \rho_M(\mathbf{E}) = 0,$ we have $\rho_{ave}^2(\mathbf{L}) = (n_2^2 - 1)^2\rho_{ave}^2(\mathbf{B}) / (n^2 - 1)^2 = 0.00002$ and $\rho_M(\mathbf{L}) = (n_2^2 - 1)\rho_M(\mathbf{B}) / (n^2 - 1) = 0.0191.$

17.2.4.4 Existence of Orthogonal Latin Hypercubes

A problem, of at least theoretical importance, in the study of orthogonal Latin hypercubes is to determine the maximum number k^* of columns in an orthogonal Latin hypercube of a given run size n . Theorem 17.4 tells us that $k^* = 1$ if n is 3 or $n = 4m + 2$ for any nonnegative integer m and $k^* \geq 2$ otherwise. Lin et al. (2010) obtained a stronger result.

TABLE 17.8
Design Matrix of \mathbf{B}_0 in Example 17.7

-15	15	-13	13	-5	-13	5	3	-1	5	-7	5	-9	-9	5
-13	-15	-3	3	7	3	15	-11	13	-5	7	-13	-7	-3	-3
-11	-9	-5	-11	-15	13	-5	11	-9	9	9	3	-5	-1	-11
-9	-1	9	-15	-11	1	-1	-13	5	-1	-15	7	1	3	15
-7	1	-7	7	15	15	-13	9	-5	-13	-3	-1	-1	7	13
-5	13	11	-5	9	-7	-3	-9	-13	11	13	-9	-3	13	1
-3	-5	13	15	-9	-9	-11	1	7	-9	15	11	9	1	-1
-1	-11	3	-7	11	-15	13	15	-7	-3	-9	9	7	9	-5
1	3	-9	-3	-1	-5	-15	-1	11	3	-11	-15	15	5	-15
3	-3	15	11	3	9	1	-7	-15	1	-13	-3	3	-15	-9
5	9	7	-1	5	11	9	13	15	15	5	1	11	-7	9
7	7	-1	-13	13	-1	-7	-5	9	-7	3	15	-13	-11	-13
9	5	-11	-9	-7	-3	7	-3	-11	-15	11	-7	13	-13	7
11	11	5	5	-13	7	11	5	3	-11	-5	-5	-11	15	-7
13	-7	-15	9	1	5	3	-15	-3	13	1	13	5	11	3
15	-13	1	1	-3	-11	-9	7	1	7	-1	-11	-15	-5	11

Theorem 17.10 *The maximum number k^* of factors for an orthogonal Latin hypercube of $n = 16m + j$ runs has a lower bound given in the following:*

- (i) $k^* \geq 6$ for all $n = 16m + j$ where $m \geq 1$ and $j \neq 2, 6, 10, 14$.
- (ii) $k^* \geq 7$ for $n = 16m + 11$ where $m \geq 0$.
- (iii) $k^* \geq 12$ for $n = 16m, 16m + 1$ where $m \geq 2$.
- (iv) $k^* \geq 24$ for $n = 32m, 32m + 1$ where $m \geq 2$.
- (v) $k^* \geq 48$ for $n = 64m, 64m + 1$ where $m \geq 2$.

Theorem 17.10 provides a general lower bound on the maximum number k^* of factors for an orthogonal Latin hypercube of n runs. We now summarize the results on the best lower bound on the maximum number k^* in an OLH(n, k^*) from all existing approaches for $n \leq 256$. Table 17.9 lists the best lower bound on the maximum number k^* in an OLH(n, k^*) for $n \leq 24$. These values except the case $n = 16$ were obtained by Lin (2008) through an algorithm. For $n = 16$, Steinberg and Lin (2006) obtained an orthogonal Latin hypercube with 12 columns. Table 17.10 reports the best lower bound on the maximum number k^* in

TABLE 17.9

Best Lower Bound k on the Maximum Number k^* of Factors in OLH(n, k^*) for $n \leq 24$

n	4	5	7	8	9	11	12	13	15	16	17	19	20	21	23	24
k	2	2	3	4	5	7	6	6	6	12	6	6	6	6	6	6

TABLE 17.10

Best Lower Bound k on the Maximum Number k^* of Factors in OLH(n, k^*) for $n > 24$

n	k	Reference	n	k	Reference
25	12	Lin et al. (2009)	144	24	Lin et al. (2010)
32	16	Sun et al. (2009)	145	12	Lin et al. (2010)
33	16	Sun et al. (2009)	160	24	Lin et al. (2010)
48	12	Lin et al. (2010)	161	24	Lin et al. (2010)
49	24	Lin et al. (2009)	169	84	Lin et al. (2009)
64	32	Sun et al. (2009)	176	12	Lin et al. (2010)
65	32	Sun et al. (2009)	177	12	Lin et al. (2010)
80	12	Lin et al. (2010)	192	48	Lin et al. (2010)
81	50	Lin et al. (2009)	193	48	Lin et al. (2010)
96	24	Lin et al. (2010)	208	12	Lin et al. (2010)
97	24	Lin et al. (2010)	209	12	Lin et al. (2010)
112	12	Lin et al. (2010)	224	24	Lin et al. (2010)
113	12	Lin et al. (2010)	225	24	Lin et al. (2010)
121	84	Lin et al. (2009)	240	12	Lin et al. (2010)
128	64	Sun et al. (2009)	241	12	Lin et al. (2010)
129	64	Sun et al. (2009)	256	248	Steinberg and Lin (2006)

an OLH(n, k^*) for $24 < n \leq 256$ as well as the corresponding approach for achieving the best lower bound.

17.3 Other Space-Filling Designs

Section 17.2 discussed various Latin hypercube designs that are suitable for computer experiments. A Latin hypercube design does not have repeated runs and each of its factors has as many levels as the run size. Bingham et al. (2009) argued that it is absolutely unnecessary to have the same number of levels as the run size in many practical applications. They proposed the use of orthogonal and nearly orthogonal designs for computer experiments, where each factor is allowed to have repeated levels. This is a rich class of orthogonal designs, including two-level orthogonal designs and orthogonal Latin hypercubes as special cases. This section reviews the concept and constructions of orthogonal designs. We also review another class of orthogonal designs provided by Moon et al. (2011). Other classes of space-filling designs that do not fall under Latin hypercube designs are low-discrepancy sequences and uniform designs. Both types of designs originate from the field of numerical analysis and give rise to attractive space-filling designs. We provide a brief account of low-discrepancy sequences and review various measures of uniformity in uniform designs.

17.3.1 Orthogonal Designs with Many Levels

Consider designs of n runs for k factors each of s levels, where $2 \leq s \leq n$. For convenience, the s levels are chosen to be centered and equally spaced; one such choice is $-(s-1)/2, -(s-3)/2, \dots, (s-3)/2, (s-1)/2$. Such a design is denoted by $D(n, s^k)$ and can be represented by an $n \times k$ design matrix $D = (d_{ij})$ with entries from the earlier set of s levels. A Latin hypercube of n runs for k factors is a $D(n, s^k)$ with $n = s$.

Let $A = (a_{ij})$ be an $n_1 \times k_1$ matrix with entries $a_{ij} = \pm 1$ and D_0 be a $D(n_2, s^{k_2})$. Bingham et al. (2009) constructed the $(n_1 n_2) \times (k_1 k_2)$ design

$$D = A \otimes D_0. \quad (17.14)$$

If A is column orthogonal, then design D in (17.14) is orthogonal if and only if D_0 is orthogonal. This provides a powerful way to construct a rich class of orthogonal designs for computer experiments, as illustrated by Example 17.8.

Example 17.8

Let D_0 be the orthogonal Latin hypercube OLH(16, 12) constructed by Steinberg and Lin (2006). The construction method in (17.14) gives a series of orthogonal designs of $16m$ runs for $12m$ factors by letting A be a Hadamard matrix of order m , where m is an integer such that a Hadamard matrix of order m exists.

Higher order orthogonality and near orthogonality of D in (17.14) were also discussed in Bingham et al. (2009). They considered two generalizations of the method (17.14).

Let D_j be a $D(n_2, s^{k_2})$, for each $j = 1, \dots, k_1$. One generalization is

$$D = (a_{ij}D_j) = \begin{bmatrix} a_{11}D_1 & a_{12}D_2 & \dots & a_{1k_1}D_{k_1} \\ a_{21}D_1 & a_{22}D_2 & \dots & a_{2k_1}D_{k_1} \\ \vdots & \vdots & \ddots & \vdots \\ a_{n_11}D_1 & a_{n_12}D_2 & \dots & a_{n_1k_1}D_{k_1} \end{bmatrix}. \tag{17.15}$$

The following results study the orthogonality of design D in (17.15).

Theorem 17.11 *Let A be column orthogonal. We have that*

- (i) $\rho_M(D) = \max\{\rho_M(D_1), \dots, \rho_M(D_{k_1})\}$;
- (ii) $\rho_{ave}^2(D) = w[\rho_{ave}^2(D_1) + \dots + \rho_{ave}^2(D_{k_1})]/k_1$, where $w = (k_2 - 1)/(k_1k_2 - 1)$;
- (iii) D in (17.15) is orthogonal if and only if D_1, \dots, D_{k_1} are all orthogonal.

The generalization (17.15) constructs designs with improved projection properties (Bingham et al. 2009). The research on orthogonal designs was further pursued by Georgiou (2011) who proposed an alternative construction method and obtained many new designs.

Another class of orthogonal designs is Gram-Schmidt designs constructed by Moon et al. (2011). A Gram-Schmidt design for n observations and k inputs is generated from an $n \times k$ Latin hypercube design $D = (d_{ij}) = (d_1, \dots, d_k)$ as follows:

Step 1: Center the j th column of D to have mean zero:

$$v_j = d_j - \sum_{i=1}^n d_{ij}/n, \quad \text{for } j = 1, \dots, k.$$

Step 2: Apply the Gram-Schmidt algorithm to v_1, \dots, v_k from Step 1 to form orthogonal columns:

$$u_j = \begin{cases} v_1, & j = 1; \\ v_j - \sum_{i=1}^{j-1} \frac{u_i v_j}{\|u_i\|^2} u_i, & j = 2, \dots, k, \end{cases}$$

where $\|u_i\|$ represents l_2 norm of u_i .

Step 3: Scale u_j from Step 2 to the desired range and denote the resulting column by x_j . Set $X = (x_1, \dots, x_k)$.

Any two columns of design X constructed via the three steps earlier have zero correlation.

17.3.2 Low-Discrepancy Sequences and Uniform Designs

Many problems in various fields such as quantum physics and computational finance require calculating definite integrals of a function over a multi dimensional unit cube. It is very common that the function may be so complicated that the integral cannot be obtained analytically and precisely, which calls for numerical methods of approximating the integral.

Recall the numerical integration problem discussed in Section 17.2.1. The quantity $\hat{\mu}$ in (17.3) is used to approximate μ in (17.2). Respecting the common notations, we use s to denote the number of factors and $\chi = [0, 1]^s$ the design region in this section. Let $\mathcal{P} = (x_1, \dots, x_n)$ be a set of n points in χ . The bound of the integration error is given by Koksma–Hlawka inequality:

$$|\mu - \hat{\mu}| \leq V(f)D^*(\mathcal{P}), \quad (17.16)$$

where $V(f)$ is the variation of f in the sense of Hardy and Krause, and $D^*(\mathcal{P})$ is the star discrepancy of the n points \mathcal{P} (Weyl 1916) described in the following. Motivated by the fact that the Koksma–Hlawka bound in (17.16) is proportional to the star discrepancy of the points, different methods for generating points in χ with as small a star discrepancy as possible have been proposed. Such methods are referred to as quasi–Monte Carlo methods (Niederreiter 1992).

For each $\mathbf{x} = (x_1, \dots, x_s)$ in χ , let $J_{\mathbf{x}} = [0, \mathbf{x})$ denote the interval $[0, x_1) \times \dots \times [0, x_s)$, $N(\mathcal{P}, J_{\mathbf{x}})$ denote the number of points of \mathcal{P} falling in $J_{\mathbf{x}}$, and $\text{Vol}(J_{\mathbf{x}})$ be the volume of interval $J_{\mathbf{x}}$. The *star discrepancy* $D^*(\mathcal{P})$ of \mathcal{P} is defined by

$$D^*(\mathcal{P}) = \max_{\mathbf{x} \in \chi} \left| \frac{N(\mathcal{P}, J_{\mathbf{x}})}{n} - \text{Vol}(J_{\mathbf{x}}) \right|. \quad (17.17)$$

A sequence S of points in χ is called a *low-discrepancy sequence* if its first n points have

$$D^*(\mathcal{P}) = O(n^{-1}(\log n)^s),$$

where $O(\cdot)$ is big O notation. As a comparison, if the set \mathcal{P} is chosen by the Monte Carlo method, that is, x_1, \dots, x_n are independent random samples from the uniform distribution, then $D^*(\mathcal{P}) = O(n^{-1/2})$, which is considered too slow in many applications (Niederreiter 2012).

Construction of low-discrepancy sequences is a very active research area in the study of quasi–Monte Carlo methods. There are many constructions available; examples are the good lattice point method, the good point method, Halton sequences, Faure sequences, and (t, s) -sequences. For a comprehensive treatment of low-discrepancy sequences, see Niederreiter (1992). Here we provide a brief account of two popular and most widely studied methods, (t, s) -sequences and uniform designs.

17.3.2.1 (t, m, s) -Nets and (t, s) -Sequences

The definitions of (t, m, s) -nets and (t, s) -sequences require a concept of elementary intervals. An *elementary interval* in base b is an interval E in $[0, 1)^s$ of the form

$$E = \prod_{i=1}^s \left[\frac{a_i}{b^{d_i}}, \frac{a_i + 1}{b^{d_i}} \right) \quad (17.18)$$

with integers a_i and d_i satisfying $d_i \geq 0$ and $0 \leq a_i < b^{d_i}$. For $i = 1, \dots, s$, the i th axis of an elementary interval has length b^{-d_i} , and thus, an elementary interval has a volume $b^{-\sum_{i=1}^s d_i}$.

For integers $b \geq 2$ and $0 \leq t \leq m$, a (t, m, s) -net in base b is a set of b^m points in $[0, 1)^s$ such that every elementary interval in base b of volume b^{t-m} contains exactly b^t points. For given values of b, m , and s , a smaller value of t leads to a smaller elementary interval and thus a set of points with better uniformity. Consequently, a smaller value of t in (t, m, s) -nets in base b is preferred.

An infinite sequence of points $\{x_n\}$ in $[0, 1)^s$ is a (t, s) -sequence in base b if for all $k \geq 0$ and $m > t$, the finite sequence $x_{kb^m+1}, \dots, x_{(k+1)b^m}$ forms a (t, m, s) -net in base b . Example 17.9 illustrates both concepts.

Example 17.9

Consider a $(0, 2)$ -sequence in base 2. Its first 8 points form a $(0, 3, 2)$ -net in base 2 and are displayed in Figure 17.5 with $t = 0, m = 3, s = 2$. There are four types of elementary intervals in base 2 of volume 2^{-3} with (d_1, d_2) 's in (17.18) being $(0, 3), (3, 0), (1, 2)$, and $(2, 1)$. Figure 17.5a through d shows a $(0, 3, 2)$ -net in base 2 when elementary intervals are given by $(d_1, d_2) = (0, 3), (d_1, d_2) = (3, 0), (d_1, d_2) = (1, 2)$, and $(d_1, d_2) = (2, 1)$, respectively. Note that in every elementary interval of the form

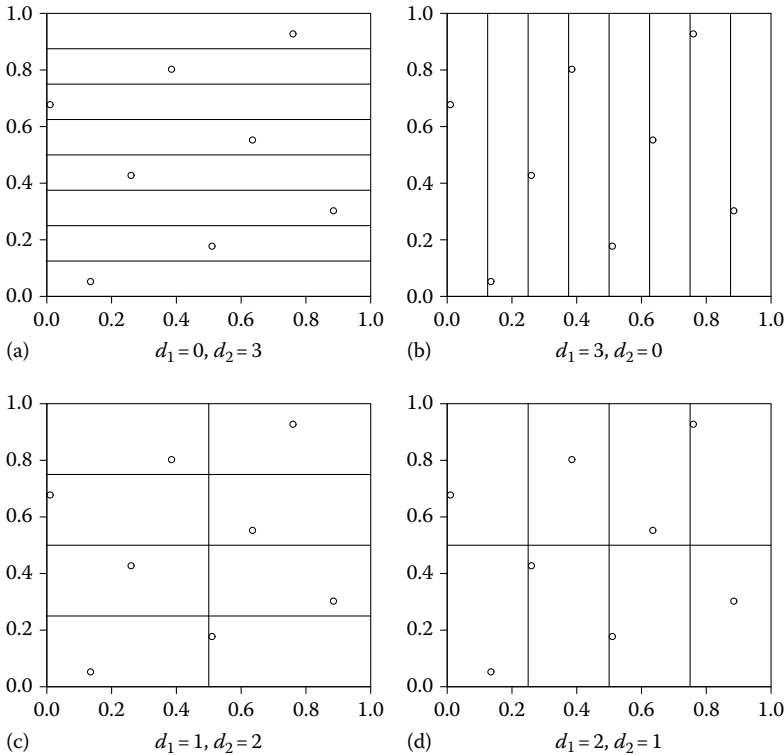


FIGURE 17.5 A $(0, 3, 2)$ -net in base 2 seen using four types of elementary intervals.

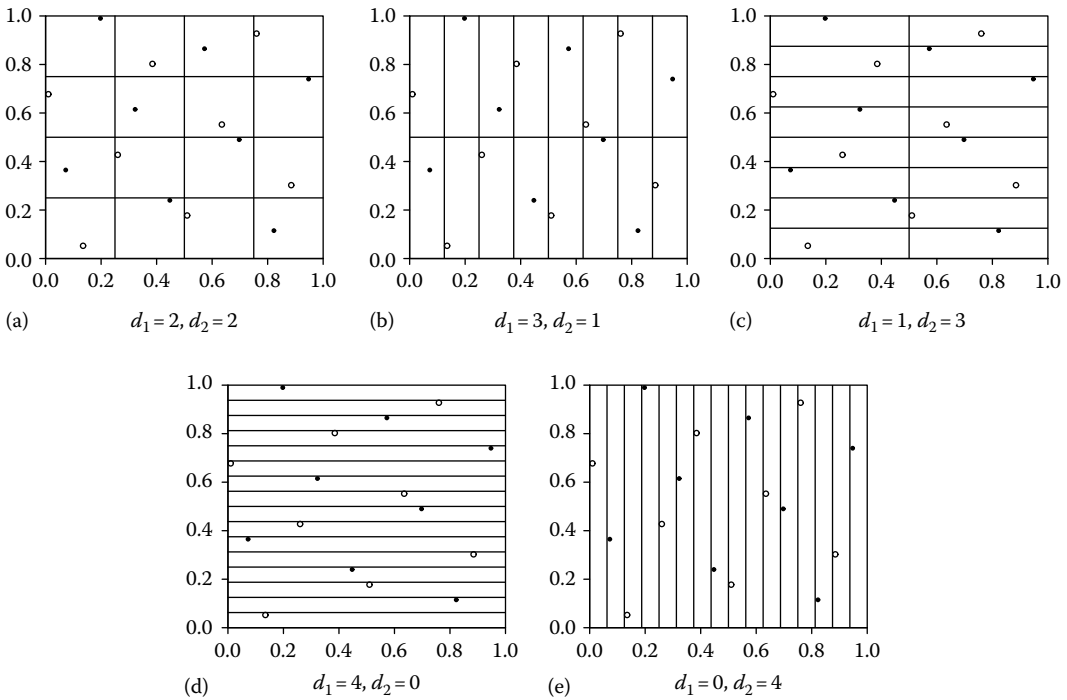


FIGURE 17.6 A $(0,4,2)$ -net in base 2 seen using (a–e) five types of elementary intervals; the first and second 8 points are represented by \circ and \bullet .

$$\left[\frac{a_1}{2^{d_1}}, \frac{(a_1 + 1)}{2^{d_1}} \right) \times \left[\frac{a_2}{2^{d_2}}, \frac{(a_2 + 1)}{2^{d_2}} \right), \quad 0 \leq a_i < 2^{d_i}, \quad i = 1, 2,$$

there is exactly one point. The next 8 points in this $(0,2)$ -sequence in base 2 also form a $(0,3,2)$ -net in base 2. The totality of all 16 points is a $(0,4,2)$ -net in base 2. Analogous to Figure 17.5, Figure 17.6(a) through (e) exhibit the $(0,4,2)$ -net in base 2 when elementary intervals are given by all (d_1, d_2) s that satisfy $d_1 + d_2 = m = 4$.

A general theory of (t, m, s) -nets and (t, s) -sequences was developed by Niederreiter (1987). Some special cases of (t, s) -sequences are as follows. Sobol’ sequences (Sobol’ 1967) are (t, s) -sequences in base 2. Faure sequences (Faure 1982) are $(0, s)$ -sequences in base q where q is a prime with $s \leq q$. Niederreiter sequences (Niederreiter 1987) are $(0, s)$ sequences in base q where q is a prime or a prime power with $s \leq q$. Niederreiter–Xing sequences (Niederreiter and Xing 1996) are (t, s) -sequences in base q for some certain t where q is a prime or a prime power with $s > q$. For constructions of all these sequences, we refer the readers to Niederreiter (2008). Results on existing (t, s) -sequences are available in Schürer and Schmid (2010).

17.3.2.2 Uniform Designs

Motivated by the Koksma–Hlawka inequality in (17.16), Fang (1980) and Wang and Fang (1981) introduced uniform designs, and by their definition, a uniform design is a set of

design points with the smallest discrepancy among all possible designs of the same run size. One choice of discrepancy is the star discrepancy in (17.17). More generally, one can use the L_p discrepancy,

$$D_p(\mathcal{P}) = \left[\int_{\mathcal{X}} \left| \frac{N(\mathcal{P}, J_{\mathbf{x}})}{n} - \text{Vol}(J_{\mathbf{x}}) \right|^p d\mathbf{x} \right]^{1/p},$$

where $N(\mathcal{P}, J_{\mathbf{x}})$ and $\text{Vol}(J_{\mathbf{x}})$ are defined as in (17.17). Two special cases of the L_p discrepancy are the L_∞ discrepancy, which is the star discrepancy, and the L_2 discrepancy. While the L_∞ discrepancy is difficult to compute, the L_2 discrepancy is much easier to evaluate because of a simple formula given by Warnock (1972):

$$D_2(\mathcal{P}) = 2^{-s} - \frac{2^{1-s}}{n} \sum_{i=1}^n \prod_{l=1}^s (1 - x_{il}^2) + \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \prod_{l=1}^s [1 - \max(x_{il}, x_{jl})],$$

where x_{il} is the setting of the l th factor in the i th run, $i = 1, \dots, n$ and $l = 1, \dots, s$.

The L_p discrepancy aims to achieve uniformity in the s -dimensional design space. Designs with the smallest L_p discrepancy do not necessarily perform well in terms of projection uniformity in low dimensions. Hickernell (1998) proposed three new measures of uniformity: the symmetric L_2 discrepancy (SL_2), the centered L_2 discrepancy (CL_2), and the modified L_2 discrepancy (ML_2). They are all defined through

$$D_{mod}(\mathcal{P}) = \sum_{u \neq \emptyset} \int_{\mathcal{X}_u} \left| \frac{N(\mathcal{P}_u, J_{\mathbf{x}_u})}{n} - \text{Vol}(J_{\mathbf{x}_u}) \right|^2 du, \tag{17.19}$$

where \emptyset represents the empty set, u is a nonempty subset of the set $\{1, \dots, s\}$, $|u|$ denotes the cardinality of u , \mathcal{X}_u is the $|u|$ -dimensional unit cube involving the coordinates in u , \mathcal{P}_u is the projection of the set of points \mathcal{P} on \mathcal{X}_u , $J_{\mathbf{x}_u}$ is the projection of $J_{\mathbf{x}}$ on \mathcal{X}_u , $N(\mathcal{P}_u, J_{\mathbf{x}_u})$ denotes the number of points of \mathcal{P}_u falling in $J_{\mathbf{x}_u}$, and $\text{Vol}(J_{\mathbf{x}_u})$ represents the volume of $J_{\mathbf{x}_u}$. The symmetric L_2 discrepancy chooses $J_{\mathbf{x}}$ such that it is invariant if x_{il} is replaced by $1 - x_{il}$, $i = 1, \dots, n$ and $l = 1, \dots, s$, and it has the formula

$$(SL_2(\mathcal{P}))^2 = \left(\frac{4}{3}\right)^s - \frac{2}{n} \sum_{i=1}^n \prod_{l=1}^s (1 + 2x_{il} - 2x_{il}^2) + \frac{2^s}{n^2} \sum_{i=1}^n \sum_{j=1}^n \prod_{l=1}^s (1 - |x_{il} - x_{jl}|).$$

The centered L_2 discrepancy chooses $J_{\mathbf{x}}$ such that it is invariant under the reflections of \mathcal{P} around any hyperplane with the l th coordinate being 0.5. Let \mathcal{A}^s denote the set of 2^s vertices of the unit cube \mathcal{X} and $\mathbf{a} = (a_1, \dots, a_s) \in \mathcal{A}^s$ be the closest one to \mathbf{x} . The centered L_2 discrepancy takes $J_{\mathbf{x}}$ in (17.19) to be

$$\{\mathbf{d} = (d_1, \dots, d_s) \in \mathcal{X} \mid \min(a_j, x_j) \leq d_j < \max(a_j, x_j), j = 1, \dots, s\}.$$

TABLE 17.11

Uniform Designs $U(6; 3^2)$ and $U(6; 6^2)$

$U(6; 3^2)$		$U(6; 6^2)$	
1	1	1	3
2	2	2	5
3	3	3	1
1	3	4	6
2	1	5	2
3	2	6	4

The formula for the centered L_2 discrepancy is given by

$$\begin{aligned}
 (CL_2(\mathcal{P}))^2 &= \left(\frac{13}{12}\right)^2 - \frac{2}{n} \sum_{i=1}^n \prod_{l=1}^s \left(1 + \frac{1}{2}|x_{il} - 0.5| - \frac{1}{2}|x_{il} - 0.5|^2\right) \\
 &\quad + \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \prod_{l=1}^s \left(1 + \frac{1}{2}|x_{il} - 0.5| + \frac{1}{2}|x_{jl} - 0.5| - \frac{1}{2}|x_{il} - x_{jl}|\right).
 \end{aligned}$$

The modified L_2 discrepancy takes $J_{\mathbf{x}} = [\mathbf{0}, \mathbf{x})$ and has the formula

$$(ML_2(\mathcal{P}))^2 = \left(\frac{4}{3}\right)^s - \frac{2^{1-s}}{n} \sum_{i=1}^n \prod_{l=1}^s (3 - x_{il}^2) + \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \prod_{l=1}^s [2 - \max(x_{il}, x_{jl})].$$

For other discrepancy measures such as the wrap-around discrepancy, see Fang et al. (2006).

Finding uniform designs based on a discrepancy criterion is an optimization problem. However, searching uniform designs in the entire unit cube is computationally prohibitive for large designs. Instead, it is convenient to find uniform designs within a class of *U-type designs*. Suppose that each of the s factors in an experiment has q levels, $\{1, \dots, q\}$. A *U-type design*, denoted by $U(n; q^s)$, is an $n \times s$ matrix in which the q levels in each column appear equally often. Table 17.11 displays a $U(6; 3^2)$ and a $U(6; 6^2)$. For $q = n$, uniform *U-type designs* can be constructed by several methods such as the good lattice method, the Latin square method, the expanding OA method, and the cutting method (Fang et al. 2006). For general values of q , optimization algorithms have been considered, such as simulated annealing, genetic algorithm, and threshold accepting (Bohachevsky et al. 1986; Winker and Fang 1997). For more detailed discussions on the theory and applications of uniform designs, see Fang et al. (2000) and Fang and Lin (2003).

17.4 Concluding Remarks

We have provided an expository account of the constructions and properties of space-filling designs for computer experiments. Research in this area remains active and will continue to thrive. Recently, a number of new directions have been pursued. He and Tang (2013)

introduced strong OAs and associated Latin hypercubes, while Tang et al. (2012) studied uniform fractional factorial designs. Research has also been conducted to take advantage of many available results from other design areas such as factorial design theory, one such work being multilayer designs proposed by Ba and Joseph (2011). Another important direction is to develop methodology for the design regions in which input variables have dependency or constraints; see Draguljic et al. (2012) and Bowman and Woods (2013) for more details.

References

- Audze, P. and Eglais, V. (1977), New approach to planning out of experiments, *Problems of Dynamics and Strength*, 35, 104–107.
- Ba, S. and Joseph, V. R. (2011), Multi-layer designs for computer experiments, *Journal of the American Statistical Association*, 106, 1139–1149.
- Bingham, D., Sitter, R. R., and Tang, B. (2009), Orthogonal and nearly orthogonal designs for computer experiments, *Biometrika*, 96, 51–65.
- Bohachevsky, I. O., Johnson, M. E., and Stein, M. L. (1986), Generalized simulated annealing for function optimization, *Technometrics*, 28, 209–217.
- Bowman, V. E. and Woods, D. C. (2013), Weighted space-filling designs, *Journal of Simulation*, 7, 249–263.
- Carnell, R. (2009), lhs: Latin hypercube samples, R package version 0.5. <http://CRAN.R-project.org/package=lhs>.
- Chen, R., Hsieh, D., Hung, Y., and Wang, W. (2013), Optimizing Latin hypercube designs by particle swarm, *Statistics and Computing*, 23, 664–676.
- Cioppa, T. M. and Lucas, T. W. (2007), Efficient nearly orthogonal and space-filling Latin hypercubes, *Technometrics*, 49, 45–55.
- Conway, J. H., Sloane, N. J. A., and Bannai, E. (1999), *Sphere Packings, Lattices, and Groups*, Vol. 290, New York: Springer Verlag.
- Dey, A. and Mukerjee, R. (1999), *Fractional Factorial Plans*, New York: Wiley.
- Draguljic, D., Dean, A. M., and Santner, T. J. (2012), Noncollapsing space-filling designs for bounded nonrectangular regions, *Technometrics*, 54, 169–178.
- Erkut, E. (1990), The discrete p-dispersion problem, *European Journal of Operational Research*, 46, 48–60.
- Fang, K. T. (1980), The uniform design: Application of number-theoretic methods in experimental design, *Acta Mathematicae Applicatae Sinica*, 3, 363–372.
- Fang, K. T., Li, R., and Sudjianto, A. (2006), *Design and Modeling for Computer Experiments*, Vol. 6, Boca Raton, FL: Chapman & Hall/CRC Press.
- Fang, K. T. and Lin, D. K. J. (2003), Uniform experimental designs and their applications in industry, *Handbook of Statistics*, 22, 131–170.
- Fang, K. T., Lin, D. K. J., Winker, P., and Zhang, Y. (2000), Uniform design: Theory and application, *Technometrics*, 42, 237–248.
- Faure, H. (1982), Discrépance de suites associées a un système de numération (en dimension s), *Acta Arith*, 41, 337–351.
- Forrester, A., Sobester, A., and Keane, A. (2008), *Engineering Design via Surrogate Modelling*, Chichester, U.K.: Wiley.
- Georgiou, S. D. (2009), Orthogonal Latin hypercube designs from generalized orthogonal designs, *Journal of Statistical Planning and Inference*, 139, 1530–1540.
- Georgiou, S. D. (2011), Orthogonal designs for computer experiments, *Journal of Statistical Planning and Inference*, 141, 1519–1525.

- Grosso, A., Jamali, A., and Locatelli, M. (2009), Finding maximin Latin hypercube designs by iterated local search heuristics, *European Journal of Operational Research*, 197, 541–547.
- He, Y. and Tang, B. (2013), Strong orthogonal arrays and associated Latin hypercubes for computer experiments, *Biometrika*, 100, 254–260.
- Hedayat, A., Sloane, N. J. A., and Stufken, J. (1999), *Orthogonal Arrays: Theory and Applications*, New York: Springer Verlag.
- Hickernell, F. J. (1998), A generalized discrepancy and quadrature error bound, *Mathematics of Computation*, 67, 299–322.
- Holland, J. H. (1975), *Adaptation in Natural and Artificial Systems*, Ann Arbor, MI: University of Michigan Press.
- Iman, R. L. and Conover, W. J. (1982), Distribution-free approach to inducing rank correlation among input variables, *Communications in Statistics–Simulation and Computation*, 11, 311–334.
- Jin, R., Chen, W., and Sudjianto, A. (2005), An efficient algorithm for constructing optimal design of computer experiments, *Journal of Statistical Planning and Inference*, 134, 268–287.
- Johnson, M. E., Moore, L. M., and Ylvisaker, D. (1990), Minimax and maximin distance designs, *Journal of Statistical Planning and Inference*, 26, 131–148.
- Joseph, V. R. and Hung, Y. (2008), Orthogonal-maximin Latin hypercube designs, *Statistica Sinica*, 18, 171–186.
- Joseph, V. R., Hung, Y., and Sudjianto, A. (2008), Blind kriging: A new method for developing metamodels, *Journal of Mechanical Design*, 130, 31102.
- Kuhfeld, W. F. (2009), Orthogonal arrays, Website courtesy of SAS Institute. Cary, NC.
- Leary, S., Bhaskar, A., and Keane, A. (2003), Optimal orthogonal-array-based Latin hypercubes, *Journal of Applied Statistics*, 30, 585–598.
- Leatherman, E., Dean, A., and Santner, T. (2014), Designs for computer experiments that minimize the weighted integrated mean square prediction error, Tech. Rep. 875, Columbus, OH: The Ohio State University.
- Li, W. W. and Wu, C. F. J. (1997), Columnwise-pairwise algorithms with applications to the construction of supersaturated designs, *Technometrics*, 39, 171–179.
- Liefvendahl, M. and Stocki, R. (2006), A study on algorithms for optimization of Latin hypercubes, *Journal of Statistical Planning and Inference*, 136, 3231–3247.
- Lin, C. D. (2008), New development in designs for computer experiments and physical experiments, PhD thesis, Vancouver, British Columbia, Canada: Simon Fraser University.
- Lin, C. D., Bingham, D., Sitter, R. R., and Tang, B. (2010), A new and flexible method for constructing designs for computer experiments, *Annals of Statistics*, 38, 1460–1477.
- Lin, C. D., Mukerjee, R., and Tang, B. (2009), Construction of orthogonal and nearly orthogonal Latin hypercubes, *Biometrika*, 96, 243–247.
- McKay, M. D., Beckman, R. J., and Conover, W. J. (1979), A comparison of three methods for selecting values of input variables in the analysis of output from a computer code, *Technometrics*, 21, 239–245.
- Melissen, J. B. M. (1997), Packing and covering with circles, PhD thesis, Utrecht, the Netherlands: Utrecht University.
- Moon, H., Dean, A. M., and Santner, T. J. (2011), Algorithms for generating maximin orthogonal and Latin hypercube designs, *Journal of Statistical Theory and Practice*, 5, 81–98.
- Morris, M. D. and Mitchell, T. J. (1995), Exploratory designs for computer experiments, *Journal of Statistical Planning and Inference*, 43, 381–402.
- Mukerjee, R. and Wu, C. F. (2006), *A Modern Theory of Factorial Designs*, Springer Verlag, New York.
- Niederreiter, H. (1987), Point sets and sequences with small discrepancy, *Monatshefte für Mathematik*, 104, 273–337.
- Niederreiter, H. (1992), Random number generation and Quasi-Monte Carlo methods, *SIAM CBMS-NSF Regional Conference Series in Applied Mathematics*. Philadelphia, PA.
- Niederreiter, H. (2008), Nets, (t, s) -sequences, and codes, *Monte Carlo and Quasi-Monte Carlo Methods 2006*, 1, 83–100.

- Niederreiter, H. (2012), Low-discrepancy simulation, In *Handbook of Computational Finance*, Eds. Duan, J. C., Härdle, W., and Gentle, J. E. Springer: Berlin Heidelberg, pp. 703–729.
- Niederreiter, H. and Xing, C. (1996), Low-discrepancy sequences and global function fields with many rational places, *Finite Fields and Their Applications*, 2, 241–273.
- Owen, A. B. (1992), A central limit theorem for Latin hypercube sampling, *Journal of the Royal Statistical Society. Series B*, 54, 541–551.
- Owen, A. B. (1994), Controlling correlations in Latin hypercube samples, *Journal of the American Statistical Association*, 89, 1517–1522.
- Pang, F., Liu, M. Q., and Lin, D. K. J. (2009), A construction method for orthogonal Latin hypercube designs with prime power levels, *Statistica Sinica*, 19, 1721–1728.
- Patterson, H. D. (1954), The errors of lattice sampling, *Journal of the Royal Statistical Society. Series B*, 16, 140–149.
- Sacks, J., Welch, W. J., Mitchell, T. J., and Wynn, H. P. (1989), Design and analysis of computer experiments, *Statistical Science*, 4, 409–423.
- Santner, T. J., Williams, B. J., and Notz, W. (2003), *The Design and Analysis of Computer Experiments*, New York: Springer Verlag.
- Schürer, R. and Schmid, W. C. (2010), MinT-Architecture and applications of the (t, m, s) -net and OOA database, *Mathematics and Computers in Simulation*, 80, 1124–1132.
- Shewry, M. C. and Wynn, H. P. (1987), Maximum entropy sampling, *Journal of Applied Statistics*, 14, 165–170.
- Sobol', I. M. (1967), On the distribution of points in a cube and the approximate evaluation of integrals, *USSR Computational Mathematics and Mathematical Physics*, 7, 86–112.
- Stein, M. (1987), Large sample properties of simulations using Latin hypercube sampling, *Technometrics*, 29, 143–151.
- Steinberg, D. M. and Lin, D. K. J. (2006), A construction method for orthogonal Latin hypercube designs, *Biometrika*, 93, 279–288.
- Sun, F., Liu, M. Q., and Lin, D. K. J. (2009), Construction of orthogonal Latin hypercube designs, *Biometrika*, 96, 971–974.
- Sun, F., Liu, M. Q., and Lin, D. K. J. (2010), Construction of orthogonal Latin hypercube designs with flexible run sizes, *Journal of Statistical Planning and Inference*, 140, 3236–3242.
- Tang, B. (1993), Orthogonal array-based Latin hypercubes, *Journal of the American Statistical Association*, 88, 1392–1397.
- Tang, B. (1994), A theorem for selecting OA-based Latin hypercubes using a distance criterion, *Communications in Statistics—Theory and Methods*, 23, 2047–2058.
- Tang, B. (1998), Selecting Latin hypercubes using correlation criteria, *Statistica Sinica*, 8, 965–978.
- Tang, Y., Xu, H., and Lin, D. K. J. (2012), Uniform fractional factorial designs, *Annals of Statistics*, 40, 891–907.
- van Dam, E. R., Husslage, B., den Hertog, D., and Melissen, H. (2007), Maximin Latin hypercube designs in two dimensions, *Operations Research*, 55, 158–169.
- Viana, F. A. C., Venter, G., and Balabanov, V. (2010), An algorithm for fast optimal Latin hypercube design of experiments, *International Journal for Numerical Methods in Engineering*, 82, 135–156.
- Wang, Y. and Fang, K. T. (1981), A note on uniform distribution and experimental design, *KeXue TongBao*, 26, 485–489.
- Warnock, T. T. (1972), Computational investigations of low-discrepancy point sets, in Zaremba S.K. (Ed.), *Applications of Number Theory to Numerical Analysis*, Academic Press, New York, 106, 319–343.
- Weyl, H. (1916), Über die Gleichverteilung von Zahlen mod. eins, *Mathematische Annalen*, 77, 313–352.
- Winker, P. and Fang, K. T. (1997), Application of threshold accepting to the evaluation of the discrepancy of a set of points, *SIAM Journal on Numerical Analysis*, 34, 2028–2042.
- Wu, C. F. J. and Hamada, M. S. (2011), *Experiments: Planning, Analysis, and Optimization*, Wiley series in probability and statistics, Hoboken, NJ: John Wiley & Sons.
- Yang, J. and Liu, M. Q. (2012), Construction of orthogonal and nearly orthogonal Latin hypercube designs from orthogonal designs, *Statistica Sinica*, 22, 433–442.

- Ye, K. Q. (1998), Orthogonal column Latin hypercubes and their application in computer experiments, *Journal of the American Statistical Association*, 93, 1430–1439.
- Ye, K. Q., Li, W., and Sudjianto, A. (2000), Algorithmic construction of optimal symmetric Latin hypercube designs, *Journal of Statistical Planning and Inference*, 90, 145–159.
- Zhu, H., Liu, L., Long, T., and Peng, L. (2011), A novel algorithm of maximin Latin hypercube design using successive local enumeration, *Engineering Optimization*, 1, 1–14.

18

Design for Sensitivity Analysis

William Becker and Andrea Saltelli

CONTENTS

18.1	Introduction.....	627
18.1.1	Black-Box Perspective.....	628
18.1.2	Types of Problem.....	629
18.2	Variance-Based Sensitivity Indices.....	632
18.3	Monte Carlo Estimation of Sensitivity Indices.....	635
18.3.1	Input Distributions and Sampling.....	636
18.3.2	Steps for Estimating Sensitivity Indices.....	637
18.3.3	Example.....	640
18.4	FAST and the Random Balance Design.....	643
18.4.1	Sampling Designs and Transformation of Variables.....	643
18.4.2	Calculation of Sensitivity Index Estimates.....	645
18.4.3	Extended FAST.....	647
18.4.4	Random Balance Designs.....	648
18.5	Screening Designs.....	649
18.5.1	Winding Stairs Design.....	650
18.5.2	Measures.....	651
18.5.3	Radial Design.....	652
18.5.4	Example.....	655
18.5.5	Discussion.....	657
18.6	Emulators.....	658
18.6.1	High-Dimensional Model Representation.....	659
18.6.2	Gaussian Processes.....	662
18.6.3	Custom Sampling for Emulators.....	664
18.7	Scatter Plot Smoothing.....	665
18.8	Conclusions and Discussion.....	668
	References.....	669

18.1 Introduction

Sensitivity analysis is the study of how uncertainty in the output of a model can be apportioned to different sources of uncertainty in the model input (Saltelli et al. 2004). Sometimes the term is also used to indicate simply the quantification of the uncertainty in the model's prediction, although strictly speaking, this is the closely related discipline of *uncertainty analysis*. In general, sensitivity analysis is used to test the robustness of

model-based inference, that is, how much the results of the model depend on the assumptions made in its construction and in particular on the specification of model input values. In engineering and risk analysis, sensitivity analysis mostly involves an exploration of the multidimensional space of the input variables.

Sensitivity analysis may also take slightly different meanings dependent on the field: in econometrics, sensitivity analysis has been advocated first in the form of *extreme bounds analysis*, measuring the sensitivity of regressor coefficients to the omission or inclusion of other regressors in a regression model (Leamer 1985, 2010). A form of derivative-based sensitivity analysis is also used to check the sensitivity of regression models to misspecification (Magnus 2007). In engineering, *design sensitivity analysis* uses the gradient of the error function between a model output and experimental measurements to estimate unknown model parameters, such as the stiffness parameters in a structural model (Tortorelli and Michaleris 1994). A succinct review of sensitivity analysis methods for use in impact assessment, that is, in relation to models used for policy, is in Saltelli and D’Hombres (2010). In this chapter, however, the focus will be on sensitivity analysis in the context of uncertainty in the inputs and outputs of a model.

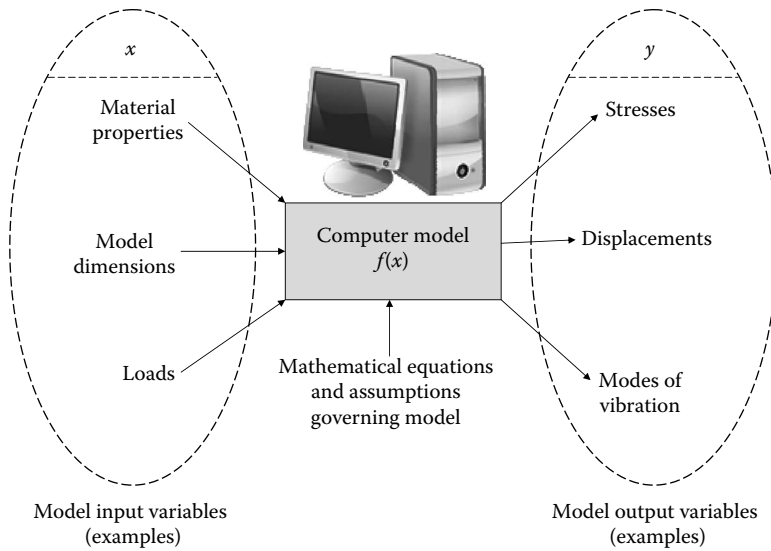
Very often, in chemistry, physics, biology and so on, one sees sensitivity analysis performed by changing one input at a time, the so-called one-at-a-time (OAT) design. This practice is not recommended because it only examines the sensitivity of model inputs at nominal (average) values and does not allow the possibility of exploring the model response at other values of input variables, which could produce very different results if the model is nonlinear (Saltelli and Annoni 2010). Instead, current best practice involves designs based on a multidimensional exploration of the space of the input variables, as in classic experimental design.

An important point to note from the start is that sensitivity analysis does not typically examine *structural uncertainty*, which is the uncertainty due to the model’s approximation to reality. As such, the results of a sensitivity analysis, which relate to uncertainty in the input variables, are conditional on the model. Structural uncertainty can often represent the largest source of uncertainty, so it is important to recognise that a sensitivity analysis is only half the story in a thorough analysis of uncertainty. Techniques for managing structural uncertainty will not be addressed here, but two approaches are noted with some references as a starting point: First, *model ensemble averaging* (Tebaldi and Knutti 2007; Rougier et al. 2013) which uses the results of a number of different but plausible models to approximate a *distribution over models* (an application to climate modelling can be found in Murphy et al. 2009) and second, an approach which considers the *discrepancy* between a single calibrated model and the true observed value (Kennedy and O’Hagan 2001; Strong et al. 2012). A discussion of model uncertainty in a wider context can be found in Saltelli and Funtowicz (2013).

18.1.1 Black-Box Perspective

In this chapter, the term *model* refers to a computer program which represents a mathematical construct built to simulate some physical, economic or other *real-world* process—examples could be models to predict climate change, engineering models to analyse the response of a component under loading or economic models forecasting the behaviour of markets.

Since any numerical model has quantifiable inputs and outputs, it is helpful in sensitivity analysis to consider it from the *black-box* perspective (see Figure 18.1). This views the model

**FIGURE 18.1**

Black-box view of an engineering computer model, with typical model inputs and outputs.

as a function $f(x)$ of k inputs, where $x = \{x_i\}_{i=1}^k$. The model will typically return a large number of output quantities, but in this chapter, for simplicity it shall be assumed that the output is a scalar y , such that $y = f(x)$. Note that although the x and y will often appear as random variables, they will always be expressed in lower case. Importantly, the models in this chapter will be assumed to be deterministic, such that f is fixed, and the *structure* of the model does not contain random components. The uncertainty in the output therefore is due uniquely to randomness in x .

Although the function (model) f is known in the sense that it represents a computer program based on mathematical equations, it will generally be complex enough as to be only accessible via simulation (i.e. not analytically tractable). Therefore in practice, all sensitivity analysis approaches involve sampling the inputs a number of times according to an experimental design, evaluating the model for each selected input vector, and estimating useful properties from the resulting outputs/data. With this in mind, it is useful to think of a particular set of model input values x as a *point* in a k -dimensional hyperspace, which is bounded by the maximum and minimum values of each input variable.

18.1.2 Types of Problem

There is no *one-size-fits-all* solution in sensitivity analysis, due to the fact that each problem has its own unique characteristics and challenges, such as a large number of model inputs, model nonlinearities, correlations or other relationships between model inputs and limitations in CPU time to name but a few. As such, there exist a great number of approaches that are designed for use under particular circumstances. Here a (non-exhaustive) taxonomy is attempted of some of the most common settings encountered in sensitivity analysis, and the tools that are available for each. One of the first defining features is whether or not the model is actually accessible to the analyst (the person performing the sensitivity analysis)—two main cases arise in this respect:

Case 1: The analyst can run the model. In other words, the model is available to the analyst, such that it can be evaluated at chosen input values and the corresponding results recorded. A design can be specified in this case where, for example, n model input points $\{x_1, x_2, \dots, x_n\}$ are selected in the k -dimensional input space, to obtain corresponding model outputs $\{y_1, y_2, \dots, y_n\}$. In this case, the sample of the input space is customarily generated without correlation among the input variables, although designs for correlated inputs are also available (Jacques et al. 2006; Xu and Gertner 2008; Li et al. 2010; Kucherenko et al. 2012). The output y could represent, for example, some modelled property of an engineering design such as an aeroplane wing or of a natural system such as groundwater flow through a geologic region.

Case 2: The sample points are given and the analyst can neither control their positioning nor generate additional points. Such data might come either from measurements or experiments or from a design that is not specifically intended for sensitivity analysis. The form of the model could be unknown, and the input variables could be correlated with one another in the sample. To give a simple example, y could be the Human Development Index computed over k countries and the x_i could be the indicators used in the construction of the index (Paruolo et al. 2013). In this case one cannot generate additional points/countries.

In Case 1 (when the design points can be specified), the best approach to performing a sensitivity analysis is determined by the cost of the model runs required to perform the analysis. In this context, *cost* refers to the total computational time required to evaluate the model at all the sample points, which is the product of the total number of model runs and the time required for each run. Since complex models can take minutes, hours or longer to evaluate for a single input point, it is not always feasible to sample a large number of input points (see, e.g. Becker et al. 2011; Batterbee et al. 2011). The strategies available for case I are as follows:

Case 1A: For *cheap* models (for which a single model evaluation will take a matter of seconds or less), a fully fledged quantitative sensitivity analysis can be performed using Monte Carlo estimators, estimating all k *first-order indices* and all k *total order indices* directly from model output values (see Section 18.3). This approach requires a large number of sample points (typically hundreds or thousands per input variable), but is preferred where possible since all sensitivity indices can be estimated with an accuracy related to the number of sample points. Furthermore, no assumptions are required about the functional form of the model (apart from that the model $f(x)$ is square-integrable, though this should not be a limitation in the vast majority of cases).

Case 1B: For expensive models, a design based on Fourier analysis can be used to compute all first-order indices at a cost which is weakly dependent on the number of input variables (see Section 18.4). The cost is of the order of some hundreds of model simulations. Alternatively, a space-filling design can be used in conjunction with an emulator (see Section 18.6 and Chapter 16). Although computationally cheaper, both of these approaches introduce a data-modelling problem which involves making assumptions about the functional form of the model, such as smoothness and continuity. Additionally, the use of certain types of emulators becomes rapidly infeasible as the number of sample points and the dimensionality of the problem increase.

Case 1C: In the case where the model is computationally expensive *and* one has many input variables, a set of methods known as *screening* can be applied to sort variables into

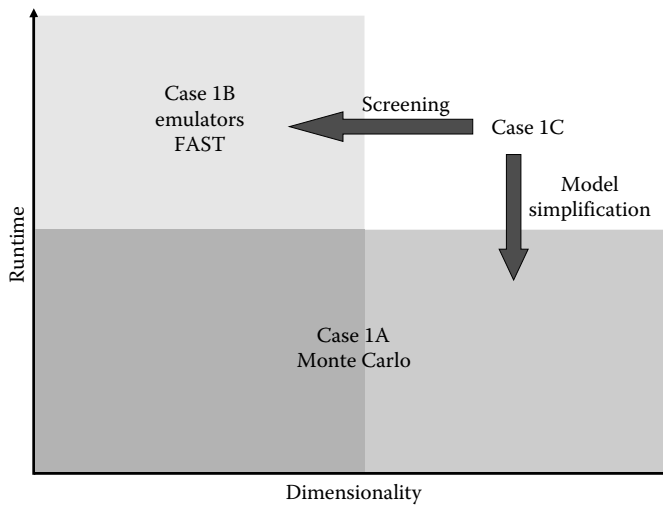


FIGURE 18.2

Case 1 sensitivity analysis problems based on dimensionality of model and run time.

influential and non-influential groups: this is known as *factor fixing* in some literature (Saltelli et al. 2008). A more detailed sensitivity analysis based on Monte Carlo or the use of emulators can then be applied on the set of influential variables, while non-influential variables are kept fixed (see Section 18.5).

A summary of Cases 1A–1C is shown in Figure 18.2: Monte Carlo methods scale well with dimensionality but need many runs per variable, so are not appropriate for expensive models. Emulators can deal with expensive models, but only for a limited number of input variables. In the case where the model has many input variables and is expensive to run, screening methods can help to reduce dimensionality to bring the problem in reach of emulators. Another possibility is to re-examine the model itself to see whether its runtime can be reduced, therefore allowing a Monte Carlo analysis. Although this is not a method of sensitivity analysis, it is an important consideration, since most models are not built with sensitivity analysis in mind and therefore are not necessarily optimised for speed. However, this requires access to the model and a deep understanding of it, which is not always possible.

For Case 2, when data are given, two approaches are considered in this chapter:

1. Use an *emulator* either to generate additional points and then perform a Monte Carlo sensitivity analysis or to directly estimate sensitivity indices from the available data (depending on the method—see Section 18.6).
2. Estimate directly the k first-order indices by univariate regression on the sorted model evaluations y_1, y_2, \dots, y_n (see Section 18.7). In effect, this involves making one-dimensional scatter plots of y against each x_i and then fitting (nonlinear) trend curves. In simple problems, even a visual inspection of scatter plots may be useful for qualitative analysis.

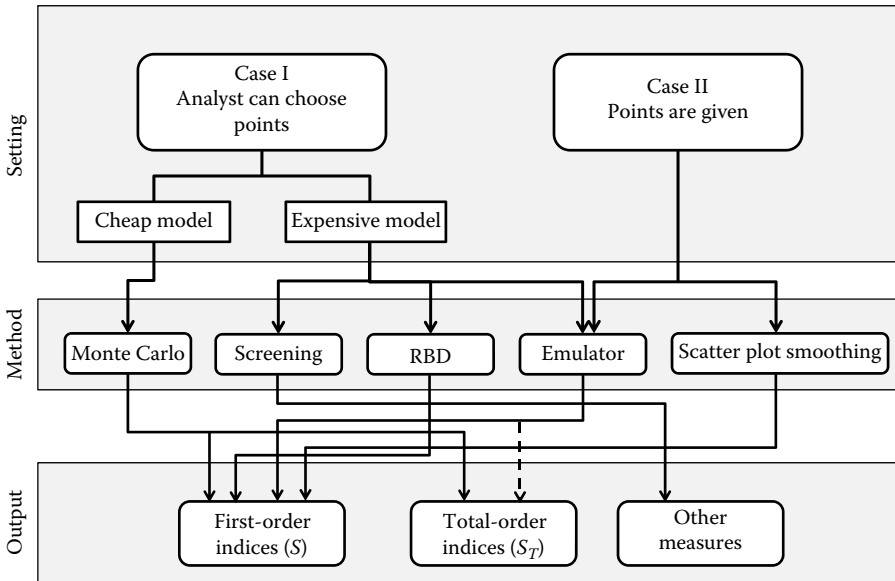


FIGURE 18.3 Various approaches to sensitivity analysis: when they can be used and what they produce. The dotted line here indicates that S_T can only be estimated with certain types of emulators, and the acronym *RBD* stands for *random balance designs*.

The various approaches discussed here and the context in which they can be applied are summarised in Figure 18.3. Note that in this chapter, only three measures of sensitivity are proposed:

1. First-order sensitivity index (see Section 18.2)
2. Total-order sensitivity index (see Section 18.2)
3. Elementary effects (see Section 18.5)

The following section gives a brief description of variance-based sensitivity analysis which underpins measures 1 and 2 earlier, after which in Section 18.3, Monte Carlo numerical procedures will be described for estimating them (this is the domain of Case 1A). In Section 18.4, an alternative way of estimating the same measures is described, using an approach based on the Fourier series. Section 18.5 deals with screening approaches (see Case 1C), while Section 18.6 explains the concepts of emulation and Section 18.7 scatter plot smoothing to deal with computationally expensive models (Case 1B), or the case where points are given (Case 2). Finally, some concluding remarks are given in Section 18.8.

18.2 Variance-Based Sensitivity Indices

Many measures of sensitivity have been proposed in the literature. For example, a well-known measure is to regress the data against each input variable x_i and take the coefficients of determination R_i^2 as measures of sensitivity. An obvious drawback of this is that

linear regression (linear in x) can only meaningfully interpret a linear model response (*model response* refers to the effect on the model output of changing its inputs). While this approach can be extended by more sophisticated forms of regression, it is preferable not to rely on any modelling of a functional relationship between y and x , since unwanted assumptions would thus be introduced. Variance-based approaches to measuring sensitivity, which consider the uncertainty in the model inputs and output from a probabilistic perspective, have become very popular in recent years, since they allow for highly nonlinear model responses, and account for variations in the output over the full input space.

In variance-based sensitivity analysis, it is required that the uncertainty in each model input x_1, x_2, \dots, x_k is characterised by known probability distributions $p_1(x_1), p_2(x_2), \dots, p_k(x_k)$ and furthermore that these distributions are independent of one another, such that $\prod_{i=1}^k p_i(x_i) = p(x)$. In fact this first step may often pose major challenges to the analyst because it is rarely the case that enough information exists to characterise the input distributions to a great degree of confidence, and correlations are not uncommon in many problems. This chapter will not however address these difficulties since the focus is on the statistical aspects of sensitivity analysis as commonly practiced, although some brief discussion of the practical aspects of sensitivity analysis is given at the end of Section 18.8.

Given known and independent input distributions then, a useful sensitivity measure for a given input variable x_i is

$$\text{Var}_{x_i}[E_{x_{\sim i}}(y \mid x_i)]. \quad (18.1)$$

The meaning of the inner E operator is the expected value of the model output y taken over all possible values of variables other than x_i (i.e. over $x_{\sim i}$) while keeping x_i fixed (the conditional mean). The outer Var is the variance taken over all possible values of x_i .

The associated normalised sensitivity measure, known as a *first-order sensitivity coefficient*, is defined as

$$S_i = \frac{\text{Var}_{x_i}[E_{x_{\sim i}}(y \mid x_i)]}{\text{Var}(y)}. \quad (18.2)$$

The measure S_i gives the fraction of model output variance which is caused by the input x_i alone, averaged over variations in all other input variables. Formula (18.2) has a long history, the foundations having been laid by Karl Pearson in the form of the *product-moment correlation coefficient* in 1895 (Pearson 1895) (a further discussion of this is given in Section 18.7). However, the use of partial variances (i.e. (18.2)) as measures of sensitivity in computer models appears to have been first proposed in 1973 in the context of the Fourier amplitude sensitivity test (FAST) (Cukier et al. 1973), which uses a Fourier series representation of the model (see Section 18.4). The idea was further developed by Ilya Sobol' in 1990 who also introduced an approach to estimate S_i by the Monte Carlo method, thereby bypassing the need to use a Fourier series approximation (Sobol' 1993) (see Section 18.3).

The numerator of S_i is in fact the first term in a variance decomposition whereby the unconditional model output variance $\text{Var}(y)$ is decomposed as the sum of a set

of conditional variances of first, second, . . . , up to the k th order (Sobol' 1993). Such a decomposition holds only if the input variables x_i are independent, in which case,

$$\text{Var}(y) = \sum_i V_i + \sum_i \sum_{j>i} V_{ij} + \cdots + V_{1,2,\dots,k}, \tag{18.3}$$

where $V_i = \text{Var}_{x_i}[E_{x_{\sim i}}(y|x_i)]$, $V_{ij} = \text{Var}_{x_i, x_j}[E_{x_{\sim ij}}(y|x_i, x_j)] - \text{Var}_{x_i}[E_{x_{\sim i}}(y|x_i)] - \text{Var}_{x_j}[E_{x_{\sim j}}(y|x_j)]$, and so on for the higher order terms. The terms in (18.3) derive from an analogous functional decomposition of $f(x)$ into orthogonal functions of increasing dimensionality:

$$f(x) = f_0 + \sum_i f_i(x_i) + \sum_i \sum_{j>i} f_{ij}(x_i, x_j) + \cdots + f_{1,2,\dots,k}(x_1, x_2, \dots, x_k) \tag{18.4}$$

where $f_0 = E(y)$, $f_i = E_{x_{\sim i}}(y|x_i) - f_0$, and $f_{ij} = E_{x_{\sim ij}}(y|x_{ij}) - f_i - f_j - f_0$. Taking the variance of (18.4) gives the variance decomposition in (18.3), noting that, for example, $\text{Var}_{x_i}(f_i) = \text{Var}_{x_i}[E_{x_{\sim i}}(y|x_i) - E(y)] = \text{Var}_{x_i}[E_{x_{\sim i}}(y|x_i)]$. A discussion of the importance of the first-order terms $f_i = E_{x_{\sim i}}(y|x_i) - E(y)$ is returned to in Section 18.7.

Dividing all terms in (18.3) by $\text{Var}(y)$ gives

$$\sum_i S_i + \sum_i \sum_{j>i} S_{ij} + \cdots + S_{1,2,\dots,k} = 1, \tag{18.5}$$

where the S_i are the first-order sensitivity coefficients defined in (18.2) and higher order terms are generalisations of these to multiple inputs. For example, S_{ij} measures the variance due to the interaction between x_i and x_j , additional to the variance caused by each input alone. A knowledge of all the sensitivity indices in (18.5) gives a detailed picture of how each input contributes to the uncertainty of the model output and the interactions between inputs in the model. Note that the case where $\sum_{i=1}^k S_i = 1$ is known as an *additive* model, in which there are no interactions between model inputs. This is however rarely the case in complex models, for which reason the calculation of the first-order indices alone is not usually sufficient.

In the ideal case then, one would like to know all sensitivity indices of all orders in (18.5). Due to computational limitations however, estimating all terms in (18.3) is often impractical for larger k given that they number $2^k - 1$ in total. For this reason, a measure known as the *total-order sensitivity index*, S_T , may be estimated, which measures the total effect of an input, including its first-order effect and interactions of any order (Homma and Saltelli 1996):

$$S_{Ti} = 1 - \frac{\text{Var}_{x_{\sim i}}[E_{x_i}(y | x_{\sim i})]}{\text{Var}(y)} = \frac{E_{x_{\sim i}}[\text{Var}_{x_i}(y | x_{\sim i})]}{\text{Var}(y)} \tag{18.6}$$

where $x_{\sim i}$ denotes the vector of all variables but x_i . In $E_{x_{\sim i}}[\text{Var}_{x_i}(y | x_{\sim i})]$, the inner variance, the scalar output of interest, is taken over all possible values of x_i while keeping $x_{\sim i}$ fixed, while the output expectation E is taken over all possible values $x_{\sim i}$ (Homma and Saltelli 1996).

One can see that $E_{x_i}(y | x_{\sim i})$ is the main effect of $x_{\sim i}$, and therefore $\text{Var}_{x_{\sim i}}[E_{x_i}(y | x_{\sim i})]$ is the variance caused by the main effects and interactions of all the variables and sets of variables *not* involving x_i . The remaining variance, $\text{Var}(y) - \text{Var}_{x_{\sim i}}[E_{x_i}(y | x_{\sim i})]$, is the

variance due to all terms in the decomposition (18.3) including x_i , that is, the variance of its main effect and all interactions of any order involving x_i , giving

$$S_{Ti} = S_i + \sum_{j>i} S_{i,j} + \sum_{l>j>i} S_{i,j,l} + \cdots + S_{1,2,\dots,k}. \quad (18.7)$$

As an example, consider a function of three input variables, $f(x_1, x_2, x_3)$. The standardised variance decomposition in (18.5) would in this case consist of

$$S_1 + S_2 + S_3 + S_{1,2} + S_{1,3} + S_{2,3} + S_{1,2,3} = 1. \quad (18.8)$$

In this case, the S_{Ti} can be expressed as the sum of any indices involving the index i :

$$\begin{aligned} S_{T1} &= S_1 + S_{1,2} + S_{1,3} + S_{1,2,3} \\ S_{T2} &= S_2 + S_{1,2} + S_{2,3} + S_{1,2,3} \\ S_{T3} &= S_3 + S_{1,3} + S_{2,3} + S_{1,2,3} \end{aligned} \quad (18.9)$$

from which one can note that in general $\sum_{i=1}^k S_{Ti} \neq 1$, unless all the interaction terms are zero. Observe also that the S_{Ti} could in fact be calculated from evaluating and summing all component sensitivity indices as in (18.9), but in practice, this is rarely done, since it involves a much higher computational effort—for this reason, the expression in (18.6) is usually used as the basis for estimation unless one is particularly interested in the precise nature of the interactions.

In the next section, the design and estimation procedures for the cases detailed in Section 24.1 are described.

18.3 Monte Carlo Estimation of Sensitivity Indices

Monte Carlo estimation of sensitivity indices is generally considered as the preferred approach to sensitivity analysis where possible, since it makes no assumptions about the functional form of the model (unlike emulators and FAST—see Sections 18.4 and 18.6). It is however only possible under the circumstances of Case 1, that is, when the analyst has full control over the placement of input points and possibly thousands of model runs can be executed without difficulty. Monte Carlo estimation involves sampling the model at a large number of points in the input space using random or *quasi-random* numbers as a basis. In this section, the use of quasi-random numbers is described, specifically the LP_τ sequences of Sobol' (Sobol' 1967, 1976) (also known simply as Sobol' sequences) coupled with a Monte Carlo design described in Section 18.3.2 (Saltelli 2002; Saltelli et al. 2010). In the following, therefore, the focus will be on the Sobol' sequence. However, the approaches described are also valid with random numbers and other low-discrepancy sequences—see Niederreiter (1992) for a summary of many common approaches.

18.3.1 Input Distributions and Sampling

It is assumed here that all random variables x_1, x_2, \dots, x_k are sampled uniformly in the k -dimensional unit hypercube \mathcal{X} :

$$x \in \mathcal{X} : \mathcal{X} = [0, 1]^k. \tag{18.10}$$

Different distributions can easily be generated by mapping the points in (18.10) onto the desired probability density function (uniform, normal, log-normal, etc.). This involves the use of the inverse cumulative distribution function of the variable of interest (also known as the *quantile function*), which allows uniformly distributed points in $[0, 1]$ to be transformed into points distributed as required (Saltelli et al. 2008). An example is shown in Figure 18.4 in which a set of equally spaced points are transformed into normally distributed points (equally spaced points are used here rather than random sampling to more clearly illustrate the transformation).

The Monte Carlo estimators presented in the following section rely on the use of random or quasi-random numbers—in particular, the approach recommended in this chapter is to use the Sobol’ sequence. The Sobol’ sequence and other quasi-random number sequences (also known as *low-discrepancy sequences*) are fixed sequences of numbers which are designed to fill hypercubes as uniformly as possible—in the context of sensitivity analysis, they can be used as a list of model input values that explore the model response with a high efficiency. Figure 18.5 shows a comparison of the Sobol’ sequence against random (strictly speaking, *pseudo-random*) numbers—observe the clusters and large *holes* in the random design compared to the relatively well-spaced points in the Sobol’ design.




FIGURE 18.4 Generating normally distributed points using equally spaced points and the inverse cumulative distribution of the normal distribution. Vertical and horizontal lines illustrate the mapping of sample points from one distribution to another.




FIGURE 18.5 128 points in two-dimensional space: (a) random numbers and (b) Sobol' sequence.

The use of quasi-random sequences is motivated by their good space-filling properties; these sequences outperform both pseudo-random Monte Carlo sampling and Latin hypercube sampling (LHS) in the estimation of multidimensional integrals (Sobol' and Kucherenko 2005)—for details on LHS, see Chapter 17. Recent extensive testing with a large variety of functions spanning different degrees of dimensionality, linearity and additivity has demonstrated their suitability for sensitivity analysis (Kucherenko et al. 2011). An additional desirable property of Sobol' sequences when compared to LHS is that with the former, additional points can be added sequentially to the analysis until a desired target accuracy is achieved (note that the points follow increasingly fine divisions of the input space—see the first four rows of Figure 18.6 for an example of the first four points in a six-dimensional Sobol' sequence). With LHS, the sample size cannot be extended once the analysis is performed, without starting again from the beginning, because the positioning of all points is dependent on the sample size. Sobol' sequences can be generated using freely available software both in FORTRAN and MATLAB (see European Commission 2012).

18.3.2 Steps for Estimating Sensitivity Indices

The steps needed to estimate a full set of first-order and total-order sensitivity indices via the Monte Carlo method are as follows (see Figure 18.6 for an illustration of the construction of the matrices):

1. Generate n points of a $2k$ -dimensional Sobol' sequence as in Figure 18.6, such that it is arranged in an $n \times 2k$ matrix with each row giving the coordinates of each point in the sequence. Call the first k -column submatrix A (i.e. the first k columns) and the remaining k -column submatrix B . The generic coordinates of A and B can be indicated, respectively, as $x_{ji}^{(a)}$ and $x_{ji}^{(b)}$, where the index i runs from 1 to k (the number of input variables) and the index j runs from 1 to n , the number of rows.*

* Note that in this chapter, although j is in general an index over samples, it is also sometimes used to index over input variables, for example, in the variance decomposition in (18.3). The meaning should however be clear given the context.




FIGURE 18.6

Construction of the A , B and A_B^i matrices, using the Sobol' LP_{τ} sequence with $k = 3$ and $N = 4$. Grey columns correspond to those taken from the matrix B .

2. Generate additional k matrices $A_B^i, i = 1, 2, \dots, k$, such that the i th matrix is entirely composed of coordinates from A except for its i th column, which is the i th column of B . A total of $k + 2$ sets of coordinates (matrices) have thus been generated.
3. Evaluate the computer model for each of the $n(k + 2)$ input vectors generated as the rows of each of the matrices A, B and $A_B^i, i = 1, 2, \dots, k$.
4. Compute the sample mean \hat{f}_0 of output associated with rows from both matrices of quasi-random points A and B combined, that is, using $\hat{f}_0 = \frac{1}{2n} \sum_{j=1}^n (f(A)_j + f(B)_j)$ where, for example, $f(A)_j$ indicates values of y computed from running the model f using the input values given by row j of matrix A . The unconditional sample variance is also calculated using the unbiased estimator $\widehat{\text{Var}}(y) = \frac{1}{2n-1} \sum_{j=1}^n [(f(A)_j - \hat{f}_0)^2 + (f(B)_j - \hat{f}_0)^2]$.
5. To estimate S_i (see (18.2)), one needs first to estimate $\text{Var}_{x_i}[E_{x_{\sim i}}(y | x_i)]$. Denoting $\text{Var}_{x_i}[E_{x_{\sim i}}(y | x_i)] = V_i$, model outputs associated with coordinates from A, B and A_B^i are used in the following estimator (Saltelli 2002; Sobol' et al. 2007; Saltelli et al. 2010):

$$\hat{V}_i = \frac{1}{n} \sum_{j=1}^n f(B)_j \left(f(A_B^i)_j - f(A)_j \right). \tag{18.11}$$

A rationale for estimator (18.11) is given below. S_i is estimated by dividing (18.11) by the sample variance $\widehat{\text{Var}}(y)$.

6. For S_{Ti} , one needs first to estimate $E_{x_{\sim i}}[\text{Var}_{x_i}(y | x_{\sim i})]$ (see (18.6)). Letting $E_{x_{\sim i}}[\text{Var}_{x_i}(y | x_{\sim i})] = V_{Ti}$, this can be estimated using model evaluations from the couple A and A_B^i (Jansen 1999):

$$\hat{V}_{T_i} = \frac{1}{2n} \sum_{j=1}^n \left(f(A)_j - f(A_B^i)_j \right)^2, \tag{18.12}$$

with a similar meaning of symbols as earlier. Again, S_{T_i} is estimated by dividing (18.12) by the sample variance $\widehat{\text{Var}}(y)$.

Note that each matrix A_B^i is used twice for estimating sensitivity indices associated with x_i , once to compute \hat{S}_i and once to compute \hat{S}_{T_i} . A derivation of estimators (18.11) and (18.12) can be found in Saltelli et al. (2008, 2010)—these designs are also called *substituted column sampling* (Morris et al. 2008). One can notice that the estimators make use of sums of products of model output values and that in each product the two function values being multiplied by one another have some symmetry. In the case of \hat{S}_i , the two function values $f(B)_j$ and $f(A_B^i)_j$ have identical values for coordinate x_i , whereas in the case of \hat{S}_{T_i} the two function values $f(A)_j$ and $f(A_B^i)_j$ have identical values for all coordinates except x_i . Take the case of \hat{S}_i for illustration: if x_i is influential, then the two function values being multiplied, $f(B)_j$ and $f(A_B^i)_j$, will be correlated, such that high values will tend be multiplied by high values and low values by low values. The resulting sum of these products will tend to be greater than the sum of the products of $f(B)_j$ and $f(A)_j$ (the two terms of which are uncorrelated), giving a value of \hat{S}_i greater than zero. In contrast, if x_i is non-influential, high and low values of $f(B)_j$ and $f(A_B^i)_j$ will be randomly coupled, resulting in an estimation of S_i which will tend to zero.

To see where the estimators (18.11) and (18.12) come from, refer back to (18.2) and (18.6). In the case of S_i , the numerator $\text{Var}_{x_i}[E_{x_{\sim i}}(y|x_i)]$ can be expressed as

$$\text{Var}_{x_i}[E_{x_{\sim i}}(y|x_i)] = \int E_{x_{\sim i}}^2(y|x_i) dx_i - \left(\int E_{x_{\sim i}}(y|x_i) dx_i \right)^2 \tag{18.13}$$

using the variance identity $\text{Var}(y) = E(y^2) - E^2(y)$. The second term in (18.13) reduces to $E^2(y)$ (since $E[E(y|x_i)] = E(Y)$) which is denoted as f_0^2 (refer back to (18.4)). The first term can be written as the following:

$$\begin{aligned} \int E_{x_{\sim i}}^2(y|x_i) dx_i &= \int E_{x_{\sim i}}(y|x_i) E_{x_{\sim i}}(y|x_i) dx_i \\ &= \int \left(\int \int f(x_{\sim i}, x_i) f(x'_{\sim i}, x_i) dx_{\sim i} dx'_{\sim i} \right) dx_i \\ &= \int \int f(x_{\sim i}, x_i) f(x'_{\sim i}, x_i) dx dx'_{\sim i}. \end{aligned} \tag{18.14}$$

Now one can see that the integral in (18.14) can be estimated by Monte Carlo integration using the first product $f(B)_j f(A_B^i)_j$ of (18.11).

Finally, it is worth briefly examining the experimental design generated in Steps 1 and 2 of this section. In the description given, matrices have been used to facilitate the programming of this procedure. However, one should note that any row of the sample matrix

A_j is simply a (quasi) random point in the input space, and the corresponding row $(A_B^i)_j$ describes a point which has the same coordinates, except for the i th input variable, which takes its value $x_{ji}^{(b)}$ from the j th row and i th column of B . Considering the set of points $\{A_j, (A_B^1)_j, (A_B^2)_j, \dots, (A_B^k)_j\}$, one can see that this subset of the design forms a *star* in the input space, with a centre point A_j and each subsidiary point $(A_B^i)_j$ a step away in the x_i direction. The example design given in Figure 18.6 is plotted for illustration in Figure 18.7. One can see that the design is nothing more than a number of OAT designs replicated at various locations in the input space. However, by performing multiple OAT experiments, one can begin to understand the *global* behaviour of the model—that is to say, the sensitivity of the model averaged over the full input space. Note that screening methods also use replicated OAT designs—(see Section 18.5). A shortcoming of the use of the Sobol’ sequence is in fact evident in Figure 18.7—one can see that the top left *star* is missing a step in the x_2 direction, and another has no steps at all. Going back to the design in Figure 18.6, the reason can be understood: the coordinate values in the Sobol’ sequence tend to repeat, which results in some instances where a coordinate value is substituted with the same number, resulting in $A_j = A_{B,j}^i$. These duplicates can however be accounted for, for example, by excluding them from the design when running the model (to avoid unnecessary runs), then adding the $f(A_B^i)_j = f(A_j)$ values where necessary. A further discussion on this point in the context of screening is given in Section 18.5.3.

18.3.3 Example

To show how the Monte Carlo estimators described earlier perform at different values of n , consider a simple polynomial example:




FIGURE 18.7 Sample design used for estimating sensitivity indices—example points as in Figure 18.6.

$$y = 3x_1^2 + 2x_1x_2 - 2x_3, \tag{18.15}$$

where the coefficients have been chosen quite arbitrarily. To illustrate the behaviour of the function with respect to its inputs, Figure 18.8 shows the scatter plots of y against x_1 , x_2 and x_3 , using random uniform sampling over $[0, 1]^3$. It is evident that x_1 has quite a strong, slightly nonlinear effect on y . Variable x_2 has apparently quite a weak effect (there is little discernable trend), whereas x_3 has a slight negative effect. These trends are clearly reflected in the coefficients of (18.15)—of course, normally one would not have the coefficients of an analytical equation to examine. The analytical values of S_i and S_{Ti} are given in columns 3 and 5 of Table 18.1.

To estimate the sensitivities of the variables, a Sobol' design is created in three dimensions, assuming uniform distributions for x_1 , x_2 and x_3 for simplicity, and estimators (18.11) and (18.12) are used. The only choice is what value of n , the number of sample points, to use. Given that the Sobol' sequence allows sequential addition of new points, one can start with a small number of points, then gradually increase until numerical convergence is observed. Figure 18.9 shows the convergence of these measures with n ranging from 8 to

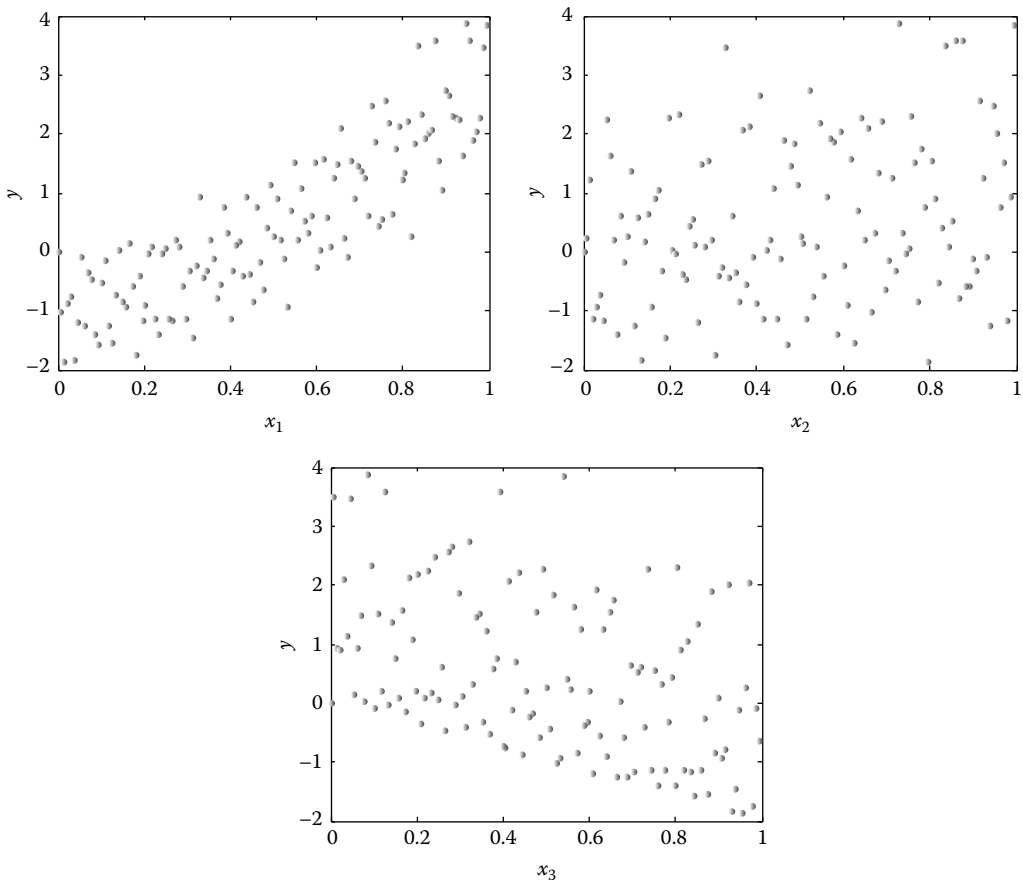


FIGURE 18.8 Scatter plots of the variables in the test equation (18.15).

TABLE 18.1

Monte Carlo Estimates and Analytical Values of S_i and S_{T_i} of Polynomial Function with $n = 128$

Variable	\hat{S}_i (MC)	S_i (Analytic)	\hat{S}_{T_i} (MC)	S_{T_i} (Analytic)
x_1	0.7517	0.7568	0.7781	0.7720
x_2	0.0503	0.0456	0.0604	0.0608
x_3	0.1870	0.1824	0.1829	0.1824




FIGURE 18.9

Convergence of the S_i and S_{T_i} of the polynomial equation (18.15) with increasing n . Lines represent, from top to bottom, x_1 , x_3 and x_2 , respectively.

1024. It is evident that the estimators converge quite quickly to an accurate estimate of the sensitivity indices; even at the lowest n , the variables are already correctly sorted, and at $n \geq 128$, the indices have converged to two decimal places. For most applications of sensitivity analysis, this would be sufficient accuracy. Table 18.1 shows the results at $n = 128$ compared to analytical values. Note that since there is a weak interaction between x_1 and x_2 , the S_{T_i} of these variables is slightly higher than their respective S_i values, due to the fact that both S_{T1} and S_{T2} additionally include the interaction effect $S_{1,2}$ (refer back to (18.8) and (18.9) to see why). The value of $S_{1,2}$ is not estimated here, though it can be deduced from the table, noticing that x_3 does not interact with any variables since $S_3 = S_{T3}$ (and hence $S_{1,3} = S_{2,3} = S_{1,2,3} = 0$), therefore $S_{1,2} = S_{T2} - S_2 = S_{T3} - S_3$.

Despite the flexibility of Monte Carlo estimators, one should remember that the cost is $n(k + 2)$ model runs (see again Figure 18.6)—that is, in the previous example the total number of model runs required was $128 \times 5 = 640$. While this is fine for fast models, for large models which are slower to run, it may be impractical. In the following sections, some alternative approaches are discussed that have lower computational requirements.

18.4 FAST and the Random Balance Design

The FAST (Cukier et al. 1973, 1978), which was actually proposed around 20 years before the Monte Carlo estimators described in the previous section, uses a transformation of the input variables of the function (model) to represent a multivariate function as a periodic function of a single *frequency variable*, s , which allows the function to be analysed using the Fourier series. The transformation of each input variable into s uses a unique characteristic sampling frequency which allows the contribution of each input to be assessed using the tools of Fourier analysis, which give analytical expressions for variance-based sensitivity indices (based on the Fourier series approximation). The advantage, compared to the Monte Carlo method, is that the integrals required to calculate the sensitivity indices, mean and variance (which are k -dimensional, see, e.g. (18.14)) can be expressed as univariate integrals with respect to s . Thus, a full set of S_i can be estimated from a single FAST sample, which means that the computational cost can be lower. However, the FAST approach relies on using the Fourier series to approximate the model output, which requires assumptions of smoothness, and furthermore uses truncated series to estimate sensitivity indices, which introduces estimation bias. A hybrid approach combines the concept of FAST with *random balance designs* (RBDs), a form of experimental design first proposed by Satterthwaite (1959)—this is described in Section 18.4.4. In the following, a description of the transformation functions is given in Section 18.4.1, followed by the estimation of sensitivity indices in Section 18.4.2. A recent overview of FAST literature can be found in Xu and Gertner (2011).

18.4.1 Sampling Designs and Transformation of Variables

In order to apply the tools of Fourier analysis to the model, each input variable x_i is transformed into a *periodic function* of a single variable, s , in the following way:

$$x_i = G(\sin(\omega_i s)), \quad i = 1, 2, \dots, k, \quad (18.16)$$

where s is a variable in $[-\pi, \pi]$, G is a specified transformation function, and ω_i is an integer. The effect of this transformation is that uniformly sampling s within its range (i.e. taking equally spaced values) results in oscillations in the corresponding sampled values of each x_i over its respective range. Figure 18.10 shows this effect on the following three transformation functions:

$$x_i = a_i \exp(b_i \sin(\omega_i s)) \quad (18.17)$$

$$x_i = a_i (1 + b_i \sin(\omega_i s)) \quad (18.18)$$

$$x_i = \frac{1}{2} + \frac{1}{\pi} \arcsin(\sin(\omega_i s)) \quad (18.19)$$

proposed, respectively, by Cukier et al. (1973), Koda et al. (1979) and Saltelli et al. (1999) as functions that are intended to approximate uniform sampling over the sample space \mathcal{X} . Note, however, that functions (18.17) and (18.18) give higher densities of points at the edges of the sample space; therefore, for truly uniform sampling, (18.19) is preferred. For non-uniform input distributions, other transformation functions would be necessary. The parameters a and b can be altered to give different ranges over x_i , for example, in Figure

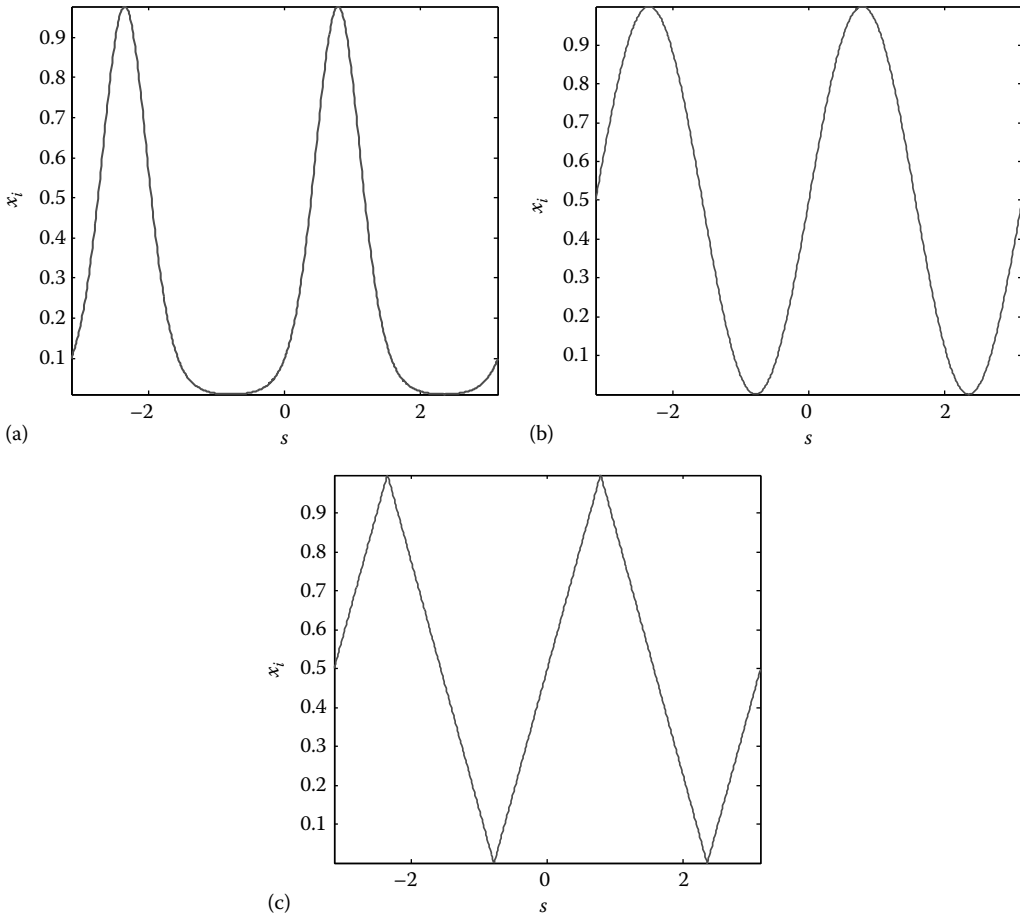


FIGURE 18.10 Examples of transformations from s to x_i : (a) Equation 18.17, (b) Equation 18.18 and (c) Equation 18.19.

18.10, they have been adjusted to give $x_i \in [0, 1]$. Note also that by varying ω_i , one can control the number of oscillations over the range of x_i ; therefore, ω_i represents the *search frequency* of each variable.

When all input variables are determined as functions of s , sampling uniformly over s produces samples along a *search curve* over the input space of the x_i . Figure 18.11 shows as an example the search curve produced for two input variables, using the triangular transformation given in (18.19), with $\omega_1 = 1$ and $\omega_2 = 4$. Notice that taking evenly spaced values of s over $[-\pi, \pi]$ results in values of x_1, x_2 that oscillate once over the range of x_1 and four times over the range of x_2 .

Clearly, a desirable property of the search curve is that it should be able to explore the input space as efficiently as possible, in other words to generate a space-filling design. Given a choice of transformation function, the extent to which the sampled points along the search curve fill the input space is dependent on the choices of the ω_i and n , the number of sample points. To illustrate this, Figure 18.12 shows the points generated by transforming s into two input variables with three different sets of search frequencies. Although in all three cases n is the same, the space-filling properties of the three curves are very different.




FIGURE 18.11

A two-dimensional example of a FAST search curve using a triangular basis function (18.19) and with $\omega_1 = 1$ and $\omega_2 = 4$.

In particular, when the two frequencies share common factors (such as in Figure 18.12a, where $\omega_1 = 10$ and $\omega_2 = 20 = 2\omega_1$), the input space is explored very poorly because the curve repeats itself with a period of $\pi/5$. In contrast, in Figure 18.12c, where $\omega_1 = 10$ and $\omega_2 = 21$, the only common factor is 1, which results in a better search curve because the points are unique over the whole range of s . The problem of frequency selection is not as simple as simply choosing, for example, large prime number values of ω_i , because more sample points are then required to adequately represent the search curve due to the Nyquist criterion—this is explained in a little more detail in the following section. In fact, the choice of the ω_i requires a balance between higher frequency and lower sample size. Sets of ω_i that optimise the space-filling properties of the search curve, for given dimensionality and sample size, can be found in Schaibly and Shuler (1973).

18.4.2 Calculation of Sensitivity Index Estimates

Given the transformation of variables, $f(x)$ is now expressed as a function $f(s)$ which is periodic over 2π . As such it can be expanded as a Fourier series:

$$f(s) = \frac{A_0}{2} + \sum_{r=1}^{\infty} (A_r \sin(rs) + B_r \cos(rs)) \tag{18.20}$$

with coefficients given as,

$$A_r = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(s) \cos(rs) ds$$

$$B_r = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(s) \sin(rs) ds, \tag{18.21}$$

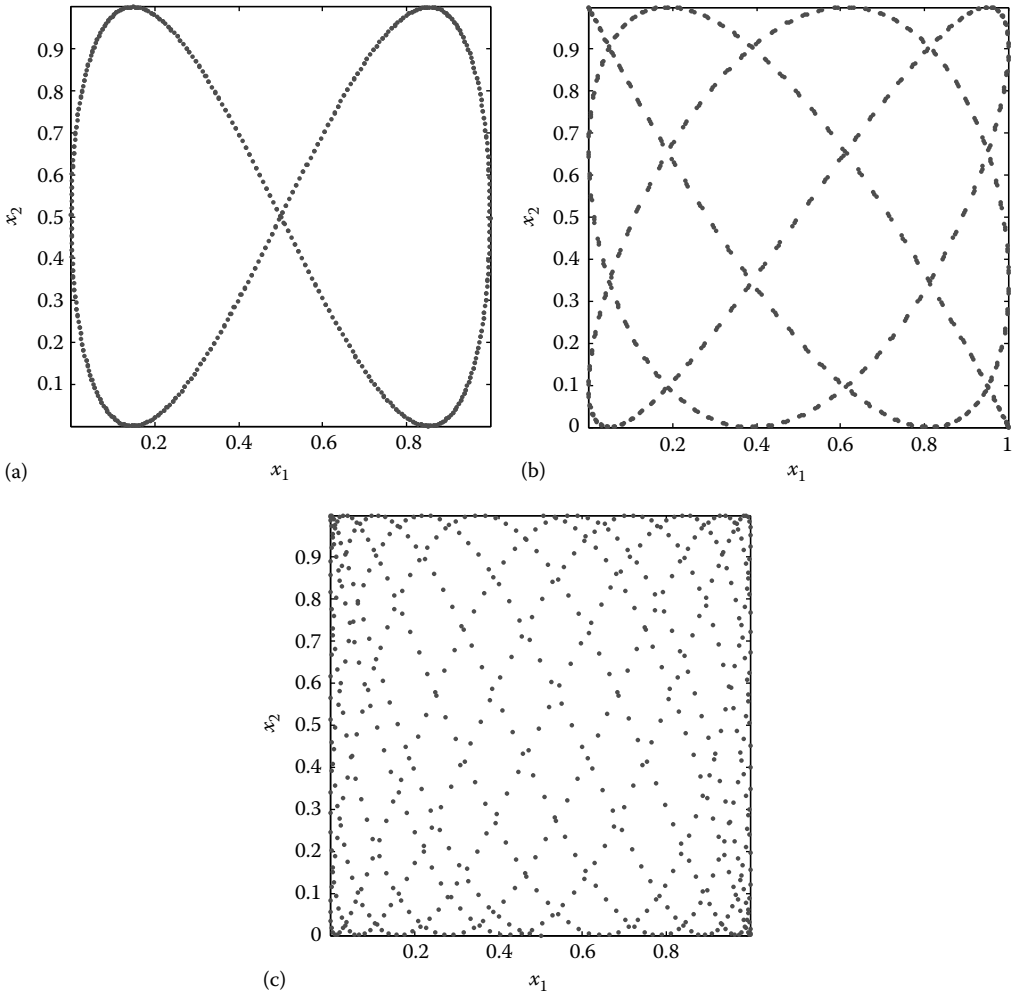


FIGURE 18.12

Search curves in two-dimensional space using transformation function (18.18): (a) $\omega_1 = 10, \omega_2 = 20$; (b) $\omega_1 = 15, \omega_2 = 21$; and (c) $\omega_1 = 10, \omega_2 = 21$.

which can also be viewed as the *amplitudes* of $f(s)$ at a given frequency r , or in other words, the contribution of frequency r to the function. This provides the basis for FAST—the coefficients of the Fourier terms at $r = \omega_i$ can be interpreted as a measure of sensitivity, because if the function output has a strong component at frequency ω_i , this implies that it is strongly affected by input x_i . The corresponding estimators for the coefficients in (18.21) are:

$$\hat{A}_r = \frac{1}{n} \sum_{j=1}^n f(s_j) \cos(rs_j)$$

$$\hat{B}_r = \frac{1}{n} \sum_{j=1}^n f(s_j) \sin(rs_j). \tag{18.22}$$

The use of the Fourier series allows expressions for the variance and partial variances of the output y . The expression for $\text{Var}(y)$ can be given in terms of the coefficients of the Fourier series as follows (Koda et al., 1979):

$$\begin{aligned}\text{Var}(y) &= \frac{1}{2\pi} \int_{-\pi}^{\pi} (f^2(s)ds - E(y)^2) ds \\ &= 2 \sum_{r=1}^{\infty} (A_r^2 + B_r^2).\end{aligned}\tag{18.23}$$

In practice, this expression is truncated to a maximum frequency R and is therefore an approximation to the variance. Similarly, the partial variance of y , $\text{Var}_{x_i}[E_{x_{\sim i}}(y|x_i)]$ (as discussed in Section 18.2), can be expressed in terms of the coefficients that correspond to the frequency ω_i and its multiples (harmonics) $p\omega_i$, where p is a positive integer (Koda et al. 1979):

$$\text{Var}_{\omega_i} = 2 \sum_{p=1}^{\infty} (A_{p\omega_i}^2 + B_{p\omega_i}^2),\tag{18.24}$$

which will again be truncated to a maximum order $p = M$, that is, the M th harmonic (multiple) of ω_i . The calculation of first-order sensitivity indices proceeds by noting that $S_i = V_{\omega_i}/\text{Var}(y)$. A problem with (18.24) is that of interference: for higher values of p , there will inevitably exist some p th harmonic of ω_i that is the same frequency as a q th harmonic of ω_j (the sampling frequency of another input variable x_j), such that $p\omega_i = q\omega_j$. This means that the coefficients of this frequency would be counted in both S_i and S_j , resulting in estimation bias. However, the amplitudes of higher harmonics in (18.24) generally decrease, so that if the ω_i are carefully chosen, any interferences up to the truncation order will be minimal.

As with Monte Carlo methods, the precision of the FAST estimates increases as n , the number of sample points, increases. For FAST, n should be at least $2R\omega_{\max} + 1$, where ω_{\max} is the highest frequency considered in the estimators (18.23) and (18.24), a limit which is imposed by the Nyquist criterion (Nyquist 1928). The required n can therefore become quite high as k increases, since higher frequencies are required to avoid interference. FAST additionally suffers from a number of sources of bias—first, that the Fourier series used to approximate the model must be truncated and, second, that for any given set of ω_i , there will be points in the input space that can never be sampled no matter how many sample points are used. Although the Fourier series is known to converge to any periodic function that is square-integrable (Carleson 1966) (the same property as the Monte Carlo estimators), the existence of bias suggests that the Monte Carlo method should generally be preferred when possible. Some further discussion on bias in FAST can be found in Xu and Gertner (2011) and Tissot and Prieur (2012).

18.4.3 Extended FAST

An extension to FAST was proposed in Saltelli et al. (1999), which additionally allows the estimation of the total effect indices discussed in Section 18.2. It proceeds by the observation that the set of frequencies *not* in the set of search frequencies and their harmonics

(i.e. $\omega \notin \{p\omega_i\}$; $p = 1, 2, 3, \dots$; $i = 1, 2, \dots, k$) contains information on the residual variance that is unaccounted for by first-order effects, which means that it contains information on the variance due to interactions of any order between variables (see again (18.3)). In the configuration described earlier, however, there is no obvious way to attribute this residual variance to interactions of particular inputs.

The proposal (known as *extended FAST*) is therefore to use a high-frequency ω_i for the i th variable and then to assign a set of low frequencies to all the remaining variables, such that their frequencies and harmonics $\{p\omega_j\}_{p=1}^M, j \neq i$, will be lower than $\omega_i/2$. The result is that the information about the output variance explained by the $x_{\sim i}$ variables, including all interactions between them, is isolated in the frequencies below $\omega/2$. By summing the variances from all these frequencies, an estimator for $\text{Var}_{x_{\sim i}}[E_{x_i}(y | x_{\sim i})]$ is obtained, which can be directly used to estimate S_{T_i} —see (18.6).

The obvious drawback to this method is that, whereas the estimation of all the S_i can be performed with one search curve, to estimate all the S_{T_i} requires k search curves. However, the approach has still been shown to be of at least comparable efficiency to the Monte Carlo approach given in Section 18.3 (Saltelli et al. 1999), although the sources of bias discussed in the previous section mean that the Monte Carlo method may still be preferable in some cases.

18.4.4 Random Balance Designs

Since different frequencies must be used to investigate each variable, the computational cost of FAST quickly rises with k (albeit less than with the Monte Carlo approach), because higher frequencies are required to avoid interferences between harmonics. Due to the Nyquist criterion, this requires more sample points. The RBD approach to FAST (Tarantola et al. 2006) circumvents this to some extent by using a single frequency ω for all inputs. With no further modifications here, the points would be very poorly distributed in the sample space—in fact they would be limited to a straight line (see Figure 18.13). The key to RBD, which was first proposed as an approach to experimental design back in 1959 by




FIGURE 18.13

A two-dimensional example of an RBD sample before and after scrambling of coordinates, with $\omega_1 = \omega_2 = 1$.

Satterthwaite (1959) and Budne (1959), is to then take random permutations of the coordinates of these points. As an example, seven points for two input variables are generated using transformation (18.18). Scrambling involves independently randomly permuting columns as follows:

x_1	x_2		x_1	x_2
0.50	0.50		0.50	0.59
0.18	0.18		0.18	0.18
0.14	0.14		0.14	0.77
0.45	0.45	\implies	0.45	0.14
0.77	0.77		0.77	0.45
0.91	0.91		0.91	0.50
0.59	0.59		0.59	0.91

where the first two columns represent the sample before scrambling and the last two show the sample after scrambling. These points are also illustrated in Figure 18.13—notice that the scrambled points fill the input space quite well because the design is very similar to a Latin hypercube (see Chapter 17).

In order to calculate sensitivity indices, the model is run at the points given by the RBD, and then for a given input i , the points are sorted into increasing order with respect to x_i , which recreates a periodic function of x_i . Then the first-order sensitivity indices can be calculated using (18.22) through (18.24) in the same way as the standard FAST method. The advantage of RBD is that, since the same frequency can be used for all inputs (which can be low, e.g. $\omega = 1$), the number of sample points required is less than conventional FAST. However, the drawback is that when estimating the effect of x_i , the random scrambling of the input variables generates random noise in the signal of x_i , resulting in a biased estimation of S_i . A bias-corrected estimator for sensitivity indices of any order has in fact been proposed by Tissot and Prieur (2012), where a more detailed discussion on bias in RBD can also be found. Readers might also want to refer back to the original extensive discussion on bias and other issues surrounding RBD (in the general context) after Satterthwaite's original paper (Youden et al. 1959). RBD has also been extended to compute total effect indices—readers are referred to Mara (2009) for more details.

18.5 Screening Designs

In the case where the user has access to the model, but it is expensive to run and there are a large number of input variables (Case 1C, see Section 18.1.2), screening methods offer a computationally efficient way of identifying influential and non-influential variables. In this setting, it is not possible to use emulators (see Section 18.6) because they become increasingly expensive to fit as the number of model inputs increases, so are not suitable for high-dimensionality problems. Typically, screening methods are used in the *factor fixing* setting, that is, to eliminate uninfluential variables before applying a more informative analysis to the remaining set (thus reducing the dimensionality—see again Figure 18.2), although they can also be used as a sensitivity analysis in their own right. The reason a user might prefer a variance-based measure, however, is because screening measures do not have a clear interpretation in terms of contribution the output variance, in contrast to S_i and S_{Ti} .

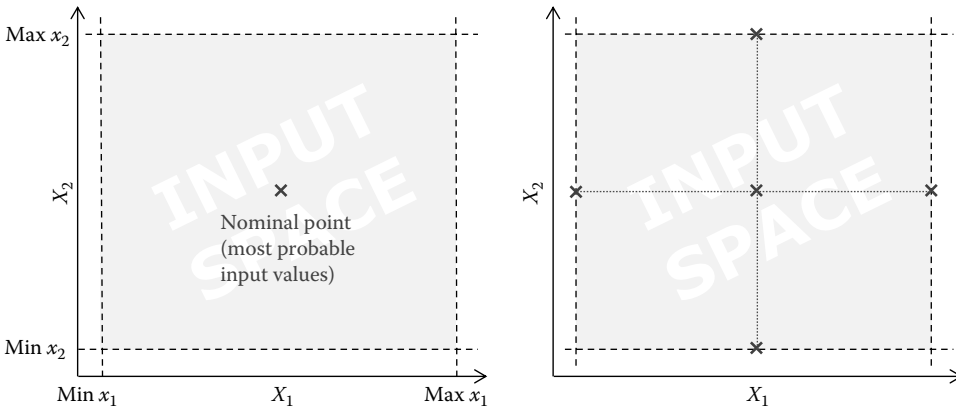


FIGURE 18.14 Basic OAT design in two dimensions; red crosses indicate design points.

The most common screening approach, suggested by Morris (1991), is an extension of a basic method of sensitivity analysis, the OAT design. Basic OAT designs simply involve varying one input while keeping all other inputs fixed at nominal (mean) values (see Figure 18.14). While this gives some information about the behaviour of the model, it is a very limited analysis because the sensitivity of input variables is only seen at the nominal value of the remaining inputs, and the extremities of the input space (i.e. the corners in Figure 18.14) are not explored. Morris’s approach, which will be called the *elementary effects method*, overcomes this problem by performing a number of OAT designs at random locations in the input space, rather than being restricted to nominal values. In this way nonlinearities and interactions can be accounted for by observing how the sensitivity of an input variable varies when moving about the input space.

18.5.1 Winding Stairs Design

The elementary effects design begins by dividing the range of each input into M equally spaced intervals (M being chosen by the analyst), such that the input space is divided into a grid of points. The design then proceeds by selecting a random point from this grid as a starting point, then moving in steps of Δ in each coordinate direction, where Δ is a predetermined multiple of $1/(M - 1)$. This design was stated in general terms by Morris (1991), but one implementation, known as the *winding stairs design* in some literature (Campolongo et al. 2011), is given here as an example. The design proceeds as follows:

1. A random point on the grid is selected as the first design point.
2. The first input, x_1 , is varied by an amount Δ_j , which is a randomly chosen multiple of the intervals in the grid, while keeping all other inputs fixed. This is the second design point.
3. Using the previous design point as a starting point, the next input, x_2 , is varied by the same amount Δ_j , keeping all other inputs fixed.




FIGURE 18.15
A trajectory screening design in three dimensions, with $n = 5$.

4. Step 3 is repeated for all k inputs.
5. Steps 1–4 are repeated for $j = 1, 2, \dots, n$ (n is selected by the user).

This design results in n trajectories, each consisting of $k + 1$ points (see Figure 18.15 for an example in three dimensions with $n = 5$). Because each trajectory is built by moving just one variable at a time, it is essentially a form of an OAT design. However, in the elementary effects design, there are n OAT designs in each x_i direction, at different points in the input space, which is much more informative than a single OAT design. The efficiency of the winding stairs design can be improved slightly by using the last point of the j th trajectory as the starting point of the $(j + 1)$ th trajectory. This would then form a single continuous trajectory with cost $nk + 1$ instead of $n(k + 1)$ in the case of the design stated here in Steps 1–5, although this may come at the expense of exploring the input space less thoroughly. Other variations of these OAT designs are discussed by Morris (1991), including *clustered* designs which can use fewer sample points to calculate the same number of elementary effects—readers are referred to the original article for more details.

18.5.2 Measures

Let $x_j^{(i)}$ and $x_j^{(i')}$ be, respectively, a point in the input space and a point that differs from $x_j^{(i)}$ only in the value of x_i . The point $x_j^{(i')}$ will therefore be the preceding point to $x_j^{(i)}$ in the five steps just described. Sensitivity is then estimated for the i th input using the n elementary effects $\{\xi_{ji}\}_{j=1}^n$, where

$$\xi_{ji} = \frac{f(x_j^{(i')}) - f(x_j^{(i)})}{\Delta_j} = \frac{f(x_j^{(i')}) - f(x_j^{(i)})}{|x_{ji}^{(i')} - x_{ji}^{(i)}|} \tag{18.25}$$

where j is the index over trajectories. The first measure of sensitivity for the i th input is thus estimated as the mean of the ξ_{ji} :

$$\hat{\mu}_i = \frac{1}{n} \sum_j^n \xi_{ji}. \tag{18.26}$$

A further useful measure of nonlinearity and interaction is given by the variance $\hat{\sigma}_i^2$ of the elementary effects, which is estimated as follows:

$$\hat{\sigma}_i^2 = \frac{1}{n-1} \sum_j^n (\xi_{ji} - \hat{\mu}_i)^2. \tag{18.27}$$

The logic here is that if the response of the output to a given input were perfectly linear, the elementary effects would be identical anywhere in the input space (and hence $\hat{\sigma}_i^2$ would be zero); in the nonlinear case, the opposite would be true.

A drawback with the sensitivity measure given in (18.26) is that if the main effect of an input is non-monotonic, the average of the elementary effects may be close to zero even though, individually, they may be significant positive or negative values. The result is that the measure $\hat{\mu}_i$ could potentially miss influential variables (although one would observe a high value of $\hat{\sigma}_i^2$). A modified measure $\hat{\mu}_i^*$, proposed in Campolongo et al. (2007), suggests the use of the mean of the absolute values of the elementary effects, that is,

$$\hat{\mu}_i^* = \frac{1}{n} \sum_j^n |\xi_{ji}|. \tag{18.28}$$

By using $\hat{\mu}_i$, $\hat{\mu}_i^*$ and $\hat{\sigma}_i^2$ in conjunction, one can assemble a picture of the strength and nature of the effect of each input at a low computational cost.

18.5.3 Radial Design

A drawback of the winding stairs design is that there is no guarantee that the trajectories are well spaced and that the input space has been well explored given the number of runs. A glance at the design in Figure 18.15 shows that points can sometimes be close to one another, therefore inefficiently exploring the input space. An alternative implementation of this design uses a so-called *radial* configuration based on Sobol’s LP_τ sequence to achieve a screening design with better-spaced trajectories (Campolongo et al. 2011). This design is in fact almost exactly the same as that used in variance-based sensitivity analysis but will be repeated here for clarity.

To construct the radial design, an LP_τ sequence of n points in $2k$ dimensions is generated and written as an $n \times 2k$ array (k being the number of model inputs). Let the first k columns be called the *baseline points*, that is, n points in k dimensions which will be denoted as a matrix A , with rows (individual points) $\{A_j\}_{j=1}^n$. The remaining k columns are called the n *auxiliary* points in a matrix B of the same size, with rows $\{B_j\}_{j=1}^n$. For a given baseline

point A_j and auxiliary point B_j , a radial configuration of $k + 1$ points is constructed as the following:

$$\begin{array}{c}
 A_{j,1}, A_{j,2}, A_{j,3}, \dots, A_{j,k} \\
 B_{j,1}, A_{j,2}, A_{j,3}, \dots, A_{j,k} \\
 A_{j,1}, B_{j,2}, A_{j,3}, \dots, A_{j,k} \\
 A_{j,1}, A_{j,2}, B_{j,3}, \dots, A_{j,k} \\
 \vdots \\
 A_{j,1}, A_{j,2}, A_{j,2}, \dots, B_{j,k}
 \end{array}$$

where, for example, $A_{j,1}$ is the first coordinate of A_j , the j th row of A . This configuration is repeated for $j = 1, 2, \dots, n$, resulting in n sets of $k + 1$ points: a total of $n(k + 1)$ model runs altogether. Notice that each set of $k + 1$ points (as shown earlier) defines a *star* with A_j as its centre and the other points at steps in each coordinate direction, defined by the coordinates of B_j . Figure 18.17 shows an example of the radial design in three dimensions, with $n = 5$, and Figure 18.16 shows the construction using the A and B matrices. Comparing Figures 18.15 and 18.17, one can see the strengths of each design: the spacing between trajectories/stars is better in the radial design because the centre points are guaranteed to be well spaced, being drawn from the Sobol' sequence. On the other hand, the spacing of points within each trajectory/star is arguably better with the winding stairs design, since points are allowed to differ in more than one variable. Which of these properties is more

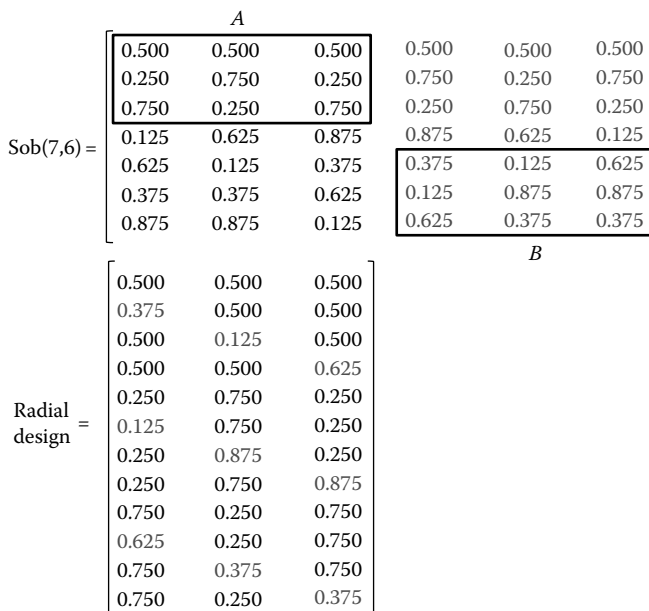


FIGURE 18.16 Construction of the first three points of the radial design based on the Sobol' sequence (design shown in Figure 18.17). Values from the B matrix are shown in grey.




FIGURE 18.17
A radial screening design in three dimensions, with $n = 5$.

important, and whether there is perhaps an advantage in combining the two designs, is left as an open question.

Since in Sobol' sequences the values of coordinates tend to repeat (see the example of the sequence in Figure 18.6), it is recommended that the baseline point A_j is not paired with B_j , but rather with a different row; otherwise, there could be no perturbation in certain dimensions (this problem is also mentioned in Section 18.3). It has been suggested that pairing A_j and $B_{j+\delta}$, where $\delta = 4$ gives good results (Campolongo et al. 2011), although there is no reason not to consider higher values of δ or to program the algorithm to skip any replicated values. In any case, this *row shift* means that for a design of n radial configurations, one needs a Sobol' sequence of $n + \delta$ points in $2k$ dimensions. This still has a computational cost of $n(k + 1)$ runs because the first δ rows of B and the last δ rows of A are discarded. Indeed, this *row-shift* strategy could also be applied to the estimators for the variance-based sensitivity indices discussed in Section 18.3, but is not generally deemed to be an issue since when estimating S_{Ti} , one is in the domain of Case 1A (a cheap computer model that can be run thousands of times) where an extra few model runs make little difference, whereas in the screening setting (Case 1C), one has to conserve model runs as much as possible.

The elementary effects method is often used in the factor fixing setting (identifying input variables that have little or no effect on the model output), with the possible intention of estimating sensitivity indices on the remaining set of *important* variables via Monte Carlo methods. An advantage therefore of using the radial design, noted in Campolongo et al. (2011), is that since it is effectively the same as the Monte Carlo design, the points can be re-used as the basis of a Monte Carlo design for estimating sensitivity indices. Alternatively (and perhaps more realistically), one might want to fit an emulator to the reduced set of input variables, since the computational cost of Monte Carlo methods is likely still prohibitive even after screening. In that case, the OAT-type designs discussed in this section (both winding stairs and radial) are wasteful since if they are projected onto the reduced-dimensionality subspace of the set of screened inputs, for every input that is discarded,




FIGURE 18.18 Collapse of radial design when projected onto a subspace: (a) a radial design in three dimensions; (b) the same design projected onto the subspace of x_1 and x_3 .

n model runs are lost. To see why this is, consider the example in Figure 18.18. A radial design is built in three dimensions, the elementary effects method is applied, and x_2 is found to be unimportant. It is desired to examine the remaining inputs (x_1 and x_3) in more detail using an emulator. When the design is projected into the subspace of x_1 and x_3 , the points that differ only by a step in the x_2 direction, a total of n points, one for each *star*, are now duplicate points. A recent approach that overcomes this problem uses a set of simplex designs to estimate screening measures (Pujol 2009), which have the property that all design points are retained after screening out unimportant variables.

18.5.4 Example

To show a simple example, the design shown in Figure 18.17 is used to estimate the μ_i , μ_i^* and σ_i^2 measures on a test function which is defined as

$$y = 3x_1^2 + 2x_1x_2 - 4\pi \sin(x_3). \tag{18.29}$$

Figure 18.19 shows scatter plots of the function values against each input, using random uniform sampling from $[0, 1]^3$ to give a visual idea of the behaviour of the function. The screening design is run with $n = 5$ and $k = 3$ (exactly as in Figure 18.17); a cost of 20 model runs. Figure 18.20 shows the results of the screening analysis. Input x_2 is clearly the least influential by any measure, and is relatively linear, having a small value of $\hat{\sigma}_i^2$. Similarly, x_1 is quite linear (having also a small value of $\hat{\sigma}_i^2$) but is judged to be more important by its $\hat{\mu}_i$ and $\hat{\mu}_i^*$ measures. In both these cases, $\hat{\mu}_i = \hat{\mu}_i^*$, which indicates monotonicity. For x_3 , one can observe a much higher $\hat{\sigma}_i^2$ value, indicating a strong nonlinearity. Finally, the fact that $\hat{\mu}_3 \neq \hat{\mu}_3^*$ indicates that the model is non-monotonic with respect to x_3 . These results are reflected by the scatter plots in Figure 18.19.

Note that the screening measures do not have an interpretation with regard to the variance of the output (as compared to the S_i and S_{T_i} measures from Section 18.2), but they

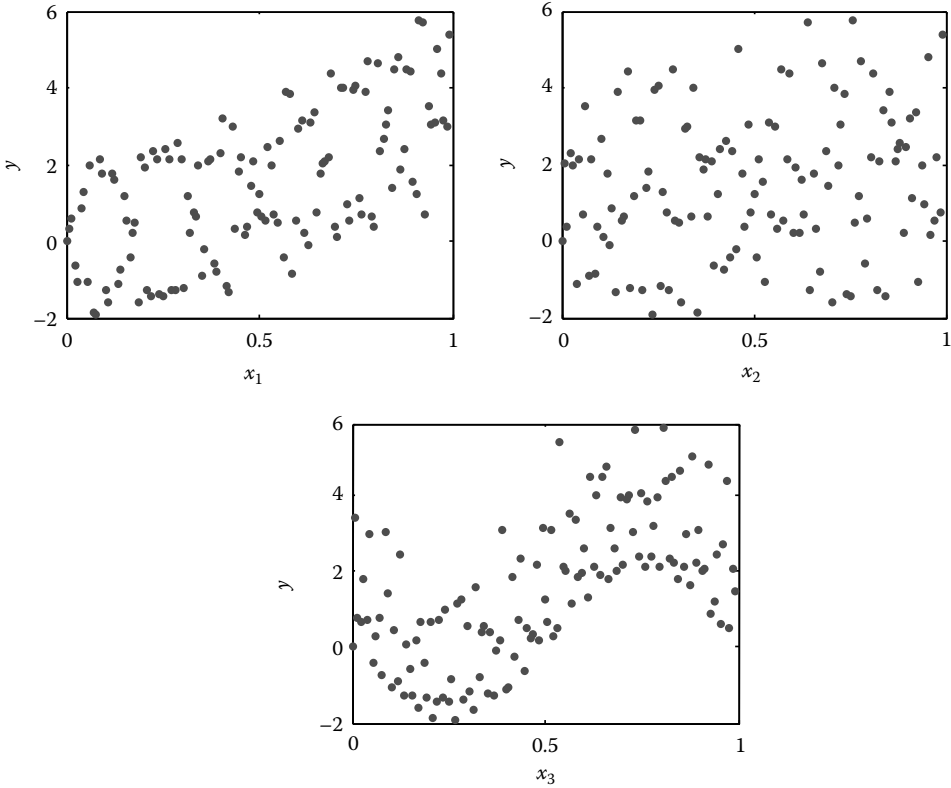


FIGURE 18.19

Scatter plots for test function (18.29). Note that these are generated using random uniform sampling from $[0, 1]^3$ —these are not the same design points used in the screening design.




FIGURE 18.20

Estimates of screening measures μ_i, μ_i^* and σ_i^2 plotted for test function (18.29).

allow the user to sort between influential and unimportant variables. Consider also that the example given is trivial, since screening is generally for use with high-dimensional problems, but even with 20 runs, the order and to some extent the magnitude of importance of each variable can be distinguished with these relatively simple measures, as well as information regarding the linearity and monotonicity.

18.5.5 Discussion

It is interesting to consider briefly the similarity of the estimator of μ^* in (18.28) and the estimator for S_{Ti} given in (18.12). Let \mathbf{x}_j be the j th point of a random or quasi-random sequence in the input space, and let $\mathbf{x}_j^{(i)}$ be a point that differs from \mathbf{x}_j only by the i th variable. Let $E_{\mathbf{x}_{\sim i}}[\text{Var}_{\mathbf{x}_i}(y | \mathbf{x}_{\sim i})]$, the numerator of S_{Ti} , also be denoted as V_{Ti} as before. Then the estimators for μ_i^* and V_{Ti} can be stated in similar terms as follows:

$$\begin{aligned} \hat{\mu}_i^* &= \frac{1}{n} \sum_{j=1}^n \frac{|f(\mathbf{x}_j^{(i)}) - f(\mathbf{x}_j)|}{|x_{ji}^{(i)} - x_{ji}|} \\ \hat{V}_{Ti} &= \frac{1}{2n} \sum_{j=1}^n |f(\mathbf{x}_j^{(i)}) - f(\mathbf{x}_j)|^2 \end{aligned} \tag{18.30}$$

where the estimator for μ_i^* has been written slightly differently compared to (18.25) and (18.28) since in the radial design all k elementary effects in each *star* use the same centre point, \mathbf{x}_j . Notice that both measures rely on averages of model outputs from multiple OAT designs conducted at different points in the input space. The only differences (up to a proportional constant) are that $\hat{\mu}_i^*$ uses the absolute value of $(f(\mathbf{x}_j^{(i)}) - f(\mathbf{x}_j))$, whereas \hat{S}_{Ti} uses the square, and $\hat{\mu}_i^*$ also incorporates the information about the distance between $\mathbf{x}_j^{(i)}$ and \mathbf{x}_j , which is instead discarded by \hat{S}_{Ti} .

Another very similar related measure of sensitivity that has been recently the subject of interest, under the heading of *derivative-based global sensitivity measures* (DGSMs), is the integral of squared derivatives, that is, a measure $v_i = \int_{\mathcal{X}} (\partial f / \partial x_i)^2 dx$ (Kucherenko et al. 2009). Using the notation just defined, its estimator is stated as

$$\hat{v}_i = \frac{1}{n} \sum_{j=1}^n \frac{|f(\mathbf{x}_j^{(i)}) - f(\mathbf{x}_j)|^2}{|x_{ji}^{(i)} - x_{ji}|}, \tag{18.31}$$

with the difference that the $\mathbf{x}_j^{(i)}$ and \mathbf{x}_j points are defined to be very close to each other to provide as close an approximation as possible to the partial derivative at \mathbf{x}_j . Some further observations on the similarities of these three measures, and further information on DGSM, are found in Sobol and Kucherenko (2009).

Connections can also be found between the radial design (Figure 18.17) and the Monte Carlo design for estimating S_i and S_{Ti} (Figure 18.7) and their respective design matrices (Figures 18.6 and 18.16). One can see that the two designs are conceptually identical, with the only difference being that the screening design here uses a shift of the \mathbf{B} matrix to avoid replacing a coordinate with an identical value. Recalling the discussion on the similarity of

the estimators of μ_i^* and S_{Ti} in the previous section, it is evident that the elementary effects method and the Monte Carlo estimation of sensitivity indices have much in common.

Finally, in some cases, the number of input variables may be so large or the cost of running the model so high that even the designs here may be too expensive. In that case, one option available to the analyst is to use grouped designs, where inputs are grouped and designs are built with respect to each group rather than to each input. In that case, one would have a measure of sensitivity for each group but be ignorant of the relative sensitivity of inputs inside each group. This concept is discussed more in Watson (1961), Morris (1991), Campolongo et al. (2007) and Moon et al. (2012).

18.6 Emulators

In Case 1B (the model is expensive to run but the dimensionality is not too high) and Case 2 (points are *given* and the analyst cannot specify new points—refer back to Section 18.1.2) a general method is presented here which adopts a data-modelling approach. The concept is to fit an *emulator* (a relatively simple mathematical function, also known as a *metamodel*) to the data, which behaves in the same way as the model itself. The emulator can then be used to estimate the model output at any point in the input space, allowing analytical or Monte Carlo estimation of the S_i and S_{Ti} (see Section 18.3) at a considerably lower computation cost. An alternative approach is to project the data onto a single axis x_i (i.e. create k one-dimensional scatter plots) and attempt to infer the main effect $E(y|x_i)$ using a smoothing regression approach, for example, kernel regression. This latter approach is discussed in Section 18.7.

The central idea of emulation is to find some relatively simple function η (the emulator) such that it closely approximates the output of the model f at any point in the input space, that is, $\eta(x) \approx f(x)$, $\forall x \in \mathcal{X}$. If η is considerably cheaper to evaluate at a given x than the original model but produces sufficiently similar results for any point in the input space, then it can be used to generate a very large number of estimated model output values, for example, at the points specified by a Monte Carlo design for estimating S_i and S_{Ti} (see Section 18.3). Even better, if η is analytically tractable, it can be used to calculate sensitivity indices analytically, because (18.2) and (18.6) can be expressed as integrals which can be solved if $\eta(x)$ is sufficiently simple. This then bypasses Monte Carlo methods altogether and the associated approximation error of numerical integration.

The four steps associated with building an emulator are as follows:

1. Select a type of emulator, η , that is appropriate for emulating the model $f(x)$. Options could be as simple as a linear regression, to more complex methods such as Gaussian processes (GPs), neural networks or smoothing splines. A brief description of some emulators follows in this section, but for a more detailed treatment, one should refer to one of the many books on the subject, such as Bishop (2006).
2. Sample (run) the model f at appropriate points in the input space to provide *training data* for the emulator. This will be a set of n points $\{x_1, x_2, \dots, x_n\}$ and n corresponding model outputs $\{y_1, y_2, \dots, y_n\}$ where the location of the inputs is usually chosen to be a space-filling design or *optimal design* (see Section 18.6.3). Emulators can however be fit to any sample of data (thus can be used for Case 2

- problems) although in areas of the input space with very few training data they will be less accurate.
3. Use the training data to estimate any parameters or hyperparameters associated with the model (this is called *training* the emulator).
 4. Check the accuracy of the estimator using methods such as cross-validation (an approach which involves training the emulator on a subset of the training data and then comparing the predictions of the emulator with the known model runs in the remaining set of data).

A large number of different emulators are available—see, e.g. Bishop (2006); a comparison of some of these methods in the context of sensitivity analysis can be found in Storlie and Helton (2008). Once an appropriate emulation method is selected, the accurate estimation of parameters can be achieved provided that a sufficiently large sample of training data (the given points) is available. How large this sample needs to be is dependent on the type of emulator and the complexity of the model f and increases dramatically with the number of input variables (the so-called curse of dimensionality). For this reason, emulator-based approaches are best suited to situations with few input variables—perhaps fewer than thirty, depending on the emulator—which demand fewer model evaluations for training data. Higher dimensionality problems can sometimes be brought into the reach of emulators by a precursory screening analysis (see Section 18.5) to reduce the number of variables (refer back to Figure 18.2). While there are many emulators available in the literature, the following sections give some brief information on only a small subset to give a flavour of the field.

18.6.1 High-Dimensional Model Representation

A significant proportion of emulation methods rely on a technique known as high-dimensional model representation (HDMR), which seeks to approximate the model by performing a functional decomposition of f into orthogonal terms and then truncating the series (Rabitz and Alis 1999; Li et al. 2002). This has already been given in (18.4) and is restated as

$$\begin{aligned}
 f(\mathbf{x}) &= f_0 + \sum_i f_i(x_i) + \sum_i \sum_{j>i} f_{i,j}(x_i, x_j) + \cdots + f_{1,2,\dots,k}(x_1, x_2, \dots, x_k) \\
 &\approx f_0 + \sum_i f_i(x_i) + \sum_i \sum_{j>i} f_{i,j}(x_i, x_j)
 \end{aligned}
 \tag{18.32}$$

where the series is truncated after the second-order terms, based on the (empirical) observation that the effect of third-order and higher interactions on the model output is negligible in many models. The task then remains to find suitable orthogonal functions to approximate the $f_i(x_i)$ and $f_{i,j}(x_i, x_j)$ (some approaches are discussed in Li et al. (2002)). Orthogonality is satisfied for any two components $f_u(\mathbf{x}_u)$, $f_v(\mathbf{x}_v)$ of the functional decomposition in (18.32) when $\int_{\mathcal{X}} f_u(\mathbf{x}_u) f_v(\mathbf{x}_v) d\mathbf{x} = 0$, where $u, v \subseteq \{1, 2, \dots, k\}$ and, for example, $\mathbf{x}_u = \{x_i\}_{i \in u}$. The advantage of HDMR is that it alleviates to some extent the problem of dimensionality, since a model with many inputs can be approximated by a sum of one- and two-dimensional terms which are relatively easy to fit.

One method for approximating the terms in the functional decomposition (18.32) which is relatively simple is to use a series of orthogonal polynomial functions (Draper and Smith 1981; Li et al. 2002; Sudret 2008). These take the form

$$\begin{aligned}
 f_i(x_i) &= \sum_{r=1}^{\infty} \alpha_r^{(i)} \phi_r(x_i) \\
 f_{i,j}(x_i, x_j) &= \sum_{p=1}^{\infty} \sum_{q=1}^{\infty} \beta_{p,q}^{(ij)} \Phi_{p,q}(x_i, x_j)
 \end{aligned}
 \tag{18.33}$$

where ϕ are terms from a suitable series of orthogonal polynomials and α and β are their corresponding coefficients. Of the many series of orthogonal polynomials, it is typical to use either the Legendre or the Hermite types (Bayin 2006). Clearly, it is necessary to truncate the infinite series of each of the equalities in (18.33) to a certain order M , which is usually done by discarding terms after the third order based on a heuristic assumption that higher orders are negligible (Li et al., 2002) or, in some cases, by sequentially adding terms and re-estimating coefficients until estimated sensitivity indices appear to converge (Zuniga et al. 2013). This gives the approximated functions as follows:

$$\begin{aligned}
 f_i(x_i) &\approx \sum_{r=1}^M \hat{\alpha}_r^{(i)} \phi_r(x_i) \\
 f_{i,j}(x_i, x_j) &\approx \sum_{p=1}^M \sum_{q=1}^M \hat{\beta}_{p,q}^{(ij)} \Phi_{p,q}(x_i, x_j)
 \end{aligned}
 \tag{18.34}$$

where the estimates of the coefficients α and β are obtained by minimising the squared difference between the terms in (18.34) and their counterparts in the HDMR decomposition one at a time. This minimisation problem can be expressed as a series of integrals which can be solved by Monte Carlo integration (Li et al. 2002). The HDMR emulator can then be built by assembling the terms from (18.34) into the truncated representation in (18.32). However, due to the simplicity of the orthogonal polynomials, it is possible to derive analytical expressions for S_i and $S_{i,j}$, which are as follows:

$$\hat{S}_i = \frac{\sum_{r=1}^{\infty} (\hat{\alpha}_r^{(i)})^2}{\widehat{\text{Var}}(y)}
 \tag{18.35}$$

$$\hat{S}_{i,j} = \frac{\sum_{p=1}^{\infty} \sum_{q=1}^{\infty} (\hat{\beta}_{p,q}^{(ij)})^2}{\widehat{\text{Var}}(y)}
 \tag{18.36}$$

where the summations are again truncated to the same orders used in (18.34). S_{T_i} may be approximated from the sum of S_i and its second-order interactions, based on the same assumption used in (18.33) that interactions above the second order do not contribute significantly to the model output (and therefore by extension to $\widehat{\text{Var}}(y)$). The variance term $\widehat{\text{Var}}(y)$ can be calculated either from an analytical expression similar to those for the sensitivity indices, or from the original sample data. The drawbacks to the use of orthogonal

polynomials are that the HDMR representation must be truncated, then a further truncation is necessary of the polynomial series. Higher order polynomials can be used, but this requires the estimation of more coefficients, each of which has an associated Monte Carlo error, and tends to increase the error rather than diminish it (Li et al. 2002).

A somewhat related method to HDMR with orthogonal polynomials is called *polynomial chaos expansion*, which is an approach which involves representing the random variable of the model output, y , as a function of other random variables (the model inputs x) via a polynomial expansion in orthogonal polynomials. The details of this method will not be explored here, but the reader is referred to Sudret (2008) for details on the methodology in the context of sensitivity analysis.

Another more complex approach for approximating functions that has been used with considerable success is the use of smoothing splines (Ratto and Pagano 2010). Smoothing splines are most commonly used on univariate data but can be extended to the multivariate case by the use of HDMR decompositions. In the following short introduction, the univariate case is therefore described for simplicity. A smoothing spline model assumes that the underlying function to be emulated is continuous and has a continuous first derivative and further that the second derivative is square-integrable. The smoothing spline estimate arises from considering the function g that minimises the following:

$$\frac{1}{n} \sum_{j=1}^n \{y_j - g(x_j)\}^2 + \lambda \int_0^1 \{g''(x)\}^2 dx. \quad (18.37)$$

The first term in this *trade-off* expression is simply the sum of squared errors between the training data and the emulator g —if the function were to pass through every data point (exact interpolation), this term would be zero. The second term expresses the integral of the square of the second derivative of g , which is a global measure of roughness, and λ is a *tuning parameter* that controls the weighting between the two terms. Overall therefore, the expression summarises the trade-off between interpolation and model simplicity. The solution to this minimisation problem can be shown to be a natural cubic spline, with knots (joins between the local cubic functions) at each of the data points. Natural cubic splines are simply local cubic polynomial functions between each data point and the next, with the constraints that the global function is continuous and the first derivative is continuous at knots. Since multivariate spline emulators rely on HDMR decompositions, however, estimates of total effect indices S_{Ti} are difficult since the HDMR series is truncated (usually neglecting third-order and higher interactions), so higher order interactions that may contribute to S_{Ti} are not accounted for. A common approximation is therefore to assume that $S_{Ti} \approx S_i + \sum_{j=1}^k S_{i,j}$.

The basis functions of splines as described earlier are not mutually orthogonal and therefore cannot be used in their standard form as the components of an HDMR emulator (Li et al. 2002). However, a class of spline methods known as *smoothing spline ANOVA models* allows the construction of mutually orthogonal spline basis functions via the reproducing kernel Hilbert space approach, which is discussed extensively in Gu (2002).

An extension of multivariate splines, known as adaptive component selection and shrinkage operator (ACOSSO), uses a modified version of (18.37) that uses norms rather than square-norms and also includes the integral of the first derivative of g (Lin and Zhang 2006). The result is that terms in the HDMR decomposition that contribute very little to the

function are eliminated, resulting in a simpler and more computationally efficient emulator. This approach has also been combined with state-dependent parameter regression, as described in Ratto et al. (2007). MATLAB® scripts for performing sensitivity analysis with this approach can be found in European Commission (2012).

18.6.2 Gaussian Processes

Another emulator that is widely used in sensitivity and uncertainty analysis is a GP, otherwise known as *kriging*. GPs are widely used in the machine learning community as a sophisticated form of nonlinear regression and classification (Rasmussen and Williams 2006). In short, a GP is a *distribution over functions*, that is, the random variable of the distribution is a function rather than a single number or fixed-length vector. Instead of returning only a point estimate \hat{y} for any given input point x (as in a standard regression), the GP returns a specification for a Gaussian probability distribution (a mean and a variance). Figure 18.21 shows an example of a simple one-dimensional GP fitted to a few points from a sine wave. Notice that at any value of x the output y is estimated by the GP as a predicted mean value (which forms a curve over the range of x) and a variance, here plotted as plus/minus two standard deviations of the predictive distribution.

GPs are fitted to training data following a Bayesian procedure. A prior distribution over functions is specified, which is then conditioned on the training points to give a posterior distribution over functions. Figure 18.21 shows an example of this process applied to data from a noisy sine wave—Figure 18.21a shows samples from the prior distribution over functions, and Figure 18.21b shows the posterior distribution over functions after conditioning on the training data.

The specification for the GP includes a number of hyperparameters. These may be estimated, for example, by plug-in estimators, for example, by using maximum likelihood




FIGURE 18.21

GP examples: (a) samples from the prior distribution over functions and (b) the GP conditioned on randomly generated points from a sine wave. The dotted line is the underlying sine function, the solid line is the fitted mean of the GP, and the grey region represents 95% prediction intervals.

estimation—this effectively treats the hyperparameters as known. Alternatively, one can assign prior distributions to hyperparameters and then integrate them out (marginalise them) either analytically or using Markov chain Monte Carlo methods—(see Bishop 2006 for more information on Bayesian inference). In this *full Bayesian* framework, the confidence intervals of the GP also include the uncertainty in the estimation of the hyperparameters.

GPs have the advantage that they do not invoke the HDMR assumption (of neglecting higher order interactions; see (18.32)) and as result can be used to estimate sensitivity indices of all orders, including the S_{Ti} . A further useful property of the GP is that, given certain assumptions, estimated sensitivity indices can be calculated analytically (Oakley and O’Hagan 2004; Becker et al. 2012). Another extremely useful property is that, since the GP is a probabilistic emulator, sensitivity indices can be stated with confidence intervals which account for the uncertainty in the fit of the emulator (and also the uncertainty in the hyperparameter estimation, when Bayesian inference is used). However, perhaps the main weakness of GPs is that the cost of training scales poorly with n , the number of training data, since it involves the inversion of a $n \times n$ covariance matrix at a cost of the order of n^3 . This limitation impacts on the dimensionality of the problems to which GPs can be applied, since more training data are required as the dimensionality increases. Added to the fact that the estimation of hyperparameters becomes increasingly expensive as the dimensionality increases, GPs tend to encounter problems emulating models with $k > 30$ inputs or so, depending on computational resources (this is however the case with all emulators, to some extent). A number of techniques are being developed to alleviate this problem (see Rasmussen and Williams (2006) for some examples). Some fairly recent additions in the field of GPs with respect to sensitivity analysis include a method of automatically screening out unimportant variables using the correlation parameter in the covariance function (see Linkletter et al. 2006, based on Welch et al. 1992), a method based on screening variables in groups using GPs (Moon et al. 2012) and the use of multiple GPs divided by decision trees to allow for discontinuous responses (Gramacy and Taddy 2010; Becker et al. 2013) (available as an R package). A very good general online resource on many aspects of GPs, emulation and sampling can be found at the *Managing Uncertainties in Complex Models (MUCM) Toolbox* (2013).

As an example of how an emulator may reduce the number of model runs required for sensitivity analysis, consider again the simple polynomial from (18.15). Using 128 points of the Sobol’ sequence over the unit cube, a GP was trained (i.e. the hyperparameters of the mean and covariance functions were estimated using the training data) and estimated sensitivity indices inferred analytically from the resulting posterior distribution. Table 18.2 shows the results. The GP is achieving accuracies of three or more decimal places on only 128 points—recall that the Monte Carlo estimator, for a similar level of accuracy, requires several thousands of runs per variable; therefore, the GP is at least an order of magnitude more efficient. However, the GP and other emulators are only as good as their

TABLE 18.2

Comparison of S_i and S_{Ti} Estimates from a GP Regression against Analytical Values

Variable	\hat{S}_i (GP)	S_i (Analytic)	\hat{S}_{Ti} (GP)	S_{Ti} (Analytic)
x_1	0.7566	0.7568	0.7715	0.7720
x_2	0.0456	0.0456	0.0605	0.0608
x_3	0.1829	0.1824	0.1830	0.1824

fit to the data: here the polynomial function is a smooth, *well-behaved* function, which is an easy data-modelling problem. For data that are heteroscedastic, discontinuous or of varying smoothness, the emulators are likely to be much less reliable. Additionally, they scale poorly with dimensionality. However, in the cases where the model gives a relatively smooth output (which appears to be the majority of physical models), emulators can offer a powerful solution to the problem of analysing uncertainty and sensitivity for computationally expensive models. For more detailed information on GPs, please refer to Rasmussen and Williams (2006) and Chapter 16.

Overall, there is no “best” emulator available. The approach will depend on computational resources, sample size and model dimension, amongst other things. Both Bishop (2006) and Storlie and Helton (2008) are recommended as background reading. Furthermore, it is essential to test the fit of any emulator by methods such as cross-validation.

18.6.3 Custom Sampling for Emulators

Although the points were considered as *given* in the discussion earlier, it can happen that the analyst has the possibility to design their own set of training data for fitting an emulator. This can be the case when, for example, the analyst only has a small number of input variables but a very computationally expensive model (Case 1B in Section 18.1.2). In this scenario, it makes sense to go directly to an emulator approach, since pure Monte Carlo would be too expensive and screening too inaccurate.

Experimental designs for fitting emulators can be divided into two categories—space-filling designs and model-based designs. In the former, the design is constructed to fill the sample space as evenly as possible, which is to say that points should be as far apart from each other as possible. The reasoning for this is that first, it is required to capture the behaviour of the model over the whole input space with as few points as possible. Second, assuming that the output of the model is deterministic and smooth with respect to its inputs, little information can be gained by having points close to each other (since the outputs will be very similar). For this reason, purely random sampling is not an efficient design.

For a general-purpose design for fitting an emulator, a space-filling design such as the Sobol’ sequence discussed in Section 18.3 is a good choice. Sobol’ designs have a low-discrepancy property that avoids *clumping* of points and allow the sequential addition of new points. For other space-filling designs, the reader is referred to Chapter 17 or Niederreiter (1992).

An even more sophisticated approach is to use *optimal design*. In this approach, the design is constructed so as to optimise some emulator-dependent criterion of interest, thus tailoring the design to the requirements of the emulator. For example, a popular criterion, called *D-optimality*, is to select design points which minimise the variance of the estimators of the emulator parameters. Another way to select design points is to minimise the maximum variance of the emulator prediction at any given input point (*G-optimality*). Note however that for emulators based on linear models or stationary GP, these designs are *not* dependent on the output of the model, only on the form of the emulator and the values of its parameters. See Section 2.3 for more information on *D-optimality*.

A further approach to building designs for emulators is to construct them sequentially. One starts with an initial design with a few input points (e.g. a space-filling or optimal design) and runs the model at these points. The emulator is fitted to this training data

(i.e. its hyperparameters are estimated), then the next input point is selected as the point which optimises some criterion of interest. An example, for a GP, would be to choose the point which returns the highest predictive variance. This new point is then run through the model and used to re-estimate the hyperparameters. The procedure is repeated, choosing each time the point with the highest predictive variance, until a design with a sufficient number of points is created. The advantage of the sequential approach (known as adaptive design) is twofold—first, the output values will influence the optimum placement of new points, so knowledge of previous points will produce a more effective design than one that is constructed in one go. Second, by proceeding in small steps, one can generate exactly the required number of points to reach some level of accuracy of interest, perhaps measured by cross-validation. More information on adaptive sampling can be found in Gramacy and Lee (2009).

The theory of model-based designs is a large field of research that is beyond the remit of this chapter; therefore, the reader is referred to Atkinson et al. (2007) for a good general resource. There is also a strong interest in Bayesian approaches to optimal design; a review can be found in Chaloner and Verdinelli (1995).

18.7 Scatter Plot Smoothing

A useful approach for estimating first-order sensitivity indices with *given* data is based on one-dimensional nonlinear smoothing regression. From a computational point of view, this method is less vulnerable to the curse of dimensionality, although it cannot be used to calculate the total-order sensitivity indices S_{T_i} .

A first visual indication of the effects of input variables can be gained by making k plots of x_i against y (see Figure 18.8). If the data show any kind of trend (or shape) with respect to x_i , this indicates that x_i has some effect on the output. Indeed, the effect of x_i on the output is described by the curve $E_{x_{-i}}(y|x_i)$ —in other words, the expected value of the model output if we were to fix x_i at some value. Over the range of x_i , $E_{x_{-i}}(y|x_i)$ is equivalent to a moving weighted average of the points in the x_i against y plot. As long as the $x_j, j = 1, 2, \dots, n$, have been randomly drawn from their specified distribution $p(x)$, S_i can be estimated by taking the variance of the y values of this curve (since $S_i = \text{Var}_{x_i}[E_{x_{-i}}(y|x_i)]$).

To estimate such a moving average, it is a matter of using any of a number of smoothing regression approaches. Kernel regression has been already used for sensitivity analysis (Paruolo et al. 2013). As with any smoothing regression, the data are modelled as

$$y = m(x_i) + \epsilon \quad (18.38)$$

where $m(x_i)$ is the smoothed curve (ideally equivalent to $E_{x_{-i}}(y|x_i)$) and ϵ is an independent error term with mean 0 and a variance that may be dependent on x_i . Although most nonlinear regression approaches assume a fixed variance over x_i , the smoothing curves that result when this assumption does not hold are still valid, albeit less efficient. However, nonlinear regression that accounts for heteroscedasticity is still a field of active research; therefore, this chapter does not venture into this territory and readers are referred to a discussion in Ruppert et al. (2003) as a starting point for further information.

In the kernel regression setting, $m(x_i)$ is typically chosen to be either a *local mean* or *local linear* kernel. The local mean estimator first proposed by Nadaraya (1964) and Watson (1964), for a single input variable x_i , is expressed as

$$\hat{m}(x_i) = \frac{\sum_{j=1}^n w(x_{ji} - x_i; h)y_j}{\sum_{j=1}^n w(x_{ji} - x_i; h)} \tag{18.39}$$

where w is a weighting function and h is a tuning parameter. The weighting function typically gives the strongest weight to points close to x_i , which reflects the belief that the closer the two points are to each other in x_i , the more likely they are to have similar values in y . A commonly used function that fulfils this requirement is a Gaussian density function with standard deviation h . The local linear estimator is expressed in a similar fashion—see Bowman and Azzalini (1997) for details—and is generally regarded as preferable to the local mean, due to its improved properties near the edges of the data cloud. In all cases, the smoothing parameter h can be optimised by cross-validation. Following the simple polynomial example from (18.15), Figure 18.22 shows an illustration of local linear kernel regression applied to scatter plots of y against each x_i . The resulting estimates of sensitivity are given in Table 18.3.

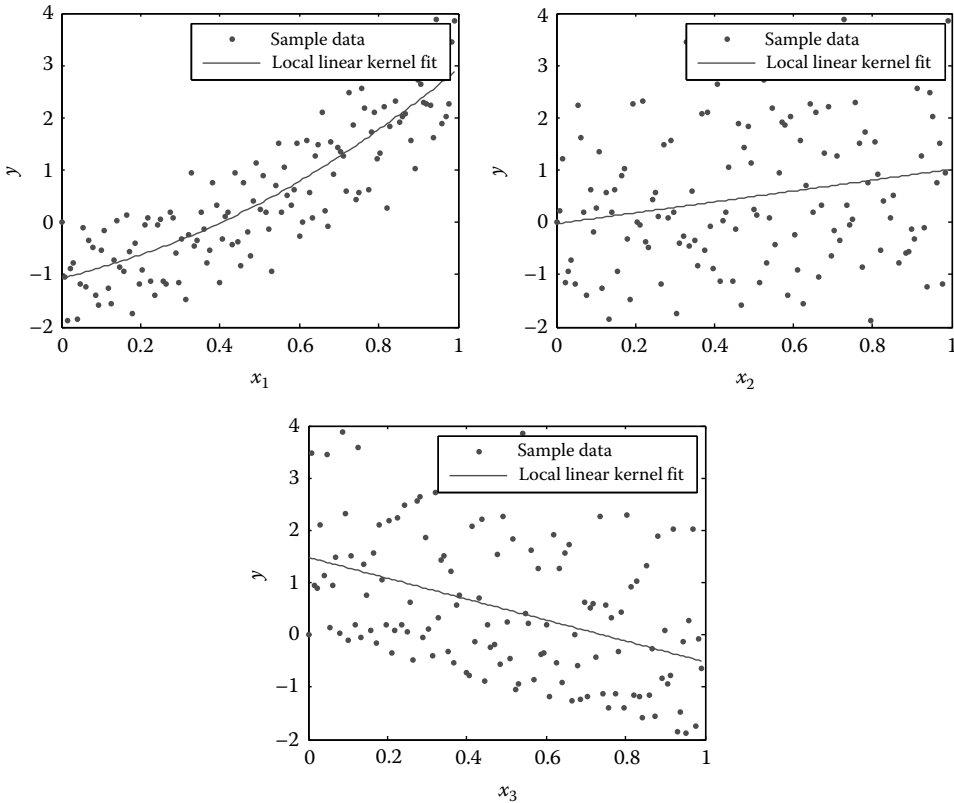


FIGURE 18.22 Local linear kernel regression applied to the polynomial function.

TABLE 18.3

Comparison of S_i Estimates from Local-Linear Kernel Regression of Polynomial Function against Analytical Values

Variable	\hat{S}_i (Kernel)	S_i (Analytic)
x_1	0.7345	0.7568
x_2	0.0494	0.0456
x_3	0.1821	0.1824

In order to estimate sensitivity from the fitted kernel regression, the estimated variance of the conditional expectation can be calculated using a standard variance identity (the domain of the expected value is explicitly stated as a subscript here for clarity):

$$\text{Var}_{x_i}\{E_{x_{\sim i}}(y|x_i)\} = E_{x_i}\{E_{x_{\sim i}}(y|x_i)^2\} - E_{x_i}\{E_{x_{\sim i}}(y|x_i)\}^2. \tag{18.40}$$

Then denoting $V_i = \text{Var}_{x_i}[E_{x_{\sim i}}(y | x_i)]$ as before, V_i is estimated by the smoothed curve $\hat{m}(x_i)$ to give

$$\hat{V}_i = E_{x_i}\{\hat{m}(x_i)^2\} - E_{x_i}\{\hat{m}(x_i)\}^2. \tag{18.41}$$

To evaluate (18.41), one simply calculates a large number of values of $\hat{m}(x_i)$ at different locations and uses the standard estimator for sample variance for each term.

Note here that since the expectation of a random variable A is defined as $\int A p(A) dA$, the expected values in (18.40) should be evaluated with respect to $p(x_i)$, the probability distribution of x_i . The simplest way of doing this is to make kernel predictions at the same x_i values as the training data. A more sophisticated approach would be to estimate $p(x_i)$ with kernel density estimation or a similar technique. In both cases, the practitioner should ensure that the given data have been sampled with respect to the correct underlying distributions—this may not necessarily be the case if the points have come from an optimisation process, for example.

While the examples here have focused on kernel smoothing, this is by no means the only viable approach to estimating $E_{x_{\sim i}}(y|x_i)$. The problem is essentially an emulation/regression problem in one dimension, which can be tackled by any number of methods such as smoothing splines (see, e.g. Ruppert et al. 2003 or earlier in Section 18.6) or GP (Rasmussen and Williams 2006). Even basic linear regression will provide a good estimate if the data are sufficiently linear. Good references on parametric and nonparametric regression can be found in Bishop (2006) and Ruppert et al. (2003).

The estimation of S_i using smoothing techniques is very closely related to the well-known *coefficient of determination*, R^2 . The *standard* R^2 is the square of the Pearson correlation coefficient and measures linear correlation between two variables. The nonlinear R^2 measure, which measures correlation when nonlinear regression is used, can be stated in general terms as $R^2_{\text{nonlin}} = SS_{\text{reg}}/SS_{\text{tot}}$, where $SS_{\text{reg}} = \sum_j (\hat{m}(x_{i,j}) - \bar{y})^2$ and $SS_{\text{tot}} = \sum_j (y_j - \bar{y})^2$. It can be seen therefore as the ratio of the variance explained by the regression of y on x_i and the total variance of y . Therefore, if the function $m(x_i)$ is equal to $E_{x_{\sim i}}(y|x_i)$, the nonlinear R^2 measure is exactly equivalent to S_i .

Finally, it should be pointed out that while the idea of reducing a multidimensional problem to a series of one-dimensional problems is very appealing from a computational point of view, it is not a *silver bullet* solution. The estimation is dependent on a good

approximation of $E_{x_{\sim i}}(y|x_i)$, which can be difficult to obtain depending on the data and smoothing method used. Moreover, as the dimensionality of the problem increases, trends in scatter plots can be increasingly biased and/or less precise due to variation in other dimensions.

18.8 Conclusions and Discussion

In this chapter, a number of *best practice* methods have been outlined that address many of the various situations that can confront an analyst. It is neither claimed nor intended by the authors that an exhaustive review of all methods has been addressed here, but the reader should have found here enough material to apply or adapt to a practical case, with references to other material to direct the reader to more in-depth descriptions.

It is however useful to summarise some areas not covered by the chapter. What did this chapter leave out? An incomplete list is as follows:

- Gradient-based sensitivity measures (mentioned in Section 18.5.1) use a measure similar to μ_i^* given in (18.28), which is the integral of the squared partial derivative as a measure of sensitivity, which has been shown to have a relationship with S_{Ti} —(see Sobol' and Kucherenko 2009; Kucherenko et al. 2009).
- Polynomial chaos expansions, briefly mentioned in Section 18.6, constitute a significant branch of emulator-based sensitivity analysis but were omitted from the chapter due to space limitations and with the intention of focusing more on the sampling aspect of sensitivity analysis. More information can be found in Sudret (2008).
- Moment independent methods form a class of sensitivity analysis approaches that examine how fixing an input variable modifies the entire empirical probability distribution function of y —(see Borgonovo 2007). The rationale behind these methods is that variance is but one of several possible moments that can indicate sensitivity. Such methods focus on measures of *distance* between the unconditional distribution of y and the conditional distribution of $y|x_i$.
- A sensitivity analysis approach based on plots of how the sample mean varies as successively greater quantiles of the distribution x_i are included is known as the *contribution to the sample mean* (Bolado-Lavin et al. 2009), with a similar method based on the contribution of the sample variance (Tarantola et al. 2012).
- Correlated input variables: variance-based sensitivity analysis (the variance decomposition given in (18.3)) is based on the assumption that the probability distributions of input variables are independent. Quite often, however, a modeller encounters the situation where this assumption is not valid. In such cases, other methods must be considered, although this is a relatively immature field of research and (arguably) no one *definitive* approach exists. Some examples of methods in the literature include (1) a method based on linear regression (Xu and Gertner 2008), (2) a method based on decomposition of the input covariance matrix (Li et al. 2010), (3) a copula-based approach (Kucherenko et al. 2012) and (4) an approach based on grouping inputs into sets which are independent of one another but variables may be correlated inside each group (Jacques et al. 2006).

This chapter has in particular stressed the fact that a design should be explorative, in contrast to approaches that rely on a single OAT experiment, a practice which is so often seen in the literature but so poor at exploring nonlinear non-additive problems (Saltelli and Annoni 2010). It should always be noted, however, that the designs that have been described here only explore the uncertainty due to the input variables of the model. This ignores the extremely significant *structural* uncertainty that results in the model's simplification of reality, which is extremely difficult to quantify.

All the sensitivity analysis techniques here assume some knowledge of the probability distributions of the input variables, from simple ranges in the elementary effects method to full probability distributions in variance-based sensitivity analysis. As any practitioner of sensitivity and uncertainty analysis knows, finding reliable information on distributions of input variables can be the most challenging task in the whole process. Some literature gives details on eliciting distributions from experts (O'Hagan et al. 2006), although this in itself assumes that there are knowledgeable experts available to consult. This issue is left as an open problem.

Lastly, every sensitivity analysis is a case apart (since every model has its idiosyncrasies), and the choice of which technique to apply is highly dependent on factors such as the model runtime, dimensionality, linearity and whether the analyst can access the model and run it at chosen input values or not. The outline that has been presented here should however be helpful in devising a suitable approach to sensitivity analysis for a wide variety of models.

References

- Atkinson, A., Donev, A. and Tobias, R. (2007), *Optimum Experimental Designs, with SAS*, Vol. 34, Oxford University Press, New York.
- Batterbee, D., Sims, N., Becker, W., Worden, K. and Rowson, J. (2011), Computational model of an infant brain subjected to periodic motion simplified modelling and Bayesian sensitivity analysis, *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine* **225**(11), 1036–1049.
- Bayin, S. (2006), *Mathematical Methods in Science and Engineering*, John Wiley & Sons, Hoboken, NJ.
- Becker, W., Oakley, J., Surace, C., Gili, P., Rowson, J. and Worden, K. (2012), Bayesian sensitivity analysis of a nonlinear finite element model, *Mechanical Systems and Signal Processing* **32**, 18–31.
- Becker, W., Rowson, J., Oakley, J., Yoxall, A., Manson, G. and Worden, K. (2011), Bayesian sensitivity analysis of a model of the aortic valve, *Journal of Biomechanics* **44**(8), 1499–1506.
- Becker, W., Worden, K. and Rowson, J. (2013), Bayesian sensitivity analysis of bifurcating nonlinear models, *Mechanical Systems and Signal Processing* **34**(1), 57–75.
- Bishop, C. (2006), *Pattern Recognition and Machine Learning*, Springer, New York.
- Bolado-Lavin, R., Castaings, W. and Tarantola, S. (2009), Contribution to the sample mean plot for graphical and numerical sensitivity analysis, *Reliability Engineering & System Safety* **94**(6), 1041–1049.
- Borgonovo, E. (2007), A new uncertainty importance measure, *Reliability Engineering & System Safety* **92**(6), 771–784.
- Bowman, A. and Azzalini, A. (1997), *Applied Smoothing Techniques for Data Analysis: The Kernel Approach with S-Plus Illustrations*, Vol. 18, Oxford University Press, New York.
- Budne, T. A. (1959), The application of random balance designs, *Technometrics* **1**(2), 139–155.
- Campolongo, F., Cariboni, J. and Saltelli, A. (2007), An effective screening design for sensitivity analysis of large models, *Environmental Modelling and Software* **22**(10), 1509–1518.
- Campolongo, F., Saltelli, A. and Campolongo, J. (2011), From screening to quantitative sensitivity analysis. A unified approach, *Computer Physics Communication* **182**(4), 978–988.

- Carleson, L. (1966), On convergence and growth of partial sums of Fourier series, *Acta Mathematica* **116**(1), 135–157.
- Chaloner, K. and Verdinelli, I. (1995), Bayesian experimental design: A review, *Statistical Science* **10**, 273–304.
- Cukier, R. I., Fortuin, C., Schuler, K. E., Petschek, A. G. and Schaibly, J. (1973), Study of the sensitivity of coupled reaction systems to uncertainties in rate coefficients. I theory, *The Journal of Chemical Physics* **59**, 3873–3878.
- Cukier, R., Levine, H. and Shuler, K. (1978), Nonlinear sensitivity analysis of multiparameter model systems, *Journal of Computational Physics* **26**(1), 1–42.
- Draper, N. R. and Smith, H. (1981), *Applied Regression Analysis*, John Wiley & Sons, New York.
- European Commission (2012), Sensitivity analysis software—Joint Research Centre, <https://ec.europa.eu/jrc/en/scientific-tool/simlab-software-sensitivity-and-uncertainty-analysis>. Accessed: January 27, 2015.
- Gramacy, R. B. and Lee, H. K. (2009), Adaptive design and analysis of supercomputer experiments, *Technometrics* **51**(2), 130–145.
- Gramacy, R. B. and Taddy, M. A. (2010), Categorical inputs, sensitivity analysis, optimization and importance tempering with tgp version 2, an R package for treed Gaussian process models, *Journal of Statistical Software* **33**(6), 1–48.
- Gu, C. (2002), *Smoothing Spline ANOVA Models*, Springer, New York.
- Homma, T. and Saltelli, A. (1996), Importance measures in global sensitivity analysis of nonlinear models, *Reliability Engineering & System Safety* **52**(1), 1–17.
- Jacques, J., Lavergne, C. and Devictor, N. (2006), Sensitivity analysis in presence of model uncertainty and correlated inputs, *Reliability Engineering & System Safety* **91**(10), 1126–1134.
- Jansen, M. (1999), Analysis of variance designs for model output, *Computer Physics Communications* **117**, 35–43.
- Kennedy, M. C. and O’Hagan, A. (2001), Bayesian calibration of computer models, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **63**(3), 425–464.
- Koda, M., Mcrae, G. J. and Seinfeld, J. H. (1979), Automatic sensitivity analysis of kinetic mechanisms, *International Journal of Chemical Kinetics* **11**(4), 427–444.
- Kucherenko, S., Feil, B., Shah, N. and Mauntz, W. (2011), The identification of model effective dimensions using global sensitivity analysis, *Reliability Engineering & System Safety* **96**(4), 440–449.
- Kucherenko, S., Rodriguez-Fernandez, M., Pantelides, C. and Shah, N. (2009), Monte carlo evaluation of derivative-based global sensitivity measures, *Reliability Engineering & System Safety* **94**(7), 1135–1148.
- Kucherenko, S., Tarantola, S. and Annoni, P. (2012), Estimation of global sensitivity indices for models with dependent variables, *Computer Physics Communications* **183**(4), 937–946.
- Leamer, E. E. (1985), Sensitivity analyses would help, *The American Economic Review* **75**(3), 308–313.
- Leamer, E. E. (2010), Tantalus on the road to asymptopia, *Journal of Economic Perspectives* **24**(2), 31–46.
- Li, G., Rabitz, H., Yelvington, P. E., Oluwole, O. O., Bacon, F., Kolb, C. E. and Schoendorf, J. (2010), Global sensitivity analysis for systems with independent and/or correlated inputs, *The Journal of Physical Chemistry A* **114**(19), 6022–6032.
- Li, G., Wang, S.-W., and Rabitz, H. (2002), Practical approaches to construct RS-HDMR component functions, *The Journal of Physical Chemistry A* **106**(37), 8721–8733.
- Li, G., Wang, S.-W., Rabitz, H., Wang, S. and Jaffé, P. (2002), Global uncertainty assessments by high dimensional model representations (HDMR), *Chemical Engineering Science* **57**(21), 4445–4460.
- Lin, Y. and Zhang, H. (2006), Component selection and smoothing in smoothing spline analysis of variance models, *Annals of Statistics* **34**(5), 2272–2297.
- Linkletter, C., Bingham, D., Hengartner, N., Higdon, D. and Ye, K. (2006), Variable selection for Gaussian process models in computer experiments, *Technometrics* **48**(4), 478–490.

- Magnus, J. R. (2007), Local sensitivity in econometrics, *Measurement in Economics: A Handbook*, ed. M. Boumans, pp. 295–319, Academic Press.
- Managing Uncertainties in Complex Models (MUCM) Toolbox* (2013), <http://mucm.aston.ac.uk/MUCM/MUCMToolkit/>. Accessed: January 27, 2015.
- Mara, T. A. (2009), Extension of the RBD-FAST method to the computation of global sensitivity indices, *Reliability Engineering and System Safety* **94**(8), 1274–1281.
- Moon, H., Dean, A. M. and Santner, T. J. (2012), Two-stage sensitivity-based group screening in computer experiments, *Technometrics* **54**(4), 376–387.
- Morris, M. (1991), Factorial sampling plans for preliminary computational experiments, *Technometrics* **33**(2), 161–174.
- Morris, M. D., Moore, L. M. and McKay, M. D. (2008), Using orthogonal arrays in the sensitivity analysis of computer models, *Technometrics* **50**(2), 205–215.
- Murphy, J. M., Sexton, D. M. H., Jenkins, G. J., Booth, B. B. B., Brown, C. C., Clark, R. T., Collins, M., Harris, G. R., Kendon, E. J., Betts, R. A., Brown, S. J., Howard, T. P., Humphrey, K. A., McCarthy, M. P., McDonald, R. E., Stephens, A., Wallace, C., Warren, R., Wilby, R., Wood, R. A. (2009), UK climate projections science report: Climate change projections, Met Office, Hadley Centre, Exeter.
- Nadaraya, E. (1964), On estimating regression, *Teoriya Veroyatnostei i ee Primeneniya* **9**(1), 157–159.
- Niederreiter, H. (1992), *Quasi-Monte Carlo Methods*, John Wiley & Sons, Philadelphia, PA.
- Nyquist, H. (1928), Certain topics in telegraph transmission theory, *American Institute of Electrical Engineers* **47**(2), 617–644.
- Oakley, J. and O'Hagan, A. (2004), Probabilistic sensitivity analysis of complex models: A Bayesian approach, *Journal of the Royal Statistical Society B* **66**, 751–769.
- O'Hagan, A., Buck, C. E., Daneshkhah, A., Eiser, J. R., Garthwaite, P. H., Jenkinson, D. J., Oakley, J. E. and Rakow, T. (2006), *Uncertain Judgements: Eliciting Experts Probabilities*, John Wiley & Sons, Chichester, U.K.
- Paruolo, P., Saisana, M. and Saltelli, A. (2013), Ratings and rankings: Voodoo or science?, *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **176**(3), 609–634.
- Pearson, K. (1895), Notes on regression and inheritance in the case of two parents, *Proceedings of the Royal Society of London* **58**, 240–242.
- Pujol, G. (2009), Simplex-based screening designs for estimating metamodels, *Reliability Engineering and System Safety* **94**(7), 1156–1160.
- Rabitz, H. and Aliş, Ö. F. (1999), General foundations of high-dimensional model representations, *Journal of Mathematical Chemistry* **25**(2–3), 197–233.
- Rasmussen, C. and Williams, C. (2006), *Gaussian Processes for Machine Learning*, The MIT Press, Cambridge, MA.
- Ratto, M. and Pagano, A. (2010), Using recursive algorithms for the efficient identification of smoothing spline anova models, *ASTA Advances in Statistical Analysis* **94**(4), 367–388.
- Ratto, M., Pagano, A. and Young, P. (2007), State dependent parameter metamodeling and sensitivity analysis, *Computer Physics Communications* **177**, 863–876.
- Rougier, J., Goldstein, M. and House, L. (2013), Second-order exchangeability analysis for multimodel ensembles, *Journal of the American Statistical Association* **108**(503), 852–863.
- Ruppert, D., Wand, M. and Carroll, R. (2003), *Semiparametric Regression*, Vol. 12, Cambridge University Press, New York.
- Saltelli, A. (2002), Making best use of model evaluations to compute sensitivity indices, *Computer Physics Communications* **145**(2), 280–297.
- Saltelli, A. and Annoni, P. (2010), How to avoid a perfunctory sensitivity analysis, *Environmental Modelling and Software* **25**(12), 1508–1517.
- Saltelli, A., Annoni, P., Azzini, I., Campolongo, F., Ratto, M. and Tarantola, S. (2010), Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index, *Computer Physics Communications* **181**(2), 259–270.

- Saltelli, A. and D'Hombres, B. (2010), Sensitivity analysis didn't help. A practitioner's critique of the Stern review, *Global Environmental Change* **20**(2), 298–302.
- Saltelli, A. and Funtowicz, S. (2013), When all models are wrong, *Issues in Science and Technology* **30**(2), 79–85.
- Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M. and Tarantola, S. (2008), *Global Sensitivity Analysis—The Primer*, John Wiley & Sons, Hoboken, NJ.
- Saltelli, A., Tarantola, S., Campolongo, F. and Ratto, M. (2004), *Sensitivity Analysis in Practice. A Guide to Assessing Scientific Models*, John Wiley & Sons, Hoboken, NJ.
- Saltelli, A., Tarantola, S. and Chan, K. (1999), Quantitative model-independent method for global sensitivity analysis of model output, *Technometrics* **41**(1), 39–56.
- Satterthwaite, F. (1959), Random balance experimentation, *Technometrics* **1**(2), 111–137.
- Schaibly, J. H. and Shuler, K. E. (1973), Study of the sensitivity of coupled reaction systems to uncertainties in rate coefficients. II Applications, *The Journal of Chemical Physics* **59**, 38–79.
- Sobol', I. M. (1967), On the distribution of points in a cube and the approximate evaluation of integrals, *USSR Computational Mathematics and Mathematical Physics* **7**(4), 86–112.
- Sobol', I. M. (1976), Uniformly distributed sequences with an additional uniform property, *USSR Computational Mathematics and Mathematical Physics* **16**(5), 236–242.
- Sobol', I. M. (1993), Sensitivity estimates for nonlinear mathematical models, *Mathematical Modeling and Computational Experiment* **1**(4), 407–414.
- Sobol', I. M. and Kucherenko, S. (2009), Derivative based global sensitivity measures and their link with global sensitivity indices, *Mathematics and Computers in Simulation* **79**(10), 3009–3017.
- Sobol', I. M. and Kucherenko, S. S. (2005), On global sensitivity analysis of quasi-Monte Carlo algorithms, *Monte Carlo Methods and Applications* **11**(1), 83–92.
- Sobol', I. M., Tarantola, S., Gatelli, D., Kucherenko, S. S. and Mauntz, W. (2007), Estimating the approximation error when fixing unessential factors in global sensitivity analysis, *Reliability Engineering & System Safety* **92**(7), 957–960.
- Storlie, C. and Helton, J. (2008), Multiple predictor smoothing methods for sensitivity analysis: Description of techniques, *Reliability Engineering and System Safety* **93**(1), 28–54.
- Strong, M., Oakley, J. E. and Chilcott, J. (2012), Managing structural uncertainty in health economic decision models: A discrepancy approach, *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **61**(1), 25–45.
- Sudret, B. (2008), Global sensitivity analysis using polynomial chaos expansions, *Reliability Engineering and System Safety* **93**(7), 964–979.
- Tarantola, S., Gatelli, D. and Mara, T. (2006), Random balance designs for the estimation of first order global sensitivity indices, *Reliability Engineering and System Safety* **91**(6), 717–727.
- Tarantola, S., Kopustinskas, V., Bolado-Lavin, R., Kaliačka, A., Ušpuras, E. and Vaišnoras, M. (2012), Sensitivity analysis using contribution to sample variance plot: Application to a water hammer model, *Reliability Engineering and System Safety* **99**, 62–73.
- Tebaldi, C. and Knutti, R. (2007), The use of the multi-model ensemble in probabilistic climate projections, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **365**(1857), 2053–2075.
- Tissot, J.-Y. and Prieur, C. (2012), Bias correction for the estimation of sensitivity indices based on random balance designs, *Reliability Engineering and System Safety* **107**, 205–213.
- Tortorelli, D. A. and Michaleris, P. (1994), Design sensitivity analysis: Overview and review, *Inverse Problems in Engineering* **1**(1), 71–105.
- Watson, G. (1961), A study of the group screening method, *Technometrics* **3**(3), 371–388.
- Watson, G. (1964), Smooth regression analysis, *Sankhyā: The Indian Journal of Statistics, Series A* **26**(4), 359–372.
- Welch, W. J., Buck, R. J., Sacks, J., Wynn, H. P., Mitchell, T. J. and Morris, M. D. (1992), Screening, predicting, and computer experiments, *Technometrics* **34**(1), 15–25.
- Xu, C. and Gertner, G. (2011), Understanding and comparisons of different sampling approaches for the Fourier amplitudes sensitivity test (FAST), *Computational Statistics and Data Analysis* **55**, 184–198.

- Xu, C. and Gertner, G. Z. (2008), Uncertainty and sensitivity analysis for models with correlated parameters, *Reliability Engineering and System Safety* **93**(10), 1563–1573.
- Youden, W., Kempthorne, O., Tukey, J. W., Box, G. and Hunter, J. (1959), Discussion of the papers of Messrs. Satterthwaite and Budne, *Technometrics* **1**(2), 157–184.
- Zuniga, M. M., Kucherenko, S. and Shah, N. (2013), Metamodelling with independent and dependent inputs, *Computer Physics Communications* **184**(6), 1570–1580.

19

Expected Improvement Designs

William I. Notz

CONTENTS

19.1	Introduction.....	676
19.1.1	Goals.....	676
19.1.2	Modeling.....	676
19.1.3	Prediction.....	679
19.1.4	Design.....	680
19.2	Expected Improvement as an Approach to Designing Computer Experiments.....	682
19.2.1	Example: Global Minimization.....	683
19.2.2	General Approach.....	685
19.3	Global Optimization.....	686
19.3.1	Schonlau, Schonlau, and Jones' Method Continued.....	686
19.3.2	Batch Sequential Design.....	688
19.3.3	Global Optimization for Control and Environmental Inputs.....	688
19.3.4	Multiple Responses.....	691
19.3.4.1	Constrained Global Optimization.....	692
19.3.4.2	Pareto Optimization.....	694
19.3.4.3	Outline of the Algorithm.....	696
19.3.4.4	0–1 Improvement.....	698
19.3.4.5	Keane's Distance-Based Improvement.....	699
19.3.4.6	Pareto Improvement.....	699
19.3.4.7	Hypervolume Improvement.....	700
19.3.5	GASP Model with Measurement Error.....	701
19.3.5.1	Multiple Simulators of Differing Fidelity.....	703
19.4	Contour Estimation.....	706
19.5	Percentile Estimation.....	707
19.5.1	Confidence Interval: Based Criterion.....	708
19.5.2	Hypothesis Testing–Based Criterion.....	709
19.6	Global Fit.....	711
19.6.1	Emulator Maturity.....	711
19.7	Questions for Further Study.....	712
19.7.1	Additional Improvement Criteria and Expected Improvement Algorithms.....	712
19.7.2	Suite of Test Functions.....	713
19.7.3	Numerical Issues.....	714
19.8	Concluding Remarks.....	714
	References.....	714

19.1 Introduction

In the study of complex physical and engineering processes, it is sometimes impossible, infeasible, or very expensive to observe the actual systems. Examples of such processes include climate change, spread of infectious diseases, and car crashes—(see Jones et al. 1998 and Santner et al. 2003). In such situations, we can sometimes use computer codes (simulators) to mimic the process that we desire to investigate. Values of the input variables are supplied to the computer code, and the code produces the corresponding values of the output variable. To study the nature of the relationship between the input variables and the output variable, it would be ideal if we could observe the code output for all possible values of the input variables. Often, however, the computer code itself is very complex and each run of the simulator turns out to be highly expensive or very time-consuming. This limits the number of times we can run the simulator. In what follows, we assume the computer code runs slowly and that the number of times we can observe the simulator is “small.”

Experimenting with the simulator rather than the actual physical or engineering process the simulator is intended to approximate is called a *computer experiment*. A computer experiment differs from a traditional physical experiment in a number of ways. For one, the output from a computer experiment is (usually) assumed to be deterministic. That is, repeated runs of the simulator at the same settings of the input variables will result in identical values of the output. Thus, random errors are (usually) not a concern in a computer experiment. Also, in a physical experiment, all the factors that affect the response may not be known. In a computer experiment, the output comes from a simulator; hence, all the factors that affect the output are known. Owing to these differences, the principles of blocking, randomization, and replication that are of great importance in the traditional designs of experiments are not relevant to computer experiments.

19.1.1 Goals

There are a variety of goals in the area of computer experiments, such as code validation (i.e., ensuring that the code is a reasonable approximation to the physical process), calibration (finding values of certain *tuning* parameters in the code that produce output that best approximates physical data), the investigation of the overall input–output relationship (finding a predictor, or *emulator*, that provides good overall fit to the output), uncertainty analysis, sensitivity analysis, global optimization, determination of all values of the inputs that produce a given value of the output, and estimation of percentiles of the distribution of the output for a given distribution of the inputs. See Santner et al. (2003) and Bayarri et al. (2007) for further discussion.

19.1.2 Modeling

We use a statistical model to estimate the input–output relationship, so that the fitted model (or emulator) can serve as a relatively inexpensive substitute for the complex and expensive computer simulator. Since we are dealing with output from computer code, the responses are considered to be deterministic (but see Section 19.3.5 where we consider simulators that

produce output with noise). Due to the lack of random error, Sacks et al. (1989) propose that any uncertainty in the code output $y(\mathbf{x})$ at any d -dimensional input $\mathbf{x} = (x_1, x_2, \dots, x_d)'$ can be assumed to be a result of systematic departure from a regression model; so the input-output relationship can be modeled by the following equation:

$$Y(\mathbf{x}) = \sum_{j=1}^k \beta_j f_j(\mathbf{x}) + Z(\mathbf{x}), \tag{19.1}$$

where $\sum_{j=1}^k \beta_j f_j(\mathbf{x})$ is the model for the mean and $Z(\mathbf{x})$ is the systematic departure (see Sacks et al. 1989). We assume that $Z(\mathbf{x})$ is a stochastic process with $E(Z(\mathbf{x})) = 0$, $\text{Var}(Z(\mathbf{x})) = \sigma_z^2$, and that the covariance between $Z(\mathbf{x}_1)$ and $Z(\mathbf{x}_2)$ is given by

$$\text{cov}(Z(\mathbf{x}_1), Z(\mathbf{x}_2)) = \sigma_z^2 R(\mathbf{x}_1, \mathbf{x}_2; \lambda),$$

where $R(\mathbf{x}_1, \mathbf{x}_2; \lambda)$ denotes the correlation between $Z(\mathbf{x}_1)$ and $Z(\mathbf{x}_2)$ and this correlation may depend on one or more parameters, denoted here by λ . It is popular in computer experiments to further assume $Z(\mathbf{x})$ is a stationary Gaussian stochastic process (GASP), with unknown variance, σ_z^2 , so that

$$\begin{aligned} \text{cov}(Z(\mathbf{x}_i), Z(\mathbf{x}_j)) &= \sigma_z^2 R(\mathbf{x}_i, \mathbf{x}_j; \lambda) = \sigma_z^2 R(\mathbf{x}_i - \mathbf{x}_j; \lambda), \\ &\text{for } \mathbf{x}_i, \mathbf{x}_j \in \mathcal{X}, \end{aligned}$$

and $\mathcal{X} \subset \mathbb{R}^d$ is the experimental or design region. Here \mathbb{R}^k denotes k -dimensional Euclidean space. We refer to this model as the GASP model.

In practice, it is not uncommon to use a special case of the GASP model, modeling the responses as realizations of a constant mean stationary GASP, that is,

$$Y(\mathbf{x}) = \beta + Z(\mathbf{x}) \quad \text{for } \mathbf{x} \in \mathcal{X}, \tag{19.2}$$

where β is the unknown overall mean and $Z(\mathbf{x})$ is a zero-mean stationary GASP with unknown variance, σ_z^2 . We refer to this model as the *GASP model with constant mean*. Empirical evidence (see Sacks et al. 1989) shows that any simplification in the usual regression part, that is, $\sum_{j=1}^k \beta_j f_j(\mathbf{x})$, is compensated for by the systematic departure part, that is, $Z(\mathbf{x})$. This, typically, holds when we are using an interpolator and the true responses are realizations of a stationary process.

At this point, when using the GASP model with constant mean, we have three unknowns in our model: the overall mean β , the process variance σ_z^2 , and the correlation function $R(\cdot, \cdot; \lambda)$. Because we assume that $Z(\mathbf{x})$ is a stationary stochastic process, the correlation

between two points $x_i = (x_{i1}, x_{i2}, \dots, x_{id})'$ and $x_j = (x_{j1}, x_{j2}, \dots, x_{jd})'$ is a function of x_i and x_j through $x_i - x_j$. Also, we assume that the product correlation rule is used. Thus,

$$R(x_i, x_j; \lambda) = R(x_i - x_j; \lambda) = \prod_{t=1}^d R_t(x_{it} - x_{jt}; \lambda_t),$$

where $\lambda = (\lambda'_1, \dots, \lambda'_d)'$ with λ_t the correlation parameters for the t th input variable and $R_t(\cdot; \lambda_t)$, itself a correlation function defined on the set of real numbers \mathbb{R} and dependent on λ_t . In general, λ_t will be different for different coordinates of x . However, if we believe our GASP model is *isotropic*, that is, the one-dimensional correlation structure is the same in every direction, then $\lambda_1 = \dots = \lambda_d$. For an example of the use of an isotropic model, see Welch et al. (1992).

When specifying the correlation structure, we require $R(x, x) = 1$ and that $R(\cdot; \cdot; \lambda)$ is positive semidefinite, that is,

$$\sum_{i,j=1}^n c_i c_j R(x_i, x_j; \lambda) \geq 0,$$

for all finite n , all x_1, \dots, x_n , and all real c_1, \dots, c_n .

For a stationary process with a product correlation structure, one very popular correlation function, called the power exponential correlation function, takes the form

$$R_t(x_{it} - x_{jt}; \theta_t, p_t) = \exp(-\theta_t |x_{it} - x_{jt}|^{p_t}), \tag{19.3}$$

where $0 < p_t \leq 2$ and $\theta_t \in (0, \infty)$. Here $\lambda_t = (\theta_t, p_t)'$. For the special case of $p_t = 2$ for all t , this corresponds to the Gaussian correlation function. Taking $p_t = 1$ for all t gives the exponential correlation function. As θ_t increases, the dependence between two sites decreases but does not go to zero. See Sacks et al. (1989) for an application with this correlation function.

Another popular correlation function is the cubic correlation function that takes the form

$$\begin{aligned} R_t((x_{it}, x_{jt}); \theta_t) &= R_t((x_{it} - x_{jt}); \theta_t) \\ &= 2 \left(1 - \frac{h_{ijt}}{\theta_t}\right)^3 1_{\left(\frac{\theta_t}{2} < h_{ijt} < \theta_t\right)} \\ &\quad + \left[1 - 6 \left(\frac{h_{ijt}}{\theta_t}\right)^2 + 6 \left(\frac{h_{ijt}}{\theta_t}\right)^3\right] 1_{\left(h_{ijt} < \frac{\theta_t}{2}\right)}, \end{aligned} \tag{19.4}$$

where $h_{ijt} = |x_{it} - x_{jt}|$, that is, the distance between the i th and j th observation on the t th dimension, for $i, j = 1, \dots, n$ and $t = 1, \dots, d$, and $1_{(\cdot)}$ is the indicator function: if A is true, then $1_A = 1$; otherwise, $1_A = 0$. The θ_t 's are the correlation parameters (here $\lambda_t = \theta_t$); small θ_t 's imply that only data with very similar values of x are used in prediction. Points (in the t th dimension) that are farther apart than θ_t units have zero correlation.

Under this setup, the parameters that need to be estimated are β , σ_z^2 , and the correlation parameters $\lambda = (\lambda_1, \dots, \lambda_d)'$. Several approaches are possible such as maximum likelihood (ML), restricted ML, and Bayesian approaches. Bayesian methods are becoming increasingly popular, and for a discussion of a fully Bayesian approach, see, for example, Higdon et al. (2003).

For purposes of this article, we use ML estimation to estimate the parameters. The ML estimates can be written as

$$\hat{\beta} = \hat{\beta}(\lambda) = (\mathbf{1}'_n \mathbf{R}^{-1} \mathbf{1}_n)^{-1} \mathbf{1}'_n \mathbf{R}^{-1} \mathbf{y}, \tag{19.5}$$

and

$$\hat{\sigma}_z^2 = \hat{\sigma}_z^2(\lambda) = \frac{1}{n} (\mathbf{y} - \mathbf{1}_n \hat{\beta})' \mathbf{R}^{-1} (\mathbf{y} - \mathbf{1}_n \hat{\beta}), \tag{19.6}$$

where $\mathbf{1}_n$ is an $n \times 1$ vector of 1s, \mathbf{Y} denotes all the available responses for the current design, and \mathbf{R} is the $n \times n$ matrix whose i, j th entry is $R(x_i, x_j; \lambda)$, the correlation between the response at the design sites x_i and x_j . We substitute $\hat{\beta}$ and $\hat{\sigma}_z^2$ in the likelihood and, as in Jones et al. (1998), use the profile likelihood to find estimates of the correlation parameters, $\lambda_1, \dots, \lambda_d$, and thus an estimate of \mathbf{R} .

The GASP model assumes stationarity of the stochastic process. Realizations of a stationary GASP with a correlation function such as the Gaussian or cubic correlation function will display a certain regularity in terms of variation in the values of $Y(x)$ and the *frequency* of variation (sometimes referred to as the *roughness* of the realization, and a parameter controlling this frequency of variation, such as the θ_t in (19.3), is sometimes called a roughness parameter). For example, when $d = 1$, the realizations remind one of the behavior of a periodic function with constant amplitude and frequency. The amplitude is determined by the process variance σ_z^2 . If, for example, the amplitude increases as x increases, this suggests that the process variance σ_z^2 depends on x and hence the process is nonstationary. The frequency, or roughness, is determined by the rate at which the correlation decreases as a function of the distance $|x_1 - x_2|$ between two points x_1 and x_2 . The frequency is small (and the realization appears to be relatively smooth) when the correlation between two points x_1 and x_2 decreases slowly as $|x_1 - x_2|$ increases. The frequency is high (and the realization appears to be “rougher”) when the correlation between two points x_1 and x_2 decreases rapidly as $|x_1 - x_2|$ increases. In the limit as the correlation between any two points $x_1 \neq x_2$ goes to 0, the realization looks very rough or *noisy*. If the frequency depends on the location of x_1 and x_2 rather than just $|x_1 - x_2|$, so that some regions exhibit high frequency variation and others low frequency variation, this suggests the process is nonstationary. In theory, the GASP model would seem to be ill suited for modeling functions that look nonstationary, and a challenging research area is investigating nonstationary models that are computationally tractable.

19.1.3 Prediction

Consider any untried input setting \mathbf{x}^* . Then, in the GASP model with constant mean, the best linear unbiased predictor (BLUP) of the response at \mathbf{x}^* is given by

$$\hat{Y} = \hat{Y}(\mathbf{x}^*) = \hat{\beta} + \mathbf{r}' \mathbf{R}^{-1} (\mathbf{y} - \mathbf{1}'_n \hat{\beta}), \tag{19.7}$$

where $\hat{\beta}$ is as in (19.5) and is the usual generalized least-squares estimator of β and $r = r(x^*) = [R(x^*, x_1; \lambda), \dots, R(x^*, x_n; \lambda)]'$ is the vector of correlations between x^* and the observed input sites. The mean squared prediction error (MSPE), denoted by $s^2(x^*)$, when β is unknown but σ_z^2 and λ are known, is given by

$$s^2(x^*) = \sigma_z^2 \left(1 - r' R^{-1} r + \frac{(1 - \mathbf{1}'_n R^{-1} r)^2}{\mathbf{1}'_n R^{-1} \mathbf{1}_n} \right). \tag{19.8}$$

In practice r and R are not known, and their estimates are used instead. In such situations, the BLUP is called the empirical BLUP (EBLUP). The behavior of the EBLUP depends on the correlation function one assumes. For the power exponential correlation function with $0 < p < 2$, the EBLUP is continuous but not differentiable. For the Gaussian correlation function ($p = 2$ in the power exponential correlation function), the EBLUP is infinitely differentiable. This would be desirable if one knows that the physical process being modeled by the computer simulator is smooth. The cubic correlation function is known to produce a predictor that is a cubic spline interpolator (see Currin et al. 1991). Compared with the Gaussian correlation function, the cubic correlation function seems less prone to numerical problems (perhaps because for small values of the correlation parameter, the correlation matrix will have many zero entries).

One reason for the popularity of the GASP model and the EBLUP as a statistical predictor (or emulator) of output from a computer simulator is that the EBLUP interpolates the data. In other words, for points at which we have observed the simulator (and have used in constructing the EBLUP), the EBLUP predicts the observed values. For deterministic simulators that yield the same output for repeated evaluations at a given set of inputs, it seems reasonable to insist that a statistical predictor be an interpolator.

Although one might argue that one should not use a stationary GASP model for a simulator that produces output that appears nonstationary, this may not be a serious problem in practice. This is because the EBLUP is an interpolator and with sufficient data should be a good approximation to the actual output, even if the output does not look stationary. See Lam and Notz (2008) for an example where the EBLUP based on the GASP model with constant mean does a good job of approximating a function that looks nonstationary. In our experience, if the points at which the simulator is observed are selected carefully, nonstationary-looking functions can be reasonably well approximated by the EBLUP even with moderate sample sizes.

19.1.4 Design

For computer codes that can be run a relatively small number of times, the inputs at which to observe the code must be chosen carefully. The selection of the input settings at which the computer code is to be observed is the design of the computer experiment. If we observe the code at the set $S_n = \{x_1, x_2, \dots, x_n\}$ of n points in \mathcal{X} , we will refer to S_n as the design. What strategies have been proposed for selecting designs in computer experiments?

One strategy is to select a design that spreads points “evenly” throughout \mathcal{X} . Such designs are sometimes called *space-filling* designs. Examples include Latin hypercube designs, uniform designs, maximin distance designs, minimax distance designs, and designs used in quasi-Monte Carlo methods. Intuitively, the idea is that the shape

of the true response surface (the functional form of the simulator output as a function of the inputs) is unknown and that important features of the response surface may occur anywhere in \mathcal{X} . Thus, it is important to observe the output in all regions of \mathcal{X} . It may also be desirable that projections of the design onto subspaces of \mathcal{X} remain space filling. Thus, if the output is found to be insensitive to some inputs, the design when projected onto the subspace of \mathcal{X} that does not include these inputs will be space filling for purposes of fitting a response surface that does not include these inputs. See Chapter 17 for more about Latin hypercube and space-filling designs.

Another strategy for selecting a design is that it be orthogonal or nearly orthogonal. Orthogonality is usually applied in a subclass of space-filling designs. For example, one might select an orthogonal design within the set of Latin hypercube designs. Orthogonal designs themselves are often space-filling, so the property of orthogonality may be compatible with various space filling criteria. In the traditional linear model, orthogonality allows independent estimates of model effects, and this may be useful in computer experiments if the goal is to conduct a sensitivity analysis.

In general, the selection of a design should depend on (1) the goals of the experiment and (2) the statistical model used to analyze the data from the experiment. In traditional optimal design, this is accomplished through a design criterion. A design criterion is a function from the space of designs to the positive real numbers. Typically, the smaller the value of this function, the better the design for achieving the goals of the experiment, and a design achieving the minimum value of this function is called an *optimal design*. An example of a popular design criterion in regression is the *D*-optimality criterion.

Many popular designs for computer experiments, such as space-filling designs, are “fixed” designs in the sense that the design fixes in advance all the runs (sets of inputs at which the computer code is to be observed). However, for designs based on criteria such as entropy, MSPE, or *D*-optimality, knowledge of the unknown parameters (e.g., the correlation parameters) is required in order to evaluate the criterion. One approach for such designs is to select a locally optimal design, namely, a design that is optimal under the design criterion for specific values of the unknown parameters. A second, related, approach is to place a prior distribution on the unknown parameters and use the expected value of the criterion with respect to this distribution. A third possibility is to adopt a sequential design strategy. Based on a small initial design, estimate the values of the unknown parameters and then select observations for the next stage based on these estimates. This can be repeated several times, each time updating estimates of the unknown parameters. A sequential approach has other advantages, allowing one to use what one learns about the true response surface to guide the selection of the next design point(s). In regression, this might include modifying the order of the response surface one is fitting and selecting additional points to facilitate fitting this new model. In the context of computer experiments, this is important. The actual response is the result of running a computer simulation. The statistical model we use (here we assume this is the GASP model) is simply a device for developing a useful statistical predictor (in particular, a predictor that is a smooth interpolator of the data). A sequential approach allows one to sample more heavily from areas of interest, for example, if the true response is nonstationary, regions in which the response varies more rapidly might benefit from a more dense sample.

We agree with the following comments in Box et al. (1978):

Scientific research is a process of guided learning. The object of statistical methods is to make that process as efficient as possible. (p. 1)

In exploring a functional relationship it might appear reasonable at first sight to adopt a comprehensive approach in which the entire range of every factor was investigated. The resulting design might contain all combinations of several levels of all factors. However, when runs can be made in successive groups, this is an inefficient way to organize experimental programs. The situation relates to the paradox that the best time to design an experiment is after it is finished, the converse of which is that the worst time is at the beginning, when the least is known. If the entire experiment was designed at the outset, the following would have to be assumed known: (1) which variables were the most important, (2) over what ranges the variables should be studied, (3) in what metrics the variables and responses should be considered (e.g., linear, logarithmic, or reciprocal scales), and (4) what multivariable transformations should be made (perhaps the effects of variables x_1 and x_2 would be most simply expressed in terms of their ratio x_1/x_2 and their sum $x_1 + x_2$).

The experimenter is least able to answer such questions at the outset of an investigation but gradually becomes more able to do so as a program evolves.

All the above arguments point to the desirability of a sequence of moderately sized designs and reassessment of the results as each group of experiments becomes available. (p. 303)

As a rough general rule, not more than one quarter of the experimental effort (budget) should be invested in a first design. (p. 304)

It is possible to develop sequential versions of space-filling designs. For certain run sizes and numbers of inputs, one can occasionally add points to a Latin hypercube design in such a way that the result is also a Latin hypercube. Also, one can develop algorithms to add points to an existing space-filling design so that the result *optimizes* some space filling criterion. Although these are sequential designs, this is not what we suggest. We advocate the use of sequential designs for computer experiments in conjunction with a quantitative design criterion consistent with the goals of the experiment.

19.2 Expected Improvement as an Approach to Designing Computer Experiments

Several criteria have been used as *optimality* criteria for selecting designs in computer experiments. These include entropy (Shewry et al. 1987) and maximum and integrated MSPE (Sacks et al. 1989, Lim et al. 2002). These can be used to generate either fixed or sequential designs.

Another approach for generating designs sequentially is the use of an *improvement criterion*. We assume that the computer experiment has a well-defined goal, for example, finding the global optimum of the output of the computer simulator (assuming the output is univariate and real valued). Given n runs of the simulator, we wish to determine the inputs at which to perform the next run. An improvement criterion is a real-valued function of the inputs that reflects the *improvement* that will be achieved in reaching our goal if a given set of inputs is used in the next run. Larger values are better. Ideally, the next run is taken at those inputs that maximize this improvement. Typically, an improvement criterion is a function of $Y(\mathbf{x})$ based on some statistical model for the simulator output, such as the GASP model. Because $Y(\mathbf{x})$ is a random variable, one takes the expected value of the

improvement function given the data and given estimates of the unknown model parameters. Hence, the actual criterion used for selecting the next run is based on the expectation of the improvement criterion, called the *expected improvement*.

For any given goal, there may be many improvement criteria that are sensible. Improvement criteria are analogous to optimality criteria in optimal experimental design. For example, in classical regression, *A*-optimality, *D*-optimality, and *E*-optimality are reasonable design criteria if the goal is to estimate all the regression coefficients with small standard errors. One way of formulating an improvement criterion is the following. Given n runs of the simulator, imagine fitting a statistical predictor $\hat{Y}(x)$ to the data. Based on this predictor and the goals of the experiment, where would you next like to observe the simulator? Presumably, certain regions appear more promising than others based on the response surface defined by the predictor $\hat{Y}(x)$. An attempt to quantify what values of $\hat{Y}(x)$ correspond to promising sites x at which to observe the simulator can lead one to an improvement criterion. For example, if the goal is to find the minimum of the simulator, small values of $\hat{Y}(x)$, in particular values smaller than any observed so far, would be potential sites at which to next observe the simulator. If the goal is overall global fit, perhaps regions where $\hat{Y}(x)$ is changing rapidly (has large absolute value of the derivative) would be potential sites at which to observe the simulator.

The notion of expected improvement was first discussed in Mockus et al. (1978) and Journel and Huijbregts (1989) and more fully developed in Schonlau (1997), Schonlau et al. (1998), and Jones et al. (1998). To further clarify the concept of expected improvement, we look more closely at the development in Schonlau (1997), Schonlau et al. (1998), and Jones et al. (1998).

19.2.1 Example: Global Minimization

In Schonlau (1997), Schonlau et al. (1998), and Jones et al. (1998), the goal is to determine the inputs for which the simulator output is a global minimum. The model for the output is the GASP model with constant mean. We assume that we have observations of the output from the simulator for n runs $\mathbf{y}^n = (y(x_1), \dots, y(x_n))'$ at the n inputs $\{x_1, \dots, x_n\}$ in \mathcal{X} . We refer to these n runs (both \mathbf{y}^n and $\{x_1, x_2, \dots, x_n\}$) as the *training data* and wish to know at what input $x_{n+1} \in \mathcal{X}$ to make the next run of the simulator.

To find $x_{\min} = \arg \min_{x \in \mathcal{X}} y(x)$, where $y(x)$ is an unknown function, we need an improvement criterion that reflects how much improvement we achieve toward finding the global minimum by taking the next observation at x^* , given the training data. The approach developed by Schonlau (1997), Schonlau et al. (1998), and Jones et al. (1998) assumes a Gaussian prior for $y(x)$, say $Y(x)$, of the form on the right-hand side of (19.2) with process variance σ_z^2 and a uniform prior distribution for the parameter β . Let \mathbf{Y}^n denote the prior associated with the vector of outputs \mathbf{y}^n . Their approach is based on the fact that

$$[Y(x) \mid \mathbf{Y}^n = \mathbf{y}^n] \sim N\left(\hat{Y}(x), s^2(x)\right), \tag{19.9}$$

where $\hat{Y}(x)$ is the EBLUP and $s^2(x)$ is the MSPE of the BLUP given in (19.8). The correlation parameters of both quantities are estimated by ML or restricted ML.

Their notion of *improvement* is defined as follows. Let

$$y_{\min}^n = \min_{i=1,\dots,n} y(x_i)$$

denote the minimum output that has been observed after the n runs of the computer code. Consider a *potential* site x^* at which to evaluate the code. Compared with the current smallest *known* minimum value of $y(\cdot)$, define the amount of improvement in $y(\cdot)$ at x to be zero if $y(x^*) \geq y_{\min}^n$, that is, $y(x^*)$ provides no improvement over y_{\min}^n . Similarly, if $y(x^*) < y_{\min}^n$, the amount of improvement at x^* is defined to be the difference $y_{\min}^n - y(x^*)$. Hence, in principle, the improvement possible at x^* is

$$\text{Improvement at } x^* = \begin{cases} y_{\min}^n - y(x^*), & y_{\min}^n - y(x^*) > 0 \\ 0, & y_{\min}^n - y(x^*) \leq 0 \end{cases}.$$

We say “in principle” because $y(x^*)$ is *unknown*, although y_{\min}^n is known from the training data. However, we have an idea of $y(x^*)$ from its prior $Y(x^*)$ (and a better one from the updated prior—the posterior of $Y(x^*)$ given Y^n in (19.9)). Hence, a probabilistically based improvement criterion can be defined by

$$I_n(x^*) = \begin{cases} y_{\min}^n - Y(x^*), & y_{\min}^n - Y(x^*) > 0 \\ 0, & y_{\min}^n - Y(x^*) \leq 0 \end{cases} \tag{19.10}$$

for $x^* \in \mathcal{X}$. The random $I_n(x^*)$ depends solely on the random quantity $Y(x^*)$. We summarize the amount of improvement possible at each potential input site x^* by its *expected improvement* with respect to the posterior distribution of $[Y(x^*)|Y^n]$.

It is straightforward to show that the *conditional expected improvement* satisfies $E\{I_n(x^*) | Y^n\} = 0$ for x^* in the input training data. This result coincides with our intuition that for deterministic simulators, there is no benefit in recomputing the output at sites x^* where we know $y(x^*)$. If x^* is not in the training data, some algebra shows that

$$\begin{aligned} & E\{I_n(x^*) | Y^n\} \\ &= s(x^*) \left\{ \frac{y_{\min}^n - \widehat{Y}(x^*)}{s(x^*)} \Phi\left(\frac{y_{\min}^n - \widehat{Y}(x^*)}{s(x^*)}\right) + \phi\left(\frac{y_{\min}^n - \widehat{Y}(x^*)}{s(x^*)}\right) \right\}, \end{aligned} \tag{19.11}$$

where $\Phi(\cdot)$ and $\phi(\cdot)$ are the $N(0, 1)$ distribution and density function, respectively. By examining the terms in (19.11), we see that the posterior expected improvement is “large” for those x^* having either a predicted value at x^* that is much smaller than the best minimum

computed so far, that is, $\widehat{Y}(\mathbf{x}^*) \ll y_{\min}^n$, or having much uncertainty about the value of $y(\mathbf{x}^*)$, that is, when $s(\mathbf{x}^*)$ is large. Candidate inputs \mathbf{x}^* are judged attractive if there is high probability that their predicted output is below the current observed minimum *and/or* there is a large uncertainty in the predicted output.

Another way to think about this is to imagine constructing a prediction interval for $\widehat{Y}(\mathbf{x}^*)$. If the lower bound of this interval is less than y_{\min}^n , then there is a some possibility that improvement could be achieved at \mathbf{x}^* . The lower bound will be less than y_{\min}^n if either $\widehat{Y}(\mathbf{x}^*) < y_{\min}^n$ or $s(\mathbf{x}^*)$ is sufficiently large. This raises the question of whether something like

$$I_n(\mathbf{x}^*) = \begin{cases} y_{\min}^n - (Y(\mathbf{x}^*) - \alpha s(\mathbf{x}^*)), & y_{\min}^n - (Y(\mathbf{x}^*) - \alpha s(\mathbf{x}^*)) > 0 \\ 0, & y_{\min}^n - (Y(\mathbf{x}^*) - \alpha s(\mathbf{x}^*)) \leq 0 \end{cases}$$

would be a reasonable variation on the improvement criterion, where α is a user-defined parameter.

Our next run of the computer simulation would be at the \mathbf{x}_{n+1} that maximizes $E\{I_n(\mathbf{x}) | \mathbf{Y}^n\}$ as given in (19.11). Once we obtain the output $y(\mathbf{x}_{n+1})$, we can redefine the training data to be $\mathbf{y}^{n+1} = (y(\mathbf{x}_1), \dots, y(\mathbf{x}_{n+1}))'$ at the $n + 1$ inputs $\{\mathbf{x}_1, \dots, \mathbf{x}_{n+1}\}$, recompute the expected improvement based on this training data, and determine the \mathbf{x}_{n+2} that maximizes the expected improvement. If the EBLUP is used for $\widehat{Y}(\mathbf{x}^*)$, one could reestimate the EBLUP at each iteration, or after a small block of iterations. This determines an algorithm for generating a design sequentially. This algorithm can begin with a randomly chosen single point or, more commonly, with an initial design of n_0 runs, for example, an n_0 -run space-filling design.

In examples discussed in Schonlau (1997), Schonlau et al. (1998), and Jones et al. (1998), this algorithm for generating a sequential design based on expected improvement performs quite well. Designs display both *local* and *global* search characteristics. By local search we mean that the algorithm tends to select points near the location of minima, eventually clustering around these minima. By global search we mean that the algorithm occasionally selects points in regions that have not been previously visited. This behavior is consistent with our comments following (19.11). This appears in other settings where improvement criteria have been used and we will discuss this more in subsequent sections.

19.2.2 General Approach

The use of improvement criterion and expected improvement for generating designs in computer experiments has several components. First, we must have a statistical model for the output of the computer experiment. This may be the GASP model, but one could use other models, for example, a standard response surface model. Second, we must have some well-defined goal for the computer experiment, such as finding the global optimum of the output. We examine other goals later in this article. Third, we must have an improvement criterion that is related to our goal. Larger values are better in our example, but the improvement criterion in (19.10) can be transformed into a criterion for finding the global maximum by replacing y_{\min}^n by a corresponding y_{\max}^n and then multiplying by -1 . Fourth, we require training data. In practice, the training data are obtained from some fixed design

such as a Latin hypercube design. However, one can implement a fully sequential procedure that begins with a randomly selected point (without training data, all points in \mathcal{X} provide the same improvement) and then use the improvement criterion for selecting subsequent points. In the author's experience, using a small space-filling initial design has performed better than a fully sequential implementation, but this needs to be explored further. Fifth, we require a stopping rule. This could simply be to stop after a total of N runs, or to stop after some criteria (e.g., the improvement does not exceed some minimal threshold) are met. We consider some stopping criteria in the next section.

These components form the basis for an algorithm for generating a design. In practice, the algorithm will need to be implemented numerically. We next consider several settings in which sequential design via expected improvement has been employed. Our primary purpose is to discuss the types of improvement criteria that have been proposed to meet a variety of goals for a computer experiment. We also hope that this inspires readers to develop new improvement criteria. Thus, in the interest of space, in some cases, we simply present the improvement criterion and refer readers to the appropriate references for details.

19.3 Global Optimization

As mentioned previously, global optimization of an unknown function (or a simulator with a single response) appears to be the first setting to employ expected improvement for generating a design. This use of expected improvement is surprisingly effective for locating the inputs at which an unknown function attains its global optimum. Examples discussed in Schonlau (1997), Schonlau et al. (1998), and Jones et al. (1998) are compelling and include functions with multiple optima. As we mentioned previously, points selected using expected improvement cluster around the locations of the optima with occasional excursions into regions that have not been previously observed, displaying both local and global search characteristics. In particular, the method avoids being trapped in a local optimum. The success of expected improvement in finding the location of a global optimum of an unknown function suggests that using expected improvement criteria might be an effective approach for experimental design for computer experiments in other settings.

The method in Schonlau (1997), Schonlau et al. (1998), and Jones et al. (1998) has been extended to many other settings for finding the location of a global optimum. In fact, global optimization appears to be the most popular goal for implementing improvement criteria. We discuss several applications in this section.

19.3.1 Schonlau, Schonlau, and Jones' Method Continued

The method of Schonlau (1997), Schonlau et al. (1998), and Jones et al. (1998) can be stated as an algorithm with a stopping rule. Starting with a space-filling design, the expected improvement algorithm updates the current training data $S_n = \{x_1, \dots, x_n\}$ as follows.

Given a specified absolute tolerance ϵ_a , if

$$\max_{x^* \in \mathcal{X}} E\{I_n(x^*) \mid Y^n\} < \epsilon_a, \quad (19.12)$$

then stop and predict x_{\min} by an input site $\widehat{x}_{\min} \in \{x_1, \dots, x_n\}$ satisfying

$$y(\widehat{x}_{\min}) = \min_{i=1, \dots, n} y(x_i).$$

Otherwise, select the point $x_{n+1} \in \mathcal{X}$ to maximize $E\{I_n(x) | \mathbf{Y}^n\}$. Set $S_{n+1} = S_n \cup \{x_{n+1}\}$, $\mathbf{Y}^{n+1} = ((\mathbf{Y}^n)', y(x_{n+1}))'$, and increment n . Continue with the next update.

Among the obvious modifications of this algorithm are the following. At stopping, one can predict x_{\min} to be the minimizer of the predictor $\widehat{Y}(\cdot)$ (based on the data available). The absolute stopping criterion in (19.12) can be replaced by the relative stopping criterion

$$\frac{\max_{x \in \mathcal{X}} E\{I_n(x) | \mathbf{Y}^n\}}{|y_{\min}^n|} < \epsilon_r,$$

where ϵ_r is a specified relative tolerance. In either case of the stopping criterion, the idea is that the algorithm should continue until the maximum possible improvement is small.

In implementing their algorithm, Schonlau (1997), Schonlau et al. (1998), and Jones et al. (1998) use the power exponential correlation structure of (19.3) for the Gaussian process $Y(\cdot)$. They estimate the correlation parameters θ_t and p_t by the method of ML. Upon completion of each update step, the correlation parameters of the GASP model can optionally be updated. The updating procedure can be computationally expensive, particularly for large designs. With this in mind, Schonlau et al. (1998) discuss a modification of their algorithm in which a correlation parameter is not updated after each iteration of the algorithm. Specifically, given a current design of size n and q iterates to be added, Schonlau et al. (1998) recommend updating the $s(x)$ coefficient in (19.11) after each iterate, but not updating the $s(x)$ term in the expressions $(y_{\min}^n - \widehat{Y}(x))/s(x)$. This heuristic forces all previously sampled inputs to be avoided, including the previous iterates of the current stage, as $s(\cdot)$ is 0 at these inputs. The EBLUP and MSPE are updated subsequent to the correlation parameters at the completion of each stage. If \mathcal{X} is finite, the expected improvement algorithm will converge to the global minimum under the assumption that ϵ_a (or ϵ_r) = 0. Schonlau (1997) demonstrates the effectiveness of this algorithm for a suite of test problems where ϵ_a (or ϵ_r) > 0.

Schonlau (1997) and Schonlau et al. (1998) consider a generalization of the expected improvement criterion in which the improvement in (19.10) is replaced by

$$I_n^g(x) = \begin{cases} (y_{\min}^n - Y(x))^g, & \text{if } Y(x) < y_{\min}^n \\ 0, & \text{otherwise} \end{cases} \tag{19.13}$$

for some $g \in \{0, 1, 2, \dots\}$. Larger values of g are associated with a more global search. Provided $I_n(x) \geq 1$, one can show $I_n^{g_1}(x) \geq I_n^{g_2}(x)$ for each input x when $g_1 \geq g_2$. Therefore, greater weight will be placed on the tails of the conditional distribution of $Y(\cdot)$ given \mathbf{Y}^n for larger values of g , so that the global potential for large improvement is given increased quantitative importance. The quantity $E\{I_n(\cdot) | \mathbf{Y}^n\}$ in the stopping rule for the expected improvement algorithm is replaced by $E\{I_n^g(\cdot) | \mathbf{Y}^n\}^{1/g}$ so that the tolerance limits ϵ_a and ϵ_r have approximately the same interpretation for any g . Schonlau et al. (1998) provide recursive relations for computing $E\{I_n^g(x) | \mathbf{Y}^n\}$.

Schonlau (1997) mentions that $g = 1$ may be undesirable and that larger g that produce a more global search may be preferable. To support this, Schonlau (1997) shows in an example that $g = 2$ and $g = 5$ outperform $g = 1$.

19.3.2 Batch Sequential Design

Suppose we have $m > 1$ computers available, each of which is able to run the computer simulation. In this case, it makes sense to run the simulations sequentially in stages of *batches* of m simultaneous runs instead of sequentially using a one-run-at-a-time algorithm. Schonlau (1997) extends the approach discussed in the previous section to computer experiments performed in batches of m simultaneous runs, one run for each of the m computers. Given training data consisting of n runs of the simulator, one would like to find the next m design points $\mathbf{x}_{n+1}, \mathbf{x}_{n+2}, \dots, \mathbf{x}_{n+m} \in \mathcal{X}$ at which to run these computers. For the GASP model and assuming the goal is to find the global minimum, Schonlau (1997) proposes selecting these m points to maximize

$$E_{Y_1} E_{Y_2} \dots E_{Y_m} (I_n^m(\mathbf{x}_{n+1}, \mathbf{x}_{n+2}, \dots, \mathbf{x}_{n+m})) \tag{19.14}$$

where Y_1, Y_2, \dots, Y_m are the random variables corresponding to the m points $\mathbf{x}_{n+1}, \mathbf{x}_{n+2}, \dots, \mathbf{x}_{n+m}$ and the m -step improvement I_n^m over y_{\min}^n is defined as

$$I_n^m(\mathbf{x}_{n+1}, \mathbf{x}_{n+2}, \dots, \mathbf{x}_{n+m}) = \max(0, y_{\min}^n - Y_1(\mathbf{x}_{n+1}), y_{\min}^n - Y_2(\mathbf{x}_{n+2}), \dots, y_{\min}^n - Y_m(\mathbf{x}_{n+m})). \tag{19.15}$$

This m -step improvement function is a natural extension of (19.10). However, analytically and computationally, this is a much harder problem than the one-at-a-time global minimization algorithm, because calculating the expected improvement involves multiple integrals with normal densities. Rather than computing these integrals numerically, Schonlau (1997) proposes some simplifications. He then demonstrates how these perform in an example (finding the global minimum of the logarithm of the Goldstein–Price function, a function of two inputs, in batches of $m = 10$ points). Once again, the use of expected improvement selects points both locally (near the true minimum) and globally. An initial design of $n_0 = 21$ is used and 13 batches of points are added, resulting in a total of 151 runs. Compared with the one-at-a-time expected improvement algorithm applied to the same example with a similar starting design, the *batch* method used more runs (151 vs. 127 for the one-at-a-time algorithm) to achieve the same degree of accuracy. However, selecting points in batches produced a design with a greater spread of points. Schonlau (1997) does not mention whether there were significant differences in overall computational times for the two methods.

Recently, there has been considerable work on computing the expected improvement in (19.14) more efficiently. See, for example, Ponweiser et al. (2008), Taddy et al. (2009), Roustant et al. (2012), and Viana et al. (2013). Roustant et al. (2012) include R code.

19.3.3 Global Optimization for Control and Environmental Inputs

In some applications, it is useful to distinguish *two* types of inputs to a simulator. One type of input variable that we may encounter is a *control variable*. If the output of the computer

experiment is some performance measure of a product or process, then control variables are those variables that can be set by an engineer or scientist to *control* the actual product or process being simulated. Some authors use the terms *engineering variables* or *manufacturing variables* rather than control variables. We use the generic notation \mathbf{x}_c to denote control variables. Control variables are present in physical experiments as well as in many computer experiments.

A second type of variable that can be present in both computer and physical experiments is an *environmental variable*. In general, environmental variables affect the output $y(\cdot)$ but depend on (uncontrolled) conditions present at the time a response is measured on the product or process. Environmental variables are sometimes called *noise variables*. We use the notation \mathbf{x}_e to denote the vector of environmental variables for a given problem. In practice, we typically regard environmental variables as random with a distribution that is known or unknown. To emphasize situations where we regard the environmental variables as random, we use the notation \mathbf{X}_e .

An example of both types of variables might be found in a manufacturing setting where the speed at which a machine is run and the ambient humidity in the shop affect some response. Assuming the machine speed can be set by the user, machine speed would be a control variable. Assuming the humidity in the shop cannot be controlled, humidity would be viewed as an environmental variable.

Suppose \mathbf{x} consists of both control and environmental components, that is, $\mathbf{x} = (\mathbf{x}_c, \mathbf{x}_e)$. For the goal of optimization, we may wish to determine the values of the control variables that, on average (averaged over the values of the noise variables), optimize the output.

In particular, suppose also that the environmental variables have a *known* probability distribution, which is specified either by the probability mass function $w_j = P\{\mathbf{X}_e = \mathbf{x}_{e,j}\}$ on n_e support points $\{\mathbf{x}_{e,j}\}_j$ or by the probability density function $w(\cdot)$. Then the quantities of interest are

$$\mu(\mathbf{x}_c) \equiv \sum_{j=1}^{n_e} w_j y(\mathbf{x}_c, \mathbf{x}_{e,j}) \quad \left(\text{or} \equiv \int y(\mathbf{x}_c, \mathbf{x}_e) w(\mathbf{x}_e) d\mathbf{x}_e \right), \quad (19.16)$$

which is the mean of $y(\cdot)$ with respect to the distribution of the \mathbf{X}_e variables.

Williams et al. (2000) extend the expected improvement algorithm of Schonlau et al. (1998) and Jones et al. (1998) to settings involving both control and environmental inputs. For discrete environmental variable distributions as used in the left-hand side of (19.16), the goal of their revised algorithm was to find a control variable input that minimized the mean function $\mu(\cdot)$ for applications with “expensive” $y(\cdot)$ outputs. The mean $\mu(\mathbf{x}_c)$ is the output $y(\mathbf{x}_c, \mathbf{x}_{e,j})$ averaged, with weights $\{w_j\}$, over these n_e values. The function $\mu(\mathbf{x}_c)$ inherits the prior process defined by

$$M(\mathbf{x}_c) \equiv \sum_{j=1}^{n_e} w_j Y(\mathbf{x}_c, \mathbf{x}_{e,j}) \quad (19.17)$$

where $Y(\mathbf{x}_c, \mathbf{x})$ has the Gaussian prior used by Schonlau et al. (1998) and Jones et al. (1998).

Let $(\mathbf{x}_{c,i}, \mathbf{x}_{e,i}), 1 \leq i \leq n$, denote a generic n -point experimental design. Then $\{\mathbf{x}_{c,i}\}$ denotes the control variable portion of this design. Williams et al. (2000) extend the improvement criterion (19.10) to this setting by defining their improvement criterion to be

$$I_n(\mathbf{x}_c) = \begin{cases} M_{\min}^n - M(\mathbf{x}_c), & \text{if } M_{\min}^n - M(\mathbf{x}_c) > 0 \\ 0 & \text{otherwise} \end{cases},$$

where

$$M_{\min}^n = \min_{i=1, \dots, n} M(\mathbf{x}_{c,i}).$$

In contrast to the known value of y_{\min}^n used in (19.10), the quantity M_{\min}^n must be treated as a random variable with a prior distribution because (19.17) is never directly calculated. This fact makes estimating the expected improvement more cumbersome than in Schonlau et al. (1998) and Jones et al. (1998), as will subsequently be discussed. In outline, the algorithm to minimize $\mu(\cdot)$ consists of the following steps:

1. Select an initial experimental design $S_n = \{(\mathbf{x}_{c,i}, \mathbf{x}_{e,i}); 1 \leq i \leq n\}$, and compute the vector of model outputs \mathbf{Y}^n at each design point.
2. Estimate the vector of correlation parameters via ML or restricted ML, under the Bayesian framework, by the mode of their joint posterior distribution.
3. Select $\mathbf{x}_{c,n+1}$ to maximize the expected improvement,

$$\mathbf{x}_{c,n+1} = \arg \max_{\mathbf{x}_c \in \mathcal{X}_c} E\{I_n(\mathbf{x}_c) \mid \mathbf{Y}^n\},$$

where \mathcal{X}_c denotes the domain of the control variables and $\mathcal{X} = \mathcal{X}_c \times \mathcal{X}_e$, where \mathcal{X}_e denotes the domain of the environmental variables.

4. Select $\mathbf{x}_{e,n+1}$ corresponding to $\mathbf{x}_{c,n+1}$ to minimize the MSPE given \mathbf{Y}^n ,

$$\mathbf{x}_{e,n+1} = \arg \min_{\mathbf{x}_e \in \mathcal{X}_e} E \left\{ \left[\widehat{M}^{n+1}(\mathbf{x}_{c,n+1}) - M(\mathbf{x}_{c,n+1}) \right]^2 \mid \mathbf{Y}^n \right\},$$

where $\widehat{M}^{n+1}(\cdot)$ is the conditional mean of $M(\cdot)$ based on the data \mathbf{Y}^n and the latent observation $Y(\mathbf{x}_{c,n+1}, \mathbf{x}_e)$.

5. Determine if the algorithm should be stopped. If not, set $S_{n+1} = S_n \cup (\mathbf{x}_{c,n+1}, \mathbf{x}_{e,n+1})$, calculate the output $y(\cdot)$ at $(\mathbf{x}_{c,n+1}, \mathbf{x}_{e,n+1})$, increment n to $n + 1$, and go to step 2 (correlation parameter estimation). Otherwise, the global minimizer in the control variable space is predicted to be the global minimizer of the EBLUP based on S_n .

The mean squared error of prediction criterion for environmental variable selection, required in step 4 of the algorithm, has a computationally convenient closed-form expression.

The expected improvement required in step 3 cannot be expressed in closed form because M_{\min}^n is not known. However, because the posterior of M_{\min}^n given \mathbf{Y}^n is known, the expected improvement can be approximated using the following Monte Carlo method. Let $\mathbf{M}^n = (M(\mathbf{x}_{c,1}), \dots, M(\mathbf{x}_{c,n}))$, and let (m_c, v_c) denote the mean and scale parameters of the tractable conditional distribution of $M(\cdot)$ given \mathbf{Y}^n and \mathbf{M}^n .

- (a) Generate r Monte Carlo samples M^n from the conditional distribution $[M^n | Y^n]$.
- (b) For each sampled M^n , compute the conditional expectation of $I_n(\cdot)$ given Y^n and M^n from the formula

$$E \{ I_n(x_c) | Y^n, M^n \} = (M^n_{\min} - m_c) T_{2n-1} \left(\frac{M^n_{\min} - m_c}{\sqrt{v_c}} \right) + \frac{1}{2(n-1)} \left[(2n-1)\sqrt{v_c} + \frac{(M^n_{\min} - m_c)^2}{\sqrt{v_c}} \right] t_{2n-1} \left(\frac{M^n_{\min} - m_c}{\sqrt{v_c}} \right),$$

where $T_v(\cdot)$ and $t_v(\cdot)$ denote the univariate t cumulative distribution and density functions, respectively, and average the results.

- (c) Average the r conditional expectations calculated in step b.

This approximation method allows one to see that, in this case, the expected improvement criterion balances local and global improvements as did (19.11) in the Schonlau et al. (1998) algorithm. The first term in $E\{I_n(x_c) | Y^n, M^n\}$ is “large” when $m_c \ll M^n_{\min}$, in other words, when the predicted objective at x_c suggests the potential for large local improvement. The second term is “large” when $\sqrt{v_c}$ is large relative to $|M^n_{\min} - m_c|$, necessarily when there is a high degree of prediction uncertainty at the candidate x_c . The global search component of this algorithm is provided by the capacity to select control sites with the potential for large improvement. One significant computational advantage of using the Monte Carlo method, (a) through (c) earlier, to calculate $E\{I_n(x_c) | Y^n\}$ is that a single Monte Carlo sample of M^n suffices to determine the expected improvement for all choices of x_c because the distribution $[M^n | Y^n]$ is independent of x_c .

The correlation parameter estimation in step 2 can be extremely time consuming, particularly for “large” experimental designs and/or high-dimensional inputs. A sensible modification of this algorithm is to update the correlation parameters frequently in the initial iterations and reduce the update rate as additional design points are added. In this way, the correlation parameters are repeatedly updated at the stage of the algorithm when they are least stable and the most substantial learning about the response surface occurs. As the response surface in the region of the optimum becomes predicted more accurately, correlation parameter updates become less necessary.

Because the correlation parameters are reestimated at each stage, the sequence of maximum expected improvements need not be monotone decreasing. The stopping rules recommended in Williams et al. (2000) are therefore based on observing a sequence of “small” maximum expected improvements relative to the history of such improvements established as the algorithm progresses. For example, moving averages and ranges of groups of observed expected improvements can be tracked, and the algorithm stopped, when these statistics reach a problem-specific threshold established relative to their initial observed values. The stopping criteria should ideally be met by two or more successive values of the moving average and range, suggesting stabilization of the expected improvement sequence.

19.3.4 Multiple Responses

Many computer experiments involve simulations that produce multiple outputs. Suppose the computer simulator generates m responses $y_1(\cdot), \dots, y_m(\cdot)$, all functions of the

same inputs. Let $\mathbf{y}(\cdot) = (y_1(\cdot), \dots, y_m(\cdot))$. $\mathbf{y}(\cdot)$ be an m -dimensional vector valued function ($m > 1$) from \mathcal{X} to \mathcal{Y} , where $\mathcal{Y} \subseteq \mathbb{R}^m$. To model multivariate output from our simulator, we need a multivariate model that specifies a joint distribution for the $y_i(\cdot)$. The simplest model assumes independence. In this case, we view $\mathbf{y}(\cdot) = (y_1(\cdot), \dots, y_m(\cdot))$ as the realization of a random process $\mathbf{Y}(\cdot) = (Y_1(\cdot), \dots, Y_m(\cdot))$ and assume $Y_1(\cdot), \dots, Y_m(\cdot)$ follow independent GASP models.

A model that assumes dependence is

$$Y_i(\mathbf{x}) = \rho_i Y_1(\mathbf{x}) + \gamma_i(\mathbf{x}), \quad i = 2, \dots, k + 1, \tag{19.18}$$

where $\gamma_i(\cdot)$ is independent of $Y_1(\cdot)$. This model says that $Y_2(\cdot), \dots, Y_m(\cdot)$ are each associated with $Y_1(\cdot)$ (positively or negatively) plus an independent refinement.

There are several other models that do not assume independence. We do not discuss these here, but see Higdon (2002), Gelfand et al. (2004), Banerjee and Johnson (2006), Banerjee et al. (2008), Conti and O’Hagan (2010), and Fricker et al. (2013) for a variety of dependence models.

The goal is to optimize $\mathbf{y}(\cdot)$ over the input space \mathcal{X} . Ideally, we would like to find a single \mathbf{x} that minimizes each $y_i(\mathbf{x})$, $1 \leq i \leq m$. In general, the same \mathbf{x} will not optimize all the $y_i(\mathbf{x})$ simultaneously. Instead, one attempts to find a set of compromise or *trade-off* solutions. One compromise solution is to optimize $y_1(\cdot)$ subject to $y_2(\cdot), \dots, y_m(\cdot)$, each taking values in reasonable intervals expressed as constraints. Another compromise solution is to define a partial ordering on $\mathbf{y}(\cdot)$ and find the inputs that produce outputs that are not dominated by any other possible set of outputs. We discuss both of these approaches.

In what follows we assume that *optimization* means minimization. It is straightforward to modify what we say to include maximization or a mixture of minimization and maximization.

19.3.4.1 Constrained Global Optimization

Schonlau et al. (1998) propose a version of the expected improvement algorithm for use in problems of constrained optimization. Here we assume that the computer output consists of multiple responses $y_1(\cdot), \dots, y_m(\cdot)$. The goal is to minimize $y_1(\mathbf{x})$ subject to \mathbf{x} satisfying $m - 1$ constraints $l_i \leq y_i(\cdot) \leq u_i$ for $i = 2, \dots, m$.

Schonlau et al. (1998) propose using the improvement function

$$I_{c,n}^g(\mathbf{x}) = \begin{cases} (y_{\min}^n - Y_1(\mathbf{x}))^g & \text{if } Y_1(\mathbf{x}) < y_{\min}^n \text{ and } l_i \leq Y_i(\mathbf{x}) \leq u_i \\ & \text{for } i = 2, \dots, m \\ 0 & \text{otherwise} \end{cases} \tag{19.19}$$

so that any constraint violation leads to zero improvement. As usual, y_{\min}^n is the minimum of the responses observed on the current experimental design and $Y_i(\cdot)$ denotes the stochastic process prior for the i th output function. Schonlau et al. (1998) assume that the objective and constraint processes were *mutually independent*; under this assumption, the conditional expected improvement is given by

$$E \left\{ I_{c,n}^g(\mathbf{x}) \mid \mathbf{Y}_1^n, \dots, \mathbf{Y}_m^n \right\} = E \left\{ I_n^g(\mathbf{x}) \mid \mathbf{Y}_1^n \right\} \\ \times P \left\{ l_2 \leq Y_2(\mathbf{x}) \leq u_2 \mid \mathbf{Y}_2^n \right\} \cdots P \left\{ l_m \leq Y_m(\mathbf{x}) \leq u_m \mid \mathbf{Y}_m^n \right\},$$

where conditionally given the observed data Y_i^n , $Y_i(x)$ has a Gaussian distribution, which is the analog of (19.9).

This methodology can be extended to accommodate correlated response functions. The expanded statistical model requires a valid cross-correlation structure. The additional model complexity increases the computational burden involved in fitting the statistical model, with possible benefits including increased flexibility in modeling the physical system and increased efficiency in finding a constrained optimum. Next, we describe an algorithm introduced by Williams et al. (2010) that uses correlated outputs for a more complicated problem involving control and environmental variables.

Williams et al. (2010) use the spatial autoregressive model in (19.18). This model says that each constraint function is associated with the objective function (positively or negatively) plus an independent refinement.

They apply this spatial autoregressive model to the bivariate ($m = 2$) response constrained optimization problem where the objective ($i = 1$) and constraint ($i = 2$) functions $\mu_i(\cdot)$ are given as in (19.16), with the corresponding output priors defined by

$$M_i(x_c) \equiv \sum_{j=1}^{n_e} w_j Y_i(x_c, x_{e,j}).$$

Their goal was to minimize $\mu_1(\cdot)$, subject to an upper bound u_2 on $\mu_2(\cdot)$. The technique employed for solving this problem extends without difficulty to lower bound or two-sided constraints.

The improvement function used by Williams et al. (2010) is a modified version of (19.19) that accounts for the focus on integrated outputs,

$$I_{c,n}(x) = \begin{cases} M_{\min}^n - M_1(x_c) & \text{if } M_1(x_c) < M_{\min}^n \text{ and } M_2(x_c) \leq u_2 \\ 0 & \text{otherwise} \end{cases}$$

where

$$M_{\min}^n = \min \left\{ M_1(x_{c,i}) : 1 \leq i \leq n \quad \text{and} \quad M_2^{0.05}(x_{c,i}) \leq u_2 \right\}$$

with $M_2^{0.05}(\cdot)$ denoting the lower fifth quantile function of the distribution of $M_2(\cdot)$. Small values of $M_1(\cdot)$ on the current experimental design are not included in the computation of the improvement function when the distribution of $M_2(\cdot)$ shows that there is only small probability that the constraint is satisfied (here, no more than 5%). This restriction eliminates the scenario involving a global minimum in $\mu_1(\cdot)$, at which the constraint on $\mu_2(\cdot)$ is violated, from being included in the computation of expected improvement. Inclusion of such points results in an expected improvement function of approximately zero throughout the control variable space, and the constrained optimum would never be found.

The algorithm proceeds in a fashion similar to the single-objective optimization algorithm of Williams et al. (2000) described in Section 19.3.3: obtain an initial experimental design, estimate the correlation parameters, determine the next control variable site by maximizing the conditional expected improvement, determine the corresponding environmental variable site by minimizing the conditional mean squared error of prediction (as in step 4 of the algorithm in Section 19.3.3), generate responses at the newly chosen experimental design site, and iterate until the stopping criteria are met. Once the algorithm is

stopped, the constrained optimizer is predicted by solving the constrained optimization problem with the EBLUPs of $M_1(\cdot)$ and $M_2(\cdot)$. The correlation parameters can be intermittently estimated as the algorithm progresses, substantially reducing computation time while sacrificing little in terms of prediction accuracy.

In applications where the simulator does not generate $y_1(\cdot)$ and $y_2(\cdot)$ simultaneously, the efficiency of the constrained optimization algorithm can be improved by utilizing an additional statistical criterion to select which one of the outcomes should be generated at each stage, rather than generating both outcomes. This feature is implemented in conjunction with choosing the environmental variable site for the subsequent iteration of the algorithm. Let $\mathbf{Y}_1^{n_1}$ and $\mathbf{Y}_2^{n_2}$ denote the vectors of outcomes calculated on the current experimental designs of n_1 and n_2 sites for the first and second outcomes, respectively, and let n denote the current total number of experimental design sites. Note that n may be less than $n_1 + n_2$ if the two responses are in some cases generated at the same input. Set $\mathbf{Y}^n = (\mathbf{Y}_1^{n_1}, \mathbf{Y}_2^{n_2})$, $\mathbf{Y}_e^j = (Y_j(\mathbf{x}_{c,n+1}, \mathbf{x}_e), \mathbf{Y}^n)$ and let $\widehat{M}_i^j(\mathbf{x}_c)$ denote the conditional mean of $M_i(\mathbf{x}_c)$ given \mathbf{Y}_e^j . The control variable site that maximizes conditional expected improvement given the current n -point experimental design is designated $\mathbf{x}_{c,n+1}$. For $j = 1, 2$, define the MSPE function by

$$\begin{aligned} MSPE_j(\mathbf{x}_e) = E \left\{ \left[\widehat{M}_1^j(\mathbf{x}_{c,n+1}) - M_1(\mathbf{x}_{c,n+1}) \right]^2 \mid \mathbf{Y}^n \right\} \\ + P \{ M_2(\mathbf{x}_{c,n+1}) > u_2 \mid \mathbf{Y}^n \} E \left\{ \left[\widehat{M}_2^j(\mathbf{x}_{c,n+1}) - M_2(\mathbf{x}_{c,n+1}) \right]^2 \mid \mathbf{Y}^n \right\}. \end{aligned}$$

For each outcome j , calculate

$$\mathbf{x}_{e,j}^* = \arg \min_{\mathbf{x}_e \in \mathcal{X}_e} MSPE_j(\mathbf{x}_e).$$

If $MSPE_1(\mathbf{x}_{e,1}^*) \leq MSPE_2(\mathbf{x}_{e,2}^*)$, augment the experimental design corresponding to the first outcome with $(\mathbf{x}_{c,n+1}, \mathbf{x}_{e,1}^*)$ and compute $y_1(\cdot)$ at this location; otherwise, augment the experimental design corresponding to the second outcome with $(\mathbf{x}_{c,n+1}, \mathbf{x}_{e,2}^*)$ and compute $y_2(\cdot)$ at this location. This criterion selects the next outcome to be generated based on minimizing a weighted sum of prediction errors in the objective and constraint functions. The error in predicting the constraint function is weighted by the probability that the constraint is exceeded given the current data; if this probability is low, then constraint satisfaction is currently determined with certainty and minimizing prediction error in the constrained outcome is unnecessary. On the other hand, if constraint satisfaction is less certain, the need to reduce prediction error in the objective function is balanced against the desire to reduce prediction uncertainty in the constraint function.

19.3.4.2 Pareto Optimization

For purposes of multiple response optimization, we define partial orderings of vectors. For two inputs \mathbf{x}_1 and \mathbf{x}_2 , we say \mathbf{x}_1 *weakly dominates* \mathbf{x}_2 ($\mathbf{x}_1 \succeq \mathbf{x}_2$) if $y_i(\mathbf{x}_1) \leq y_i(\mathbf{x}_2)$ for all $i = 1, \dots, m$. If at least one inequality is strict, then \mathbf{x}_1 is said to *dominate* \mathbf{x}_2 ($\mathbf{x}_1 \succ \mathbf{x}_2$). Similarly, for two elements $\mathbf{y}(\mathbf{x}_1)$ and $\mathbf{y}(\mathbf{x}_2)$ in \mathcal{Y} , we say $\mathbf{y}(\mathbf{x}_1)$ *weakly dominates* $\mathbf{y}(\mathbf{x}_2)$ ($\mathbf{y}(\mathbf{x}_1) \succeq \mathbf{y}(\mathbf{x}_2)$) if $y_i(\mathbf{x}_1) \leq y_i(\mathbf{x}_2)$ for all $i = 1, \dots, m$. If at least one inequality is strict, then $\mathbf{y}(\mathbf{x}_1)$ is said to *dominate* $\mathbf{y}(\mathbf{x}_2)$ ($\mathbf{y}(\mathbf{x}_1) \succ \mathbf{y}(\mathbf{x}_2)$). An input vector $\mathbf{x} \in \mathcal{X}$ is *Pareto optimal* if and only if there is no $\mathbf{x}' \in \mathcal{X}$ such that $\mathbf{x} < \mathbf{x}'$. (Such \mathbf{x} are also referred to as *nondominated inputs*. The image

$\mathbf{y}(x)$ of a nondominated input is sometimes referred to as a *nondominated output*.) The set of all Pareto optimal points in \mathcal{X} is referred to as the *Pareto set*. Denote this set as $\mathbb{P}_{\mathcal{X}}$. The corresponding image of $\mathbb{P}_{\mathcal{X}}$ in \mathcal{Y} is referred to as the *Pareto front*, and it is denoted $\mathbb{P}_{\mathcal{Y}}$. In multiresponse optimization, the goal is to find (or, more realistically, approximate) the possibly uncountable sets $\mathbb{P}_{\mathcal{Y}}$ and $\mathbb{P}_{\mathcal{X}}$.

Given two solutions (approximations) to $\mathbb{P}_{\mathcal{Y}}$ or $\mathbb{P}_{\mathcal{X}}$, how do we compare them? Several metrics exist for comparing solutions to multiresponse optimization problems. These metrics rely on the notion of an approximation set and a generalization of dominance for sets of points, rather than just single points, in the input space \mathcal{X} or the output space \mathcal{Y} .

Let A be a set of points in the output space. A is an *approximation set* if none of its elements dominate each other. In a one-dimensional setting ($m = 1$), an approximation set would just be our best estimate of the true minimum of some function. In this situation, it is easy to compare two competing minima, because of the inherent ordering of the real numbers \mathbb{R} . In a multiresponse problem, though, this is much more complicated, and we must rely on the notion of dominance to compare sets.

A set A in the output space *weakly dominates* the set B in the output space ($A \succeq B$) if every vector in B is weakly dominated by some vector in A . Set A *dominates* B ($A \succ B$) if every vector in B is dominated by a vector in A . These terms can be defined analogously for sets of points in the input space. Obviously, an approximation set B that is weakly dominated or dominated by another approximation set A would be considered an inferior approximation to the Pareto front.

Unfortunately, for any two approximation sets A and B , neither one need dominate the other. To help determine how “good” an approximation set is, researchers studying multiresponse optimization have introduced several *Pareto set approximation quality indicators*. A detailed survey of this topic can be found in Zitzler et al. (2008).

Pareto set approximation quality indicators play two important roles. First, they serve as metrics for comparing various approaches to approximating the Pareto front of a computer experiment with multiple outputs. Second, many proposed approaches for the multiresponse optimization of computer experiments are based on these Pareto set approximation quality indicators. Here we only discuss one of these quality indicators, the hypervolume indicator. This quality indicator was first introduced in Zitzler and Thiele (1998) and remains one of the most popular quality indicators in the multiresponse optimization literature. To calculate the hypervolume indicator for a set A in the objective space, we first need to define a reference point \mathbf{v} , which is a point weakly dominated by all vectors in \mathcal{Y} . Define the hypervolume indicator of A as

$$I_H(A, \mathbf{v}) = \int_{\mathcal{Y}} 1_{\{\mathbf{y} \mid \mathbf{y} \succeq \mathbf{v}, A \succeq \{\mathbf{y}\}\}} d\mathbf{y}.$$

Basically, it is the Lebesgue measure of the set of points in the output space that are dominated by A but also dominate some reference point \mathbf{v} .

The hypervolume indicator’s main advantage is that it is *strictly monotonic* with respect to dominance. This means that if we have two approximation sets A and B and $A \succ B$, then $I_H(A, \mathbf{v}) > I_H(B, \mathbf{v})$. While there are many Pareto set approximation indicators that are monotonic, that is, $A \succeq B$, then $I_H(A, \mathbf{v}) \geq I_H(B, \mathbf{v})$, Zitzler et al. (2008) state that the hypervolume indicator and a related indicator, the weighted hypervolume indicator, are the only known Pareto set approximation quality indicators that are *strictly* monotonic.

However, there are three disadvantages of the hypervolume indicator: One, it is dependent on the scaling of the various outputs. Two, it requires some additional problem knowledge, as it requires one to specify a suitable upper bound \mathbf{v} on the output space. Three, it has a high computational cost. According to Fonseca et al. (2006), the best known algorithms for calculating the hypervolume indicator have running times that are exponential in m , the number of response functions. Therefore, for even moderate-sized output spaces, one must usually resort estimating the hypervolume indicator via Monte Carlo methods.

We now generalize the notion of *expected improvement* as given in Schonlau et al. (1998). First, we replace, y_{\min}^n , the current minimum among n function evaluations, by the current Pareto front, denoted $\mathbb{P}_{\mathbf{y}}^n$, which is the set of nondominated outputs among the set of n function evaluations of the m -dimensional function $\mathbf{y}(\cdot)$. Next, we must define some sort of improvement function to quantify how much an output $\mathbf{y}(x)$ improves upon $\mathbb{P}_{\mathbf{y}}^n$. Possible improvement functions have been proposed by Emmerich et al. (2006), Keane (2006), Bautista (2009), and Svenson (2011).

Before discussing specific improvement criteria, we describe a general algorithm for solving our multiresponse optimization.

19.3.4.3 Outline of the Algorithm

We assume that the m -dimensional function $\mathbf{y}(\cdot)$ is a realization of an m -variate GASP model $\mathbf{Y}(\cdot)$. For simplicity, we also assume independent GASP models for the $Y_i(\cdot)$. The algorithm is as follows:

1. Evaluate $\mathbf{y}(\cdot)$ at an initial space-filling design $S_n = (x_1, \dots, x_n) \subset \mathcal{X}$. Let $\mathbf{y}^{m,n} = (\mathbf{y}'(x_1), \dots, \mathbf{y}'(x_n))'$.
2. Estimate the unknown parameters using ML or REML based on the prior distribution associated with $\mathbf{y}^{m,n} = (\mathbf{y}'(x_1), \dots, \mathbf{y}'(x_n))'$.
3. Calculate the current Pareto set $\mathbb{P}_{\mathcal{X}}^n$ and current Pareto front $\mathbb{P}_{\mathbf{y}}^n$. These are the set of nondominated inputs among all inputs in S_n and all nondominated outputs among all function evaluations in $\mathbf{y}^{m,n}$, respectively.
4. Find x^{n+1} by maximizing a given *quality improvement* measure $QI(x)$. The function $QI(\cdot)$ will depend on the specific improvement criterion $I^*(\mathbf{y}(x))$, that is, the multivariate generalization of $I(x)$ in (19.10). In the following, we present two ways that $QI(x)$ has been defined in the literature. Note that $QI(x)$ is more commonly referred to as the *expected improvement* at x , but in the literature, it has not always been defined as the expectation of the selected improvement criterion. Hence, we use here the notation $QI(\cdot)$ to emphasize that it is not necessarily an expectation.
5. Evaluate $\mathbf{y}(x^{n+1})$. If our budget has been exhausted or a stopping criterion has been met, terminate the algorithm. Otherwise, repeat steps 2 through 5 with $\mathbf{y}^{m,n+1} = (\mathbf{y}'(x_1), \dots, \mathbf{y}'(x_n), \mathbf{y}'(x_{n+1}))'$.

Two methods have been proposed in the literature for defining $QI(x)$ in terms of a given improvement function $I^*(\mathbf{y}(x))$:

- **Method 1:** Define the quality improvement associated with $I^*(\cdot)$ to be

$$QI(x) = E [I^*(Y(x)) | Y^{m,n} = \mathbf{y}^{m,n}]. \tag{19.20}$$

$QI(\cdot)$ is an expectation of the posterior predictive distribution, and thus it is natural to refer to this improvement measure as an expected improvement function.

- **Method 2:** Define the quality improvement associated with $I^*(\cdot)$ to be

$$QI(x) = P(Y(x) \in R_n | Y^{m,n} = \mathbf{y}^{m,n}) I^*(\bar{Y}(x)) \tag{19.21}$$

where

$$R_n = \{ \mathbf{y} : \mathbf{y} \not\leq \mathbf{z}, \forall \mathbf{z} \in \mathbb{P}_y^m \} \tag{19.22}$$

is the region of the output space that is nondominated by the current Pareto front and $\bar{Y}(x)$ is the mean of the distribution of $[Y(x) | Y^{m,n} = \mathbf{y}^{m,n}, Y(x) \in R_n]$, which can be expressed as

$$\bar{Y}(x) = \frac{E \left[Y(x) 1_{[Y(x) \in R_n]} | Y^{m,n} = \mathbf{y}^{m,n} \right]}{P(Y(x) \in R_n | Y^{m,n} = \mathbf{y}^{m,n})}. \tag{19.23}$$

Notice that $\bar{Y}(x)$ is an m -dimensional vector.

Method 1 is the more natural multivariate extension of the improvement criterion given in (19.10) and the corresponding expected improvement in (19.11). We define an improvement function and compute its expectation conditional on the function evaluations we have already observed. In practice, one attempts to find analytic expression for (19.20). If this is not possible, one can turn to Monte Carlo methods by simulating a large number of draws from the distribution of $[Y(x) | Y^{m,n} = \mathbf{y}^{m,n}]$, evaluating the improvement function at each draw, and averaging these values. The strong law of large numbers will guarantee convergence to (19.20) at any given x .

Method 2 is less natural. First and foremost, (19.21) is a *function* of expectations, but generally is not an expectation. Therefore, referring to it as an *expected improvement function* can be misleading. However, it plays a role similar to the expected improvement in (19.11), so we will refer to it as an *expected improvement*.

For output with only a single response, (19.21) is equivalent to (19.10). In particular, when $m = 1$,

$$\begin{aligned}
 & E \left[I(y(x)) | \mathbf{Y}^{n,1} = \mathbf{y}^{n,1} \right] \\
 &= E \left[\left(y^{\min} - Y(x) \right) 1_{[Y(x) - y^{\min} < 0]} | \mathbf{Y}^{n,1} = \mathbf{y}^{n,1} \right] \\
 &= y^{\min} E \left[1_{[Y(x) - y^{\min} < 0]} | \mathbf{Y}^{n,1} = \mathbf{y}^{n,1} \right] \\
 &\quad - E \left[Y(x) 1_{[Y(x) - y^{\min} < 0]} | \mathbf{Y}^{n,1} = \mathbf{y}^{n,1} \right] \\
 &= E \left[1_{[Y(x) - y^{\min} < 0]} | \mathbf{Y}^{n,1} = \mathbf{y}^{n,1} \right] \\
 &\quad \times \left(y^{\min} - E \left[\frac{Y(x) 1_{[Y(x) - y^{\min} < 0]}}{E \left[1_{[Y(x) - y^{\min} < 0]} \right]} | \mathbf{Y}^{n,1} = \mathbf{y}^{n,1} \right] \right) \\
 &= P \left(Y(x) \not\geq y^{\min} \mid \mathbf{Y}^{n,1} = \mathbf{y}^{n,1} \right) I(\bar{Y}(x)).
 \end{aligned}$$

Aside from its equivalence to the usual single-objective expected improvement function, there are three reasons why Method 2 is considered. First, it has made its way into the literature. Keane (2006) essentially uses Method 2 to define $QI(x)$ in a proposed expected improvement algorithm to be discussed shortly. Second, in situations where we must calculate expectations via Monte Carlo integration and $I^*(\cdot)$ is expensive to compute, using Method 1 can be computationally demanding because $I^*(\cdot)$ must be computed for every single draw from the predictive distribution. For example, the improvement criteria defined by Emmerich et al. (2006) require a calculation of the hypervolume indicator that has a very high computational cost for larger m . Finally, the most compelling reason to study Method 2 is that it can work surprisingly well in practice and, when paired with certain improvement functions, can even outperform Method 1.

To keep notation simple, we will use the simplified expressions QIM1 and QIM2 for Method 1 and Method 2, respectively.

19.3.4.4 0–1 Improvement

Keane (2006) introduced the simplest of all multiresponse expected improvement algorithms, the *probability of improvement*. Essentially, we sequentially add inputs by maximizing $P(Y(x) \in R_n | \mathbf{Y}^{m,n} = \mathbf{y}^{m,n})$, and so we are using the improvement criterion

$$I_{\mathbb{P}\mathbb{I}}(\mathbf{y}(x)) = 1_{[\mathbf{y}(x) \in R_n]}$$

where $1_{[\mathbf{y}(x) \in R_n]}$ is the indicator function and equals 1 if $\mathbf{y}(x) \in R_n$ and 0 otherwise. We will refer to $I_{\mathbb{P}\mathbb{I}}(\cdot)$ as the *0–1 improvement criterion*, as it returns a 0 if $\mathbf{y}(x)$ is dominated by the current Pareto front and a 1 otherwise. Notice that, for this particular improvement criterion, QIM1 and QIM2 are equivalent. For the case where $m = 2$ and the outputs are modeled as independent Gaussian processes, Keane (2006) derives a closed-form expression for $P(Y(x) \in R_n | \mathbf{Y}^{m,n} = \mathbf{y}^{m,n})$.

While relatively simple, the probability of improvement is not a very effective expected improvement. Keane (2006), Forrester et al. (2007), and Bautista (2009) point out, based on numerical studies, that the probability of improvement tends to add outputs that are “clumped” together. Essentially, this particular choice for $QI(\mathbf{x})$ does not explore the input space very efficiently. This is not surprising, since the improvement criterion $I_{PI}(\mathbf{Y}(\mathbf{x}))$ only takes into account whether or not an improvement occurs and ignores the magnitude of the improvement.

19.3.4.5 Keane’s Distance-Based Improvement

Keane (2006) also introduces a distance-based expected improvement algorithm in the case where $m = 2$, with the goal of balancing exploration of the input space and exploitation of the surrogate approximation in the search for nondominated sets. This approach employs a QIM2-based $QI(\mathbf{x})$. While only presented for 2D outputs, it is not difficult (at least in principle) to generalize for arbitrary $m \geq 2$. Let

$$\begin{aligned}
 QI(\mathbf{x}) &\equiv P(\mathbf{Y}(\mathbf{x}) \in R_n | \mathbf{Y}^{m,n}) \\
 &\times \min_{\mathbf{x}_i \in \mathbb{P}_n^m} \sqrt{\sum_{k=1}^m (\bar{Y}_k(\mathbf{x}) - y_k(\mathbf{x}_i))^2} \\
 &= \min_{\mathbf{x}_i \in \mathbb{P}_n^m} \sqrt{\sum_{k=1}^m \left(E \left[(Y_k(\mathbf{x}) - y_k(\mathbf{x}_i)) 1_{[\mathbf{Y}(\mathbf{x}) \in R_n]} | \mathbf{Y}^{m,n} \right] \right)^2}. \tag{19.24}
 \end{aligned}$$

The improvement function implicit in this formulation is

$$I_{\mathbb{K}}(\mathbf{y}(\mathbf{x})) = \begin{cases} \min_{\mathbf{x}_i \in \mathbb{P}_n^m} \sqrt{\sum_{k=1}^m (y_k(\mathbf{x}) - y_k(\mathbf{x}_i))^2} & \text{if } \mathbf{y}(\mathbf{x}) \in R_n, \\ 0 & \text{otherwise.} \end{cases}$$

Just like $P(\mathbf{Y}(\mathbf{x}) \in R_n | \mathbf{Y}^{m,n} = \mathbf{y}^{m,n})$, Keane (2006) derives a closed-form expression for $\bar{Y}(\mathbf{x})$ in the two-response, independent-output case.

This expected improvement can substantially outperform the probability of improvement, as the distance-based improvement criterion is larger for outputs that are farther from the current Pareto front. Therefore, the magnitude of improvement is taken into consideration. This encourages the sequentially added points to be more spread out in the output space than when using the probability of improvement. However, $I_{\mathbb{K}}(\mathbf{y}(\mathbf{x}))$ has some theoretical disadvantages that are discussed in Svenson (2011). Also, Svenson (2011) shows that $I_{\mathbb{K}}(\mathbf{y}(\mathbf{x}))$ is relatively inefficient when compared to many other methods.

19.3.4.6 Pareto Improvement

Bautista (2009) introduces the *EMAX* algorithm, which follows the general algorithm outlined previously and, in particular, the QIM1 interpretation of $QI(\mathbf{x})$. The Pareto improvement criterion is explicitly defined as

$$I_{\mathbb{P}}(\mathbf{y}(\mathbf{x})) = - \max_{\mathbf{x}_i \in \mathbb{P}_{\mathcal{X}}^n} \min_{j=1, \dots, m} (y_j(\mathbf{x}) - y_j(\mathbf{x}_i)).$$

Bautista (2009) shows that $I_{\mathbb{P}}(\mathbf{y}(\mathbf{x}))$ is an attractive improvement criterion because it allows one to easily distinguish whether or not a given point in the output space is non-dominated. If $I_{\mathbb{P}}(\mathbf{y}(\mathbf{x})) > 0$, then $\mathbf{y}(\mathbf{x})$ is not dominated by any vectors in $\mathbb{P}_{\mathcal{Y}}^n$. If $I_{\mathbb{P}}(\mathbf{y}(\mathbf{x})) < 0$, then $\mathbf{y}(\mathbf{x})$ is dominated by at least one vector in $\mathbb{P}_{\mathcal{Y}}^n$. If $I_{\mathbb{P}}(\mathbf{y}(\mathbf{x})) = 0$, then $\mathbf{y}(\mathbf{x})$ is weakly dominated by at least one vector in $\mathbb{P}_{\mathcal{Y}}^n$. This improvement criterion is based on the modified maximin fitness function

$$\text{fitness}(\mathbf{y}(\mathbf{x})) = \max_{\mathbf{x}_i \in \mathbb{P}_{\mathcal{X}}^n} \min_{j=1, \dots, m} (y_j(\mathbf{x}) - y_j(\mathbf{x}_i)) \quad (19.25)$$

presented in Balling (2003), which was originally introduced as a component of a multiresponse evolutionary algorithm. One can interpret the Pareto improvement (and the closely related maximin fitness function) as a distance metric from an output space vector to the Pareto front.

A potential drawback of this particular improvement criterion, discussed in Svenson (2011), is that it is not equal to zero in the currently dominated region, and one consequence of this definition is that the improvement criterion for single-response optimization employed in Jones et al. (1998) and Schonlau (1997) is not a special case of the Pareto improvement criterion.

This issue can be easily resolved by introducing the *maximin improvement function*, which Svenson (2011) defines as

$$I_{\mathbb{M}}(\mathbf{y}(\mathbf{x})) = - \max_{\mathbf{x}_i \in \mathbb{P}_{\mathcal{X}}^n} \min_{j=1, \dots, m} (y_j(\mathbf{x}) - y_j(\mathbf{x}_i)) \mathbf{1} \left[- \max_{\mathbf{x}_i \in \mathbb{P}_{\mathcal{X}}^n} \min_{j=1, \dots, m} (y_j(\mathbf{x}) - y_j(\mathbf{x}_i)) < 0 \right].$$

See Svenson (2011) for more discussion.

19.3.4.7 Hypervolume Improvement

Emmerich et al. (2006) incorporated the hypervolume indicator into an improvement criterion. The idea is to choose the next input \mathbf{x} in our sequential design so that we get the largest possible expected increase in the hypervolume indicator. To do this, we define the *hypervolume improvement criterion* as

$$I_{\mathbb{H}}(\mathbf{y}(\mathbf{x})) = \begin{cases} 0 & \text{if } \mathbf{y}(\mathbf{x}) \leq \mathbb{P}_{\mathcal{Y}}^n \text{ or } \mathbf{y}(\mathbf{x}) \not\leq \mathbf{v}, \\ I_{\mathbb{H}}(\{\mathbf{y}(\mathbf{x})\} \cup \mathbb{P}_{\mathcal{Y}}^n, \mathbf{v}) - I_{\mathbb{H}}(\mathbb{P}_{\mathcal{Y}}^n, \mathbf{v}) & \text{otherwise,} \end{cases}$$

for some suitable reference point \mathbf{v} .

Svenson (2011) discusses additional improvement functions, the Gaussian weighted hypervolume improvement function, and the completeness indicator improvement function, based on the Gaussian weighted hypervolume and the completeness quality indicators. In addition, he discusses the QIM1 and QIM2 implementations of these and other improvement functions for Pareto optimization.

For more discussion of Pareto optimization than we give here, see Coello et al. (2006). See Svenson (2011) for a detailed discussion of improvement functions for Pareto optimization

as well as performance comparisons of improvement functions. Pareto optimization has been extensively studied and here we hope to give a sense of how improvement criteria have been utilized.

19.3.5 GASP Model with Measurement Error

The GASP model in (19.1) can be modified to allow for measurement or random error. This might occur in computer simulations with a noise component. Assume we have n runs of our simulator. Let $\tilde{Y}^i(\mathbf{x})$ denote the output for the i th run of the simulator at input \mathbf{x} . In this case, the GASP model for run i at input \mathbf{x} becomes

$$\tilde{Y}^i(\mathbf{x}) = \sum_{j=1}^k \beta_j f_j(\mathbf{x}) + Z(\mathbf{x}) + \epsilon^i(\mathbf{x}), \tag{19.26}$$

where β_j , f_j , and Z have the same values as given previously and $\epsilon^i(\mathbf{x})$ is normally distributed with mean 0 and variance δ^2 independent of $Z(\mathbf{x})$. We assume $\epsilon^i(\cdot)$ and $\epsilon^j(\cdot)$, corresponding to runs $i \neq j$, are independent, and thus, the covariance between $\tilde{Y}^i(\mathbf{x}_1)$ and $\tilde{Y}^j(\mathbf{x}_2)$ is given by

$$\text{cov}(\tilde{Y}^i(\mathbf{x}_1), \tilde{Y}^j(\mathbf{x}_2)) = \sigma_z^2 R(\mathbf{x}_1, \mathbf{x}_2) + \delta^2 1_{i=j},$$

where $1_{i=j}$ is the indicator function and equals 1 if $i = j$; otherwise, it equals 0.

We wish to predict the output $Y(\mathbf{x})$ of a hypothetical “noiseless” version of the simulator, which we model as

$$Y(\mathbf{x}) = \sum_{j=1}^k \beta_j f_j(\mathbf{x}) + Z(\mathbf{x}) \tag{19.27}$$

with

$$\text{cov}(Y(\mathbf{x}_1), \tilde{Y}^j(\mathbf{x}_2)) = \sigma_z^2 R(\mathbf{x}_1, \mathbf{x}_2).$$

If we restrict to the GASP model with constant mean for both the noisy and noiseless cases, the EBLUP $\hat{Y}(\mathbf{x})$ and corresponding (empirical) MSPE described in (19.7) and (19.8) and the discussion following are still valid for predicting $Y(\mathbf{x}^*)$, provided we replace $\sigma_z^2 \mathbf{R}$ by \mathbf{V} , where \mathbf{V} is the $n \times n$ matrix of covariances between the responses at the observed design sites, and we replace $\sigma_z^2 \mathbf{r}$ by $\mathbf{v} = \mathbf{v}(\mathbf{x}^*) = [\text{cov}(Y(\mathbf{x}^*), \tilde{Y}^1(\mathbf{x}_1)), \dots, \text{cov}(Y(\mathbf{x}^*), \tilde{Y}^n(\mathbf{x}_n))]'$, the vector of covariances between $Y(\mathbf{x}^*)$ and the $\tilde{Y}^i(\cdot)$ at the observed input sites. In addition to modeling output from computer simulations with noise, some authors advocate using this GASP model with random error (sometimes referred to as a GASP model with a nugget effect) because it has certain numerical advantages (\mathbf{R} can sometimes be nearly singular, whereas this will not be the case with \mathbf{V}) and because it accommodates the use of correlation functions that produce very smooth predictors, even if the computer simulation is not actually a smooth function.

Huang et al. (2006b) consider the goal of optimization and consider the expected improvement and algorithm of Jones et al. (1998) in the GASP model with constant mean and with random error. They use independent diffuse priors for the constant mean β as well as for σ_z^2, δ^2 , and the correlation parameters. Let $Y_p(x)$ denote the posterior distribution of the output of the hypothetical noiseless simulator as a function of the input x , and let $\hat{Y}(x)$ be its posterior mean and $s^2(x)$ its posterior variance. They use an augmented expected improvement criterion at a potential site x^* :

$$E(I(x^*)) = (E\{\max(\hat{Y}(x_0) - Y_p(x^*), 0)\}) \times \left(1 - \frac{\delta}{\sqrt{s^2(x^*) + \delta^2}}\right), \tag{19.28}$$

where x_0 denotes the current *effective best solution*. The current effective best solution is defined to be the point among all those that have been observed so far that maximizes the utility function

$$u(x) = -\hat{Y}(x) - cs(x), \tag{19.29}$$

where c is a user-specified constant. The constant c reflects the relative importance of a predicted small value of the hypothetical noiseless simulator at x versus a small prediction standard deviation. Another way to interpret this utility is that we seek the point x_0 that minimizes the upper end point of a *prediction interval* of the form $\hat{Y}(x) \pm cs(x)$ for a given c . Taking $c = 0$ corresponds to the *utility* used by Jones et al. (1998). Huang et al. (2006b) use $c = 1$.

The second term on the right side of (19.28) is a factor intended to account for diminishing returns of additional replicates as the prediction becomes more accurate. It is equal to the proportion of the reduction in the posterior standard deviation after a new replicate is added. It approaches 1 when the variance of the random errors approaches 0.

As in Jones et al. (1998), one can show that the conditional expected improvement in (19.28) given the past data and estimates of the correlation parameters is

$$E(I(x^*)) = s(x^*) \left\{ \frac{\hat{Y}(x_0) - \hat{Y}(x^*)}{s(x^*)} \Phi \left(\frac{\hat{Y}(x_0) - \hat{Y}(x^*)}{s(x^*)} \right) + \phi \left(\frac{\hat{Y}(x_0) - \hat{Y}(x^*)}{s(x^*)} \right) \right\} \times \left(1 - \frac{\delta}{\sqrt{s^2(x^*) + \delta^2}}\right) \tag{19.30}$$

where, as previously, Φ and ϕ are the standard normal probability cumulative distribution and density functions, respectively.

Not surprisingly, this approach works very well in the examples discussed in Huang et al. (2006b). Forrester and Bressloff (2006) independently propose a similar approach and demonstrate that it is effective in an airfoil example.

Picheny et al. (2013) also consider optimization in computer experiments with noise. They propose an expected improvement criterion based on quantiles. The goal is similar to the modified goal in Huang et al. (2006b). Based on training data from n observations x_1, \dots, x_n , they seek the point x_0 that minimizes

$$q_n(x) = \hat{Y}(x) + \Phi^{-1}(\gamma)s(x) \tag{19.31}$$

for a given quantile $\gamma \in [0.5, 1)$. Notice that this looks similar to $u(\mathbf{x})$ in (19.29). The improvement at a point \mathbf{x}^* is then defined to be

$$I(\mathbf{x}^*) = \min_{x \in \{x_1, \dots, x_n\}} q_n(\mathbf{x}) - q_{n+1}(\mathbf{x}^*). \tag{19.32}$$

Replacing $\widehat{Y}(\cdot)$ with $Y_p(\cdot)$ in the (19.31) and taking expectations yields the expected improvement

$$EI(\mathbf{x}^*) = \left(q_n^{\min} - \widehat{Y}(\mathbf{x}^*) \right) \Phi \left(\frac{q_n^{\min} - \widehat{Y}(\mathbf{x}^*)}{s(\mathbf{x}^*)} \right) + s(\mathbf{x}^*) \phi \left(\frac{q_n^{\min} - \widehat{Y}(\mathbf{x}^*)}{s(\mathbf{x}^*)} \right), \tag{19.33}$$

where

$$q_n^{\min} = \min_{x \in \{x_1, \dots, x_n\}} q_n(\mathbf{x})$$

and $\widehat{Y}(\cdot)$ and $s^2(\cdot)$ are the appropriate conditional mean and variance (see the appendix of Picheny et al. 2013 for details). Properties, performance, and some variations of this expected improvement are discussed in Picheny et al. (2013), and the criterion performs well.

19.3.5.1 Multiple Simulators of Differing Fidelity

Huang et al. (2006a) consider the situation in which there are m computer simulations available to observe of differing accuracy or *fidelity*. Those of low fidelity are cheaper (run more quickly) than those of higher fidelity; thus, for a give budget, one can observe the lower-fidelity simulations more frequently than the higher fidelity simulations. On the other hand, the higher fidelity simulations more accurately represent some physical process of interest, so the higher fidelity simulations are more informative than the lower fidelity simulations. Denote observations from these simulators in increasing order of fidelity by $\tilde{Y}_1(\mathbf{x}), \tilde{Y}_2(\mathbf{x}), \dots, \tilde{Y}_m(\mathbf{x})$. Let $Y_1(\mathbf{x}), Y_2(\mathbf{x}), \dots, Y_m(\mathbf{x})$ denote hypothetical “noiseless” versions of these simulator outputs.

For $l = 1, 2, \dots, m$, define

$$\zeta_l(\mathbf{x}) = \beta_l + Z_l(\mathbf{x}), \tag{19.34}$$

$$Y_l(\mathbf{x}) = Y_{l-1}(\mathbf{x}) + \zeta_l(\mathbf{x}), \tag{19.35}$$

and

$$\tilde{Y}_l(\mathbf{x}) = Y_l + \epsilon_l(\mathbf{x}) \tag{19.36}$$

where we take $Y_0(\mathbf{x}) = 0$. The $Z_l(\mathbf{x})$ are mean 0 stationary GASP with $\text{var}[Z_l(\mathbf{x})] = \sigma_l^2$ and are assumed to be all independent of each other. The $\epsilon_l(\mathbf{x})$ are normally distributed with mean 0 and variance δ_l^2 , independent of each other and of the $Z_l(\mathbf{x})$. A Gaussian correlation structure (as in (19.3) with $p = 2$) is assumed for each $Z_l(\cdot)$. Each $Y_l(\cdot)$ is modeled as a refinement (through the addition of $\zeta_l(\cdot)$) of the previous (lower fidelity) one.

This model is inspired by Kennedy and O'Hagan (2001). However, the goal in Kennedy and O'Hagan (2001) is calibration. They consider the situation in which one has data from both a computer simulator and the actual physical process being simulated. The output from the simulator is modeled by a stationary GASP. The model for the physical process is the sum of a term corresponding to the model for the simulator output, a term representing model uncertainty (arising from the fact that the computer simulator only approximates the physical process), and a term representing random error. Kennedy and O'Hagan (2001) construct a predictor for observations from the physical process based on data from both the computer simulator and the physical process. In Kennedy and O'Hagan (2001), the model for the data from the physical process is a refinement of the model for the simulator output and was the basis for the model in Huang et al. (2006b).

The goal in Huang et al. (2006a) is to find the optimum (minimum) of the highest-fidelity hypothetical noiseless simulator. Huang et al. (2006a) use independent diffuse priors for the β_l , σ_l^2 , δ_l^2 , and the correlation parameters. Let $Y_{l,p}(x)$ denote the posterior distribution of the output of hypothetical noiseless simulator l as a function of the input x , and let $\hat{Y}_l(x)$ be its posterior mean and $s_l^2(x)$ its posterior variance. Using an improvement criterion similar to that in Jones et al. (1998), Huang et al. (2006a) show that the conditional expected improvement in $Y_{m,p}(\cdot)$, based on the model for the highest level of fidelity, at a potential site x^* as compared to the current minimum $\hat{Y}_m(x_0)$ (assumed to occur at x_0), is

$$E\{\max(\hat{Y}_m(x_0) - Y_{m,p}^p(x^*), 0)\} \\ = s_l(x^*) \left\{ \frac{\hat{Y}_m(x_0) - \hat{Y}_m(x^*)}{s_m(x^*)} \Phi \left(\frac{\hat{Y}_m(x_0) - \hat{Y}_m(x^*)}{s_m(x^*)} \right) + \phi \left(\frac{\hat{Y}_m(x_0) - \hat{Y}_m(x^*)}{s_m(x^*)} \right) \right\}. \quad (19.37)$$

In order to modify this improvement criterion to select both the next site x^* at which to observe a simulator and the site at which simulator (which level of fidelity) to use, Huang et al. (2006a) multiply $E\{\max(\hat{Y}_m(x_0) - Y_{m,p}(x^*), 0)\}$ by three terms. The first of these is the correlation between $Y_{m,p}(x^*)$ and $Y_{l,p}(x^*)$:

$$\alpha_1(x, l) = \text{corr}(Y_{m,p}(x), Y_{l,p}(x)),$$

and this favors the simulators that are most highly correlated with the highest-fidelity simulator. The second term is

$$\alpha_2(x, l) = \left(1 - \frac{\delta_l}{\sqrt{s_l^2(x) + \delta_l^2}} \right),$$

and at a given input x , this favors simulators whose random error variance δ_l^2 is small relative to the prediction variance $s_l^2(x)$. In other words, at a given input x , this favors simulators that are relatively "nonnoisy." The third term is

$$\alpha_3(l) = \frac{C_m}{C_l},$$

where C_i is the cost of observing simulator i . This favors the “cheapest” simulators. Notice that this improvement criterion requires the user to specify the error variances δ_l^2 and the costs C_l .

The expected improvement algorithm is the following:

1. Perform an initial experiment of sample size $n = n_1 + n_2 + \dots + n_m$, where n_l is the number of points in the initial design allocated to simulator l , to obtain training data

$$\tilde{y}_1(x_1), \dots, \tilde{y}_1(x_{n_1}), \tilde{y}_2(x_{n_1+1}), \dots, \tilde{y}_2(x_{n_1+n_2}), \dots, \tilde{y}_m(x_{n_1+n_2+\dots+1}), \dots, \tilde{y}_m(x_{n_1+\dots+n_m}).$$

2. Fit the $\hat{Y}_l(x)$ s and compute the $s_l(x)$. Verify, through cross validation, that the $\hat{Y}_l(x)$ s are satisfactory. If not, try transformations of the y s to improve the quality of the fits of the EBLUPs.
3. Identify the point $x_{n+1} \in \mathcal{X}$ and the simulator l that maximizes the expected improvement, and run simulator l at this point to obtain $\tilde{y}_l(x_{n+1})$.
4. Update the training data by adding x_{n+1} and $\tilde{y}_l(x_{n+1})$ to the training data and setting n to $n + 1$.
5. Iterate between steps 2 and 4 until the experimental budget has been exhausted or some stopping criterion is met.
6. Estimate the minimum of the highest-fidelity simulator from the final $\hat{Y}_m(x)$.

Huang et al. (2006a) suggest that the initial design has the following properties. The initial design for the lowest-fidelity simulator should be a Latin hypercube design. For higher fidelity simulators, the initial design should be a subset of the designs used for lower fidelity simulators. Ideally, one would like these subsets to themselves be Latin hypercubes, but this is not generally possible. If these subset designs cannot be selected to be Latin hypercubes, then choose the subset of specified size with maximal minimum distance between points.

Huang et al. (2006a) propose that instead of using the minimum of $\hat{Y}_m(\cdot)$, the current best site x_0 be the site satisfying

$$x_0 = \arg \max_{x \in \{x_1, \dots, x_n\}} [u(x)],$$

where $u(x) = -\hat{Y}_m(x) - cs_m(x)$ is a utility function that is determined by a user-specified constant c that reflects the relative importance of a predicted small value of simulator m at x versus a small prediction standard deviation. Another interpretation of this utility is that we seek the site x_0 that minimizes the upper end point of a sort of *prediction interval* of the form $\hat{Y}_m(x) \pm cs_m(x)$. Taking $c = 0$ corresponds to the *utility* (global minimum) used by Jones et al. (1998). Huang et al. (2006a) use $c = 1$.

This expected improvement algorithm is applied to several numerical examples and performs sensibly.

Forrester et al. (2007) also consider optimization in a multiple fidelity model setting. They focus on settings with two levels of fidelity (although their results can be extended to more than two levels of fidelity), a cheap and an expensive simulator, and use cokriging methods for developing predictors of the simulator output. Their methods are applied to an example involving the optimization of an aircraft wing.

19.4 Contour Estimation

In Ranjan et al. (2008), the goal is to determine the set of all inputs $\mathbf{x} \in \mathcal{X}$ for which the output $y(\mathbf{x}) = a$, where a is a user-specified constant. This set defines a contour on the response surface determined by $y(\mathbf{x})$. A GASP model with constant mean is adopted for the output, and an approach, analogous to that used by Schonlau et al. (1998) and Jones et al. (1998), as expressed in (19.9), is employed.

The improvement criterion they use to sequentially add design points to the existing design is the following:

$$I(\mathbf{x}) = \epsilon^2(\mathbf{x}) - \min \left\{ (Y(\mathbf{x}) - a)^2, \epsilon^2(\mathbf{x}) \right\} \quad (19.38)$$

or equivalently, as pointed out by Roy and Notz (2013),

$$I(\mathbf{x}) = \begin{cases} \epsilon^2(\mathbf{x}) - (Y(\mathbf{x}) - a)^2, & \text{if } Y(\mathbf{x}) \in (a - \epsilon(\mathbf{x}), a + \epsilon(\mathbf{x})) \\ 0, & \text{otherwise} \end{cases}, \quad (19.39)$$

where $\epsilon(\mathbf{x}) = \alpha s(\mathbf{x})$ for some positive constant α . There is no improvement if $Y(\mathbf{x})$ is “far away from” a , that is, more than $\epsilon(\mathbf{x})$ away from a . However, if $Y(\mathbf{x})$ is “close to” a , that is, within $\epsilon(\mathbf{x})$ of a , there is improvement and the amount of improvement is a function of how close $Y(\mathbf{x})$ is to a . In particular, the improvement is $\epsilon^2(\mathbf{x}) - (Y(\mathbf{x}) - a)^2$. The maximum possible improvement is $\epsilon^2(\mathbf{x})$.

As expressed in (19.39), the improvement function has a *confidence interval* interpretation. One sees that there is improvement if the interval $Y(\mathbf{x}) \pm \epsilon(\mathbf{x})$ for the actual value $y(\mathbf{x})$ of the computer code includes a . In other words, there is improvement if a is not an “implausible” value for $y(\mathbf{x})$.

With some algebra, one can show that the expected improvement is

$$\begin{aligned} E[I(\mathbf{x})] &= \left[(\alpha s(\mathbf{x}))^2 - (\hat{Y}(\mathbf{x}) - a)^2 \right] \left[\Phi \left(\frac{a - \hat{Y}(\mathbf{x})}{s(\mathbf{x})} + \alpha \right) - \Phi \left(\frac{a - \hat{Y}(\mathbf{x})}{s(\mathbf{x})} - \alpha \right) \right] \\ &\quad + 2(\hat{Y}(\mathbf{x}) - a)s^2(\mathbf{x}) \left[\phi \left(\frac{a - \hat{Y}(\mathbf{x})}{s(\mathbf{x})} + \alpha \right) - \phi \left(\frac{a - \hat{Y}(\mathbf{x})}{s(\mathbf{x})} - \alpha \right) \right] \\ &\quad - \int_{a - \alpha s(\mathbf{x})}^{a + \alpha s(\mathbf{x})} (y - \hat{Y}(\mathbf{x}))^2 \phi \left(\frac{y - \hat{Y}(\mathbf{x})}{s(\mathbf{x})} \right) dy, \end{aligned} \quad (19.40)$$

where $\Phi(\cdot)$ and $\phi(\cdot)$ denote the cumulative distribution function and probability density function of the standard normal distribution, respectively. Although not immediately obvious, it turns out that the expected improvement at a potential input \mathbf{x}^* has both a local search component (values of \mathbf{x}^* for which $\hat{Y}(\mathbf{x}^*)$ is close to a yield improvement) and a global search component (values of \mathbf{x}^* for which the uncertainty $s(\mathbf{x}^*)$ is large, which typically occurs at values of \mathbf{x}^* that are far from the training data, can yield improvement). The greatest improvement is achieved for points \mathbf{x}^* at which $s(\mathbf{x}^*)$ is large but for which $\hat{Y}(\mathbf{x}^*)$

is close to a . We note that the constant α in the definition of $\epsilon(x)$ determines the extent to which the search is local versus global. Larger values of α produce a more global search.

Ranjan et al. (2008) use this improvement criterion as the basis for an algorithm for contour estimation. Their algorithm is as follows:

1. Perform an initial experiment of sample size n to obtain training data.
2. Fit the EBLUP and compute $s(x)$ as in (19.7) and (19.8).
3. Identify the point $x_{n+1} \in \mathcal{X}$ that maximizes $E[I(x)]$ (shown in (19.40)), and run the computer simulator at this point to obtain $y(x_{n+1})$.
4. Update the training data by adding x_{n+1} and $y(x_{n+1})$ to the training data and setting n to $n + 1$.
5. Iterate between steps 2 and 4 until the experimental budget has been exhausted or some stopping criterion is met.
6. Estimate the desired contour $\hat{C} = \{x; \hat{Y}(x) = a\}$ from the final EBLUP.

This algorithm is illustrated by examples that show it to be effective in estimating the contour. The examples also show the improvement criterion adding points that are near the true contour (local search) and occasionally points that are in regions of high uncertainty (global search).

19.5 Percentile Estimation

Roy and Notz (2013) extend the results in Ranjan et al. (2008) to estimating percentiles. Suppose the inputs x to the computer simulator represent characteristics of members of some population. For example, x could be physical characteristics of American adults. The values of x are thus distributed in some fashion in this population. We assume this distribution is known and we can view x as a random variable. We use upper case X if we want to think of the inputs as random. The distribution of X induces a distribution on $y(X)$. It may be of interest to estimate percentiles of this distribution. For example, we may wish to estimate the median or a reasonably large value such as the 90th percentile. In fact, when the actual range of values of the output $y(x)$ is not well known, it may not be possible to specify what one means by a “large” value of $y(x)$ other than through specifying a large percentile.

Roy and Notz (2013) take as their goal the estimation of the p th percentile value, ζ_p , of this induced distribution on $y(X)$. For simplicity, X is assumed to be uniformly distributed over \mathcal{X} , but the results in Roy and Notz (2013) can be easily modified to handle any known continuous distribution. As in Ranjan et al. (2008), the GASP model with constant mean is assumed for the output and an approach as expressed in (19.9) is employed. The cubic correlation function (see (19.4)) is used to model the correlation structure.

Roy and Notz (2013) discuss two types of improvement criteria. One is a confidence interval–inspired approach and is thus motivated: given we have a fair idea about the location of the p th percentile ζ_p , we should look at inputs for which the output lies in an interval about ζ_p to refine our estimate of ζ_p . The other approach is a hypothesis testing–inspired approach, namely, we try to select design points at which the response is not significantly different from ζ_p .

For either approach, an important component is estimating ζ_p from a given set of training data. Roy and Notz (2012) use the following method:

1. Generate a large random sample from the distribution of the input variables. (They use a large maximin distance Latin hypercube sample in their examples.)
2. Use the EBLUP based on the training data to predict the values of the output for this large sample.
3. Numerically estimate ζ_p from the ordered vector of predicted values obtained in step 2.

19.5.1 Confidence Interval: Based Criterion

Given the current estimate of the p th percentile value, ζ_p , of the induced distribution of the output variable, define the improvement at any untried x to be

$$I(x) = \begin{cases} h\left(Y(x) - \zeta_p, \frac{1}{s(x)}\right), & \text{if } Y(x) \in (\zeta_p - \alpha s(x), \zeta_p + \alpha s(x)), \\ 0, & \text{otherwise} \end{cases}, \quad (19.41)$$

where $h(\cdot)$ is a decreasing function of $|Y(x) - \zeta_p|$ and $\frac{1}{s(x)}$. According to this improvement criterion, if the current estimate of the p th percentile, ζ_p , lies within αs units of the response at input site x , the improvement at that design point is set equal to $h\left(Y(x) - \zeta_p, \frac{1}{s(x)}\right)$. Else, the improvement is set equal to 0. Taking

$$h\left(Y(x) - \zeta_p, \frac{1}{s(x)}\right) = (\alpha s)^g - (Y - \zeta_p)^g$$

yields the following improvement criterion:

$$I_g(x, \alpha) = \begin{cases} (\alpha s(x))^g - (Y(x) - \zeta_p)^g, & \text{if } Y(x) \in (\zeta_p - \alpha s(x), \zeta_p + \alpha s(x)), \\ 0, & \text{otherwise} \end{cases}, \quad (19.42)$$

where ζ_p is the current estimate of the p th percentile value of the induced distribution of the output variable, g is a positive even integer, and $\alpha > 0$. It is clear that the improvement function of Ranjan et al. (2008) given by (19.39) is a special case of the improvement function in (19.42), with $\zeta_p = a$ and $g = 2$. To average the improvement over the uncertainty in the predicted response surface, we take an expectation over the distribution of the response variable, which is $N(\hat{Y}(x), s^2(x))$, conditioned on the training data and parameter estimates. Here, $\hat{Y}(x)$ is the EBLUP. The design site for which the expected improvement is maximized is added to the existing design.

Roy and Notz (2013) show that the relative amounts of local versus global search in their improvement criterion $I_g(x, \alpha)$ can be controlled by changing g and α . From (19.42) it is clear that increasing α results in a wider interval of candidate design points (leading to a more

global search) from which we may pick the most “informative”, at any iteration. The use of g to control the relative amount of local versus global search was inspired by Schonlau et al. (1998).

The expected improvement (given the training data and parameters) at any design site \mathbf{x} is given by

$$E[I_g(\mathbf{x}, \alpha)] = E \left[\left\{ ((\alpha s(\mathbf{x}))^g - (Y(\mathbf{x}) - \zeta_p)^g) \cdot 1_{\zeta_p - \alpha s(\mathbf{x}) < Y(\mathbf{x}) < \zeta_p + \alpha s(\mathbf{x})} \right\} \right],$$

where 1_{\cdot} is the indicator function: if A is true, then $1_A = 1$; otherwise, $1_A = 0$. For $g = 2$, the associated expected improvement at any design site \mathbf{x} can be written as

$$\begin{aligned} E[I_2(\mathbf{x}, \alpha)] &= [(\alpha s)^2(\mathbf{x}) - (\hat{Y}(\mathbf{x}) - \zeta_p)^2] \left[\Phi \left(\frac{\zeta_p - \hat{Y}(\mathbf{x})}{s(\mathbf{x})} + \alpha \right) - \Phi \left(\frac{\zeta_p - \hat{Y}(\mathbf{x})}{s(\mathbf{x})} - \alpha \right) \right] \\ &\quad + 2(\hat{Y}(\mathbf{x}) - \zeta_p)s^2(\mathbf{x}) \left[\phi \left(\frac{\zeta_p - \hat{Y}(\mathbf{x})}{s(\mathbf{x})} + \alpha \right) - \phi \left(\frac{\zeta_p - \hat{Y}(\mathbf{x})}{s(\mathbf{x})} - \alpha \right) \right] \\ &\quad - \int_{\zeta_p - \alpha s(\mathbf{x})}^{\zeta_p + \alpha s(\mathbf{x})} (y - \hat{Y}(\mathbf{x}))^2 \phi \left(\frac{y - \hat{Y}(\mathbf{x})}{s(\mathbf{x})} \right) dy, \end{aligned} \tag{19.43}$$

where $\Phi(\cdot)$ and $\phi(\cdot)$ denote the cumulative distribution function and probability density function of the standard normal distribution, respectively.

In general, for any positive even g , the expected improvement at any design site \mathbf{x} can be written as

$$\begin{aligned} E[I_g(\mathbf{x}, \alpha)] &= (\alpha s(\mathbf{x}))^g \left[\Phi \left(\frac{\zeta_p - \hat{Y}(\mathbf{x})}{s(\mathbf{x})} + \alpha \right) - \Phi \left(\frac{\zeta_p - \hat{Y}(\mathbf{x})}{s(\mathbf{x})} - \alpha \right) \right] \\ &\quad - \sum_{k=0}^g \binom{g}{k} (\hat{Y}(\mathbf{x}) - \zeta_p)^{g-k} \int_{\zeta_p - \alpha s(\mathbf{x})}^{\zeta_p + \alpha s(\mathbf{x})} (y - \hat{Y}(\mathbf{x}))^k \phi \left(\frac{y - \hat{Y}(\mathbf{x})}{s(\mathbf{x})} \right) dy. \end{aligned}$$

19.5.2 Hypothesis Testing–Based Criterion

Roy and Notz (2013) propose another criterion that may be used to select new input settings to be observed and sequentially added to the existing design. The inspiration for this selection criterion comes from the idea of hypothesis testing. They define the *discrepancy* between the current estimate of the p th percentile, ζ_p , and $Y = Y(\mathbf{x})$, the response (as modeled by a GASP model) at any untried input setting, \mathbf{x} , to be

$$D_\epsilon(\mathbf{x}) = \begin{cases} \frac{(Y(\mathbf{x}) - \zeta_p)^2 + \epsilon}{s^2(\mathbf{x})}, & \text{if } s(\mathbf{x}) \neq 0 \\ \infty, & \text{otherwise} \end{cases}, \tag{19.44}$$

where $\epsilon > 0$. The idea is to choose that design site at which the response is not significantly different from the p th percentile, ζ_p . One wants to select points for which the discrepancy is minimized. The negative of $D_\epsilon(\mathbf{x})$ would be an improvement criterion. It is to be noted that

if a particular design site has already been observed, the MSPE, s^2 , at that point is zero; since we are working with output from computer experiments and do not want replications, we set the discrepancy of previously observed input sites to be high.

Note that if ϵ is set equal to zero in the discrepancy function given in (19.44), the expression looks like an F -statistic. The reason for the inclusion of a positive ϵ term in the expression is simple. Suppose we have two competing design sites x_1 and x_2 , and the responses at both x_1 and x_2 are equal to the current estimate of the p th percentile, ζ_p . The value of the discrepancy function as given in (19.44) sans the ϵ term is zero for both x_1 and x_2 . The question that arises is: Which design site should we pick in such a situation? Naturally, we prefer the design site where there is more uncertainty, that is, the design site that has a higher value of s^2 . The inclusion of an $\epsilon (> 0)$ results in nonzero values of the discrepancy function; more specifically, among competing design points, those with higher values of $s^2(x)$ will have smaller values of discrepancies associated with them. Thus, selecting the design site with the least value of the discrepancy function ensures that we pick the design site with the highest mean square prediction error from a set of competing design sites.

Increasing the value of ϵ leads to a more global search. As discussed earlier, the inclusion of ϵ ensures that design sites with higher values of s^2 are selected over other candidate sites. Increasing ϵ allows us to put more weight on $s^2(x)$, relative to $(Y(x) - \zeta_p)^2$.

We have seen that conditional on the observed data and parameter estimates, the response at any selected design site x is normally distributed with mean $\hat{Y}(x)$ and variance $s^2(x)$. Thus, the expected discrepancy at any design site, x , is given by

$$E[D_\epsilon(x)] = \begin{cases} \left(\frac{\hat{Y}(x) - \zeta_p}{s(x)} \right)^2 + \frac{\epsilon}{s^2(x)} + 1, & \text{if } s(x) \neq 0 \\ \infty, & \text{otherwise} \end{cases} \quad (19.45)$$

The design site x , at which $E[D_\epsilon(x)|\mathbf{y}, \beta = \hat{\beta}, \sigma_z^2 = \hat{\sigma}_z^2, \lambda = \hat{\lambda}, \zeta_p = \hat{\zeta}_p]$ is minimized, is added to the existing design.

Regardless of whether one uses the improvement criterion or the discrepancy criterion given in (19.41) and (19.44), respectively, Roy and Notz (2013) propose the following algorithm to obtain a design that would allow one to estimate the p th percentile of the induced distribution of the response variable, where the value of p is predetermined:

1. Obtain observations on an initial design.
2. Estimate the input–output relationship using all the available data, and estimate the p th percentile, ζ_p .
3. Use the estimated input–output relationship and a criterion to choose a design site at which we will next observe the output.
4. Update the data with the new design site and the corresponding observed output; update the estimates of the input–output relationship and the estimate of the p th percentile, ζ_p .
5. Iterate steps (3) and (4) until a stopping criterion is met. If the stopping criterion is satisfied, go to step (6).
6. Use the updated data to obtain the final estimate of the p th percentile value, ζ_p , of the induced distribution of the output variable. Call it ζ_p^* .

Roy and Notz (2013) use examples to demonstrate the effectiveness of their method. They note that both the improvement criterion and the discrepancy criterion given in (19.41) and (19.44), respectively, are based on the assumption that our current estimate of the p th percentile, ζ_p , is correct. This makes it very important that the improvement criterion has a strong global search component. They mention that, in practice, it is also important to have a good initial design.

19.6 Global Fit

A very challenging problem is developing criteria for the goal of producing a predictor of the simulator output that has “good” overall fit, that is, it gives “good” predictions over all of \mathbb{X} . Popular criteria are integrated MSPE, maximum MSPE, methods based on cross validation, and entropy. See Lam and Notz (2008) for more discussion. Can one development improvement criteria that are effective for the goal of good overall fit?

Lam and Notz (2008) propose the following improvement criterion:

$$I(x) = (Y(x) - y(x_{j^*}))^2, \quad (19.46)$$

where $y(x_{j^*})$ refers to the observed output at the sampled point, x_{j^*} , that is closest (in distance) to the candidate point x . They determine this nearest sampled design point using the Euclidean distance.

The intuition behind this criterion is to place high value on *informative* regions in the domain that will help improve the global fit of the model. By informative Lam and Notz (2008) mean regions with significant variation in the response values.

The expected improvement for this criterion is

$$E(I(x)) = (\hat{Y}(x) - y(x_{j^*}))^2 + s^2(x). \quad (19.47)$$

The expected improvement in (19.47) consists of two search components—local and global. The local component of the expected improvement will tend to be large at a point where it has the largest (response) increase over its nearest sampled point. The global component is large for points with the largest prediction error as defined in (19.8), that is, points about which there is large uncertainty, and these tend to be far from existing sampled points.

Lam and Notz (2008) compare the performance of this criterion with sequential and fixed sample implementations of design criteria based on MSPE, changes in measures based on cross validation error, entropy, and a space-filling design. Their method performs well on response surfaces that look nonstationary, in particular, surfaces that have regions of high variability as well as regions that are very flat. It performs satisfactorily in examples where the response surface looks stationary.

19.6.1 Emulator Maturity

Loeppky et al. (2010) investigate strategies for adding design points in batches for the goal of improving overall global fit of the GASP model. They refer to this as *emulator*

maturity, the idea being that we seek to improve our emulator (predictor) of the physical process by adding observations from the computer simulation (and, perhaps, from the physical process itself). They consider several strategies. One is the one-at-a-time sequential strategy using the expected improvement criterion of Lam and Notz (2008). Others are batch sequential implementations of the entropy criterion, integrated MSPE, maximum MSPE, and two proposed distance-based strategies that add points in batches to maintain some degree of *space fillingness*. The surprising result is that in the examples considered by Loepky et al. (2010), entropy and the proposed distance-based strategies perform best. The reason this is surprising is that neither entropy nor the distance-based strategies make use of the observed values obtained from the computer simulator. It seems counterintuitive that ignoring data should lead to improved performance. The reasons are unclear, but the following may be relevant.

First, under the GASP model, it may be that (asymptotically) one cannot do better than a design that is space filling in some reasonable sense. So to the extent that the output from the computer simulator resembles a realization of a GASP model, space-filling designs should perform well in terms of producing predictors with good overall global fit. However, one reason why EBLUPs and other predictors based on the GASP model perform well is that they are interpolators. Even if the output of the simulator does not resemble a realization of a GASP model (e.g., the response surface produced by the simulator has distinct nonstationary features), EBLUPs will often provide good global fit with enough data. Whether theoretical results about the asymptotic performance of space-filling designs for the GASP model are valid when the true response does not resemble a realization of a GASP model is not clear. And to what extent asymptotic results provide insight into performance for smaller sample sizes is also not clear. More research is probably needed. The improvement criterion of Lam and Notz (2008) appears to perform well for response surfaces that appear nonstationary and satisfactorily when the response surface appears stationary. More extensive testing on a variety of response surfaces with differing features (some that appear to be realizations of a GASP model and some that have distinctly nonstationary features) might improve understanding of the behavior of different design criteria.

Second, it may be that there are better design criteria for obtaining good global fit than those studied by Loepky et al. (2010). A challenging research question is whether there are better improvement criteria for the goal of good overall fit of the predictor than those that currently appear in the literature.

19.7 Questions for Further Study

19.7.1 Additional Improvement Criteria and Expected Improvement Algorithms

For many of the problems we have discussed, only a single improvement criterion has been investigated. In global optimization, for example, many of the improvement criteria in the literature are based on that proposed by Schonlau (1997), Schonlau et al. (1998), and Jones et al. (1998). Are there other plausible improvement criteria, and how do they perform compared to existing criteria?

In developing new criteria, it is our opinion that improvement criteria based on simple, intuitive notions (possibly motivated by statistical considerations) of *improvement* are most appealing. The basic improvement criterion of Schonlau (1997), Schonlau et al. (1998), and

Jones et al. (1998) is an example, as are those of Ranjan et al. (2008) (using the confidence interval interpretation) and of Roy and Notz (2013) (using the confidence interval and testing interpretations). Unfortunately, in settings such as constrained optimization or Pareto optimization, criteria tend to be less intuitive simply because of the complexity of the goal.

Any criterion can be modified to incorporate additional features. For example, that used in Huang et al. (2006b) for simulators of varying fidelity multiplies a criterion analogous to that in Schonlau (1997), Schonlau et al. (1998), and Jones et al. (1998) by several extra terms to capture the effects of observing different simulators. One could multiply any criterion by indicator functions that penalize points that are too close to points already observed. Such modifications may produce more effective hybrid improvement criteria, but the resulting criteria may be less intuitively appealing.

Two problems where we believe new improvement criteria are needed are those where the goal is global fit (emulator maturity) and calibration. As far as we know, there are no improvement criteria in the literature for the problem of calibration.

Another question that has not been studied in the literature is the possibility of a sort of adaptive improvement criterion. The criteria in (19.13), (19.38), (19.41), and (19.44) all contain user-defined constants that control the trade-off between local and global fit. One can imagine an algorithm that changes these constants at each stage. For example, the constant at early stages might be chosen to favor global search but at later stages favor local search. Popular optimization techniques, such as simulated annealing, do this automatically. How to adaptively modify these constants is unclear. One possibility might be to assess the quality of the fit of the EBLUP at each stage in the algorithm. If the fit is poor, use a value of the constant that favors global fit for purposes of selecting the next input for the simulator.

Another issue is selection of good stopping criteria. Some authors propose possible stopping criteria, but in most examples the expected improvement algorithm is stopped after a certain sample size is reached.

19.7.2 Suite of Test Functions

As research on criterion-based designs, including those based on improvement functions, continues to grow, it would be helpful to develop *standards* for evaluating and comparing designs. These standards should include a variety of run sizes, a variety of input sizes (values of the dimension of the input space, d), and a variety of test functions. The latter might include test functions whose shapes resemble realizations of a (stationary) GASP model. It might include functions with shapes that are clearly nonstationary in a variety of ways (the magnitude and frequency of variation in the response surface differs in different regions of the input space \mathcal{X} , e.g., the response is relatively flat in some regions but oscillates rapidly in other regions with oscillations varying in magnitude). It might include functions reflecting different correlation structures, for example, response surfaces that can be approximated by a first-order response surface and response surfaces that look a bit more like white noise. An extensive suite of test functions would also help researchers identify the sorts of response surfaces that proposed designs work well on, as well as those that they do not work well on.

In some of the early papers on global optimization, the test functions used in examples come from the literature on optimization. These include many functions with multiple optima that are either equal or nearly so. Examples are the Branin function, the three-dimensional and six-dimensional Hartman function, the Goldstein–Price function, and the Shekel function. These are discussed in Schonlau (1997). There are several online sites that

contain extensive lists of test functions for global optimization. A list of test functions for evaluating designs for computer experiments would be welcome and would provide a standard for authors seeking to introduce new designs.

19.7.3 Numerical Issues

Using a sequential design strategy (algorithm) based on an improvement criterion requires one to write code to implement the strategy. There are many nontrivial numerical problems one faces, for example, nearly singular correlation matrices in calculating the EBLUP. We have not discussed these in detail, but some of the papers cited do address these problems. Good solutions are not always available, but are important. In assessing any improvement criterion (or any criterion for selecting a design), it is important to know that it has been implemented accurately and effectively. At each stage, has one really found the point that maximizes the expected improvement? How good is the predictor (EBLUP or otherwise) at each stage? As more points are added, it sometimes happens that the quality of the EBLUP suddenly changes (perhaps because of numerical problems) and it is suddenly a poor predictor. It would be unfortunate to reject a criterion because it appeared to be inferior to competitors, only to discover that the poor performance was due to a poor implementation.

19.8 Concluding Remarks

In the spirit of optimal design, we encourage the development of sequential designs based on optimality criteria relevant to the goals of the experiment. Improvement functions are one approach to developing such criteria. They have been extensively employed for the purpose of identifying the optimum of an unknown function or functions. They have been introduced, but to a much lesser extent, for other purposes. Perhaps the most challenging application is overall global fit. I encourage those interested in design to develop new improvement criteria in new settings. This is a promising, but largely unexplored, area of experimental design. This area will benefit from improvements in numerical methods to insure that sequential designs are implemented accurately. The area will also benefit from the development of suites of test cases for comparing designs, avoiding the danger of a researcher selecting only examples that showcase a proposed method and ignoring examples in which the proposed method appears to perform poorly.

References

- Balling, R. (2003), The maximin fitness function: A multiobjective city and regional planning, in *Evolutionary Multi-Criterion Optimization*, eds. Fonseca, C., Fleming, P., Zitzler, E., Deb, K., and Thiele, L., Springer, Berlin, Germany, pp. 1–15.
- Banerjee, S., Gelfand, A. E., Finley, A. O., and Sang, H. (2008), Gaussian predictive process models for large spatial data sets, *Journal of the Royal Statistical Society Series B*, 70, 825–848.
- Banerjee, S. and Johnson, G. A. (2006), Coregionalized single- and multiresolution spatially varying growth curve modeling with application to weed growth, *Biometrics*, 62, 864–867.

- Bautista, D. C. (2009), A sequential design for approximating the Pareto Front using the expected Pareto improvement function, PhD thesis, Department of Statistics, The Ohio State University, Columbus, OH.
- Bayarri, M. J., Paulo, R., Berger, J. O., Sacks, J., Cafeo, J. A., Cavendish, J., Lin, C.-H., and Tu, J. (2007), A framework for validation of computer models, *Technometrics*, 49, 3251–3269.
- Box, G., Hunter, W., and Hunter, J. (1978), *Statistics for Experimenters*, New York: John Wiley & Sons.
- Coello, C. A. C., Lamont, G. B., and Veldhuizen, D. A. V. (2006), *Evolutionary Algorithms for Solving Multi-Objective Problems (Genetic and Evolutionary Computation)*, Secaucus, NJ, Springer-Verlag New York, Inc.
- Conti, S. and O'Hagan, A. (2010), Bayesian emulation of complex multi-output and dynamic computer models, *Journal of Statistical Planning and Inference*, 140, 640–651.
- Currin, C., Mitchell, T. J., Morris, M. D., and Ylvisaker, D. (1991), Bayesian prediction of deterministic functions, with applications to the design and analysis of computer experiments, *Journal of the American Statistical Association*, 86, 953–963.
- Emmerich, M. T., Giannakoglou, K. C., and Naujoks, B. (2006), Single- and multiobjective evolutionary optimization assisted by Gaussian random field metamodels, *IEEE Transactions on Evolutionary Computation*, 10, 421.
- Fonseca, C. M., Paquete, L., and López-Ibáñez, M. (2006), An improved dimension-sweep algorithm for the hypervolume indicator, in *IEEE Congress on Evolutionary Computation*, IEEE Press, Piscataway, NJ, pp. 1157–1163.
- Forrester, A. and Bressloff, N. (2006), Design and analysis of “noisy” computer experiments, *AIAA Journal*, 463, 2331–2339.
- Forrester, A., Sobester, A., and Keane, A. (2007), Multi-fidelity optimization via surrogate modeling, *Proceedings of the Royal Society A*, 463, 138–154.
- Fricker, T. E., Oakley, J. E., and Urban, N. M. (2013), Multivariate Gaussian process emulators with nonseparable covariance structures, *Technometrics*, 55, 47–56.
- Gelfand, A., Schmidt, A., Banerjee, S., and Sirmans, C. (2004), Nonstationary multivariate process modeling through spatially varying coregionalization, *Test (Madrid)*, 13, 263–294.
- Higdon, D. (2002), Space and space-time modeling using process convolutions, in *Quantitative methods for Current Environmental Issues*, eds. Anderson, C., Barnett, V., Chatwin, P., and El-Shaarawi, A., Springer, London, pp. 37–54.
- Higdon, D., Lee, H., and Holloman, C. (2003), Nonstationary multivariate process modeling through spatially varying coregionalization, *Bayesian Statistics 7: Proceedings of the Seventh Valencia International Meeting*, Oxford University Press, Oxford, pp. 181–197.
- Huang, D., Allen, T. T., Notz, W. I., and Miller, R. A. (2006a), Sequential kriging optimization using multiple fidelity evaluations, *Structural and Multidisciplinary Optimization*, 32, 369–382.
- Huang, D., Allen, T. T., Notz, W. I., and Zeng, N. (2006b), Global optimization of stochastic black-box systems via sequential kriging metamodels, *Journal of Global Optimization*, 34, 441–466.
- Jones, D. R., Schonlau, M., and Welch, W. J. (1998), Efficient global optimization of expensive black-box functions, *Journal of Global Optimization*, 13, 455–492.
- Journel, A. G. and Huijbregts, C. J. (1989), *Bayesian Approach to Global Optimization*, Dordrecht, the Netherlands: Kluwer Academic Publishers.
- Keane, A. J. (2006), Statistical improvement criteria for use in multiobjective design optimization, *AIAA Journal*, 44, 879–891.
- Kennedy, M. C. and O'Hagan, A. (2001), Bayesian calibration of computer models (with discussion), *Journal of the Royal Statistical Society B*, 63, 425–464.
- Lam, C. and Notz, W. (2008), Sequential adaptive designs in computer experiments for response surface model fit, *Statistics and Applications*, 6, 207–233.
- Lim, Y. B., Sacks, J., Studden, W. J., and Welch, W. J. (2002), Design and analysis of computer experiments when the output is highly correlated over the input space, *The Canadian Journal of Statistics*, 30, 109–126.
- Loeppky, J., Moore, L., and Williams, B. (2010), Batch sequential designs for computer experiments, *Journal of Statistical Planning and Inference*, 140, 1452–1464.

- Mockus, J., Tiesis, V., and Žilinskas, A. (1978), The application of Bayesian methods for seeking the extremum, in *Towards Global Optimisation*, eds. Dixon, L. C. W. and Szego, G. P., Amsterdam, the Netherlands, Vol. 2, pp. 117–129.
- Picheny, V., Ginsbourger, D., Richet, Y., and Caplin, G. (2013), Quantile-based optimization of noisy computer experiments with tunable precision, *Technometrics*, 55, 2–13.
- Ponweiser, W., Wagner, T., and Vincze, M. (2008), Clustered multiple generalized expected improvement: A novel infill sampling Criterion for surrogate models, *IEEE Congress on Evolutionary Computation*, 12, 3515–3522.
- Ranjan, P., Bingham, D., and Michailidis, G. (2008), Sequential adaptive designs in computer experiments for response surface model fit, *Technometrics*, 50, 527–541.
- Roustant, O., Ginsbourger, D., and Deville, Y. (2012), DiceKriging, DiceOptim: Two R packages for the analysis of computer experiments by kriging-based metamodeling and optimization, *Journal of Statistical Software*, 51, 1–55.
- Roy, S. and Notz, W. I. (2013), Estimating percentiles in computer experiments: A comparison of sequential-adaptive designs and fixed designs, *Journal of Statistical Theory and Practice*, 8, 12–29.
- Sacks, J., Schiller, S., and Welch, W. (1989), Designs for computer experiments, *Technometrics*, 31, 31–47.
- Santner, T. J., Williams, B. J., and Notz, W. I. (2003), *The Design and Analysis of Computer Experiments*, New York: Springer Verlag.
- Schonlau, M. (1997), Computer experiments and global optimization, PhD thesis, Department of Statistics and Actuarial Science, University of Waterloo, Waterloo, CA.
- Schonlau, M., Welch, W. J., and Jones, D. R. (1998), Global versus local search in constrained optimization of computer models, in *New Developments and Applications in Experimental Design*, eds. Flournoy, N., Rosenberger, W. F., and Wong, W. K., Institute of Mathematical Statistics, Hayward, CA, Vol. 34, pp. 11–25.
- Shewry, M., Welch, W., and Jones, D. R. (1987), Maximum entropy sampling, *Journal of Applied Statistics*, 14, 165–170.
- Svenson, J. D. (2011), Computer experiments: Multiobjective optimization and sensitivity analysis, PhD thesis, Department of Statistics, The Ohio State University, Columbus, OH.
- Taddy, M., Lee, H., Gray, G., and Griffin, J. (2009), Bayesian guided pattern search for robust local optimization, *Technometrics*, 51, 389–401.
- Viana, F., Haftka, R., and Watson, L. (2013), Efficient global optimization algorithm assisted by multiple surrogate techniques, *Journal of Global Optimization*, 56, 669–689.
- Welch, W. J., Buck, R. J., Sacks, J., Wynn, H. P., Mitchell, T. J., and Morris, M. D. (1992), Screening, predicting, and computer experiments, *Technometrics*, 34, 15–25.
- Williams, B. J., Santner, T. J., and Notz, W. I. (2000), Sequential design of computer experiments to minimize integrated response functions, *Statistica Sinica*, 10, 1133–1152.
- Williams, B. J., Santner, T. J., Notz, W. I., and Lehman, J. S. (2010), Sequential design of computer experiments for constrained optimization, in *Statistical Modelling and Regression Structures: Festschrift in Honour of Ludwig Fahrmeir*, eds. Kneib, T. and Tutz, G., Berlin, Springer Verlag, pp. 449–472.
- Zitzler, E., Knowles, J., and Thiele, L. (2008), Quality assessment of Pareto set approximations, in *Multiobjective Optimization: Interactive and Evolutionary Approaches*, eds. Branke, J., Deb, K., Miettinen, K., and Slowinski, R., Springer, Berlin, Germany, pp. 373–404.
- Zitzler, E. and Thiele, L. (1998), Multiobjective optimization using evolutionary algorithms—A comparative case study, in *Fifth International Conference on Parallel Problem Solving from Nature (PPSN V)*, eds. Eiben, A. E., Bäck, T., Schoenauer, M., and Schwefel, H. P., Springer, Berlin, Germany, 292–301.

Section VI

Cross-Cutting Issues

20

Robustness of Design

Douglas P. Wiens

CONTENTS

20.1 Introduction.....	719
20.2 Robustness against a Misspecified Response Function.....	721
20.2.1 Minimum Bias Designs: Uniformity.....	728
20.2.2 Polynomial Models of Uncertain Degree.....	730
20.2.3 Models with Truncated Series Response Functions.....	732
20.2.4 Nonlinear Regression Models.....	733
20.2.5 Generalized Linear Models.....	734
20.3 Robustness against a Misspecified Error Structure.....	735
20.3.1 Robustness against Heteroscedasticity.....	735
20.3.2 Robustness against Dependence.....	736
20.4 Special Applications.....	739
20.4.1 Extrapolation.....	740
20.4.2 Model Selection: Discrimination and Goodness of Fit.....	741
20.4.3 Dose–Response Designs.....	742
20.4.4 Clinical Trials.....	745
20.4.5 Spatial Designs: Field and Computer Experiments.....	745
20.4.6 Mixture Experiments.....	747
20.4.7 Designs for Robust Inference.....	747
References.....	748

20.1 Introduction

Robustness of design is a vast topic – the Current Index to Statistics lists over 700 articles purporting to discuss it – and so we shall first clarify our restricted interpretation of the term. By *robustness of design* we shall mean the scenario outlined in the next few paragraphs.

An investigator anticipates planning a study that will result in a number of observations on a random variable y , whose probability distribution – often merely through its expected value – depends on a vector x of covariates *that can be set by the investigator*, hence the design. After the data are gathered, the relationship between y and x is to be assessed. This will generally involve both estimation and prediction and is often done in the context of a particular model *of which the experimenter might have only partial knowledge and in which he might have little faith* – hence the robustness requirement.

Robustness has numerous meanings in statistics. The notion appears to have been introduced by Box (1953) and was given a firm mathematical basis by Huber (1964, 1981), for whom it generally – but certainly not exclusively – meant the relative insensitivity of a statistical procedure to departures from the assumed Gaussian error distribution. In design, the usual performance measures depend on the error distribution only through the first two moments, and beyond this, the distributional shape is not so relevant. One does however have in mind a particular model to be fitted once the data are gathered. In *classical* optimal design theory, one believes explicitly that the model one fits is the correct one and measures the quality of a design through a *loss function* such as the determinant, or trace, or maximum eigenvalue of the covariance matrix, corresponding to the well-known D , A and E -optimality criteria discussed in Chapters 1 and 2. In *model robust* design theory, one instead anticipates that the model that will be fitted by the experimenter is not necessarily the true one – a simple example to bear in mind is that of fitting a straight line regression when the true response function is possibly not exactly linear in the covariate – and so the loss function highlights some more general feature such as the mean squared error (*mse*). This will of course depend on the true, rather than fitted, model, and so one seeks a design minimizing some scalar quantity summarizing the increased loss – perhaps the maximum, or average, of the *mse* over the predicted values – as the true model varies.

It is the notion of model robustness on which this chapter is primarily focussed. An experimenter plans to fit a particular model to his data while realizing that any model is at best an approximation. He seeks protection, at the design stage, from increased loss incurred by model misspecification.

Some other robustness concepts – certainly worthy of discussion, but which space limitations prevent us from discussing more fully here – are as discussed as follows:

- *Criterion robustness*: Here, one aims to find a design optimizing a mixture of criteria, for instance, discriminatory power to identify a true model among a class of competitors, and estimation efficiency at this model. In some cases – for instance, if the competing models are all only vaguely specified – this can be viewed as a sub-topic of model robustness and will be discussed in such a context. For other notions of criterion robustness – for instance, designing to minimize some mixture of the D , A and E -optimality measures applied to the covariance matrix when the model is not in doubt – a pioneering reference is Kiefer (1975).
- *Robustness of designs for comparative experiments*: The robustness issues, and techniques for handling them, are very different for block designs, factorial designs, etc. than for regression-based experiments. The robustness issues tend to revolve around effects whose estimation is precluded by the design – because of blocking, or aliasing, for instance – and the concentration is more on assessing the robustness of existing designs, and choosing among them, than on constructing optimally robust designs.

In the next two sections of this chapter, we discuss model robustness in some generality, classifying by the types of model departures of particular interest. Such a classification is quite crude, and so in Section 20.4 we look at a number of special applications in which robustness is sought. Several, evidently open, problems of potential interest are noted. Our coverage of the field is of necessity quite selective, and we apologize to those whose contributions we have had to omit.

20.2 Robustness against a Misspecified Response Function

In planning a design strategy that is to be robust against response functions other than those which will be fitted by the analyst, one must first characterize the class of alternatives. The seminal work here is by Box and Draper (1959), who considered, among others, the problem of designing for a polynomial response when the class of alternatives consists of all such functions of a given higher degree. As a simple yet motivating example of their results, suppose that one is to fit a straight line in an independent variable $x \in [-1, 1]$ when in fact the mean response is $E[y(x)] = \phi_0 + \phi_1x + \phi_2x^2$ for parameters Φ . Assume uncorrelated errors with common variance σ_ϵ^2 . With (not necessarily distinct) design points x_1, \dots, x_n , and with X denoting the model matrix with the i th row $f'(x_i) = (1, x_i)$, the model to be fitted to the vector y of observations is $E[y] = X\theta$ for some $\theta = (\theta_0, \theta_1)'$; neither θ_0 nor θ_1 need equal the corresponding parameter in the ϕ -parameterization describing the true model. The least squares estimate (*lse*) is $\hat{\theta} = (X'X)^{-1} X'y$. Define $\tau_k = \sum x_i^k/n$ and assume that $\tau_1 = \tau_3 = 0$, as, for instance, is the case if the design is symmetric. Then one finds that, under the true quadratic model, the mean vector and covariance matrix of $\hat{\theta}$ are $(\phi_0 + \tau_2\phi_2, \phi_1)'$ and $(\sigma_\epsilon^2/n) \text{diag}(1, \tau_2^{-1})$, respectively, so that the predictions $\hat{y}(x) = (1, x) \hat{\theta}$ have *mse*

$$\text{MSE}[\hat{y}(x)] = E\left[\{\hat{y}(x) - E[y(x)]\}^2\right] = \frac{\sigma_\epsilon^2}{n} \left(1 + \frac{x^2}{\tau_2}\right) + \left(\phi_2(\tau_2 - x^2)\right)^2.$$

A common measure of performance is the integrated mean squared error (*imse*) of the predictors, which in this instance is

$$\text{IMSE} = \int_{-1}^1 \text{MSE}[\hat{y}(x)] dx = \left\{ \frac{2\sigma_\epsilon^2}{n} \left(1 + \frac{1}{3\tau_2}\right) \right\} + \left\{ 2\phi_2^2 \left(\left(\tau_2 - \frac{1}{3}\right)^2 + \frac{4}{45} \right) \right\}. \tag{20.1}$$

The first term in braces in (20.1) is the integrated variance and is minimized by the design – simultaneously optimal with respect to several criteria (*D*-, *A*-, *E*-, etc.) – with half of the observations made at each of $x = \pm 1$. The second – the integrated bias, dominating the first once n is sufficiently large – is minimized if $\tau_2 = 1/3$. This can be attained in many ways, but we note that it is in particular the second moment of the (continuous) uniform distribution on $[-1, 1]$. If higher order alternatives are considered, then minimization of this second term requires correspondingly higher order moments to agree with those of the uniform distribution. A design approximating the uniform is the equally spaced design $x_i = -1 + 2(i - 1)/(n - 1)$, with $\tau_2 = (1/3) + 2/(3(n - 1))$. The conclusion that can be drawn here and that was drawn in the other cases considered is that “... the optimal design in typical situations in which both variance and bias occur is very nearly the same as would be obtained if *variance were ignored completely* and the experiment designed so as to *minimize the bias alone*” (Box and Draper 1959, p. 622).

To generalize the setup of this example and to set notation to be used throughout this chapter, suppose that one has a p -vector $f(x)$ of regressors, each element of which is a

function of q functionally independent variables $x = (x_1, \dots, x_q)'$, with x to be chosen from a design space \mathcal{X} . Then the fitted model is $E[y(x)] = f'(x)\theta$, with alternatives

$$E[y(x)] = f'(x)\theta + \psi(x), \quad (20.2)$$

for some function ψ . There is an immediate problem concerning the interpretation of θ , since in the alternate models, one might equally well write $E[y(x)] = f'(x)(\theta + \Phi) + (\psi(x) - f'(x)\Phi)$ for arbitrary Φ , whence the parameter vector is not identifiable. This is avoided by first *defining* the target parameter, for instance, by

$$\theta = \arg \min_{\eta} \int_{\mathcal{X}} (E[y(x)] - f'(x)\eta)^2 dx, \quad (20.3)$$

and then *defining*

$$\psi(x) = E[y(x)] - f'(x)\theta; \quad (20.4)$$

this leads to the orthogonality requirement

$$\int_{\mathcal{X}} f(x)\psi(x) dx = \mathbf{0}. \quad (20.5)$$

Under the very mild assumption that the matrix $A = \int_{\mathcal{X}} f(x)f'(x) dx$ be invertible – equivalent to the statement that if $c'f(x) = 0$ (a.e. $x \in \mathcal{X}$), then $c = \mathbf{0}$ – the parameter defined by (20.4) and (20.5) is unique:

$$\theta = A^{-1} \int_{\mathcal{X}} f(x) E[y(x)] dx.$$

In the motivating example given earlier, $f(x) = (1, x)'$ and (20.5) applied to quadratic alternatives forces $\psi(x) = \phi_2(1/3 - x^2)$ for some ϕ_2 .

In what follows, we identify a design, denoted ξ , with its design measure – a probability measure $\xi(dx)$ on \mathcal{X} . If n_i of the n observations are to be made at x_i , we also write $\xi_i = \xi(x_i) = n_i/n$. Define

$$M_{\xi} = \int_{\mathcal{X}} f(x)f'(x) \xi(dx), \quad (20.6a)$$

$$b_{\psi, \xi} = \int_{\mathcal{X}} f(x)\psi(x) \xi(dx), \quad (20.6b)$$

and assume that M_{ξ} is invertible. The covariance matrix of the *lse* $\hat{\theta}$ is $(\sigma_{\varepsilon}^2/n)M_{\xi}^{-1}$, the bias is $E[\hat{\theta} - \theta] = M_{\xi}^{-1}b_{\psi, \xi}$, and the general version of (20.1) is found to be

$$\text{IMSE} = \int_{\mathcal{X}} \text{MSE}[\hat{y}(x)] dx = \frac{\sigma_{\varepsilon}^2}{n} \text{trace}(AM_{\xi}^{-1}) + b'_{\psi, \xi} M_{\xi}^{-1} AM_{\xi}^{-1} b_{\psi, \xi} + \int_{\mathcal{X}} \psi^2(x) dx. \quad (20.7)$$

It is obvious from (20.7), and true even if $imse$ is not the loss function, that one must bound the influence of $\psi(\cdot)$, to complete the definition of a class Ψ of alternatives in (20.2). For straight line regression, Huber (1975) defines Ψ by (20.5) together with

$$\int_{\mathcal{X}} \psi^2(x) dx \leq \tau^2/n, \tag{20.8}$$

for a given constant τ ; this class was generalized to other scenarios by Wiens (1992). That the bound be $O(n^{-1})$ is required for a sensible asymptotic treatment based on the mse – it forces the bias of the estimates to decrease at the same rate as their standard error. Marcus and Sacks (1976), Pesotchinsky (1982) and Li and Notz (1982) instead take

$$|\psi(x)| \leq \phi(x), \tag{20.9}$$

for a specified function $\phi(\cdot)$. The resulting optimal *minimax* designs (maximize over ψ and then minimize over ξ) depend on the form of ϕ , but commonly – as in the three articles just mentioned – the design mass is concentrated on a small number of extreme points of \mathcal{X} . This precludes an investigation of the response function in the interior of \mathcal{X} and so is clearly not robust.

The class Ψ defined by (20.5) and (20.8) is not immune from criticism – it is so rich that any design with finite maximum loss is necessarily absolutely continuous, hence must be approximated in order to be implemented. That this is so is intuitively clear: the Lebesgue integrals defining Ψ may be modified on sets of Lebesgue measure zero; thus, if ξ places mass on any such set – as does any discrete measure – then one can choose ψ to be arbitrarily large there, thus exploding the elements of $b_{\psi, \xi}$. A formal proof may be found in Heo et al. (2001). However, as in Wiens (1992), “Our attitude is that an approximation to a design which is robust against more realistic alternatives is preferable to an exact solution in a neighbourhood which is unrealistically sparse.” To implement a continuous design on an interval, one might place the design points at the quantiles: $x_i = \xi^{-1}((i - 0.5)/n)$, $i = 1, \dots, n$; this empirical approximation is optimal in a minimum Kolmogorov discrepancy sense (Fang and Wang 1994, §1.2) and approaches ξ weakly as $n \rightarrow \infty$. For this case, and especially for multidimensional designs, there is a variety of such approximation methods (see, e.g., Xu and Yuen 2011). Rounding strategies for implementing discrete designs with continuous weights ξ_i are discussed by Pukelsheim and Rieder (1992).

To obtain a design, robust with respect to (20.5) and (20.8), one first maximizes (20.7) under these constraints. In contrast to the next step – minimization over the class of designs – this can be done in complete generality. Let $m(x)$ be the density of ξ , and define $H_\xi = M_\xi A^{-1} M_\xi$, $K_\xi = \int_{\mathcal{X}} f(x) f'(x) m^2(x) dx$ and

$$G_\xi = K_\xi - H_\xi = \int_{\mathcal{X}} \left[\left(m(x) I_p - M_\xi A^{-1} \right) f(x) \right] \left[\left(m(x) I_p - M_\xi A^{-1} \right) f(x) \right]' dx.$$

The matrix G_ξ is clearly positive semidefinite; assume for the moment that it is positive definite and define a function $r(x) = (\tau/\sqrt{n}) G_\xi^{-1/2} (m(x) I_p - M_\xi A^{-1}) f(x)$. We have the following identities:

1. $\int_{\mathcal{X}} r(x) r'(x) dx = (\tau^2/n) I_p$.
2. $\int_{\mathcal{X}} f(x) r'(x) dx = \mathbf{0}_{p \times p}$.

3. $\int_{\mathcal{X}} f(x)r'(x)m(x) dx = (\tau/\sqrt{n}) G_{\xi}^{1/2}$.
4. $\int_{\mathcal{X}} r(x)\psi(x) dx = (\tau/\sqrt{n}) G_{\xi}^{-1/2} \mathbf{b}_{\psi,\xi}$.

As in Wiens (1992), it follows from (1) and (2) that the class $\Psi_0 = \{\psi_{\beta}(x) = \mathbf{r}'(x)\boldsymbol{\beta} \mid \|\boldsymbol{\beta}\| = 1\}$ is a subclass of Ψ (with equality in (20.8)) that is least favourable in that the supremum, over Ψ , of (20.7) is attained by a member of Ψ_0 . To see this last point, let $\psi \in \Psi$ be arbitrary and set $\boldsymbol{\beta}_{*} = G_{\xi}^{-1/2} \mathbf{b}_{\psi,\xi} / \|G_{\xi}^{-1/2} \mathbf{b}_{\psi,\xi}\|$. By (3), (20.7) evaluated at $\psi_{\boldsymbol{\beta}_{*}}$ gives

$$\text{IMSE}_{|\psi_{\boldsymbol{\beta}_{*}}} = \frac{\sigma_{\xi}^2}{n} \text{trace} \left(\mathbf{A} \mathbf{M}_{\xi}^{-1} \right) + \frac{\tau^2 \mathbf{b}'_{\psi,\xi} \mathbf{H}_{\xi}^{-1} \mathbf{b}_{\psi,\xi}}{\|G_{\xi}^{-1/2} \mathbf{b}_{\psi,\xi}\|^2} + \frac{\tau^2}{n}. \quad (20.10)$$

The Cauchy–Schwarz inequality, followed by (4), gives

$$\frac{\tau^2}{n} \geq \left| \int_{\mathcal{X}} \psi(x) \psi_{\boldsymbol{\beta}_{*}}(x) dx \right| = \frac{\tau}{\sqrt{n}} \|G_{\xi}^{-1/2} \mathbf{b}_{\psi,\xi}\|,$$

so that the *imse* (20.10) is at least as large as that – (20.7) – evaluated at ψ .

Evaluating (20.7) at $\psi_{\boldsymbol{\beta}}$ for arbitrary $\boldsymbol{\beta}$ gives

$$\text{IMSE}_{|\psi_{\boldsymbol{\beta}}} = \frac{\sigma_{\xi}^2}{n} \text{trace} \left(\mathbf{A} \mathbf{M}_{\xi}^{-1} \right) + \frac{\tau^2}{n} \boldsymbol{\beta}' \left[G_{\xi}^{1/2} \mathbf{H}_{\xi}^{-1} G_{\xi}^{1/2} + \mathbf{I}_p \right] \boldsymbol{\beta};$$

now maximizing over $\boldsymbol{\beta}$ yields the result that $\max_{\Psi} \text{IMSE}$ is $(\sigma_{\xi}^2 + \tau^2)/n$ times

$$\mathcal{L}_{\nu}(\xi) = (1 - \nu) \text{trace} \left(\mathbf{A} \mathbf{M}_{\xi}^{-1} \right) + \nu \text{ch}_{\max} \left(\mathbf{K}_{\xi} \mathbf{H}_{\xi}^{-1} \right), \quad (20.11)$$

where $\nu = \tau^2/(\sigma_{\xi}^2 + \tau^2)$ and ch_{\max} denotes the maximum eigenvalue. If \mathbf{G} is singular, then one first perturbs it to make it nonsingular and then passes to the limit – details in Heo et al. (2001).

One is now to choose $\nu \in [0, 1]$, reflecting the relative importance to the experimenter of errors due to bias rather than to variance, and minimize $\mathcal{L}_{\nu}(\xi)$. This step is highly dependent on the form of the model being fitted. In some simple cases, it can be done analytically, using variational methods to minimize $\mathcal{L}_{\nu}(\xi)$ over $\xi' = m$ subject to various side conditions. For instance, in straight line regression – under the restriction that the design be symmetric – the maximization can be carried out in two stages. At the first stage, one imposes the conditions that m has an integral of one and a fixed second moment $\gamma = \int_{\mathcal{X}} x^2 m(x) dx$; this fixes \mathbf{M}_{ξ} and hence the first term in (20.11), and one then seeks a non-negative, symmetric function minimizing the second term subject to these side conditions. The result is a partially minimized value $\min \mathcal{L}_{\nu}(\xi|\gamma)$; at the second stage, a minimization over γ is performed. See Huber (1981) for details. For bivariate regression (without interactions) on a spherical design space, details are in Wiens (1990) – see Figure 20.1 for a plot of $m(x)$, minimizing $\mathcal{L}_{\nu}(\xi)$ when $\nu = 0.36$. We note that in both of these cases, a proof that the optimal design is *necessarily* symmetric is still outstanding.

In these and the other examples in Wiens (1990, 1992), a difficulty encountered is that it is not clear which of the eigenvalues $e_1(\xi), \dots, e_p(\xi)$ of $\mathbf{K}_{\xi} \mathbf{H}_{\xi}^{-1}$ – note that these depend




FIGURE 20.1
 Minimax design density for bivariate regression without interaction on a spherical design space.

on the design – will turn out to be the maximum one when evaluated at the final, minimizing design. The usual approach is to attempt to guess correctly which one is to be optimized and to then verify at the end that the guess was correct. There are cases however in which this approach fails. One of the simplest of these is quadratic regression ($p = 3$), where for a large range of values of ν , it happens that each design ξ_k derived to minimize $l_k(\xi) = (1 - \nu)\text{trace}(AM_{\xi}^{-1}) + \nu e_k(\xi)$ turns out to result in a maximum eigenvalue in (20.11) that is not the one that was minimized: $\text{ch}_{\max}(K_{\xi_k}H_{\xi_k}^{-1}) \neq e_k(\xi_k)$. In such cases, a possible approach, detailed in Daemi and Wiens (2012), is to find designs ξ_k minimizing the k th of the competing forms $\{l_1(\xi), \dots, l_p(\xi)\}$ of $\mathcal{L}_{\nu}(\xi)$ subject to the side condition that $e_k(\xi_k)$ exceed $e_i(\xi_k)$ for $i \neq k$; the optimal design is that for which $l_k(\xi_k) = \min \{l_1(\xi_1), \dots, l_p(\xi_p)\}$.

Shi et al. (2003) instead minimize the nondifferentiable functional $\mathcal{L}_{\nu}(\xi)$ using nonsmooth optimization methods. Another approach is to restrict to a smaller but more tractable class of designs. Heo et al. (2001) consider the class of designs with densities of the form $m_{\beta}(x) = \max(0, \sum_j \beta_j f_j(x_1^2, \dots, x_q^2))$, where $f(x) = (f_1(x), \dots, f_p(x))'$ in (20.2), and minimize (20.11) numerically over β subject to the requirement that the arguments of m be exchangeable. Design points are then chosen in such a way that the resulting design has empirical moments matching those of the optimal $m_{\beta}(x)$, to as high an order as possible. For fitting a full second-order model, including linear, quadratic and interaction terms, in two variables x_1 and x_2 with $n = 48$ and $\nu = 1/6$, this yields the design in Figure 20.2.

The situation is somewhat simpler, with little loss of generality, in a finite design space. In practice, one is very often restricted to choosing from a finite, if large, set of levels of the independent variables. Li (1984) proposed designs for straight line regression, robust against departures satisfying (20.5) and (20.9) with constant ϕ , concentrated on sets of equally spaced points in the interval of interest. The resulting minimax designs spread out their mass near the end points of this interval. This is a recurring theme – a naive yet sensible and *near optimal* method of robustifying a design is to take the replicates in the classically optimal design for the model in question and spread these out into nearby clusters. This paradigm is also exemplified in Figure 20.2.

Fang and Wiens (2000) continue the approach of Li (1984). They take an N -element design space \mathcal{X} and discretize the definition of Ψ by replacing the integrals in (20.3),




FIGURE 20.2
 Restricted minimax design for fitting a full second-order model. (From Heo, G. et al., *Can. J. Stat.*, 29, 117, 2001.)

(20.5) and (20.8) by sums over \mathcal{X} . To maximize the *imse* $N^{-1} \sum_{i=1}^N E \left[\{\hat{y}(x_i) - E[y(x_i)]\}^2 \right]$, it is convenient to temporarily transform to an orthogonal basis. Let Q_1 be an $N \times p$ matrix whose columns form an orthogonal basis for the column space of the matrix with rows $\{f'(x) \mid x \in \mathcal{X}\}$; this is computed in the QR decomposition of this matrix. Augment Q_1 by $Q_2 : N \times (N - p)$ whose columns form an orthogonal basis for the orthogonal complement of this space. Then $(Q_1 : Q_2)$ is an orthogonal matrix and $\psi = (\psi(x_1), \dots, \psi(x_N))'$ is necessarily of the form $\psi = (\tau/\sqrt{n}) Q_2 c$, where $\|c\| \leq 1$. Define $\tilde{A} = N^{-1} \sum_{i=1}^N f(x_i) f'(x_i)$, $\tilde{M}_\xi = \sum_{i=1}^N \xi_i f(x_i) f'(x_i)$ and $D_\xi = \text{diag}(\xi_1, \dots, \xi_N)$. Then the analogue of (20.7) is

$$\text{IMSE} = \frac{\sigma_\xi^2}{n} \text{trace}(\tilde{A} \tilde{M}_\xi^{-1}) + \frac{\tau^2}{nN} c' \left[Q_2' D_\xi Q_1 (Q_1' D_\xi Q_1)^{-2} Q_1' D_\xi Q_2 + I_{N-p} \right] c. \quad (20.12)$$

Carrying out the maximization over c and returning to the original notation gives a direct analogue of (20.11): in terms of $\tilde{K}_\xi = \sum_{i=1}^N \xi_i^2 f(x_i) f'(x_i)$ and $\tilde{H}_\xi = \tilde{M}_\xi \tilde{A}^{-1} \tilde{M}_\xi$, the maximum loss is $(\sigma_\xi^2 + \tau^2)/n$ times

$$\tilde{\mathcal{L}}_\nu(\xi) = (1 - \nu) \text{trace}(\tilde{A} \tilde{M}_\xi^{-1}) + \nu \text{ch}_{\max}(\tilde{K}_\xi \tilde{H}_\xi^{-1}). \quad (20.13)$$

The minimization is carried out numerically – by simulated annealing, or via a genetic algorithm – yielding *exact*, that is, integer-valued, designs. See Figure 20.3 for an example, with $N = 40, n = 20$ and $\nu = .5$; this again illustrates the aforementioned paradigm.

This summary of possible approaches is by no means exhaustive. For instance, the use of *imse* to measure loss, while attractive, is not universal – one can instead apply the usual *D*-, *G*-, *A*- or *E*-criteria to the *mse* matrix of the regression parameters, rather than to the covariance matrix alone. Pesotchinsky (1978, 1982) measures the loss via a general L^p norm of the eigenvalues of the *mse* matrix; Marcus and Sacks (1976) and Li and Notz (1982) use a weighted trace of this matrix. Wiens (1993) maximizes the coverage probability of confidence ellipsoids. Zhou (2008) assumes a finite design space and obtains minimax *D*-optimal




FIGURE 20.3
Minimax design ($n = 20$) for approximate cubic regression when bias and variance receive equal weight.

designs, with the maximum taken over the discrete version of Ψ described earlier. Adewale and Wiens (2006) *average* the *imse* (rather than maximizing it) with respect to a uniform prior on c in (20.12). Exact designs are then obtained by simulated annealing.

Various approaches, besides those described thus far, have been proposed with the aim of restricting Ψ while still hoping to capture a broad spectrum of alternatives. Fedorov et al. (1998) take $\psi(x)$ to be a linear combination of additional regressors $\{v_j(x)\}$, that is, functions other than the elements of $f(x)$. Notz (1989) places a prior distribution on possible departures from the assumed model, and Rychlik (1987) assumes a finite-dimensional Ψ constrained by (20.5) and (20.8); both go on to specialize to polynomial responses with polynomial alternatives, a class of applications that will be discussed in more detail in Section 20.2.2. Yue and Hickernell (1999) assume that ψ comes from a reproducing kernel Hilbert space. The definition of such alternatives is quite technical, and it can be difficult to see just how broad the class is; the approach is however promising in that, in the examples given, it leads to designs affording a full exploration of the design space.

Some attempts have been made to construct designs tailored to estimates that are linear in the observations, but are not necessarily *lse*'s. Karson et al. (1969) revisit the setup of Box and Draper (1959) and propose linear estimates that minimize the integrated squared bias (*isb*) resulting from fitting a low-order polynomial when the true response function contains higher order terms. For such estimates, they go on to propose designs that minimize the integrated variance. Draper and Sanders (1988) continue the design aspects of this approach and search for parsimonious designs meeting these criteria. Sacks and Ylvisaker (1984) study designs, robust against certain very broad classes Ψ of nonparametric alternatives, to be used with linear estimates. They comment that "...when Ψ is an infinite dimensional class ... it is by no means certain that the optimal choice of both design and linear estimates leads to least squares nor is there much a priori justification for adherence to least squares". Their approach is summarized in Chang and Notz (1996, Section 3.2) and extended in Tang (1993).

20.2.1 Minimum Bias Designs: Uniformity

In the development of Box and Draper (1959) described earlier, a motivating goal was the maximization of the power of a test of lack of fit (*lof*); this led in a natural way to considerations of bias as well as variance, thence to the prescription of choosing a *near uniform* design if consideration focusses on bias alone. Subsequent researchers have been led by this to investigate various optimality properties of the design with continuous uniform measure $\lambda(x)$, that is, Lebesgue measure on the design space \mathcal{X} , normed to have unit total mass. Implicit here is an assumption that the optimality will carry over to appropriate discretizations of this uniform measure; exactly how this should be done seems to be an area where further investigation could be fruitful (but see Fang and Wang 1994).

To describe the current situation, recall that the usual test of *lof* takes replicates at each of several locations, with covariates $\{x_j\}$, and compares the unstructured *full* model $E[y_j] = \mu_j$ with the *restricted* model $\mu_j = f'(x_j) \theta$ being fitted. If the class of full models is written in the form (20.2), with θ defined by (20.3), then the power of the F-test of *lof* is an increasing function of the noncentrality parameter $(n/\sigma_\epsilon^2) \mathcal{B}(\psi, \xi)$, where

$$\mathcal{B}(\psi, \xi) = \int_{\mathcal{X}} \{E[\hat{y}(x) - y(x)]\}^2 \xi(dx) = \int_{\mathcal{X}} \psi^2(x) \xi(dx) - \mathbf{b}'_{\psi, \xi} \mathbf{M}_\xi^{-1} \mathbf{b}_{\psi, \xi}.$$

Maximizing the power is equivalent to maximizing $\mathcal{B}(\psi, \xi)$. This goal is closely related to the observation that the expectation of the regression-based estimate of σ_ϵ^2 – the mean square of the residuals – is given by $\sigma_\epsilon^2 + n\mathcal{B}(\psi, \xi)/(n - p)$. We note as well that the *isb* of $\hat{y}(x)$ in the full model is, as at (20.7),

$$\text{ISB}(\psi, \xi) = \int_{\mathcal{X}} \{E[\hat{y}(x) - y(x)]\}^2 dx = \int_{\mathcal{X}} \psi^2(x) dx + \mathbf{b}'_{\psi, \xi} \mathbf{M}_\xi^{-1} \mathbf{A} \mathbf{M}_\xi^{-1} \mathbf{b}_{\psi, \xi}.$$

Now denote by Ψ^- the class given by (20.5) and (20.8) and by Ψ^+ the class that has the inequality reversed in (20.8). For testing *lof*, which we now express through the null hypothesis that $\psi \equiv 0$, it is necessary to separate the null and alternate hypotheses; hence, we restrict to Ψ^+ . In Wiens (1991), it is shown that the uniform design has a *maximin* property – it maximizes the minimum (over Ψ^+) power of the test of *lof*, with $\mathcal{B}(\psi, \lambda) = \int_{\mathcal{X}} \psi^2(x) \lambda(dx) \geq \tau_0^2 \stackrel{\text{def}}{=} \tau^2 / \int_{\mathcal{X}} dx$. The proof consists of showing that if the design ξ is non-uniform, then one can construct $\psi \in \Psi^+$ with $\mathcal{B}(\psi, \xi) \leq \tau_0^2$. The same method of proof establishes a dual, *minimax* property of λ – it minimizes the maximum (over Ψ^-) bias of the mean square of the residuals as an estimate of σ_ϵ^2 . Similarly, λ is also a minimax (over Ψ^-) design with respect to $\text{ISB}(\psi, \xi)$, with $\text{ISB}(\psi, \lambda) = \int_{\mathcal{X}} \psi^2(x) \lambda(dx)$.

Xie and Fang (2000) obtain similar optimality results in a nonparametric regression setting. They assume that $E[y(x)] = g(x)$ is to be estimated by a linear combination of the members of a set $\Psi_k = \{\psi_i\}_{i=1}^k$ of square-integrable and functionally independent but otherwise arbitrary functions and show that the maximum, over all such approximations, of *isb* is minimized by the continuous uniform design. They study examples in which Ψ_k consists of Fourier regression functions, or Haar wavelets.

Wiens (2000) works with the class Ψ^- and allows *weighted* least squares (*wls*) estimates (while continuing to assume homoscedastic errors), with weights $w(x)$. The analogues of $\mathbf{b}_{\psi, \xi}$ and \mathbf{M}_ξ at (20.6) are $\mathbf{b}_{\psi, w, \xi} = \int_{\mathcal{X}} \mathbf{f}(x) \psi(x) w(x) \xi(dx)$ and $\mathbf{M}_{w, \xi} = \int_{\mathcal{X}} \mathbf{f}(x) \mathbf{f}'(x) w(x) \xi(dx)$. It follows that if ξ has a density $k(x)$ with $w(x) k(x) \propto \lambda(x)$

(implying that $M_{w,\xi} \propto A$), then $b_{\psi,w,\xi} \equiv 0$, that is, $E[\hat{\theta} - \theta | \psi] \equiv 0$. Within the class of such unbiased designs, a member minimizing the integrated variance of $\hat{y}(x)$ is found. This integrated variance is

$$\frac{\sigma^2}{n} \int_{\mathcal{X}} \frac{1}{w(x)} dx \int_{\mathcal{X}} f'(x) A^{-1} f(x) w(x) dx,$$

minimized by weights

$$w(x) \propto \frac{1}{\sqrt{f'(x) A^{-1} f(x)}}. \tag{20.14}$$

Examples of the resulting *minimum variance unbiased (mvu)* designs and weights are shown in Figure 20.4, for approximate polynomial regression.

Figure 20.4 illustrates another recurring feature that occurs not only in robust design theory but also in optimal design theory when the model is not in doubt and only the variance is minimized. This is the similarity of the *I* and *D*-optimal designs to each other and the quite different nature of the *A*-optimal designs. (In Figure 20.4 *I*-optimality is referred to as

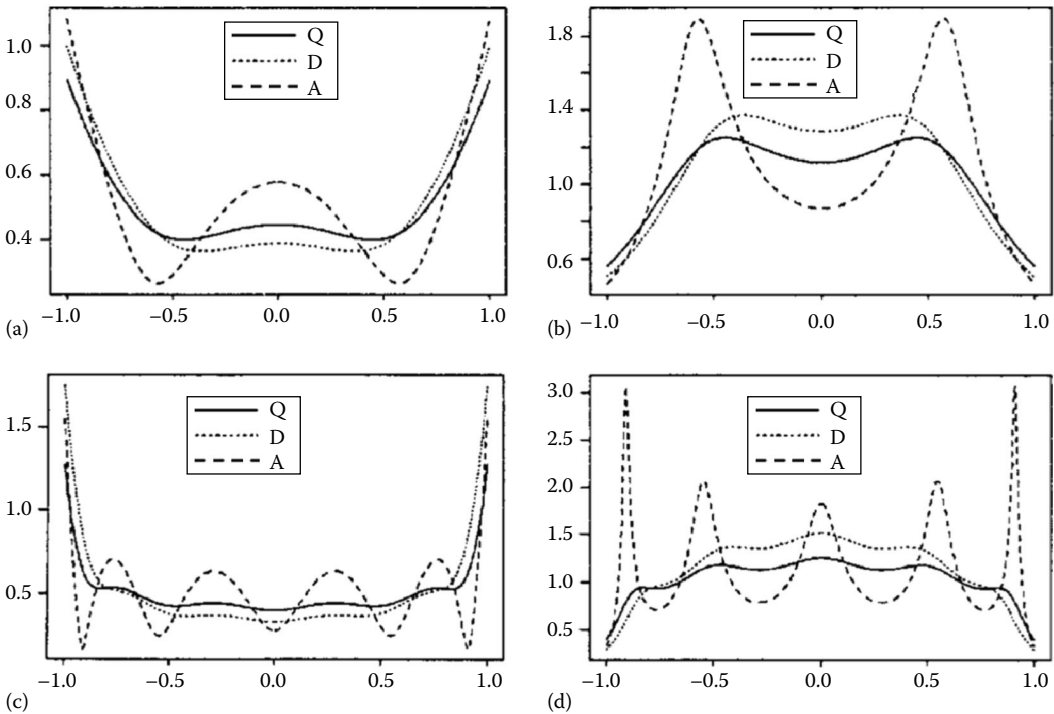


FIGURE 20.4 *mvu* design densities and weights for degree- q approximate polynomial regression and various loss functions: (a) design densities, $q = 2$; (b) weights, $q = 2$; (c) design densities, $q = 5$; (d) weights, $q = 5$. The loss functions corresponding to *Q*- ($=I$), *D*- and *A*-optimality are respectively *imse*, the determinant of the *mse* matrix and the trace of the *mse* matrix. (From Wiens, D.P., *J. Stat. Plan. Inference*, 83, 395, 2000.)

Q-optimality, a term that is now out of fashion.) This seems to be related to the fact that the former criteria are range invariant, whereas the latter is not – a linear transformation of the design space incurs not merely a linear transformation of the A -optimal design points but a change in the design weights as well. See the discussion in Wierich (1988).

Biedermann and Dette (2001) extend the results of Wiens (1991, 2000) in several directions. They introduce a function $v(x)$ whose role can be characterized by replacing (20.3) by $\theta = \arg \min_{\eta} \int_{\mathcal{X}} (E[y(x)] - f'(x)\eta)^2 v(x) dx$ and allow for heteroscedastic errors with variance $\sigma^2(x)$. Using the method of proof in Wiens (1991), they show that the minimum powers of each of several tests of *lof* are maximized by weights inversely proportional to $\sigma^2(x)$ and by a design density with $w(x)k(x) \propto v(x)$. Bischoff and Miller (2006) extend Wiens (1991) and Biedermann and Dette (2001) – they propose that one design to minimize the variance, subject to a lower bound on the power of the *lof* test.

We see from this that the unique, bias-minimizing role of the uniform design is determined, through (20.3), by the experimenter's specification of just what *bias* is to mean – indeed, Biedermann and Dette (2001, p. 223) state that "...the optimality property of a particular design with respect to Wiens (1991) maximin criterion is intimately related to the particular measure used in the definition of the set Ψ ". A comprehensive study of other measures used to define the target model would no doubt yield interesting unbiased designs.

20.2.2 Polynomial Models of Uncertain Degree

When the target response function being fitted is a polynomial in a single covariate x , a more focussed form of model robustness is to entertain only other polynomials, of different degrees than that being fitted. Then the matrices M_{ξ} under all models and the bias vector $b_{\psi, \xi}$ (with ψ containing the powers of x of interest but not being fitted) defined in Section 20.2 depend solely on the moments $c_k(\xi) = E_{\xi}[X^k]$ (hence the common use of *moment matrix* to refer to M_{ξ} in polynomial regression). Consequently, the mathematical development is closely related to the theory of moment spaces and of *canonical moments*. Roughly speaking, the canonical moment $p_{n+1}(\xi)$ describes the possible values of $c_{n+1}(\xi)$, among all distributions ξ with given values of $c_k(\xi)$ for $1 \leq k \leq n$. For a book-length treatment, see Dette and Studden (1997).

Studden (1982) considered the following problem (extending a proposal of Stigler 1971). One intends to fit a polynomial of degree r to observations at $x \in [0, 1]$, but wishes to be able to determine whether or not the coefficients β_{r+s} of x^{r+s} ($1 \leq s \leq m - r$) would be statistically significant were these coefficients to be assessed. With $f_1(x) = (1, x, \dots, x^r)'$, $f_2(x) = (x^{r+1}, \dots, x^m)'$ and $M_{ij}(\xi) = \int_0^1 f_i(x)f_j'(x)\xi(dx)$, the moment matrix for the model of degree m is partitioned as

$$M_m(\xi) = \int_0^1 \begin{pmatrix} f_1(x) \\ f_2(x) \end{pmatrix} (f_1'(x), f_2'(x)) \xi(dx) = \begin{pmatrix} r+1 & m-r \\ M_{11}(\xi) & M_{12}(\xi) \\ M_{21}(\xi) & M_{22}(\xi) \end{pmatrix} \begin{matrix} r+1 \\ m-r \end{matrix}. \quad (20.15)$$

Note that $M_{11}(\xi) = M_r(\xi)$ is the moment matrix for the fitted model of degree r . Were $\hat{\beta}_{r+1}, \dots, \hat{\beta}_m$ to be fitted, their covariance matrix would be the inverse of $M_{22.1} \stackrel{\text{def}}{=} M_{22} - M_{21}M_{11}^{-1}M_{12}$, and so an extended D -optimality criterion is to maximize $\det(M_{11}(\xi))$, subject to a lower bound

$$\det(\mathbf{M}_{22.1}) \geq \rho^{m-r} \sup_{\eta} \det(\mathbf{M}_{22.1}(\eta)), \tag{20.16}$$

for given $\rho \leq 1$. Studden (1982) phrases and solves this problem in terms of the canonical moments and then finds the design ξ_{rm} with the optimal canonical moments. For instance, ξ_{12} , for fitting a straight line with a quadratic alternative, is of the form

$$\xi = \{(0, \alpha), (1/2, 1 - 2\alpha), (1, \alpha)\}, \tag{20.17}$$

with $\alpha = (1 + \sqrt{1 - \rho}) / 4$.

Dette and Studden (1995) consider designs for polynomial regression when the experimenter seeks reasonable efficiency for all degrees $r = 1, 2, \dots, K$. They define the D -efficiency

$$eff_r(\xi) = \left(\frac{\det(\mathbf{M}_r(\xi))}{\sup_{\eta} \det(\mathbf{M}_r(\eta))} \right)^{1/r+1}. \tag{20.18}$$

They then seek designs maximizing the weighted p -norm

$$\Phi_{p,w} = \left(\sum_{r=1}^K w_r (eff_r(\xi))^p \right)^{1/p}, \quad -\infty \leq p \leq 1, \tag{20.19}$$

for nonnegative weights w_1, \dots, w_K . Again the problem is first formulated and solved in terms of canonical moments. For instance, suppose that $K = 2$ (linear or quadratic regression), $x \in [0, 1]$ and the weights are uniform: $(w_1, w_2) = (1/2, 1/2) \stackrel{def}{=} w_0$. Then the Φ_{p,w_0} -optimal design maximizing (20.19) is of the form (20.17) with α depending on p . In fact for each $p \in [-\infty, 1]$, there is $\rho \leq 1$ – recall (20.16) – such that the design of Studden (1982) is the Φ_{p,w_0} -optimal design of Dette and Studden (1995) given here. This is an instance of a more general result – that any discrete design is $\Phi_{p,w}$ -optimal for some choices of p and w – due to Dette (1991).

When the fitted polynomial response of degree m is in doubt, one might seek a high D -efficiency $eff_m(\xi)$ and simultaneously a high power for the t-test of the hypothesis that the coefficient of x^m , in a model of degree m , is zero. This latter quantity is maximized by a design maximizing the D_1 -efficiency $eff_m^{D_1}(\xi)$. In the notation leading to (20.15), with $r = m - 1$, the D_1 -efficiencies can be defined by

$$eff_m^{D_1}(\xi) = \left(\frac{\det(\mathbf{M}_{22.1}(\xi))}{\sup_{\eta} \det(\mathbf{M}_{22.1}(\eta))} \right)^{1/m-r}.$$

(In the present case, the exponent is 1 and $\mathbf{M}_{22.1}$ is a scalar; we present the D_1 -efficiency in this more general form for later reference.) Dette and Franke (2001) construct designs that maximize the minimum of $\{eff_m(\xi), eff_{m-j}^{D_1}(\xi), \dots, eff_{m+k}^{D_1}(\xi)\}$ for nonnegative j, k . Again, canonical moments play a central role in the derivations. Fang (2006) instead maximizes a weighted average of the logarithms of $eff_m(\xi)$ and $eff_{m-1}^{D_1}(\xi)$.

Rather than considering only polynomial alternatives, one might entertain alternatives with polynomial bounds. Liu and Wiens (1997) consider degree r polynomial fits with alternatives of the form (20.2) with $\psi(x) = x^{r+1}\pi(x)$, where $\pi(\cdot)$ is continuous and bounded in absolute value by a given function $\phi(\cdot)$. Here we discuss only the case of constant $\phi(\cdot)$, so that ψ is bounded by a fixed multiple of x^{r+1} . Liu and Wiens optimize with respect to several criteria, one of which is to maximize the determinant of the moment matrix $M_r(\xi)$, subject to a bound on the normalized bias: $\sup_{\psi} \mathbf{b}'_{\psi, \xi} M_r^{-1}(\xi) \mathbf{b}_{\psi, \xi} \leq d^2$. Designs optimal with respect to this criterion are termed *bounded-bias* (BB) designs, and it was shown that BB designs with $r + 1$ support points always exist. Of course these do not allow higher order models to be explored; for instance, when $r = 1$ – straight line regression – with $x \in [0, 1]$, these BB designs place equal mass at two points $\{c, 1 - c\}$, with c depending on d . Fang and Wiens (2003) employ the theory of canonical moments to extend the notion to generalized bounded bias (GBB) designs. They find that if d^2 is less than the value of $\sup_{\psi} \mathbf{b}'_{\psi, \xi_D} M_r^{-1}(\xi) \mathbf{b}_{\psi, \xi_D}$ attained by the D -optimal design ξ_D for fitting a polynomial of degree r , then the BB criterion does not determine all of the canonical moments and that more design points may then be added. They do this so as to satisfy the additional requirement of maximizing $eff_m^{D_1}(\xi)$ for a fixed $m > r$. For instance, if $r = 1$ and $m = 2$, the GBB designs on $[0, 1]$ are of the form (20.17) with $\alpha = d/2 < 1/2$. If $r = 2$ and $m = 3$, there are 4 symmetrically placed design points, including 0 and 1.

20.2.3 Models with Truncated Series Response Functions

Polynomial models sometimes arise by truncating a Taylor series expansion of the response function. This interpretation is particularly appealing when the fitted terms $f(x)$ in (20.2) constitute the first terms in a series $\sum_{i=1}^{\infty} \theta_i f_i(x)$ of orthogonal functions (over \mathcal{X} and with respect to some weight function); then $\psi(x)$ may be viewed as representing the tail of the series. Oyet and Wiens (2000, 2003) and Oyet (2002) obtain minimax designs, minimizing (20.11) and its discrete or continuous analogues for loss functions such as the trace or determinant of the *mse* matrix, when the elements of $f(x)$ are Haar wavelets or multiwavelets. For these functions and with $\mathcal{X} = [0, 1]$, $\int_{\mathcal{X}} f_i(x) f_j(x) dx = \delta_{i=j}$ and then *mvu* designs and weights assume a particularly simple form: from (20.14), we obtain design densities $k(x) \propto \|f(x)\|$ and weights $w(x) \propto 1/\|f(x)\|$. See Figure 20.5.

Practitioners using orthogonal series expansions sometimes avoid doing least squares, since the matrix whose inversion is required can be huge. Instead, *direct estimation* is




FIGURE 20.5

Minimum variance unbiased design densities for multiwavelet approximations with $N \cdot 2^{m+1}$ regressors. Each density has period $2^{-(m+1)}$; the values of N and m are (a) (2,0), (b) (2,2), (c) (3,0), and (d) (3,2). (From Oyet, A.J. and Wiens, D.P., *J. Nonparam. Stat.*, 12, 837, 2000.)

employed: if $E[y(x)] = \sum_{i=1}^p \theta_i f_i(x)$, with $\int_{\mathcal{X}} f_i(x) f_j(x) \mu(x) dx = \delta_{i=j}$ for some function $\mu(\cdot)$, then $\theta_j = \int_{\mathcal{X}} E[y(x)] f_j(x) \mu(x) dx$ and estimates are computed via a suitable discretization – $\hat{\theta} = n^{-1} \sum_{i=1}^n y(x_i) f(x_i) \mu(x_i)$. Dette and Wiens (2008) propose the modified estimate

$$\hat{\theta} = \frac{1}{n} \sum_{i=1}^n f(x_i) \frac{y(x_i) w(x_i) \mu(x_i)}{k(x_i)},$$

where $k(x)$ is the design density and weights $w(x)$ are introduced for flexibility in meeting robustness requirements. Constant weights, with $k(x) \propto \|f(x)\| \mu(x)$, yield asymptotically unbiased estimates that minimize the *imse* within the class of such unbiased estimates; otherwise, both weights and designs are chosen to minimize the maximum *imse* unconditionally. Examples are given with expansions in terms of Zernike polynomials and to spherical harmonics; in the latter case, similar results are obtained for *wls* estimation in Dette and Wiens (2009). Dette and Melas (2010) study bias-minimizing designs, with the bias arising through truncation of an expansion of the response function as a series of spline functions.

20.2.4 Nonlinear Regression Models

A problem that arises immediately in considering a design for nonlinear regression is that, by the nature of nonlinear models, the common measures of loss or efficiency depend on the unknown values of the parameters. This problem is often handled by constructing a *locally optimal* design – one that is optimal only at a particular value θ_0 of the parameter. The designer hopes that the optimality will extend to nearby parameter values. The choice of θ_0 might arise from the experimenter's prior knowledge, or perhaps as an estimate from an earlier experiment.

A mild form of robustness allows for uncertainty about the parameter values but not the form of the response function. A common approach is to first maximize the chosen loss function (or minimize an efficiency measure such as the determinant of the information matrix) over a neighbourhood of a local parameter θ_0 and to then optimize over the class of designs. In some cases, there are easily checked first-order conditions, to verify if a design is minimax. King and Wong (2000) use this approach to construct minimax *D*-optimal designs for the two-parameter logistic model and give a numerical algorithm for computing the designs; depending on the ranges of the parameters in their examples, these have up to 9 design points. Dette and Biedermann (2003) consider instead the two-parameter Michaelis–Menten model, with $E[y|x, \theta] = \theta_0 x / (\theta_1 + x)$. This model is conditionally linear in θ_0 , with the consequence that the determinant of the information matrix (computed under a normal likelihood) of the *lse*'s of θ_0 and θ_1 depends only on θ_1 . They find designs that maximize the minimum value of the corresponding *D*-efficiency, which is (20.18) with $r = 2$ and with $M_2(\xi)$ replaced by the information matrix. The resulting designs are supported on two points. Dette et al. (2003) carry out a similar program for this Michaelis–Menten model, concentrating on *E*-optimality rather than *D*-optimality; in some cases, the designs are supported on three, rather than two, points. For various exponential models, see Dette et al. (2006) and Dette and Pepelyshev (2008).

Another manner in which one can deal with this problem is through Bayesian optimality. Here the loss function is averaged, with respect to a *prior* distribution on the parameters; this averaged value is then minimized. We mention Chaloner and Larntz (1989) and Dette and Neugebauer (1997) as representative approaches. We note however that at this point,

no particular viewpoint on Bayes inference is required; the adoption of a prior may be seen merely as a convenient means to an end.

The minimax and Bayesian designs described here often have very few design points, either as a consequence of the criterion or because the mathematical difficulties preclude a search for optimal designs with more support points than regression parameters. Moreover, as emphasized by Ford et al. (1989, p. 54), "Of more serious concern is the misspecification of the model itself. Some static designs offer little or no scope for testing the validity of the assumed model. Indeed, if the model is seriously in doubt, the forms of design which we have considered may be completely inappropriate." To address this issue, Sinha and Wiens (2002) entertain a class of nonlinear models forming a neighbourhood of that which the experimenter will fit, maximize the loss over this neighbourhood and then sequentially construct minimizing designs. The model development parallels (20.2) through (20.5), but with both $f(x|\theta) = \partial E[y|x,\theta] / \partial \theta$ and $\psi(x|\theta)$ depending on the parameter; thus, robustness is being sought against alternate response surfaces tangent, at the parameter defined by (20.3), to that being fitted. Then (20.7) defines $IMSE(\xi|\psi, \theta)$ under homoscedasticity, with alternate forms applied if the design space is discrete or if the errors are heteroscedastic. Sinha and Wiens (2002) suppose that n observations have been made, defining the design ξ_n and yielding an estimate $\hat{\theta}_n$. They estimate ψ by smoothing the residuals, and choose the next design point, or batch of design points, so as to minimize $IMSE(\xi|\psi, \hat{\theta}_n)$, where now ξ denotes ξ_n augmented by the potential additional points. The result of a simulation study is shown in Figure 20.6. The fitted response function is exponential: $E[y|x,\theta] = \theta_0(1 - e^{-\theta_1 x})$; the true response function used in the simulations is either this exponential or the nearby Michaelis-Menten, both for $x \in [.5, 5]$. Comparisons are made to the uniform and locally D -optimal design (with the local parameter being based on a small initial design). See also Sinha and Wiens (2003) for the asymptotic theory of this sequential approach.

If one insists on static designs with robustness against model misspecifications of this form, then local optimality combined with a minimax approach is a possibility; Xu (2009a) finds designs minimizing $\max_{\psi} IMSE(\xi|\psi, \theta_0)$ for a local parameter θ_0 . One can instead take a Bayesian approach as described earlier, leading to the requirement that one minimize $\int_{\Theta} [\max_{\psi} IMSE(\xi|\psi, \theta)] p(\theta) d\theta$ - some details are in Karami and Wiens (2011), who use a discrete design space and show that $\max_{\psi} IMSE(\xi|\psi, \theta)$ is as at (20.13), evaluated at $f(\cdot|\theta)$. The minimization is carried out using a genetic algorithm.

20.2.5 Generalized Linear Models

Generalized linear models (GLMs) form an increasingly popular class, for which some classical design issues were surveyed by Khuri et al. (2006). In the standard setup, it is assumed that observations are made on a random variable y , with density $p(y; \alpha, \phi) = \exp\left\{\frac{\alpha y - b(\alpha)}{a(\phi)} + c(y, \phi)\right\}$. For a vector x of covariates and vector $f(x)$ of regressors, the mean response $E[y|x] = \mu(x)$ satisfies $g(\mu(x)) \stackrel{def}{=} \eta(x) = f'(x)\theta$ for a strictly monotonic, differentiable function g . Some robustness issues that might be addressed at the design stage are then misspecifications of $\eta(x)$ and uncertainties about the form of $g(\cdot)$ or $a(\cdot)$. As in nonlinear regression, a further concern is the dependence of the measures of loss on the unknown parameters. Abdelbasit and Plackett (1983) discuss robustness of locally optimal designs to poorly chosen local values in binary models. Again in a context of binary models, King and Wong (2000) propose maximin procedures, akin to those discussed in Section 20.2.4, with the min taken over a range of parameter values.

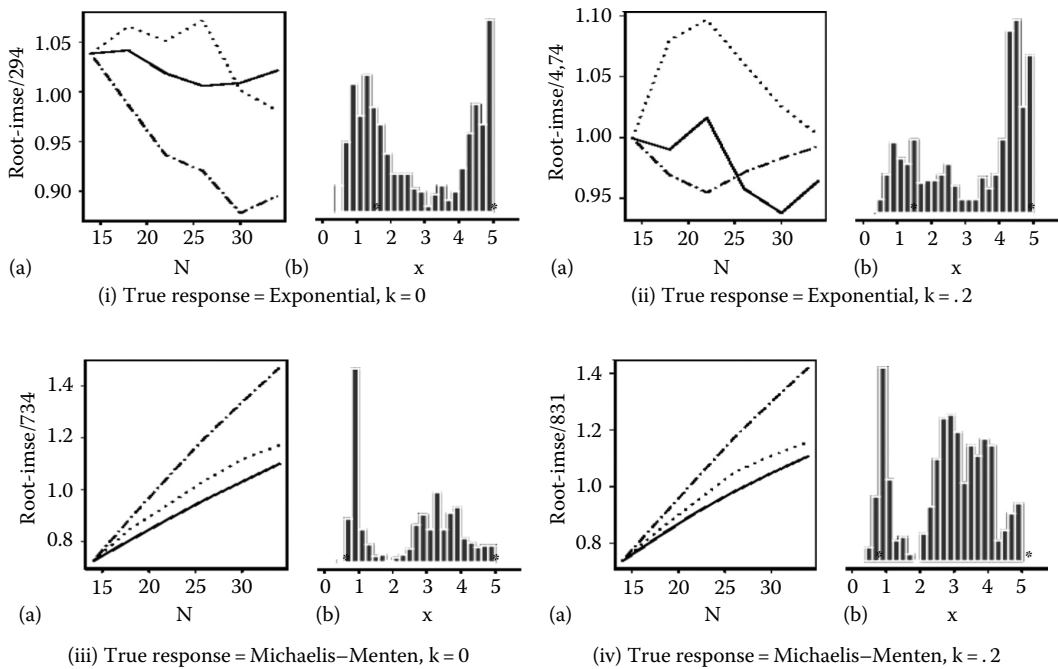


FIGURE 20.6

Sequential robust designs versus locally D -optimal and uniform designs. Fitted response is exponential; true response is either exponential or Michaelis–Menten. The fitting assumes homoscedasticity, whereas the true variance function is $\sigma^2(x) = 1 + k(x - .5)^2$ for $k = 0$ or $k = .2$ as indicated in the captions. (a) Average (over 100 sample paths) normalized $\sqrt{n} \cdot IMSE$ of robust (—), uniform (\cdots) and D -optimal ($-\cdot-\cdot-$) designs. (b) Probability histogram of all points chosen by the 100 sequential designs; asterisks are at the average sites of the D -optimal designs. (From Sinha, S. and Wiens, D. P., *Can. J. Stat.*, 30, 601.)

Woods et al. (2006) (see also Dror and Steinberg 2006) for a discussion of the computational aspects) seek *compromise* designs. These are defined as maximizers of the average of the log-determinant of the information matrix resulting from a design with a particular g , η and θ , with the average taken over various sets of these three quantities. Adewale and Wiens (2009) consider logistic models and robustness only against the form of $\eta(\cdot)$, which is allowed to vary over a class defined by discrete versions of (20.2) through (20.5). Adewale and Xu (2010) take an approach similar to that of Woods et al. (2006), but with the log-determinant replaced by *imse* and g replaced by a in the departures against which robustness is sought. In all of these articles, there is a mixture of analytic and numerical methods, with the latter being dominated by simulated annealing.

20.3 Robustness against a Misspecified Error Structure

20.3.1 Robustness against Heteroscedasticity

Wiens (1998) extends the methods of Section 20.2 to allow for heterogeneous error variances and allows the use of weights in the *lse*'s. In the minimax procedures, a maximum is taken




FIGURE 20.7

Minimax design densities (a) and weights (b) for heteroscedastic straight line regression for two values of ν as defined at (20.11): $\nu = 20/21$ (—) and $\nu = 1/51$ (· · ·). (Reprinted from Wiens, D.P., *J. Am. Stat. Assoc.*, 93, 1440, 1998. Copyright 1998 by the American Statistical Association. With permission.)

over the variance function $\sigma^2(x)$, subject to an L^2 -bound, as well as the neighbourhood of regression responses, and the minimization of the resulting generalization of (20.11) yields optimal weights as well as designs. See Figure 20.7. Montepiedra and Wong (2001) suppose that the experimenter will compute the unweighted *lse*, and they investigate whether or not the D -optimal design continues to be optimal in the presence of a particular nonconstant variance function. They give conditions under which the two-point design for straight line regression continues to be D -optimal and present examples for which the D -optimal design that addresses the heteroscedasticity has three support points.

20.3.2 Robustness against Dependence

Experimental observations may be autocorrelated – for instance, because they are gathered serially, or because spatial or other correlations enter in through the regressors. Sacks and Ylvisaker (1966, 1968) study design problems for regression in which the error process is a time series with a known covariance function that is used in the estimation process. Bickel and Herzberg (1979) and later Bickel et al. (1981) entertain the following form of robustness. It is supposed that the data will be analyzed by ordinary least squares (*ols*), tailored to a straight line fit (with or without an intercept) in the covariate t (time). The straight line model is not in doubt, but the true autocorrelation function is a mixture of a particular autocorrelation function $\rho(\cdot)$ and that under independence, namely, $(1 - \gamma) \delta_{\{t_i=t_j\}} + \gamma\rho(t_i - t_j)$ for a mixing parameter γ and observations made at times t_i, t_j . Their asymptotics lead to the conclusion that at least when $\rho(\cdot)$ is convex in $|t|$, the uniform design is very nearly optimal for regression with an intercept. For regression through the origin, an example of an exactly optimal design density is in Figure 20.8. From these cases, we are led to the conjecture that designs robust against autocorrelations of this form are qualitatively similar to designs constructed to be robust against misspecifications in the regression response but assuming that the observations will be independent.

Wiens and Zhou (1996) consider regression models in which the error process $\{\varepsilon(t)\}$ has a mixed autocorrelation function as the preceding, with $\rho(t)$ absolutely summable (leading to the existence of a spectral density) but otherwise arbitrary. Misspecified response functions (in covariates x) following (20.2) through (20.5) are allowed and the parameters are




FIGURE 20.8
 Optimal design density of Bickel and Herzberg (1979) for regression through the origin; $\gamma = .5$, $\rho(t) = \exp(-.483|t|)$. (From Bickel, P.J. and Herzberg, A.M., *Ann. Stat.*, 7, 77, 1979.)

estimated by *ols*. They show that a design ξ_* that is asymptotically minimax optimal for uncorrelated errors retains its optimality if the design points come from the measure ξ_* , or from a design tending weakly to ξ_* , and are then implemented in random order. This complements results of Wu (1981) who established the model robustness of various randomized designs for comparative experiments such as randomized complete block designs and randomized Latin square designs.

The omnibus correlation structure of the previous paragraph calls for the design points to be implemented in random order. If the class of autocorrelation models against which one seeks protection is less broad, then more precise guidance can be given. Wiens and Zhou (1997) work with models as at (20.2) through (20.5) and correlation structures as the aforementioned, expressed through autocorrelation matrices $P_\gamma = (1 - \gamma)I + \gamma P$ for $P \in \mathcal{P}$. Particular attention is paid to the classes $\mathcal{P}_1, \mathcal{P}_2$ of MA(1) structures with positive or negative lag-1 correlations, respectively. Under P_γ , the determinant of the *mse* matrix of the *ols* estimate $\hat{\theta}$ is

$$\mathcal{D}(\psi, \xi, P_\gamma) = \frac{\sigma^{2p}}{n} |M_\xi|^{-2} \left| \frac{X'P_\gamma X}{n} \right| \cdot \left\{ 1 + \frac{n}{\sigma^2} b'_{\psi, \xi} \left(\frac{X'P_\gamma X}{n} \right)^{-1} b_{\psi, \xi} \right\}.$$

With $\psi_0 = 0$ and $P_0 = I$, define the *change-of-variance function* in the direction of $P \in \mathcal{P}$ by

$$CVF(\xi, P) = \frac{\frac{d}{d\gamma} \mathcal{D}(\psi_0, \xi, P_\gamma)|_{\gamma=0}}{\mathcal{D}(\psi_0, \xi, P_0)},$$

and the *change-of-bias function* in the direction of $\psi \in \Psi$ by

$$CBF(\xi, \psi) = \frac{\frac{1}{2} \frac{d^2}{d\gamma^2} \mathcal{D}((1 - \gamma)\psi_0 + \gamma\psi, \xi, P_0)|_{\gamma=0}}{\sigma^{-2} \mathcal{D}(\psi_0, \xi, P_0)}.$$

The supremum of $CVF(\xi, P)$ over $P \in \mathcal{P}$ is termed the *change-of-variance sensitivity*, denoted $CVS(\xi, \mathcal{P})$ and that of $CBF(\xi, \psi)$ over Ψ the *change-of-bias sensitivity*, denoted $CBS(\xi, \Psi)$. A design ξ is *V-robust* if it minimizes $\mathcal{D}(\psi_0, \xi, P_0)$, that is, maximizes $\det(M_\xi)$, subject to

a bound $CVS(\xi, \mathcal{P}) \leq \alpha$, and *most V-robust* in a class of designs if α is the infimum of the CVS over this class. Similarly, a design is *B-robust* if it maximizes $\det(\mathbf{M}_\xi)$, subject to a bound $CBS(\xi, \psi) \leq \beta$, and *most B-robust* if β is the infimum of the CBS over a given class of designs. These extend notions of robustness, introduced in other contexts by Hampel et al. (1986), under which one seeks efficiency at the assumed model together with a bound on the rate at which the performance deteriorates as one moves away from this model.

It is evident that B-robust designs coincide with the BB designs of Section 20.2.2. A design is *M-robust* if it is both V-robust and B-robust and *most M-robust* if it is both most V-robust and most B-robust. The question of the existence of most M-robust designs seems to have remained open. Examples of most V-robust, most B-robust and M-robust designs are given in Wiens and Zhou (1997) for q -dimensional multiple linear regression with an intercept: $f(x) = (1, x')'$. The restriction is made to design spaces and designs for which \mathbf{M}_ξ is diagonal. For $q = 1$ and $q = 2$, see Figure 20.9. We see from this that a general guiding principle is that the design points should be robust against response function misspecifications and then implemented in an order such that there are as many sign changes as possible under \mathcal{P}_1 and as few as possible under \mathcal{P}_2 . Other examples of M-robust designs are given by Tsai and Zhou (2005).

Wiens and Zhou (1999) make similar findings, when robustness is sought against possible autocorrelations with an AR(1) structure. A method established there, to asymptotically minimize the maximum *imse*, over Ψ and over AR(1) models with lag-1 correlation $\rho < 0$, is to first generate optimal design points $\{t_j\}_{j=1}^n$ assuming independence. Then, having implemented x_1, \dots, x_{m-1} (with $x_0 = \mathbf{0}$), x_m is to be the nearest neighbour, among those t_j not yet chosen, of x_{m-1} . If instead $\rho > 0$, then the induced design $\{(-1)^m x_m\}_{m=1}^n$ is asymptotically minimax. Zhou (2001a) replaces the *imse* criterion by the trace of the *mse* matrix and obtains designs numerically for MA(q) errors ($q = 1, 2$), by a two-stage algorithm in which the design points are found by simulated annealing and an optimal order of implementation is obtained by a nested exchange algorithm. Shi et al. (2007) extend these results to $q = 3$.

Zhou (2001b) considers a modified form of the CVF. Motivated by a desire for accurate confidence intervals on linear functions $\mathbf{a}'\theta$, she defines

$$CVF_a(\xi, \mathbf{P}) = \frac{\frac{d}{d\gamma} \mathbf{a}'V(\mathbf{P}_\gamma) \mathbf{a}|_{\gamma=0}}{\mathbf{a}'V(\mathbf{P}_0) \mathbf{a}},$$

where $V(\mathbf{P}_\gamma) = \mathbf{M}_\xi^{-1} \left(\frac{\mathbf{X}'\mathbf{P}_\gamma\mathbf{X}}{n} \right) \mathbf{M}_\xi^{-1}$ is the covariance matrix of the *ols* estimate $\sqrt{n}\hat{\theta}$. It turns out, after a calculation, that

$$CVF_a(\xi, \mathbf{P}) = \frac{\mathbf{a}'(V(\mathbf{P}) - V(\mathbf{I})) \mathbf{a}}{\mathbf{a}'V(\mathbf{I}) \mathbf{a}},$$

the numerator of which is proportional to the difference in the squared widths of confidence intervals computed using a knowledge of \mathbf{P} and those based on *ols*. Zhou (2001b) goes on to evaluate $CVF_a(\xi, \mathbf{P})$ for fixed \mathbf{P} and at a given design ξ , chosen, for instance, for its efficiency or robustness under independence. Then the order of implementation of the design points is optimized, so as to minimize $|CVF_a(\xi, \mathbf{P})|$. An extended criterion studied is to minimize a sum $\sum_i |CVF_{a_i}(\xi, \mathbf{P})|$ (One might instead take a maximum over \mathbf{a} , leading




FIGURE 20.9 Most V-, most B, and M-robust designs, $n = 25$. (a,b): Most V-robust designs in $[-1, 1]$ for \mathcal{P}_1 and \mathcal{P}_2 . (c,d): Most V-robust designs in $[-1, 1] \times [-1, 1]$, with the indices of the design points plotted for \mathcal{P}_1 and \mathcal{P}_2 . (e,f): Most B-robust designs for Ψ , ordered for \mathcal{P}_1 and \mathcal{P}_2 . (g,h): M-robust designs for Ψ , ordered for \mathcal{P}_1 and \mathcal{P}_2 . In (a,e,g) the design points are at the vertices; in (b,f,g) they are at the plotted squares; in all six cases the order of implementation is as on the vertical axes. (Reprinted from Wiens, D.P. and Zhou, J., *J. Am. Stat. Assoc.*, 92, 1503, 1997. Copyright 1997 by the American Statistical Association. With permission.)

to apparently open problems.) The examples given indicate that using this optimal ordering, the coverage probabilities of confidence intervals, computed incorrectly assuming independence, deteriorate only very slightly from the nominal values.

20.4 Special Applications

In design, as in all applied research, general theory such as those presented in the previous sections must often be modified or adapted to suit specific applications. In the next few

sections, we illustrate this in the context of some particular examples in which notions of robustness of design have been found to be fruitful.

20.4.1 Extrapolation

Extrapolation poses special problems in statistics in general and design in particular, since a model that is in doubt in the design space may be even more so in the extrapolation space. Reiterating sentiments expressed by us elsewhere in this chapter, Lawless (1984) remarks that “in extrapolation problems, a slight degree of model inadequacy quickly wipes out advantages that minimum variance designs possess when the model is exactly correct”.

Huber (1975) supposes that a function f is to be observed with additive random error, at n points chosen by the experimenter from $(0, \infty)$, with the intention of predicting $f(-1)$. Assuming the use of a linear predictor, he finds designs that minimize the maximum *mse* $E \left[\left\{ \hat{f} - f(-1) \right\}^2 \right]$, with the maximum taken over all functions f that are $h + 1$ times differentiable, with a uniformly bounded $(h + 1)$ th derivative. These turn out to have $h + 1$ points of support constituting a certain set of Chebyshev points. Huang and Studden (1988) find and correct an error in Huber’s (1975) proof, while agreeing with the final result, and derive similar designs for extrapolation from $[-1, 1]$ to a single point outside of this interval. In this latter context, Hoel and Levine (1964) derive designs for extrapolating a polynomial of known degree; Spruill (1985) and Dette and Wong (1996) discuss the use of these designs, and suggest improvements, when the form or degree of the function being extrapolated is in doubt.

Fang and Wiens (2000) consider the extrapolation of approximate regression responses as at (20.2) through (20.5), but with x chosen from a finite design space \mathcal{X} to a space \mathcal{T} on which $E[y(x)] = f'(x)\theta + \psi_T(x)$ with $\int_{\mathcal{T}} \psi_T^2(x) \mu(dx) \leq \tau_T^2$ for some measure μ (typically either Lebesgue measure or counting measure). Designs minimizing the maximum, over both ψ at (20.4) and over ψ_T , of the integrated mean squared prediction error $\int_{\mathcal{T}} E \left[\left\{ \hat{y}(x) - E[y(x)] \right\}^2 \right] \mu(dx)$ are derived. The development largely parallels the derivation and minimization of (20.13); in particular the maximum is derived analytically and minimized by simulated annealing. The results complement Fang and Wiens (1999) and Wiens and Xu (2008a,b), where absolutely continuous designs are sought.

To illustrate the use of these designs, consider the following experiment of Guess et al. (1977), discussed by Hoel and Jennrich (1979). A sample of $n = 235$ animals is to receive doses, at levels x , of a carcinogen. The mean response is approximated by a cubic polynomial in $x \in \mathcal{X} = [1, 500]$; the goal is extrapolation to the region $\mathcal{T} = \{0.5\}$ where the *background effect* is to be estimated. The minimum variance design proposed by Hoel and Jennrich places 63, 125, 35 and 12 observations at $x = 1, 82.6, 342$ and 500. The Huber/Huang/Studden design discussed earlier returns design weights that must be rounded to yield implementable designs; using the rounding mechanism of Pukelsheim and Rieder (1992) results in frequencies of 232, 2, 1 and 0 at $x = 1, 125.75, 373.3$ and 500 (so that the rounded design becomes supersaturated for the cubic model). With \mathcal{X} discretized to 47 equally spaced points, the designs of Fang and Wiens (2000) are shown in Figure 20.10. Both homogeneous and heterogeneous error variances (ranging over a neighbourhood as described in Section 20.3.1 in the latter case), and *ols* and *wls* estimates, are considered. To some extent these designs again illustrate the paradigm of Section 20.2, with the replicates of the Hoel and Jennrich design spread out into clusters.




FIGURE 20.10

Minimax robust designs ($n = 235$) for extrapolation from $[1, 500]$ to $\{.5\}$. (a) Design frequencies for *ols*, homoscedasticity. (b) design for *ols*, heteroscedasticity. (c) design for *wls*, heteroscedasticity. (d) Weights for *wls*, heteroscedasticity. (Reprinted from Fang, Z. and Wiens, D.P., *J. Am. Stat. Assoc.*, 95, 807, 2000. Copyright 2000 by the American Statistical Association. With permission.)

A class of applications in which extrapolation designs are required is accelerated life testing, in which items are tested at unusually high stress levels – at which they break down quickly, if at all – in order to predict their reliability under more usual levels. This is often combined with censoring. Chernoff (1962) derives accelerated life designs assuming that the models and parameters are known exactly, while remarking that “It is not possible to overemphasize the importance of the underlying assumptions. Doubts about them would cast serious doubts on the results to be obtained from accelerated designs”. Ginebra and Sen (1998) address these reservations by constructing minimax designs, with the maximum taken over plausible ranges of the unknown model parameters. Xu (2009b) obtains designs for situations in which both acceleration and censoring are present, under departures similar to those in Fang and Wiens (2000) discussed earlier, but with a continuous design space.

20.4.2 Model Selection: Discrimination and Goodness of Fit

In the face of model uncertainty, it is natural to seek design strategies to clarify the nature of the response function. When the experimenter is choosing among a set of models whose functional forms are specified and that are perhaps even nested, then ideas similar to those in Section 20.2.2 might be applied – see, for example, Biedermann et al. (2007), where the goal is to select the correct response function from a set $\left\{ \sum_{i=1}^k a_i \exp(-\lambda_i x) ; k \geq 1 \right\}$ and to estimate the parameters $\{a_i, \lambda_i ; i = 1, \dots, k\}$ with reasonable efficiency. This is continued in Chapter 14.

Suppose, however, as is perhaps more common, that the correct model is a member of one of two only approximately known classes $\mathcal{C}_0, \mathcal{C}_1$. In \mathcal{C}_j , the response variable has density $p(y|x, \mu_j)$, parameterized by the mean response $\mu_j(x) = \eta_j(x|\theta_j) + \psi(x)$ with ψ ranging over a class Ψ_j similar to those in Section 20.2. If $\mathcal{C}_0, \mathcal{C}_1$ each contains only a single member,

that is, if each of Ψ_0 and Ψ_1 contains only the zero function, then one can discriminate between the classes by means of the Neyman–Pearson test of $H_0: p(y|x, \mu_0)$ versus $H_1: p(y|x, \mu_1)$. One might then seek a design maximizing the power of this test; this leads to maximizing $E_\xi [\mathcal{I}\{\mu_0(x), \mu_1(x)\}]$, where

$$\mathcal{I}\{\mu_0(x), \mu_1(x)\} = \int_{-\infty}^{\infty} p(y|x, \mu_1) \log \left\{ \frac{p(y|x, \mu_1)}{p(y|x, \mu_0)} \right\} dy$$

is the Kullback–Leibler (KL) divergence, measuring the information that is lost when $p(\cdot|x, \mu_0)$ is used to approximate $p(\cdot|x, \mu_1)$ – see López-Fidalgo et al. (2007). Wiens (2009) continues this approach and constructs *robust KL-optimal* designs that maximize the minimum value of $E_\xi [\mathcal{I}\{\mu_0(x), \mu_1(x)\}]$, over full classes \mathcal{C}_0 and \mathcal{C}_1 . This corresponds to discriminating between the members of a least favourable pair

$$(\mu_0^* = \eta_0(\cdot|\theta_0) + \psi_0^*(\cdot), \mu_1^* = \eta_1(\cdot|\theta_1) + \psi_1^*(\cdot)),$$

at which \mathcal{C}_0 and \mathcal{C}_1 are closest. Both static and sequential designs are studied. In the former case, the designs can roughly be described as placing clusters of points in regions where the two competing mean responses are farthest apart – as one might expect. See also O’Brien (1995) who gives a method of adding locations to a D -optimal design – typically with no more design points than parameters – in such a way as to allow testing for *lof*.

In the sequential case, this KL criterion of Wiens is modified so as to allow a transition, over time, from choosing design points solely to maximize the discriminatory power to choosing them for increased efficiency of parameter estimation in the more plausible model. This criterion of optimizing a mixture of goals is used as well in Wiens (2010), where the KL criterion is applied to testing an approximately specified model for *lof*, while simultaneously aiming for some efficiency in estimating the parameters. This is also the goal of Dette et al. (2005), who consider the EMAX class of response functions $E[y|x] = \theta_0 x^h / (\theta_1 + x^h)$ and design to test the fit of a Michaelis–Menten response ($h = 1$).

20.4.3 Dose–Response Designs

Consider an experiment in which independent binary responses are to be observed, each representing the presence or absence of a particular response to an application of a *dose* at level x . The analyst fits a model $\Pr(y = 1|x) = F_0(\alpha + \beta x)$ for some distribution function (*link*) $F_0(\cdot)$, typically the logistic or the normal. Often the goal is the estimation of ED_{100p} – the level $x_p = (F_0^{-1}(p) - \alpha)/\beta$ at which 100 p % of the population will respond. The associated design problem is to determine the optimal doses and the number of observations to be made at each. This is then an instance of the class of problems discussed in Section 20.2.5, with issues related to poor initial parameter estimates and inappropriately chosen links.

Sitter (1992) notes the lack of robustness to poor initial parameter estimates and proposes minimax procedures, with the maximum taken over a given parameter region.

Only symmetric designs are considered; this restriction is dropped by Biedermann and Dette (2005). Hedayat et al. (1997, 2002) instead derive designs that are D -optimal among those with k ($3 \leq k \leq 7$) levels and note that as k increases, the effect of poor initial guesses for α and β decreases.

Li and Wiens (2011) design to minimize the asymptotic mse of the estimate \hat{x}_p , possibly integrated over an interval of values of p , in the face of uncertainty about the link. The estimate \hat{x}_p is computed under F_0 , and the mse is then computed under an alternate link F_n and maximized over a Kolmogorov neighbourhood $\{F_n | \sup |F_n(t) - F_0(t)| \leq \tau/\sqrt{n}\}$. Designs are computed that minimize the maximized mse .

To compare some of these approaches, consider an experiment of Rosenberger and Grill (1997), the aim of which was to elicit information about the relationship between stimulus level (x) and response, by estimating quantiles of the stimulus–response curve. Subjects sequentially received *marking stimuli* (auditory marking clicks) at various levels and at random times near to those of a certain event and were then asked whether the event occurred before or after the stimulus. The response y was binary, with $y = 1$ being recorded if the subject reported that the event occurred before the stimulus. The principal goal was to estimate the median (ED_{50}) of the stimulus–response curve. A secondary goal was to design so as to allow for the estimation of other quantiles such as the lower and upper quartiles. The investigators assumed a logistic link relating $\Pr(y = 1)$ to a linear function of the stimulus level. This experiment was subsequently redesigned by Biedermann, Dette and Pepelyshev (2006, henceforth BDP), by Zhu and Wong (2000, henceforth ZW) and by Li and Wiens (2011, henceforth LW). The approaches were, respectively, maximin optimality (with the minimum over a range of parameters), Bayesian optimality and minimax robustness as described in the previous paragraph. Each group constructed one design for estimation of the median and another for estimation of the quantiles.

See Figure 20.11, where the four designs of BDP and ZW are illustrated, together with four LW designs – two assuming that the parameter τ of the Kolmogorov neighbourhood defined earlier is zero, that is, that the fitted link is exactly correct, and two using $\tau = 1$. Comparative mse (maximized, if $\tau > 0$) in these neighbourhoods are given in Table 20.1. The designs LW2, LW4 do very well in the case of quartile estimation for which they are intended; the minimax design LW4 performs well overall. Design LW1 = LW3 also does well for ED_{50} estimation, but has a rather counterintuitive concentration of mass near the centre of the design space, illustrating that efficiency or robustness for one particular feature might lead to poor performance in other instances. An indicator of the robustness of the LW designs is that in three of the four cases in which they are intended to be optimal, there is only a small increase in mse when τ changes from 0 to 1.

To do only these comparisons is unfair, since we are looking at criteria with respect to which the designs LW have been optimally tuned. Those of ZW and BDP were tuned for optimal efficiency at one or several fitted models, a natural measure of which is the root of the total variance, that is, $[\det(\mathbf{M}_\xi)]^{-1/2}$ where \mathbf{M}_ξ is the information matrix, given in the final column of Table 20.1. With respect to this measure, the latter designs fare well. The design LW1 = LW3 does poorly, while the others are quite comparable.

An intriguing extension of these ideas, which to our knowledge has not been pursued, is to allow $\Pr(y = 1)$ to depend not only on the level x but as well on covariates – perhaps in an only approximately specified manner. For instance, aircraft manufacturers attempt to estimate and then reduce the probability of an *ignition event* ($y = 1$), in response to a lightning strike at level x . The reduction in this probability is brought about through various surface coatings of the aircraft; these form the covariates.




FIGURE 20.11

Dose–response designs for estimating the median (a, c, e, and g) or the quantiles (b, d, f, and h) of the stimulus–response curve. (From Li, P. and Wiens, D.P., *J. R. Stat. Soc. (Ser. B)*, 17, 215, 2011.)

TABLE 20.1

Comparative Losses of Dose–Response Designs

Design	Max mse; $\tau = 0^a$	ED ₅₀ Estimation $\tau = 1$	Max mse; $\tau = 0^1$	ED _{25,75} Estimation $\tau = 1$	$[\det(M_\xi)]^{-1/2}$
ZW1	6.51	46.32	[3.86]	[28.33]	1.51
ZW2	[9.01]	[57.32]	5.08	49.91	1.75
BDP1	6.15	41.81	[3.74]	[26.04]	1.55
BDP2	[8.83]	[18.99]	5.12	27.85	1.91
LW1	4.01	4.03	[117.87]	[129.06]	16.48
LW2	[5.18]	[92.74]	3.49	25.06	1.65
LW3	4.01	4.03	[117.87]	[129.06]	16.48
LW4	[4.89]	[5.75]	3.75	4.37	1.93

^a When $\tau = 0$, the loss is due solely to variation at the fitted model.

Notes: 1. Square brackets denote estimation situations in which the indicated design is not intended to be appropriate.

2. Figures in bold correspond to situations in which the LW designs are intended to be optimal.

20.4.4 Clinical Trials

In their simplest form, clinical trials involve the random allocation of patients into groups, after which each receives a treatment (perhaps a control) specific to the group. An element of randomness is required, in order that the investigator remain *blinded* to the assignments. Responses are assumed to depend on the treatment and perhaps on various covariates (age, weight, etc.). Heckman (1987) assumes that for each of two treatments ($t = 0, 1$), with a single covariate, the mean response is approximately linear:

$$E[y|x, x_0] = \beta_0(x_0) + \beta_1(x_0)t + \beta_2(x_0)(x - x_0) + r(t, x_0, x),$$

for $|r(t, x_0, x)|$ bounded by a known multiple, depending on t and x_0 , of $(x - x_0)^2$. (This is her *model 2*; *model 1* has $\beta_2(x_0) \equiv 0$ and $|r(t, x_0, x)|$ is bounded by a multiple of $|x - x_0|$.) The difference in the mean effects of the treatments at $x = x_0$ is $\beta_1(x_0)$, so that one is interested in estimating this quantity. Sacks and Ylvisaker (1978) had obtained linear, minimum *mse* estimators for such models, and Heckman obtains the correspondingly optimal sequences of allocations to the two groups. Wiens (2005a) instead takes

$$E[y|x] = \beta_0 + \beta_1 t + f'(x)\theta + \psi_t(x)$$

(his *model 2*; *model 1* has $\theta \equiv \mathbf{0}$). Static and sequential allocation schemes are considered; in the former case a function $\rho(\cdot)$ is derived, and a subject who arrives with covariates x is assigned to the treatment group with probability $\rho(x)$. For this scheme, the results are qualitatively similar to those of Heckman. In the sequential case, they compare favourably with those of Atkinson (1982) for exactly linear responses when such do indeed hold, and have significantly reduced *mse* otherwise. There is however a significant cost to be paid for this robustness in terms of *balance* (see the discussions in Atkinson 1996, 2002) across the covariates.

20.4.5 Spatial Designs: Field and Computer Experiments

In spatial experiments – covered more broadly in Chapter 15 – one makes observations on units that are spatially arranged, and the spatial layout becomes important. A fairly general formulation is that the experimenter samples locations chosen from a set $\mathcal{X} = \{t_1, \dots, t_N\} \subset \mathbb{R}^d$. At these locations, he observes

$$y(t) = \mu(t; \theta) + \delta(t) + \varepsilon(t), \tag{20.20}$$

where $\mu(t; \theta)$ is a deterministic mean function possibly depending on unknown parameters θ , $\delta(t)$ represents stochastic, spatially correlated departures, and $\varepsilon(t)$ represents measurement error. The goals often focus on the prediction of functions of $X(t) = \mu(t; \theta) + \delta(t)$. Such models and prediction techniques are commonly also used in computer experiments, described in Section V; designs constructed in this context often (see in particular Chapter 17) aim for a fairly uniform exploration of the design space – itself a quite robust approach, as we have seen.

Misspecification in (20.20) can occur in a number of ways – (1) misspecifying $\mu(t; \theta)$, perhaps in ways as outlined in Section 20.2; (2) misspecifying the covariance structure of $\delta(t)$,




FIGURE 20.12

Minimax spatial designs for prediction of X_{Total} as in Section 20.4.5. (a) Isotropic correlations; a maximum is taken over a neighbourhood of $\mu(t; \theta)$ but the nominal correlation function for δ and homogeneous variance function for ε are assumed to be correct. (b) As in (a) but with anisotropic correlations. (c,d): As in (a,b), respectively, but maxima are taken as well over neighbourhoods of the correlation function for δ and over the nominal variance function of ε . (From Wiens, D. P., *Environmetrics*, 16, 205, 2005b.)

for instance, by assuming an isotropic structure when such is not the case; and (3) misspecifying the variance/covariance structure of $\varepsilon(t)$, as in Section 20.3. Wiens (2005b) constructs designs that are minimax robust against any of these three types of departures, using neighbourhoods of the nominal model to describe each, and simulated annealing to carry out the minimization. See Figure 20.12 for examples – in each case, 7 points are sampled from a 25-point grid (t_1, t_2) . The nominal response is $\mu(t; \theta) = \theta_0 + \theta_1 t_1 + \theta_2 t_2$, the nominal covariance function Gaussian: $\text{CORR}[\delta(t), \delta(t')] = \exp\{-\lambda_1 (t_1 - t'_1)^2 - \lambda_2 (t_2 - t'_2)^2\}$, isotropic if $\lambda_1 = \lambda_2$ and anisotropic otherwise, and the nominal error variances homogeneous. The aim is to predict $X_{Total} = \sum_{t \in \mathcal{X}} X(t)$ with minimum *mse*.

Spatial designs are commonly employed in field experiments, where plots, perhaps in a rectangular grid, are to be planted with various crops. The goal is the optimal – with respect to some criterion – assignment of crops to plots. Minimax designs in such situations are studied by Wiens and Zhou (2008) and Ou and Zhou (2009). An example given by Ou and Zhou has eight crops, with crop one being a *control*, arranged in a 3×8 field of plots. Possible correlation structures are $NN(\rho)$, in which each plot has nonzero correlation ρ only with its immediate neighbours, and $DG(\rho)$, in which the correlation between plots is ρ^d when d is the sum of the east-west and north-south distances between them. These are the nominal

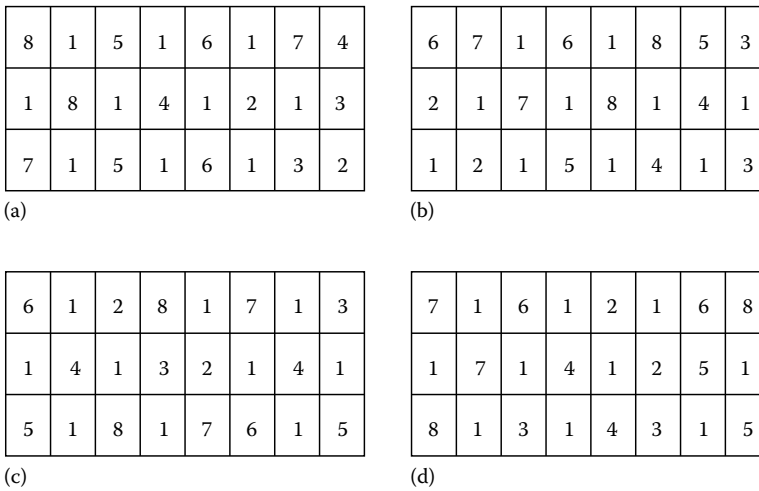


FIGURE 20.13 Minimax assignment of 8 crops to 24 fields. (a) $NN(\rho)$ correlations with *ols*, (b) $NN(\rho)$ correlations with *gls*, (c) $DG(\rho)$ correlations with *ols*, (d) $DG(\rho)$ correlations with *gls*. In all cases $\rho = .15$. (From Ou, B. and Zhou, J., *Metrika*, 69, 45, 2009.)

correlations; neighbourhood structures of these are introduced, and minimax designs are obtained assuming the use of either *ols* or generalized least squares (*gls*) using the possibly incorrect nominal correlation structure. In each case, various arrangements of replicates are compared, and it is found that, quite generally, 10 replicates of the control and 2 of all other crops is optimal for the goal of efficient estimation of the crop effects (row / column effects were not modelled). See Figure 20.13.

This brief description barely scratches the surface of applications in which spatial designers might benefit from principles of robustness. As examples, two that appear to be as yet unexplored are spatial models with discrete – perhaps binary – data, and threshold models, describing, for instance, the possible exceedance of a regulatory threshold.

20.4.6 Mixture Experiments

In mixture experiments, the covariates x_1, \dots, x_q are nonnegative and sum to one, representing the relative proportions of ingredients blended in a mixture. This is another area that is relatively unexplored by robust designers. Huang et al. (2009) find designs in the case that the experimenter is uncertain whether a first- or second-order model is appropriate. Smucker et al. (2011) extend this idea to designing so as to maximize a geometrically weighted product $\phi(\xi) = \prod_j [\det(M_\xi(f_j))]^{\omega_j}$ of determinants of information matrices arising from candidate models with regressors f_j . Exchange algorithms are derived for the maximization.

20.4.7 Designs for Robust Inference

Even the most carefully designed experiment can lose effectiveness if the responses themselves are erroneous, or outlying, or missing altogether. While these eventualities are not directly under the control of the designer, their effects can be tempered by an appropriate design.

The influence of an observation, at x_i , on the *lse* in a linear model is measured by its *leverage* – the diagonal element h_{ii} of the *hat* matrix $X(X'X)^{-1}X'$. Recall that the h_{ii} lie in $[0, 1]$, and that their average is $\bar{h} = p/n$. Huber (1975) recommends reducing, by (approximate) replication, the values of any h_{ii} that are close to 1. Box and Draper (1975) argue that one should aim for a small value of $\sum_{i=1}^n h_{ii}^2$, or equivalently of the variance of the h_{ii} . They point out that this is also related to robustness against nonnormality, an observation bolstered by Huber's (1973, p. 804) remark that the condition that $\max h_{ii} \rightarrow 0$ as $n \rightarrow \infty$ "appears to be indispensable for any reasonably simple general asymptotic theory of robust estimation".

Herzberg and Andrews (1976) and Andrews and Herzberg (1979) discuss the possibility of design *breakdown*, by which they mean that the model matrix might become rank deficient after perceived outliers have been discarded, or as a result of cases that are missing entirely. This line of investigation is continued by Akhtar and Prescott (1986); Herzberg et al. (1987); and Ahmad and Gilmour (2010). They find that the loss in information (relative to the *D*-criterion) due to a missing observation at x_i is also measured by h_{ii} , and so a design with moderate values of $\{h_{ii}\}$ is robust to missing observations.

In the literature of robust estimation, one is commonly cautioned against removing cases thought to be outliers; instead, one is encouraged to downweight them in the estimation process. This is the effect of ordinary and generalized M-estimation, each of which includes least squares as a special case. Wiens and Wu (2010) construct designs that minimize the maximum asymptotic *mse* of M-estimates of regression, with the maximum taken over the discrete versions of neighbourhoods as at (20.2) through (20.5). This *mse* depends on the error distribution and on the score functions used in the estimation, but it is found that the designs change only slowly with variations in these quantities, so that a design which is robust for *lse*, against the types of departures of concern in this chapter, will remain so, if perhaps slightly suboptimal, for M-estimation.

References

- Abdelbasit, K. M. and Plackett, R. L. (1983), Experimental design for binary data, *Journal of the American Statistical Association*, 78, 90–98.
- Adewale, A. and Wiens, D. P. (2006), New criteria for robust integer-valued designs in linear models, *Computational Statistics and Data Analysis*, 51, 723–736.
- Adewale, A. and Wiens, D. P. (2009), Robust designs for misspecified logistic models, *Journal of Statistical Planning and Inference*, 139, 3–15.
- Adewale, A. J. and Xu, X. (2010), Robust designs for generalized linear models with possible overdispersion and misspecified link functions, *Computational Statistics and Data Analysis*, 54, 875–890.
- Andrews, D. and Herzberg, A. (1979), The robustness and optimality of response surface designs, *Journal of Statistical Planning and Inference*, 3, 249–257.
- Ahmad, T. and Gilmour, S. G. (2010), Robustness of subset response surface designs to missing observations, *Journal of Statistical Planning and Inference*, 140, 92–103.
- Akhtar, M. and Prescott, P. (1986), Response surface designs robust to missing observations, *Communications in Statistics: Simulation and Computation*, 15, 345–363.
- Atkinson, A. C. (1982), Optimum biased coin designs for sequential clinical trials with prognostic factors, *Biometrika*, 69, 61–67.
- Atkinson, A. C. (1996), The usefulness of optimum experimental designs, *Journal of the Royal Statistical Society (Series B)*, 58, 59–76.

- Atkinson, A. C. (2002), The comparison of designs for sequential clinical trials with covariate information, *Journal of the Royal Statistical Society (Series A)*, 165, 349–373.
- Bickel, P. J. and Herzberg, A. M. (1979), Robustness of design against autocorrelation in time I: Asymptotic theory, optimality for location and linear regression, *The Annals of Statistics*, 7, 77–95.
- Bickel, P. J., Herzberg, A. M., and Schilling, M. F. (1981), Robustness of design against autocorrelation in time II: Optimality, theoretical and numerical results for the first-order autoregressive process, *Journal of the American Statistical Association*, 76, 870–877.
- Biedermann, S. and Dette, H. (2001), Optimal designs for testing the functional form of a regression via nonparametric estimation techniques, *Statistics and Probability Letters*, 52, 215–224.
- Biedermann, S. and Dette, H. (2005), Numerical construction of parameter maximin D -optimal designs for binary response models, *South African Statistical Journal*, 39, 127–161.
- Biedermann, S., Dette, H., and Pepelyshev, A. (2006), Some robust design strategies for percentile estimation in binary response models, *The Canadian Journal of Statistics*, 4, 603–622.
- Biedermann, S., Dette, H., and Pepelyshev, A. (2007), Optimal discrimination designs for exponential regression models, *Journal of Statistical Planning and Inference*, 137, 2579–2592.
- Bischoff, W. and Miller, F. (2006), Optimal designs which are efficient for lack of fit tests, *Annals of Statistics*, 34, 2015–2025.
- Box, G. E. P. (1953), Non-normality and tests on variances, *Biometrika*, 40, 318–335.
- Box, G. E. P. and Draper, N. R. (1959), A basis for the selection of a response surface design, *Journal of the American Statistical Association*, 54, 622–654.
- Box, G. E. P. and Draper, N. R. (1975), Robust designs, *Biometrika*, 62, 347–352.
- Chaloner, K. and Larntz, K. (1989), Optimal Bayesian experimental design applied to logistic regression experiments, *Journal of Statistical Planning and Inference*, 21, 191–208.
- Chang, Y.-J. and Notz, W. I. (1996), Model robust designs, in: *Design and Analysis of Experiments, Handbook of Statistics Vol. 13*, pp. 1055–1098; eds. Ghosh, S. and Rao, C. R., North-Holland, the Netherlands: Elsevier.
- Chernoff, H. (1962), Optimal accelerated life designs for estimation, *Technometrics*, 4, 381–408.
- Daemi, M. and Wiens, D. P. (2012), Techniques for the construction of robust regression designs, *The Canadian Journal of Statistics*, 41, 679–695.
- Dette, H. (1991), A note on robust designs for polynomial regression, *Journal of Statistical Planning and Inference*, 28, 223–232.
- Dette, H. and Biedermann, S. (2003), Robust and efficient designs for the Michaelis–Menten model, *Journal of the American Statistical Association*, 98, 679–686.
- Dette, H. and Franke, T. (2001), Robust designs for polynomial regression by maximizing a minimum of D - and D_1 -efficiencies, *The Annals of Statistics*, 29, 1024–1049.
- Dette, H., Lopez, I. M., Ortiz Rodriguez, I. M., and Pepelyshev, A. (2006), Maximin efficient design of experiment for exponential regression models, *Journal of Statistical Planning and Inference*, 136, 4397–4418.
- Dette, H. and Melas, V. B. (2010), A note on all-bias designs With applications in spline regression models, *Journal of Statistical Planning and Inference*, 140, 2037–2045.
- Dette, H., Melas, V. B., and Pepelyshev, A. (2003), Standardized maximin E -optimal designs for the Michaelis–Menten model, *Statistica Sinica*, 13, 1147–1163.
- Dette, H., Melas, V. B., and Wong, W.-K. (2005), Optimal design for goodness-of-fit of the Michaelis–Menten enzyme kinetic function, *Journal of the American Statistical Association*, 100, 1370–1381.
- Dette, H. and Neugebauer, H.-M. (1997), Bayesian D -optimal designs for exponential regression models, *Journal of Statistical Planning and Inference*, 60, 331–349.
- Dette, H. and Pepelyshev, A. (2008), Efficient experimental designs for sigmoidal growth models, *Journal of Statistical Planning and Inference*, 138, 2–17.
- Dette, H. and Studden, W. J. (1995), Optimal designs for polynomial regression when the degree is not known, *Statistica Sinica*, 5, 459–473.
- Dette, H. and Studden, W. J. (1997). *The Theory of Canonical Moments with Applications in Statistics, Probability, and Analysis*, New York: Wiley.

- Dette, H. and Wiens, D. P. (2008), Robust designs for series estimation, *Computational Statistics and Data Analysis*, 52, 4305–4324.
- Dette, H. and Wiens, D. P. (2009), Robust designs for 3D shape analysis with spherical harmonic descriptors, *Statistica Sinica*, 19, 83–102.
- Dette, H. and Wong, W.-K. (1996), Robust optimal extrapolation designs, *Biometrika*, 83, 667–680.
- Draper, N. R. and Sanders, E. R. (1988), Designs for minimum bias estimation, *Technometrics*, 30, 319–325.
- Dror, H. A. and Steinberg, D. M. (2006), Robust experimental design for multivariate generalized linear models, *Technometrics*, 48, 520–529.
- Fang, K. T. and Wang, Y. (1994), *Number-Theoretic Methods in Statistics*, New York: Chapman and Hall.
- Fang, Z. (2006), Some robust designs for polynomial regression models, *The Canadian Journal of Statistics*, 34, 623–638.
- Fang, Z. and Wiens, D. P. (1999), Robust extrapolation designs and weights for biased regression models with heteroscedastic errors, *The Canadian Journal of Statistics*, 27, 751–770.
- Fang, Z. and Wiens, D. P. (2000), Integer-valued, minimax robust designs for estimation and extrapolation in heteroscedastic, approximately linear models, *Journal of the American Statistical Association*, 95, 807–818.
- Fang, Z. and Wiens, D. P. (2003), Robust regression designs for approximate polynomial models, *Journal of Statistical Planning and Inference*, 117, 305–321.
- Fedorov, V. V., Montepiedra, G., and Nachtsheim, C. J. (1998), Optimal design and the model validity range, *Journal of Statistical Planning and Inference*, 72, 215–227.
- Ford, I., Titterton, D. M., and Kitsos, C. P. (1989), Recent advances in nonlinear experimental design, *Technometrics*, 31, 49–60.
- Ginebra, J. and Sen, A. (1998), Minimax approach to accelerated life tests, *IEEE Transactions on Reliability*, 47, 261–267.
- Guess, H., Crump, K., and Peto, R. (1977), Uncertainty estimates for low-dose extrapolation of animal carcinogenicity Data, *Cancer Research*, 37, 3475–3483.
- Hampel, F. R., Ronchetti, E., Rousseeuw, R. J., and Stahel, W. (1986), *Robust Statistics: The Approach Based on Influence Functions*, New York: Wiley.
- Heckman, N. E. (1987), Robust design in a two treatment comparison in the presence of a covariate, *Journal of Statistical Planning and Inference*, 16, 75–81.
- Hedayat, A. S., Yan, B., and Pezzuto, J. M. (1997), Modeling and identifying optimum designs for fitting dose-response curves based on raw optical density data, *Journal of the American Statistical Association*, 92, 1132–1140.
- Hedayat, A. S., Yan, B., and Pezzuto, J. M. (2002), Optimum designs for estimating ED_p based on raw optical density data, *Journal of Statistical Planning and Inference*, 104, 161–174.
- Heo, G., Schmuland, B., and Wiens, D. P. (2001), Restricted minimax robust designs for misspecified regression models, *The Canadian Journal of Statistics*, 29, 117–128.
- Herzberg, A. M. and Andrews, D. F. (1976), Some considerations in the optimal design of experiments in non-optimal situations, *Journal of the Royal Statistical Society, Series B*, 38, 284–289 (Corr: V39 p394).
- Herzberg, A. M., Prescott, P., and Akhtar, M. (1987), Equi-information robust designs: Which designs are possible? *The Canadian Journal of Statistics*, 15, 71–76.
- Hoel, P. G. and Jennrich, R. I. (1979), Optimal designs for dose response experiments in cancer research, *Biometrika*, 66, 307–316.
- Hoel, P. G. and Levine, A. (1964), Optimal spacing and weighting in polynomial prediction, *Annals of Mathematical Statistics*, 35, 1533–1560.
- Huang, M.-N. L., Hsu, H.-L., Chou, C.-J., and Klein, T. (2009), Model-Robust D - and A -optimal designs for mixture experiments, *Statistica Sinica*, 19, 1055–1075.
- Huang, M.-N. L. and Studden, W. J. (1988), Model robust extrapolation designs, *Journal of Statistical Planning and Inference*, 18, 1–24.
- Huber, P. J. (1964), Robust estimation of a location parameter, *The Annals of Mathematical Statistics*, 35, 73–101.

- Huber, P. J. (1973), Robust regression: Asymptotics, conjectures and Monte-Carlo, *The Annals of Statistics*, 1, 799–821.
- Huber, P. J. (1975), Robustness and designs, in: *A Survey of Statistical Design and Linear Models*, ed. Srivastava, J. N., North Holland, pp. 287–303.
- Huber, P. J. (1981), *Robust Statistics*, Wiley.
- Karami, H. K. and Wiens, D. P. (2014), Robust static designs for approximately specified nonlinear regression Models, *Journal of Statistical Planning and Inference*, 144, 55–62.
- Karson, M. J., Manson, A. R., and Hader, R. J. (1969), Minimum bias estimation and experimental design for response surfaces, *Technometrics*, 11, 461–475.
- Khuri, A. I., Mukherjee, B., Sinha, B. K., and Ghosh, M. (2006), Design issues for generalized linear models: A review, *Statistical Science*, 21, 376–399.
- Kiefer, J. (1975), Optimal design: Variation in structure and performance under change of criterion, *Biometrika*, 62, 277–288.
- King, J. and Wong, W.-K. (2000), Minimax D -optimal designs for the logistic model, *Biometrics*, 56, 1263–1267.
- Lawless, J. F. (1984), Some problems concerning experimental designs for extrapolation, in: *Topics in Applied Statistics*, ed. Dwivedi, T. D., Montreal, Quebec, Canada: Concordia University Press, pp. 357–366.
- Li, K. C. (1984), Robust regression designs when the design space consists of finitely many points, *The Annals of Statistics*, 12, 269–282.
- Li, K. C. and Notz, W. (1982), Robust designs for nearly linear regression, *Journal of Statistical Planning and Inference*, 6, 135–151.
- Li, P. and Wiens, D. P. (2011), Robustness of design for dose-response studies, *Journal of the Royal Statistical Society (Series B)*, 17, 215–238.
- Liu, S. X. and Wiens, D. P. (1997), Robust designs for approximately polynomial regression, *Journal of Statistical Planning and Inference*, 64, 369–381.
- López-Fidalgo, J., Tommasi, C., and Trandafir, P. C. (2007), An optimal experimental design criterion for discriminating between non-normal models, *Journal of the Royal Statistical Society B*, 69, 231–242.
- Marcus, M. B. and Sacks, J. (1976), Robust designs for regression problems, in: *Statistical Theory and Related Topics II*, eds. Gupta, S. S. and Moore, D. S., Academic Press, New York, pp. 245–268.
- Montepiedra, G. and Wong, W.-K. (2001), A new design criterion when heteroscedasticity is ignored, *Annals of the Institute of Statistical Mathematics*, 53, 418–426.
- Notz, W. (1989), Optimal designs for regression models with possible bias, *Journal of Statistical Planning and Inference*, 22, 43–54.
- O'Brien, T. E. (1995), Optimal design and lack of fit in nonlinear regression models, in: *Proceedings of the 10th International Workshop on Statistical Modelling*, Lecture Notes in Statistics, Springer-Verlag, New York, pp. 201–206.
- Ou, B. and Zhou, J. (2009), Minimax robust designs for field experiments, *Metrika*, 69, 45–54.
- Oyet, A. J. (2002), Minimax A - and D -optimal integer-valued wavelet designs for estimation, *The Canadian Journal of Statistics* 30, 301–316.
- Oyet, A. J. and Wiens, D. P. (2000), Robust designs for wavelet approximations of regression models, *Journal of Nonparametric Statistics*, 12, 837–859.
- Oyet, A. J. and Wiens, D. P. (2003), On exact minimax wavelet designs obtained by simulated annealing, *Statistics and Probability Letters*, 61, 111–121.
- Pesotchinsky, L. (1978), Φ_p -optimal second order designs for symmetric regions, *Journal of Statistical Planning and Inference*, 2, 173–188.
- Pesotchinsky, L. (1982), Optimal robust designs: linear regression in R^k , *The Annals of Statistics*, 10, 511–525.
- Pukelsheim, F. and Rieder, S. (1992), Efficient rounding of approximate designs, *Biometrika*, 79, 763–770.
- Rosenberger, W. F. and Grill, S. E. (1997), A sequential design for psychophysical experiments: An application to estimating timing of sensory events, *Statistics in Medicine*, 16, 2245–2260.

- Rychlik, T. (1987), Robust experimental design: A comment on Huber's result, *Zastosowania Matematyki*, 19, 93–107.
- Sacks, J. and Ylvisaker, D. (1966), Designs for regression problems with correlated errors, *The Annals of Statistics*, 37, 66–89.
- Sacks, J. and Ylvisaker, D. (1968), Designs for regression problems with correlated errors; Many Parameters, *The Annals of Statistics*, 39, 49–69.
- Sacks, J. and Ylvisaker, D. (1978), Linear estimation in approximately linear models, *The Annals of Statistics*, 6, 1122–1137.
- Sacks, J. and Ylvisaker, D. (1984), Some model robust designs in regression, *The Annals of Statistics*, 12, 1324–1348.
- Shi, P., Ye, J., and Zhou, J. (2003), Minimax robust designs for misspecified regression models, *The Canadian Journal of Statistics*, 31, 397–414.
- Shi, P., Ye, J., and Zhou, J. (2007), Discrete minimax designs for regression models with autocorrelated MA errors, *Journal of Statistical Planning and Inference*, 137, 2721–2731.
- Sinha, S. and Wiens, D. P. (2002), Robust sequential designs for nonlinear regression, *The Canadian Journal of Statistics*, 30, 601–618.
- Sinha, S. and Wiens, D. P. (2003), Asymptotics for robust sequential designs in misspecified regression models; mathematical statistics and applications, *Festschrift for Constance van Eeden, IMS Lecture Notes - in: Monograph Series*, eds. Moore, M., Léger, C. and Froda, S., 233–248. Institute of Mathematical Statistics, Hayward, CA.
- Sitter, R. R. (1992), Robust designs for binary data, *Biometrics*, 48, 1145–1155.
- Smucker, B. J., del Castillo, E. and Rosenberger, J. (2011), Exchange algorithms for constructing model-robust experimental designs, *Journal of Quality Technology*, 43, 1–15.
- Spruill, C. (1985), Model robustness of hoel-levine optimal designs, *Journal of Statistical Planning and Inference*, 11, 217–225.
- Stigler, S. M. (1971), Optimal experimental design for polynomial regression, *Journal of the American Statistical Association*, 66, 311–318.
- Studden, W. J. (1982), Some robust-type D -optimal designs in polynomial regression, *Journal of the American Statistical Association*, 77, 916–921.
- Tang, D.-I. (1993), Minimax regression designs under uniform departure models, *The Annals of Statistics*, 21, 434–446.
- Tsai, Y.-L. and Zhou, J. (2005), Discrete M -robust designs for regression models, *Journal of Statistical Planning and Inference*, 131, 393–406.
- Wiens, D. P. (1990), Robust, minimax designs for multiple linear regression, *Linear Algebra and Its Applications*, 127, 327 - 340.
- Wiens, D. P. (1991), Designs for approximately linear regression: Two optimality properties of uniform designs, *Statistics and Probability Letters*; 12, 217–221.
- Wiens, D. P. (1992), Minimax designs for approximately linear regression, *Journal of Statistical Planning and Inference*, 31, 353–371.
- Wiens, D. P. (1993), Designs for approximately linear regression: Maximizing the minimum coverage probability of confidence ellipsoids, *The Canadian Journal of Statistics*, 21, 59–70.
- Wiens, D. P. (1998), Minimax robust designs and weights for approximately specified regression models with heteroscedastic errors, *Journal of the American Statistical Association*, 93, 1440–1450.
- Wiens, D. P. (2000), Robust weights and designs for biased regression models: Least squares and generalized M -estimation, *Journal of Statistical Planning and Inference*, 83, 395–412.
- Wiens, D. P. (2005a), Robust allocation schemes for clinical trials with prognostic factors, *Journal of Statistical Planning and Inference*, 127, 323–340.
- Wiens, D. P. (2005b), Robustness in spatial studies II: Minimax design, *Environmetrics*, 16, 205–217.
- Wiens, D. P. (2009), Robust discrimination designs, *Journal of the Royal Statistical Society (Series B)*, 71, 805–829.
- Wiens, D. P. (2010), Robustness of design for the testing of lack of fit and for estimation in binary response models, *Computational Statistics and Data Analysis*, 54, 3371–3378.

- Wiens, D. P. and Wu, E. K. H. (2010), A comparative study of robust designs for M-estimated regression models, *Computational Statistics and Data Analysis*, 54, 1683–1695.
- Wiens, D. P. and Xu, X. (2008a), Robust prediction and extrapolation designs for misspecified generalized linear regression models, *Journal of Statistical Planning and Inference*, 138, 30–46.
- Wiens, D. P. and Xu, X. (2008b), Robust designs for one-point extrapolation, *Journal of Statistical Planning and Inference*, 138, 1339–1357.
- Wiens, D. P. and Zhou, J. (1996), Minimax regression designs for approximately linear models with autocorrelated errors, *Journal of Statistical Planning and Inference*, 55, 95–106.
- Wiens, D. P. and Zhou, J. (1997), Robust designs based on the infinitesimal approach, *Journal of the American Statistical Association*, 92, 1503–1511.
- Wiens, D. P. and Zhou, J. (1999), Minimax designs for approximately linear models with AR(1) errors, *The Canadian Journal of Statistics*, 27, 781–794.
- Wiens, D. P. and Zhou, J. (2008), Robust estimators and designs for field experiments, *Journal of Statistical Planning and Inference*, 138, 93–104.
- Wierich, W. (1988), A-optimal design measures for one-way layouts with additive regression, *Journal of Statistical Planning and Inference*, 18, 57–68.
- Woods, D. C., Lewis, S. M., Eccleston, J. A., and Russell, K. G. (2006), Designs for generalized linear models with several variables and model uncertainty, *Technometrics*, 48, 284–292.
- Wu, C.-F. (1981), On the robustness and efficiency of some randomized designs, *Annals of Statistics*, 9, 1168–1177.
- Xie, M.-Y. and Fang, K.-T. (2000), Admissibility and minimaxity of the uniform design measure in nonparametric regression model, *Journal of Statistical Planning and Inference*, 83, 101–111.
- Xu, X. (2009a), Robust designs for misspecified exponential regression models, *Applied Stochastic Models in Business and Industry*, 25, 179–193.
- Xu, X. (2009b), Robust prediction and extrapolation designs for censored data, *Journal of Statistical Planning and Inference*, 139, 486–502.
- Xu, X. and Yuen, W. K. (2011), Applications and implementations of continuous robust designs, *Communications in Statistics – Theory and Methods*, 40, 969–988.
- Yue, R.-X. and Hickernell, F. J. (1999), Robust designs for fitting linear models with misspecification, *Statistica Sinica*, 9, 1053–1069.
- Zhou, J. (2001a), Integer-valued, minimax robust designs for approximately linear models with correlated errors, *Communications in Statistics: Theory and Methods*, 30, 2139.
- Zhou, J. (2001b), A robust criterion for experimental designs for serially correlated observations, *Technometrics*, 43, 462–467.
- Zhou, J. (2008), D-optimal minimax regression designs on discrete design space, *Journal of Statistical Planning and Inference*, 138, 4081–4092.
- Zhu, W. and Wong, W.-K. (2001), Bayesian optimal designs for estimating a set of symmetrical quantiles, *Statistics in Medicine*, 20, 123–137.

21

Algorithmic Searches for Optimal Designs

Abhyuday Mandal, Weng Kee Wong, and Yaming Yu

CONTENTS

21.1 Introduction.....	755
21.2 Algorithmic Approach to Solve a Design Problem: A Motivating Example.....	756
21.2.1 Drug Discovery Problem.....	757
21.3 Background.....	758
21.4 Review of Selected Traditional Algorithms.....	761
21.4.1 Fedorov–Wynn Type of Algorithms.....	762
21.4.2 Exchange Algorithms.....	763
21.4.3 Issues with Algorithms.....	764
21.5 Alternative Algorithms for Finding Optimal Designs.....	765
21.5.1 Multiplicative Algorithms.....	767
21.5.2 Application of Multiplicative Algorithm in Pharmacology.....	768
21.5.3 Cocktail Algorithm.....	769
21.5.4 Application of a Cocktail Algorithm in Pharmacology.....	770
21.6 Metaheuristic Algorithms.....	771
21.6.1 Genetic Algorithms.....	772
21.6.2 Simulated Annealing.....	775
21.6.3 Particle Swarm Optimization.....	776
21.7 Summary.....	778
Acknowledgments.....	779
References.....	780

21.1 Introduction

Research in optimal experimental design has a long history and dates back as early as 1918 in a seminal paper by Smith (1918) and probably earlier. This chapter discusses algorithms for finding an optimal design given a statistical model defined on a given design space. We discuss background and the need for algorithms to find an optimal design for various situations. There are different types of algorithms available in the literature, and even for the same design problem, the researcher usually has several algorithms to choose from to find an optimal design. There are also algorithms that use specialized methods to find an optimal design for a very specific application. For example, Syed et al. (2011) used a mathematical programming technique to search for a D -optimal design using cyclotomic cosets. The literature on algorithms to find an optimal design for a statistical model is therefore huge and diverse.

The aim of this chapter is to give a brief overview on algorithms for finding optimal designs and to discuss selected algorithms that represent some of the current trends. We also highlight algorithms that are more widely applicable for solving different types of design problems. We discuss typical problems encountered in using an algorithmic approach to find an optimal design and present alternative algorithms from other fields that seem capable of generating efficient designs quickly and easily for any model and objective in the study. Along the way, we provide pseudocodes for some of these algorithms and illustrate how they work to solve real and current design problems in biomedicine.

In recent years, experimental costs have risen steeply at many levels, and researchers increasingly want to minimize study costs without having to sacrifice the quality of the statistical inference. Even with cost considerations aside, design issues are still very important. This is because a poor design can provide unreliable answers either because the estimates have unacceptably large variances or it provides low power for testing the main hypothesis in the scientific study. In the extreme case, when the study is so badly designed, it may not even provide an answer to the main scientific question of interest no matter how large the sample size. Thus all studies should be carefully planned at the onset. The main goal or goals have to be clearly specified in advance, along with all model assumptions, constraints, and practical concerns associated with execution and interpretation of the experiment.

Typically, a mathematical function called a design optimality criterion is formulated to reflect the objectives of the study as accurately as possible. A common criterion is D -optimality for estimating all model parameters; if only a subset of the model parameters is of interest, D_s -optimality is used. Both criteria seek to minimize the volume of the confidence region for the parameters of interest when errors are independent and normally distributed. Other commonly used design criteria are discussed in other chapters of this book. In general, given the user-selected criterion and a statistical model, the design problem is to find the optimal distinct combinations of the independent variables that define the treatments for which the outcome is to be observed and the number of replicates (or repeats) at each of these settings. Throughout the chapter, we assume the sample size n is predetermined, and so we are not dealing with a sample size determination problem.

This chapter is organized as follows. In the next section, we discuss the need for a carefully designed study in a real biomedical problem and briefly describe how design issues can be addressed using an algorithm described later in the chapter. Section 21.3 provides background material, fundamental concepts in designs, common terminology, and tools to find an optimal design as well as verify if a design is optimal. Section 21.4 reviews two popular types of algorithms: Fedorov–Wynn type of algorithms and exchange algorithms. In Section 21.5, we discuss alternative and modern algorithms for generating an optimal design, and in Section 21.6, we present metaheuristic algorithms, which have more recently been used as effective and practical tools for finding optimal designs. A summary is offered in Section 21.7.

21.2 Algorithmic Approach to Solve a Design Problem: A Motivating Example

This section provides an example of a real problem that can be solved using an algorithm discussed in this chapter. We omit technical details but provide references for the source where the solution of the problem, codes, and implementation of the actual algorithm can

be found. The motivating application we have in mind is how to efficiently design a drug discovery study.

21.2.1 Drug Discovery Problem

Identifying promising compounds from a vast collection of feasible compounds is a challenging problem in drug discovery. In this combinatorial chemistry problem, the goal is to obtain sets of reagents (or monomers) that maximize certain efficacy of a compound. However, here, the objective is to identify several “nearly best” combinations, rather than only one “best” or optimal one. In a typical problem, a core molecule is identified to which reagents are attached at multiple locations. Each attachment location may have tens or hundreds of potential monomers. Mandal et al. (2009) considered an example where a compound was created by attaching reagents to the three locations of a molecule, denoted by A, B, and C (e.g., see Figure 21.1). In this kind of application, the compound library (the set of all feasible compounds) may consist of 5 feasible substructures (monomers) at position A, 35 at position B and 250 at position C. That is, the compound library has a total of $5 \times 35 \times 250 = 43,750$ chemical compounds. Production of all these compounds for physical testing is expensive, and thus, it is desirable to select a relatively much smaller subset of the compounds with desirable properties. Once a compound is created, its physiochemical properties (namely, absorption/administration, distribution, metabolism, excretion, toxicity [ADMET]) are used to identify whether the compound is promising or not.

Using terminology in the design literature, we have here three factors (A, B, and C) in this design problem with 5, 35, and 250 levels, respectively. The response can be one of the ADMET properties of a compound, identified by a particular combination of A, B, and C. Alternatively, we can use a multiple-objective design discussed in Chapter 25 to capture the goals of the study simultaneously. The purpose of this study is to identify level combinations that can reduce, say, the toxicity of a compound, and at the same time increase, say, its absorption capability. Mandal et al. (2007) used desirability scores to reduce the dual goals to a single-objective optimization problem, where the goal was to identify a compound (i.e., a design point) x_i that will maximize the objective function ψ given by

$$\psi(x_i) = \psi(x_{i1}, \dots, x_{ip}). \quad (21.1)$$

Here, $p = 3$ and (x_{i1}, x_{i2}, x_{i3}) denote the levels of the three factors A, B, and C. Common design techniques such as fractional factorial designs, orthogonal arrays, and response surface designs have been widely used in screening studies for a long time in many industries (Dean and Lewis 2006). In this problem, the outcome or outcomes do not have a known mean structure in terms of the factors, and so they cannot be applied directly. With the




FIGURE 21.1

The core molecule of a compound with three reagent locations. (Reprinted (adapted) with permission from Mandal, A., Johnson, K., Wu, C.F.J., and Bornmeier, D., Identifying promising compounds in drug discovery: Genetic algorithms and some new statistical techniques, *J. Chem. Inf. Model.*, 47, 981–988. Copyright 2007 American Chemical Society.)

advent of technology, more complicated assumptions are increasingly required, and consequently, nonstandard design approaches are called for to solve the design problem. There is no analytical description of the optimal design, and so an optimal design has to be found using an algorithm. In this case, we applied a metaheuristic algorithm to search for an optimal design.

Metaheuristic algorithms are gaining in popularity because of their many successes in tackling high-dimensional optimization problems in other fields. The research in this field is very active; there are many book chapters and monographs in this area, including some that are updated in less than 2 years just to keep track of its rapid development; see Yang (2008, 2010) for example. Metaheuristic algorithms are based on heuristics that are, in part or all, not supported by theory. Here, we take the term *heuristic* to mean any procedure, or algorithm, that is not wholly supported by theory but appears to work well in practice. For example, in the algorithm used by Cook and Nachtshiem (1982) described just before Section 21.4.1, the authors offer a heuristic rule calling for restarting their algorithm after every 25 iterations for greater efficiency. Such rules are provided based on empirical experience and may not even apply to other problems or scenarios. The prefix meta- in metaheuristic suggests that it has a common mission and not a specific one for a particular problem. In our case, a metaheuristic algorithm for finding an optimal design means no specific feature of the design problem should greatly affect the functionality of the algorithm. For instance, the algorithm should work whether the design criterion is differentiable or not or whether the criterion is convex or not. Within broad limits, performance of the algorithm should not depend on the number of design variables, and constraints on those variables should be easily accommodated. Generally, only obvious modifications of the algorithm are required, and they will include, for example, modifying the statistical model or the design criterion or the design space. This is in contrast to, say, Fedorov-type algorithms where it is assumed that the optimality criterion is differentiable. Of course, as expected, more complicated optimization problems will require more time to solve.

Mandal et al. (2006) used a version of the genetic algorithm (GA), which is a metaheuristic algorithm, to find the optimal compounds to be created in the laboratory. In their application, $p = 3$ and x_{i1} , x_{i2} , and x_{i3} take the possible values of A, B, and C, respectively. The ψ in (21.1) is user selected and an example of such a function is given in (21.9). In that section, we illustrate how GA can be applied to identify the settings that maximize the objective function.

21.3 Background

Throughout this chapter, we assume that we are given a model, an optimality criterion, and a fixed sample size n and the problem is how to take n independent observations from the given design space in an optimal way. When a parametric approach is adopted, the i th outcome y_i is modeled by

$$y_i = f(x_i, \beta) + \epsilon_i, \quad i = 1, \dots, n. \quad (21.2)$$

Here, ϵ_i is the error incurred at the i th trial, and all errors are independent, and each is distributed as $\epsilon_i \sim N(0, \sigma^2)$. The mean response function $f(x, \beta)$ is assumed to be known, and β is the vector of unknown parameters. All outcomes are observed at user-selected

points from a specified compact design space \mathcal{X} , which can be multidimensional. The choice of the user-selected points $x_1, x_2, \dots, x_n \in \mathcal{X}$ makes up the design for the study.

Linear models are commonly used in scientific studies. They arise as a special case of (21.2) when the mean response function can be written as the product of two components with $f'(x_i, \beta) = f'(x_i)\beta$ and $f(x)$ is the regression function with p linearly independent components. In this chapter, linear models with heteroscedastic, independent, and normal errors are called general linear models and defined by

$$y_i = f'(x_i)\beta + \epsilon_i, \quad i = 1, \dots, n. \quad (21.3)$$

We follow optimal design terminology and use an efficiency function $\lambda(x)$ to incorporate the heteroscedasticity by letting $\text{Var}(\epsilon_i) = \sigma^2/\lambda(x_i)$. When the efficiency function is known, design issues can be directly addressed using a suitably transformed homoscedastic model.

Following convention, the goodness of a design is measured by the Fisher information matrix. Apart from an unimportant multiplicative constant, this matrix is obtained by first computing the expectation of the matrix of second derivatives of the log likelihood function at a single point and then averaging it over all the design points and multiplying by -1 . For nonlinear models, this matrix depends on the unknown parameters β , but not so for linear models. For example, if (21.3) holds and we have resources to take n_i observations at $x_i, i = 1, 2, \dots, k$, then the information matrix for this linear model is

$$\sum n_i f(x_i) f'(x_i),$$

apart from a multiplicative constant. When the model is nonlinear, we approximate the information matrix by replacing $f(x_i)$ by the gradient of $f(x_i, \beta)$ in the aforementioned matrix. The simplest way to construct optimal designs for nonlinear models is to assume that nominal values for β are available. After plugging the nominal values into the information matrix, the criterion no longer depends on β , and the design problem can be solved using design techniques for linear models. Because such optimal designs depend on the nominal parameter values, they are termed locally optimal. In what is to follow, the information matrix is denoted by $M(\xi, \beta)$ where ξ is the design used to collect the data. If we have a linear model, the information matrix does not depend on β , and we simply denote it by $M(\xi)$.

There are two types of optimal designs: approximate optimal designs and exact optimal designs. Approximate designs, or continuous designs, are easier to find and study than exact optimal designs. Approximate designs are essentially probability measures defined on a compact and known design space. In our setting, we assume that we are given a pre-selected sample size n , a design criterion, and a statistical model. The design questions are how to choose design points in the design space to observe the outcome in an optimal manner. For approximate designs, the optimization problem finds the optimal probability measure that then has to be rounded appropriately to an exact design before it can be implemented in practice. For example, if n is the total number of observations to be taken in the experiment and the approximate design calls for taking observations at three points x_1, x_2 , and x_3 from \mathcal{X} with weights w_1, w_2 , and w_3 , the implemented design takes nw_i observations at $x_i, i = 1, 2, 3$ such that $nw_1 + nw_2 + nw_3 = n$, and the weights are rounded so that each of the summands is an integer. Clearly, the implemented design is not unique as it depends on the rounding procedure. In contrast, an optimal exact design solves the optimization problem by finding an optimal integer variable k , the number of unique design

points, where these points x_1, x_2, \dots, x_k are from the design space and the number of observations n_1, n_2, \dots, n_k to be taken at each of these design points and subject to the constraint that they sum to n .

The main appeal of working with approximate designs is that after formulating the objective as a convex function of the information matrix, we have a convex optimization problem, and equivalence theorems from convex analysis theory can be directly used to verify whether a design is optimal or not (Kiefer 1974). In addition, algorithms are available for finding different types of optimal approximate designs. There is no general algorithm for finding an optimal exact design and no general analytical tool for confirming whether an exact design is optimal or not.

The traditional way of finding an optimal design is a mathematical derivation based on model assumptions. Frequently, tedious algebra and specialized mathematical tools that exploit special properties of the regression functions are required to determine the optimal design. For simpler problems, a closed form formula for the optimal design may be possible, but at other times only an analytical description is available. Generally, D -optimal designs are considered the easiest to find compared with other types of optimal designs. For example, if we wish to estimate only a subset of the whole set of model parameters in the mean function, formulae can be complicated. Furthermore, a D_s -optimal design for estimating selected parameters in a polynomial regression model is usually described by a system of equations whose solutions provide the canonical moments of the optimal design, and the optimal design is then recovered from the set of canonical moments in a complex manner; details are given in a monograph by Dette and Studden (1997).

Having a formula or a closed-form description of the optimal design is helpful because it facilitates studying properties of the optimal design, including its sensitivities to model assumptions. However, a purely theoretical approach to describing an optimal design analytically can be limiting in terms of scope and usability. This is because the optimal design is derived under a very strict set of assumptions, and so the theoretical results are applicable to that specific setting only. For instance, model assumptions frequently made at the onset of the study may be questionable, and it is desirable to know how robust the optimal design is when some aspects of the model assumptions are changed. Chapter 20 elaborates on this important issue and how to make a design more robust to model assumptions. As a specific case, errors are often assumed to have a known homoscedastic structure for simplicity. What is the corresponding optimal design when errors become slightly heteroscedastic? Unfortunately, the optimal design for the model with heteroscedastic errors frequently cannot be deduced from the theory used for the construction of the optimal design for the homoscedastic model because the technical justifications used in the derivation of an optimal design are usually quite involved and not applicable for another model. This is especially true for nonlinear models where the method of proof usually depends uniquely on the model and the design criterion. Consequently, analytical results stated for a particular situation are of limited use in practice where different conditions apply.

Algorithms are therefore very useful for generating different types of optimal designs in practice because they can overcome the problems just described associated with the theoretical approach. Algorithms generally do not depend on the mathematical complexities of the problem as much as analytical approaches, and at the same time, they can also apply more generally to different settings. Users provide design inputs for the problems of interest, and the algorithm converges to the optimal design or a design close to the optimum. If the criterion is differentiable, Fedorov-type algorithms have been shown to converge to the optimal design for a linear model (Fedorov 1972). If the criterion is not differentiable, we know of no algorithm that has been proven to converge to the optimal

design for a general linear or nonlinear model. An example of a class of nondifferentiable criteria is the minimax type of criteria, where the optimal design minimizes the maximum of some quantity and the minimization is over all designs associated with the defined space. For example, in response surface estimation, one may want to find a design that minimizes the maximum variance of the fitted response over a user-selected region. Another example of a minimax criterion concerns parameter estimation, where there are variances from a few estimated parameters in the model, and the goal is to find a design that minimizes the maximum of these variances. A sample of work in the construction of minimax optimal designs for linear models can be found in Wong (1992), Wong and Cook (1993), Brown and Wong (2000), and Chen et al. (2008). King and Wong (2000) constructed minimax optimal designs for the two-parameter logistic model when the nominal value of each of the two parameters was known to belong to a specified interval, and the goal was to find a minimax D -optimal design that minimized the maximal inefficiency that could arise from misspecification of the nominal values from the two intervals. Using similar ideas, Berger et al. (2000) found minimax optimal designs for the more complicated item theory response models commonly used in education research. Most recently, Chen et al. (2014) proposed a nature-inspired metaheuristic algorithm for finding minimax optimal designs in a general nonlinear regression setup.

21.4 Review of Selected Traditional Algorithms

This section provides a review of a few algorithms to find an optimal design. Because there are many algorithms in the field, we only discuss selected methods and provide a list of references from the literature.

With continued advances in technology, computational cost has decreased rapidly, and exhaustive search is becoming feasible for some problems thought to be *prohibitively large* even a decade ago. However, with the advancement of science, the demand for more complex designs has gone up as well. One example is in the construction of optimal designs for event-related fMRI studies in Chapter 25. Finding alternatives to an exhaustive search is always desirable, and some of the old algorithms, after suitable modifications, have reemerged and have been shown to be quite successful in recent years. For example, Ranjan et al. (2008) used the branch-and-bound algorithm proposed by Welch (1982) and found sequential optimal designs for contour estimation using complex computer codes. Yang et al. (2015) used Fedorov-type exchange algorithm, originally published in 1969 (see also Fedorov 1972), to obtain optimal designs for generalized linear models. There are several versions of the exchange algorithms where the search begins with a single random design, and then each design point is considered for exchange with other points. The pair of points chosen for exchange is the pair that results in maximum gain of the optimality criterion. Similarly, the “DETMAX” algorithm proposed by Mitchell (1974) may be considered an early version of an exchange algorithm. Atkinson et al. (2007) discussed variants of the exchange algorithms.

To delineate properties of the algorithms, it is helpful to compare their performances under a broad variety of settings and identify situations where some may outperform others. This is typically a hard job as one has to carefully define the scope and choose appropriate algorithms to compare using several well-defined measures of goodness. Cook and Nachtsheim (1980) compared several algorithms for constructing exact discrete D -optimal

designs, and Johnson and Nachtsheim (1983) gave guidelines for constructing exact D -optimal designs. Nguyen and Miller (1992) gave a comprehensive review of some exchange algorithms for constructing exact discrete D -optimal designs.

Applications of computer algorithms to find an optimal design are abundant in the literature, and several of them can be found in our reference list. One such example is given in Cook and Nachtsheim (1982), where they wanted to estimate uranium content in calibration standards. They assumed the underlying model is a polynomial of degree 1–6 and modified the Fedorov's algorithm described in the succeeding text to find an optimal design capable of providing maximal information on the uranium density along the uranium log. The optimal design provides useful guidelines for how the alternating sequence of thin and thick disks should be cut from the uranium–aluminum alloy log before the thin disks are used for destructive analyses.

21.4.1 Fedorov–Wynn Type of Algorithms

The Fedorov–Wynn type of algorithm is one of the earliest and most notable algorithms for finding optimal approximate designs that has enjoyed and continues to enjoy widespread popularity. In the most basic form, the algorithm requires a starting design and a stopping rule. It then proceeds by adding a point to the current design to form a new design and repeats this sequence until it meets the condition of the stopping rule. There are several modifications of the original algorithm currently in use; they may be tailored to a particular application or modified to speed up the convergence. A main reason for its popularity is that this is one of the few algorithms that can be proved to converge to the optimal design if it runs long enough. Many subsequent algorithms in the literature for finding optimal designs have features in common with the original Fedorov–Wynn algorithm.

As an illustration, we describe here the essential steps in the Fedorov–Wynn type of algorithm for finding a D -optimal approximate design for a general linear model in (21.3) where $\text{Var}(y(x))$ is $\sigma^2/\lambda(x)$ so that $\lambda(x)$ is inversely proportional to the known variance of the response at the point $x \in \mathcal{X}$. Technical details of the algorithms including proof of its convergence can be found in Fedorov (1972).

Pseudocode for Fedorov–Wynn algorithm:

1. Set $t = 0$ and choose a starting approximate design ξ_t with a nonsingular information matrix.
2. Compute its information matrix $M(\xi_t)$ and obtain its inverse $M^{-1}(\xi_t)$.
3. Determine the point x^* that maximizes $\lambda(x^*)d(x^*, \xi_t)$ over all x in the design space, where $d(x, \xi_t) = f'(x^*)M^{-1}(\xi_t)f(x^*)$ is the variance of the fitted response at the point x^* , apart from an unimportant multiplicative constant.
4. Generate a new design $\xi_{t+1} = (1 - \alpha_t)\xi_t + \alpha_t\nu_{x^*}$ where α_t 's is a preselected sequence of numbers between 0 and 1 such that its limit is 0 and its sum is infinite. Here, ν_{x^*} is the one point design supported at the point x^* .
5. If the stopping rule is not met, replace t by $t + 1$ and go to step 2.

A more specific choice for α_t is possible to increase the efficiency of the algorithm. For example, one may choose α_t to maximize the increase in the determinant of the current information matrix at each iteration. It can be shown using simple calculus that this leads to the choice of $\alpha_t = \zeta_t/[(\zeta_t + p - 1)p]$ where $\zeta_t = \lambda(x^*)d(x^*, \xi_{t-1}) - p$ and p is the number of the parameters in the mean function. In the last step, the stopping rule is user specified.

Two common examples of stopping rules are the maximum number of iterations allowed and when $\lambda(x^*)d(x^*, \xi_t) - p < \kappa$ for some small prespecified positive κ . The term $\lambda(x)d(x, \xi_t)$ appears frequently when we study D -optimal designs for heteroscedastic models and is sometimes called the weighted variance of the response at the point x using design ξ_t .

21.4.2 Exchange Algorithms

In this section, we review several exchange strategies for exact designs. Some of these are analogues of Fedorov–Wynn algorithms of the previous section for exact designs. Consider the general optimization problem where the objective function is $\psi(x_1, \dots, x_n)$ and one wishes to maximize $\psi(\cdot)$ with respect to $x_j \in \mathcal{X}$, $j = 1, \dots, n$, where \mathcal{X} is the design space. Again, let us use D -optimal designs for linear regression as an example. Suppose an observation is taken at the design point x_i , and the outcome is modeled as $y_i = f'(x_i)\beta + \epsilon_i$, where $\epsilon_i \sim N(0, \sigma^2)$ and all errors are independent. D -optimality corresponds to maximizing $\psi(x_1, \dots, x_n) = \det[f'(X)f(X)]$ where $X' = (x_1, \dots, x_n)$ and $f'(X) = (f(x_1), \dots, f(x_n))$. In other words, each design point $\{x_j\}$ corresponds to a row of the model matrix $f(X)$.

Exchange algorithms iteratively modify the current design by deleting existing design points and adding new points from the design space \mathcal{X} in an effort to increase the design criterion ψ . Multiple runs with different starting designs are often performed due to issues with local maximizers. Well-known algorithms that fall in this category include Wynn’s (1972) algorithm, Fedorov’s (1972) algorithm and its modification (Cook and Nachtsheim 1980), and k -exchange algorithms (Johnson and Nachtsheim 1983). A basic step in these algorithms is exchanging a point x_i in the current design with some $x_* \in \mathcal{X}$ where x_* is chosen such that the improvement in the objective function ψ is the greatest. They differ, however, in the choice of x_i . For Fedorov’s algorithm, x_i is chosen such that the improvement in ψ after the exchange is the greatest among all x_i . Thus, each exchange effectively performs n optimizations, one for each x_i in the current design, but only implements the best of these exchanges in the next design. Cook and Nachtsheim (1980) propose a modified Fedorov algorithm, where each iteration performs n exchanges, one for each x_i . The following are basic pseudocodes for these algorithms:

Pseudocode for Fedorov algorithm:

1. Choose the initial design $(x_1^0, x_2^0, \dots, x_n^0)$.
2. At iteration t , suppose the current design is $(x_1^t, x_2^t, \dots, x_n^t)$.
 - a. For $1 \leq j \leq n$, compute $\psi_j \equiv \max_{\mathcal{X}} \psi(x_1^t, \dots, x_{j-1}^t, x, x_{j+1}^t, \dots, x_n^t)$ where the maximization is over all $x \in \mathcal{X}$. Let x_j^* be the corresponding maximizer.
 - b. Find $j^* = \arg \max_j \psi_j$. Set $x_k^{t+1} = x_k^*$, $k = j^*$ and $x_k^{t+1} = x_k^t$, $k \neq j^*$.
3. Stop when there is no appreciable improvement in ψ .

Pseudocode for modified Fedorov algorithm:

1. Choose the initial design $(x_1^0, x_2^0, \dots, x_n^0)$.
2. At iteration t , suppose the current design is $(x_1^t, x_2^t, \dots, x_n^t)$. For $j = 1, \dots, n$ in turn, set $x_j^{t+1} = \arg \max_{\mathcal{X}} \psi(x_1^{t+1}, \dots, x_{j-1}^{t+1}, x, x_{j+1}^t, \dots, x_n^t)$ where the maximization is over all $x \in \mathcal{X}$.
3. Stop when there is no appreciable improvement in ψ .

In the k -exchange method, a subset of k current design points is chosen for exchange (replacement). These k points are often chosen so that deletion of each results in the smallest k decreases in ψ . More sophisticated algorithms include the DETMAX algorithm of Mitchell (1974). For some design criteria, these algorithms exploit special relationships between ψ before and after the exchange. For example, for D -optimality, when a design point x_i is replaced by another point $x_* \in \mathcal{X}$, the multiplicative change in ψ is given by

$$\Delta(x_i, x_*) = 1 + v(x_*, x_*) - v(x_i, x_i) + v^2(x_i, x_*) - v(x_*, x_*)v(x_i, x_i), \quad (21.4)$$

where $v(a, b) = f'(a)M^{-1}f(b)$ and M is the information matrix before the exchange, $M = f'(X)f(X)$. There is also a simple formula to compute M^{-1} after each exchange without any full matrix inversion.

Meyer and Nachtsheim (1995) proposed a cyclic coordinate-exchange algorithm and showed its effectiveness on several common design criteria such as the D -criterion and the linear criteria (i.e., linear functionals of the inverse information matrix). Rodriguez et al. (2010) used a similar strategy for G -optimal designs. In its basic form, the cyclic coordinate-exchange algorithm works as follows. Suppose each design point can be written as a vector with p coordinates, and suppose the design space \mathcal{X} is the Cartesian product of the corresponding p subspaces. Writing the design criterion as

$$\psi(x_1, \dots, x_n) \equiv \psi(x_{11}, \dots, x_{1p}, \dots, x_{n1}, \dots, x_{np}), \quad (21.5)$$

where $x_i \equiv (x_{i1}, \dots, x_{ip}) \in \mathcal{X}$, we iteratively optimize ψ over each x_{ij} , $i = 1, \dots, n$, $j = 1, \dots, p$ in turn. Thus, the method is an example of the widely used cyclic ascent algorithm. Here, each x_{ij} may be discrete or continuous. For a continuous one-dimensional x_{ij} , one can use various optimization routines such as golden-section search or parabolic interpolation (Rodriguez et al. 2010). One cycle of the algorithm consists of a sequence of np optimizations, one over each variable x_{ij} . The algorithm is stopped either when there is not much improvement in ψ after the latest cycle or when a prespecified number of cycles have been performed. As with other exchange methods, multiple runs with different starting values are recommended.

21.4.3 Issues with Algorithms

Many algorithms proposed in the literature typically proceed sequentially as follows. The user first inputs quantities that define the design problem. They typically include the statistical model, the design space, and the optimality criterion. If the model is nonlinear, the user would also have to supply nominal values for the model parameters before running the algorithm. The user then provides a (nonsingular) starting design, and the algorithm iterates until the stopping criterion is satisfied. For approximate designs, the procedure typically iterates by mixing the current design with a specially selected point to form a new design; this can be done simply by taking a convex combination of the design and the point using a sequence of weights that converges but not prematurely. An example of such a sequence of weights is to use $1/n$ at the n th iteration. For a differentiable design criterion, such as D -optimality, the selected point to introduce at each iteration is the point that maximizes the weighted variance function. The user also specifies a stopping criterion to tell the algorithm when to stop; this can be in terms of either the maximum number of iterations allowed or the minimum change in the criterion value over successive iterations. These steps are clearly exemplified in the Fedorov–Wynn type of algorithms and its

many modifications. Technical details including proof of their convergence can be found in design monographs such as Fedorov (1972), Silvey (1980), and Pázman (1986). Cook and Nachtsheim (1980) provide a review and comparisons of performance of various algorithms in different settings.

With algorithms, the main issues are proof of convergence, speed of convergence, ease of use and implementation, how applicable they are to find optimal designs for different types of problems, and how their performance compares to that of competing algorithms. Therefore, care should be exercised in the selection of an appropriate algorithm. For example, the Fedorov–Wynn type of algorithms introduces a point to the current design at each iteration, and frequently the generated design has many support points clustered around the true support points of the optimal design. Sometimes, periodic collapsing of these clusters of points to single points can accelerate convergence. Typical rules for collapsing the clusters of points into single points may be applied after a certain number of iterations, say, 100, with the expectation that this number may vary from problem to problem. If the rule for collapsing accumulated points is not appropriate, it may take a longer time to find the optimal design. Some algorithms, such as particle swarm optimization (PSO)-based algorithms described later on in the chapter, do not have this issue. However, mathematical proof of convergence of Fedorov–Wynn type of algorithms is available, but none exists for particle swarm-based algorithms.

Another issue with algorithms is that they may get *stuck* at some design and have difficulty in moving away from it in the direction of the optimal design. This can happen randomly or in part because of a poor choice of the starting design. Further, even though an algorithm has been proven to converge for some types of models, the result may not hold for others. For example, we recall the proof of convergence of Fedorov–Wynn type of algorithms was given for linear models only. When we apply them to nonlinear models or models with random effects, the algorithm may not work well.

A distinguishing feature of approximate designs is that when the criterion is a convex function of the information matrix, it is possible to verify whether the generated design is optimal or not. When the design space is low dimensional, a graphical plot can confirm optimality using convex analysis theory. The same theory also provides a lower bound for the efficiency of the generated design if it is not optimal. The case with exact optimal designs is very different. Frequently, there is no guarantee that the design generated by the algorithm is indeed the optimal one because a general tool for confirming optimality is not available. Researchers often show that their proposed algorithm works better by producing a design that is superior to those obtained from other algorithms.

21.5 Alternative Algorithms for Finding Optimal Designs

As mentioned before, there are other numerical tools for finding optimal designs. For example, Chaloner and Larntz (1989) used Nelder–Mead method to search for Bayesian A -, c -, and D -optimal designs for logistic regression models with one independent variable. Purely exchange-type algorithms are known to perform extremely poorly in this context. One can also rewrite the approximate design problem as an unconstrained optimization problem and apply the popular conjugate gradient or quasi-Newton methods to solve the resulting optimization problem (Atkinson et al. 2007). These are powerful and general optimization methods, which can be applied to find various types of optimal designs for a wide

variety of models and design criteria. However, they may not be the best ways for finding approximate designs (Yu 2011) because they can be quite slow and convergence is not guaranteed. We elaborate on this a bit in Section 21.5.1.

Another class of methods that has recently gained popularity is mathematical programming methods. Some of these methods were used as early as the 1970s but only more recently have statisticians started investigating how such methods work for finding optimal designs. Examples of such methods include semidefinite programming (SDP) that is well discussed in Vandenberghe and Boyd (1996) and semi-infinite programming (SIP) that is well discussed in Reemtsen and Ruckman (1998). In SDP, one optimizes a linear function subject to linear constraints on positive semidefinite matrices. Depending on the design criteria, this may not be an easy task. For example, it is relatively easy for A - or E -optimality, where the optimization problem can be rewritten as an SDP and then solved using standard methods such as interior point methods. Atashgah and Seifi (2009) discussed SDP for handling optimal design for multiresponse experiments. The examples provided by the authors were quite illuminating, and the approach seemed more versatile and powerful than the algorithm proposed by Chang (1997) for finding D -optimal designs for multiple response surface models constructed from polynomials. See also Papp (2012) and Duarte and Wong (2014a) who applied SDP to find optimal designs for rational polynomial models and several types of Bayesian optimal design problems, respectively. Filová et al. (2012) also applied SDP to find another type of optimal designs under a nondifferentiable criterion, namely, a design that maximizes the minimum of some measure of goodness in an experiment. Duarte and Wong (2014b) applied SIP to find minimax optimal designs for nonlinear models. Such mathematical programming tools have long been widely used in the engineering field for various optimization purposes, and it is a curiosity why such methods are not as popular in statistics as other optimization tools.

Another class of algorithms with a long history (explained in more detail in Section 21.5.1) for finding optimal designs is the class of multiplicative algorithms. Early work in this area was initiated by Titterton (1976, 1978). Recent theoretical advances and new ways to increase the speed of such algorithms have resulted in its renewed interest. In particular, Mandal and Torsney (2006) applied multiplicative algorithms to clusters of design points for better efficiency. Harman and Pronzato (2007) proposed methods to exclude nonoptimal design points so as to reduce the dimension of the problem. Dette et al. (2008) modified the multiplicative algorithm to take larger steps at each iteration but still maintain monotonic convergence. Yu (2011) combined the multiplicative algorithm, a Fedorov exchange algorithm, and a nearest neighbor exchange (NNE) strategy to form the cocktail algorithm, which appears to be much faster without sacrificing monotone convergence; see Section 21.5.3 for further discussion. See also Torsney and Martin-Martin (2009) who used multiplicative algorithms to search for optimal designs when responses are correlated. Harman (2014) proposed easy-to-implement multiplicative methods for computing D -optimal stratified designs. Here, *stratified* refers to allocating given proportions of trials to selected nonoverlapping partitions of the design space. Harman (2014) proposed two methods: one using a renormalization heuristic and the other using a barycentric algorithm.

Recently, Yang, Biedermann and Tang (2013) proposed a new iterative algorithm for computing approximate designs. At each iteration, a Newton-type algorithm is used to find the optimal weights given the current set of support points; a new support point is also added as in Fedorov's exchange method to ensure that true support points are not accidentally omitted. The method seems very promising (the authors reported computational speeds even better than the cocktail algorithm). Although the Newton steps make

the implementation more involved, their algorithm readily applies to various models and design criteria.

The next few subsections describe a couple of special algorithms that seem to have garnered increased interest of late. We provide some details, including their relationship to exchange algorithms, and also demonstrate how to use them for a few specific applications in the pharmaceutical arena.

21.5.1 Multiplicative Algorithms

Multiplicative algorithms are a class of simple procedures proposed by Titterton (1976, 1978) to find approximate designs in a discrete design space; see also Silvey et al. (1978). Unlike Fedorov–Wynn type of algorithms, a multiplicative algorithm at each iteration adjusts the whole vector $\boldsymbol{w} = (w_1, \dots, w_n)$ of design weights. Note that weights are computed for all points in the design space. Each weight is adjusted by a multiplicative factor so that relatively more weight is placed on design points whose increased weight may result in a larger gain in the objective function. Mathematically, suppose $\psi(\boldsymbol{w})$ is the objective function (design criterion) evaluated at the Fisher information matrix corresponding to the weight allocation (\boldsymbol{w}) and then each iteration of a general multiplicative algorithm can be written as

$$w_i^{(t+1)} \propto w_i^{(t)} \left(\frac{\partial \psi(\boldsymbol{w}^{(t)})}{\partial w_i} \right)^\rho, \quad i = 1, \dots, n, \quad (21.6)$$

where $\rho > 0$ and the ρ th power of the derivative serves as the multiplicative factor (other functions are also possible).

This simple and general algorithm has received considerable attention; see, for example, Titterton (1976, 1978), Silvey et al. (1978), Pázman (1986), Fellman (1989), Pukelsheim and Torsney (1991), Mandal and Torsney (2006), Harman and Pronzato (2007), Dette et al. (2008), Yu (2010), and Yu (2011). On the theoretical side, Yu (2010) derived general conditions under which the multiplicative algorithm monotonically increases the objective function, which yields stable convergence. One main advantage of multiplicative algorithms is their simplicity, as illustrated by the following pseudocode.

Pseudocode for multiplicative algorithm ($\rho > 0$):

1. Choose starting weights $\boldsymbol{w}^{(0)} = (w_1^{(0)}, w_2^{(0)}, \dots, w_n^{(0)})$ such that $w_j^{(0)} > 0$ for all $1 \leq j \leq n$.
2. At iteration t , suppose the current weights are $\boldsymbol{w}^{(t)} = (w_1^{(t)}, w_2^{(t)}, \dots, w_n^{(t)})$. Compute $\chi_j = \partial \psi(\boldsymbol{w}) / \partial w_j$ at $\boldsymbol{w} = \boldsymbol{w}^{(t)}$, and form $w_j^{(t+1)} = w_j^{(t)} \chi_j^\rho / \sum_i w_i^{(t)} \chi_i^\rho$, $1 \leq j \leq n$.
3. Iterate steps 1 and 2 until convergence.

Despite their simplicity and (in certain cases) monotonic convergence, multiplicative algorithms are often slow. In Section 21.5.3, we discuss some improvements and alternatives. Some recent multiplicative algorithms are based on the Fedorov–Wynn algorithm and its modifications.

21.5.2 Application of Multiplicative Algorithm in Pharmacology

Let us consider the following compartmental model that is commonly used in pharmacokinetics. The mean of the response variable at time x is modeled as

$$\eta(x, \theta) = \theta_3[\exp(-\theta_2x) - \exp(-\theta_1x)], \quad x \geq 0. \tag{21.7}$$

Frequently, the mean response measures the movement of the drug concentration in the target compartment in the body, say, in the liver. Common interests in such studies are to estimate the time the drug spends inside the compartment, the peak concentration in the compartment, and the time it takes for the compartment to receive the maximum concentration. Our interest here is to select the optimal number of time points and the optimal number of subjects from which responses will be taken at each of these time points.

Since this is a nonlinear model, the optimal design will depend on the parameter $\theta = (\theta_1, \theta_2, \theta_3)'$. We compute the locally D -optimal design at the prior guess $\theta = (4.29, 0.0589, 21.80)'$ (see Atkinson et al. 2007, Example 17.4). The design space is the interval $[0, 20]$ (in minutes) discretized, specifically $x \in \{x_1, x_2, \dots, x_n\}$ where $n = 200$ and $x_j = (j - 1)/10$. The D -optimality criterion corresponds to

$$\psi_D(w) = \log \det M(w), \quad M(w) = \sum_{i=1}^n w_i f(x_i, \theta) f'(x_i, \theta), \tag{21.8}$$

where the gradient vector $f'(x, \theta) \equiv \nabla_{\theta} \eta(x, \theta)$, often called the parameter sensitivity, is of length $m = 3$. We note that here the information matrix depends only on the weights since the whole space has been discretized, and the design problem reduces to just finding the optimal weights at these points. The multiplicative algorithm (21.6) (using $\rho = 1$, a common choice that has a convergence guarantee for the D criterion) takes a particularly simple form:

$$w_i^{(t+1)} = w_i^{(t)} m^{-1} f'(x_i, \theta) M^{-1} \left(w^{(t)} \right) f(x_i, \theta), \quad i = 1, \dots, n.$$

Table 21.1 displays the design points and their weights after 2,000, 10,000, and 50,000 iterations of this algorithm. Design weights less than 0.01 are omitted. Here, the optimal design should have weights 1/3 at each of $x = 0.2, 1.4$ and 18.4 . As one can readily see, the convergence of the algorithm is slow. After a large number of iterations, there is still a noticeable clustering of design weights around $x = 18.4$.

TABLE 21.1

Multiplicative Algorithm for D -Optimal Design for a Compartmental Model

x	0.2	1.3	1.4	18.0	18.1	18.2	18.3	18.4	18.5	18.6	18.7	18.8	18.9
$w^{(2,000)}$	0.33	0.01	0.32	0.02	0.03	0.04	0.04	0.05	0.04	0.04	0.03	0.02	0.01
$w^{(10,000)}$	0.33		0.33			0.03	0.07	0.10	0.08	0.03			
$w^{(50,000)}$	0.33		0.33				0.04	0.23	0.07				

21.5.3 Cocktail Algorithm

Several exchange strategies exist for finding approximate designs, similar to the case of exact designs. To describe these exchange strategies, let us define the directional derivative $d(i, \mathbf{w}) = \partial\psi((1 - \delta)\mathbf{w} + \delta\mathbf{e}_i)/\partial\delta|_{\delta=0+}$ where \mathbf{e}_i is the vector that assigns all the mass to the i th design point, $i = 1, \dots, n$, and zero to all other design points. As before, ψ is the objective function, and \mathbf{w} is the vector of design weights. We use the vertex direction method (VDM) (Fedorov, 1972), which is a simple iterative method: given the current $\mathbf{w}^{(t)}$, we first select the index i_{\max} with the maximum directional derivative, that is, $d(i_{\max}, \mathbf{w}^{(t)}) = \max_{1 \leq i \leq n} d(i, \mathbf{w}^{(t)})$, and then set the next iteration $\mathbf{w}^{(t+1)}$ as the maximizer of $\psi(\mathbf{w})$ along the line segment $\mathbf{w} = (1 - \delta)\mathbf{w}^{(t)} + \delta\mathbf{e}_{i_{\max}}$, $\delta \in [0, 1]$. This one-dimensional maximization is usually not too difficult and can often be done in closed form. Plainly, each iteration of VDM moves the vector of weights $\mathbf{w}^{(t)}$ toward the vertex of maximum directional derivative.

Similar to VDM, one can define a general exchange between two design points i and j as follows. Given the current $\mathbf{w}^{(t)}$, set the new vector $\mathbf{w}^{(t+1)}$ as the maximizer of $\psi(\mathbf{w})$ along the line segment $\mathbf{w} = \mathbf{w}^{(t)} + \delta(\mathbf{e}_j - \mathbf{e}_i)$, $\delta \in [-w_j, w_i]$. In other words, with the weights at other design points unchanged, those between i and j are allocated so that the objective function is maximized. Such exchanges can set the mass at a design point at zero in one iteration, unlike the multiplicative algorithm. For example, if $\delta = w_i$, then all the mass on i is transferred to j , and $w_i^{(t+1)} = 0$. Exchanging between two design points forms the basis of Bohning’s (1986) *vertex exchange method* (VEM), which uses a special choice of i and j at each iteration. Specifically, at iteration t , VEM chooses i_{\min} and i_{\max} such that

$$d(i_{\min}, \mathbf{w}^{(t)}) = \min \{d(i, \mathbf{w}^{(t)}) : w_i^{(t)} > 0\};$$

$$d(i_{\max}, \mathbf{w}^{(t)}) = \max \{d(i, \mathbf{w}^{(t)}) : 1 \leq i \leq n\}.$$

In other words, we exchange mass between a design point with maximum directional derivative and another point (among those that have positive mass) with minimum directional derivative.

Another strategy, known as the NNE, is critical to the cocktail algorithm of Yu (2011). At each iteration of NNE, we first determine an ordering on the set of support points (design points with positive mass). An ideal ordering should place similar design points close to each other. For example, suppose the design variable is a discretization of the interval $[0, 1]$, then an obvious choice is the natural ordering of the real numbers. (Note that design points with zero mass are excluded.) In general, there is no definite rule for choosing an ordering. In the context of D -optimal designs for regression problems, Yu (2011) advocates dynamically choosing an ordering at each iteration based on the L_1 distances between vectors of explanatory variables that the design points represent. Once an ordering is chosen, one performs pairwise exchanges between consecutive points on the list. For example, if the design space has four points $\{x_1, x_2, x_3, x_4\}$ and the current ordering is (x_3, x_1, x_4) (design point x_2 has zero mass and is excluded), then we exchange mass between x_3 and x_1 and then between x_1 and x_4 . Each exchange involves a one-dimensional maximization that improves the design.

The cocktail algorithm of Yu (2011) builds on both the multiplicative algorithm and exchange strategies. Each iteration of the cocktail algorithm is a simple concatenation of one iteration of VDM, one iteration of the multiplicative algorithm, and the full sequence

of NNE described earlier. Heuristically, NNE uses local exchanges that might complement the global nature of the multiplicative algorithm. NNE can also quickly eliminate design points (i.e., setting their masses to zero) because of the many pairwise exchanges involved. The VDM step ensures that a design point eliminated by NNE has a chance of *resurrection* if indeed it should receive positive mass in the optimal solution; this enables theoretical proofs of convergence (to the global optimum). Empirical comparisons of the cocktail algorithm with VEM as well as several off-the-shelf optimization methods (e.g., Nelder–Mead, conjugate gradient, and quasi-Newton) show that the cocktail algorithm improves upon traditional algorithms considerably in terms of computational speed, using effectively the same stopping criterion.

21.5.4 Application of a Cocktail Algorithm in Pharmacology

As an example, let us consider the compartmental model (21.7) again. In addition to D -optimality, we also consider c -optimality, which aims to minimize the variance of an estimated scalar function of the vector of model parameters θ , say, $g(\theta)$. The optimality criterion can be written as

$$\psi_c(w) = c'M^{-1}(w)c,$$

where M is as in (21.8) and $c' \equiv \nabla g(\theta)$. Because c -optimal designs often result in singular M matrices, a small positive quantity is added to the diagonals of $M(w)$ to stabilize the matrix inversion mentioned earlier. We emphasize that this is introduced merely to avoid numerical difficulties and does not correspond to what is intended by the c criterion. Two functions of interest (see Atkinson et al. 2007, Example 17.4) are the area under the curve (AUC), which is defined as $g_1(\theta) = \theta_3/\theta_2 - \theta_3/\theta_1$, and the time to maximum concentration (TMC), which is defined as $g_2(\theta) = (\log \theta_1 - \log \theta_2)/(\theta_1 - \theta_2)$. Table 21.2 displays the D -optimal design as well as the c -optimal designs for g_1 and g_2 found by the cocktail algorithm. These generally agree with Table 17.1 of Atkinson et al. (2007); slight discrepancies exist because we discretize the design space. In this example, it takes the cocktail algorithm only a few iterations to obtain the D -optimal design to the degree of accuracy as displayed

TABLE 21.2

Cocktail Algorithm for D - and c -Optimal Designs for a Compartmental Model

Criterion	x	Weight	Criterion Value
D -optimality	0.2	0.3333	7.3713
	1.4	0.3333	
	18.4	0.3333	
AUC	0.2	0.0137	2190.2
	17.5	0.1459	
	17.6	0.8404	
TMC	0.2	0.5916	0.028439
	3.4	0.3025	
	3.5	0.1059	

in Table 21.2. However, it takes considerably more iterations (hundreds or even thousands) to find the c -optimal designs.

21.6 Metaheuristic Algorithms

Recently, evolutionary algorithms have become more popular in finding optimal designs for industrial experiments. Here, we present some examples. Accelerated life testing has been used to study the degradation of reliable components of materials with high risk, such as nuclear reactors and aerospace vehicles. In order to set up a degradation testing procedure, several objectives have to be considered. For example, in such experiments, we want to have an adequate model for the degradation process to understand the relationship between degradation and failure time and, at the same time, high-quality estimates of model parameters. This leads to a highly nonlinear multiobjective optimization problem. Marseguerra et al. (2003) proposed using a GA for finding optimal designs and found nondominated solutions with two objective functions for estimating the test parameters efficiently. Similar algorithms are used for finding optimal designs in product assembly line as well. Traditional fractional factorial and response surface type of designs cannot be used when practical restrictions are imposed on factor-level combinations. Sexton et al. (2006) considered exchange algorithms and GA to compare their performances on product designs. In the two hydraulic gear pump examples and one electroacoustic transducer example they considered, the exchange algorithm performed better than the GA, although the authors noted that in the early stages of searching, GA performed better, and hence they recommended a combination of two algorithms, with GA to be used at the early stages before switching to the exchange algorithms.

GA are often used in obtaining robust parameter designs in the presence of control and noise variables, where low prediction variances for the mean response are often desirable. Goldfarb et al. (2005) used GA to obtain efficient designs for mixture processes. With the examples of soap manufacturing and another tightly constrained mixture problem, they obtained D -efficient designs. Rodriguez et al. (2009) used GA as well to construct optimal designs using a desirability score function (Harrington 1965) to combine dual objectives of minimizing the variances for the mean and slope into one objective function.

We focus on such algorithms inspired by nature, and so they are generally referred to as nature-inspired metaheuristic algorithms. The genesis of such algorithms could be based on a variety of observations from nature, for example, how ants colonize or how fish swim in large schools when they perceive a threat or birds fly as a flock in search of food. Generally, there is no mathematical proof that shows metaheuristic algorithms will converge to the optimum; that is why they are called metaheuristic in the first place! However, repeated experiences from researchers in many fields report that they frequently do find the optimum, and if they do not, these algorithms get to the proximity of the optimum quickly.

There are many metaheuristic algorithms in the optimization literature, and it is a curiosity why statisticians do not seem to explore and use more of them in their work. The most common examples in the statistical literature are simulated annealing (SA) and GA, discussed in the succeeding text. There are many newer ones, such as PSO, bat algorithm, differential evolution algorithm, and cuckoo algorithm. For space consideration, we discuss in the succeeding text only PSO, which is just beginning to appear in the statistical

literature. All are nature-inspired algorithms with different search techniques based on metaheuristics. Yang (2010) has a brief but good discussion on the other algorithms.

21.6.1 Genetic Algorithms

GA are one of the evolutionary algorithms that are very popular for finding exact optimal designs. Developed by John Holland and his colleagues at the University of Michigan (Holland 1975), these algorithms mimic Darwin's theory of evolution. The metaphors of natural selection, crossbreeding, and mutation have been helpful in providing a framework to explain how and why GA work. In our context, this translates to the heuristic that from two "good" design points, one should sometimes be able to generate an even "better" design point. The objective function to be optimized, denoted by ψ as in (21.5), can be treated as a *black box* without the mathematical requirements of continuity, differentiability, convexity, or other properties required by many traditional algorithms. Because of its simplicity, this algorithm has also gained popularity in finding optimal designs for computer experiments (see Bates et al. 2003; Liefvendahl and Stocki 2006; Crombecq and Dhaene 2010). It is an iterative algorithm that starts with a set of candidates and can explore very large spaces of candidate solutions. Convergence is usually not guaranteed but GA often yield satisfactory results.

In an attempt to understand how GA function as optimizers, Reeves and Wright (1999) considered GA as a form of sequential experimental design. (For details of sequential design algorithms, see Chapter 19.) Recently, GA have been used quite successfully in solving statistical problems, particularly for finding near-optimal designs (Hamada et al. 2001; Heredia-Langner et al. 2003, 2004). In this chapter, we will occasionally deviate from what is used in the optimization literature and present our algorithm in the context of search for optimal designs. One of the simplest version of the algorithm process is as follows:

1. *Solution representation*: For a single design point, a factor at a particular level is called a *gene*. The factor combinations, that is, the entire design point or *run*, are called a *chromosome*. Using the notation of Section 21.4.2, $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})$ is a chromosome, whereas each x_{ij} is called a gene. Note that it is possible to define the chromosome a little differently, as we will see in the Broudiscou et al. (1996) example later. Initially a large population of random candidate solutions (design points) is generated; these are then continually transformed following steps (2) and (3).
2. *Select* the best and eliminate the worst design point on the basis of a fitness criterion (e.g., the higher the better for a maximization problem) to generate the next population of design points.
3. *Reproduce* to transform the population into another set of solutions by applying the genetic operations of *crossover* and *mutation*. Different variants of crossover and mutations are available in literature.
 - a. *Crossover* : A pair of design points (chromosomes) are split at a random position, and the head of one is combined with the tail of other and vice versa.
 - b. *Mutation* : The state (i.e., level) of a randomly chosen factor is changed. This helps the search avoid being trapped at local optima.
4. *Repeat* steps (2) and (3) until some convergence criterion is met (usually no significant improvement for several iterations).

Note that in this sketched algorithm, it is more in the line of solving an optimization problem than a design problem. For example, we inherently assumed that all the designs we consider consist of only one run, which is not the case for most of our problems. But the aforementioned algorithm can be modified very easily to solve the standard design problems (e.g., finding orthogonal arrays with fixed run size). Next, we illustrate the algorithm stated earlier with an optimization problem, mainly for simplicity. After that, we will discuss common design problems where the entire design with multiple runs is taken as a chromosome, and the collection of such designs represent the population.

Let us illustrate the algorithm sketched earlier with a *black box* function provided by Levy and Montalvo (1985):

$$\begin{aligned} \psi(x_{i1}, \dots, x_{ip}) = & \sin^2 \left\{ \pi \left(\frac{x_{i1} + 2}{4} \right) \right\} + \sum_{k=1}^{p-1} \left(\frac{x_{ik} - 2}{4} \right)^2 \left\{ 1 + 10 \sin^2 \left(\pi \left(\frac{x_{ik} + 2}{4} \right) + 1 \right) \right\} \\ & + \left(\frac{x_{ip} - 2}{4} \right)^2 \left\{ 1 + \sin^2 (2\pi (x_{ip} - 1)) \right\}. \end{aligned} \tag{21.9}$$

Here, $p = 4$ and only integer values of x_{ik} 's ($0 \leq x_{ik} \leq 10$) are considered. This corresponds to an experiment with four factors each at 11 levels denoted by $0, 1, \dots, 10$. Suppose we start with a random population of size 10 given in Table 21.3. In this case, the ten runs $x_1 = (2, 1, 4, 5)'$, \dots , $x_{10} = (0, 9, 8, 7)'$ are the ten chromosomes that constitute the initial population.

Suppose that the best two design points, namely, $(10, 6, 7, 8)'$ and $(0, 9, 8, 7)$, are chosen as *parents* to *reproduce*. Then suppose crossover happens at location 2, so the new sets of design points will be $(10, 6, 8, 7)'$ and $(0, 9, 7, 8)'$ (the first two factors of the first design point are combined with the last two factors of the second design point and vice versa). Now suppose for $(10, 6, 8, 7)'$ the mutation happens at location 3 (for factor C) and the level is changed from 8 to 7. Then the resulting design point becomes $(10, 6, 7, 7)'$ with the value of Levy–Montalvo function as 26.81. Note that this value is worse than $(10, 6, 7, 8)'$ and that is not unexpected. What is expected is that, on average, when 10 new *offsprings* (design points) are generated in this fashion, some of them will be “good”. Those new 10 offsprings

TABLE 21.3
Initial Population for Levy–Montalvo Example

A	B	C	D	ψ
2	1	4	5	1.63
3	4	1	7	3.13
1	1	8	8	11.66
10	6	7	8	26.81
9	5	5	9	5.20
7	8	2	8	11.57
2	3	5	3	1.54
6	5	4	6	2.80
5	9	5	8	8.02
0	9	8	7	15.83

will constitute the second generation of the population, and the same process continues. Usually hundreds or even thousands of generations are needed to obtain a near-optimal design. Mandal et al. (2006) used an efficient modification of GA on this same example and showed that it produced good results.

Broudiscou et al. (1996) used GA for constructing D -optimal designs, in the context of an antigen/antibody test to diagnose an infection to a virus. There were six factors with a total of $7 \times 6 \times 6 \times 5 \times 3 \times 3 = 11,340$ combinations. In this example, the basic exchange algorithms discussed in Section 21.4.2 are not very efficient. Consider a design with $n = 28$ points. In Fedorov's exchange algorithm, during each iteration, a design point x_i will be replaced by the point x_* that maximizes $\Delta(x_i, x_*)$ of (21.4). Also, one has to evaluate all the pairs (x_i, x_*) . Naturally, when the total number of candidate design points is high (11,340 in this case), even with moderate run size (28 in this case), the algorithm will be slow because finding the "best" combination, at each iteration, is time-consuming. Stochastic search-based algorithms, such as GA, turn out to be very efficient in such situations. The authors reported a D -optimal design in Table 5 of their paper, for a mixed-level factorial main effects model. In this example, the model matrix \mathbf{X} has $1 + 6 + 5 \times 2 + 4 + 2 \times 2 = 25$ columns. The objective function is $\psi(x_1, \dots, x_n) = \det [f'(\mathbf{X})f(\mathbf{X})]$, as mentioned before, and our objective is to maximize ψ . In this case, the chromosomes are possible designs obtained by juxtaposing the rows of the design matrix, such as $(x'_1, x'_2, \dots, x'_{28})'$. Each chromosome has $6 \times 28 = 168$ genes.

Hamada et al. (2001) used GA to find near-optimal Bayesian designs for regression models. The objective function ψ in their case was the expected Shannon information gain of the posterior distribution. It is equivalent to choosing designs that maximize the expected Kullback–Leibler distance between the prior and posterior distributions (Chaloner and Verdinelli 1995). The reader is referred to the pseudocode given in the appendix of their paper. They also considered a design with p factors and n runs, and each factor level as a gene such that each chromosome (where it is the design) had np genes.

Kao et al. (2009) used GA to construct optimal designs for event-related fMRI studies. The authors considered multiple objectives, and ψ was defined to reflect all of them. In that application, the chromosomes can be 300–600 genes long, each gene taking 5–13 different values. See Chapter 25 for a detailed discussion of their approach.

Pseudocode for GA:

1. Generation 0: Generate M random designs, evaluate them by the objective function ψ , and order them by their fitness (ψ -values).
2. For generations $g = 1, \dots, G$, the following apply:
 - a. With probability proportional to fitness, draw with replacement $M/2$ pairs of designs to crossover—select a random cut point, and exchange the corresponding fractions of genes in paired designs to generate M offspring designs.
 - b. Randomly select $q\%$ of the genes from the M offspring designs and perform mutation to generate M new offspring designs.
 - c. Obtain the fitness scores of the new offsprings.
 - d. Repeat steps (a through c) until stopping criterion is met.

As mentioned before, there are several variants of GA. Being one of the metaheuristic algorithms, although the performance of GA in finding optimal design is, in general, not greatly affected by the functionality of the algorithm, in order to apply this algorithm, a number of parameters need to be determined beforehand. These include the population

size (M), mutation rates (q), crossover, and mutation locations. Such details of implementation, however, are essential for anyone to reproduce the results of a search, even after accounting for the randomness of the procedure. Lin et al. (2014) explored the merits of GA in the context of design of experiments and discussed some elements that should be considered for an effective implementation.

21.6.2 Simulated Annealing

While GA mimic Darwin’s theory of evolution, SA simulates the process called annealing in metallurgy where some physical properties of a material are altered by heating to above a critical temperature, followed by gradual cooling. It starts with a random design and then moves forward by replacing the current solution by another solution with the hope of finding near-optimal solutions. Introduced by Kirkpatrick et al. (1983), SA is a probabilistic global optimization technique and has two basic features. The first one is motivated by the Metropolis algorithm (Metropolis et al. 1953), in which designs that are worse than the current one are sometimes accepted in order to better explore the design space. The other one is the strategy for lowering the *temperature* (T), by which the probability of inclusion of new “bad” solutions is controlled. In the traditional formulation for a minimization problem, initially this time-varying parameter T is set at a high value in order to explore the search space as much as possible. At each step, the value of the objective function is calculated and compared with the best design at hand. In the context of optimal design search, designs with higher values of the objective function will be accepted according to the Metropolis algorithm if $e^{-\Delta D/T} > U$ where ΔD is the change of the objective function, T is the *current temperature*, and U is a uniform random number. Usually, the temperature is lowered at some specified intervals at a geometric rate. The higher the temperature value, the more unstable the system is, and the more frequently “worse” designs are accepted. This helps explore the entire design space and jump out of local optimum values of the objective function. As the process continues, the temperature decreases and leads to a steady state. The temperature T , however, should not be decreased too quickly in order to avoid getting trapped at local optima. There are various *annealing schedules* for lowering the temperature. Similar to GA, the algorithm also usually stops when no significant improvements are observed for several iterations.

Woods (2010) used SA to obtain optimal designs when the outcome is a binary variable. The objective of the study is to maximize $\psi(x_1, \dots, x_n)$ the log determinant of the Fisher information matrix that depends on the unknown values of the parameters. As in Section 21.4.2, here, x_i represents a vector (x_{i1}, \dots, x_{ip}) where x_{ij} represents the value of the j th variable in the i th trial, with the constraint that $-1 \leq x_{ij} \leq 1$ ($j = 1, \dots, p; i = 1, \dots, n$). In this case, x_{ij} has been perturbed to a new design point x_{ij}^{new} by

$$x_{ij}^{new} = \min\{1, \max[-1, x_{ij} + ud_T]\}$$

where u is a random number drawn from a uniform $U[-1, 1]$ distribution and d_T is the size of the maximum allowed perturbation at temperature T . The new (perturbed) design is compared with the original design via the calculation of the objective function ψ , and the new design is accepted if its ψ value is greater than that of the original one. Otherwise, the new design is accepted with probability $\min\left\{1, \exp\left(\frac{\psi(\xi^{new}) - \psi(\xi^{original})}{T}\right)\right\}$. Here, $\xi^{original}$ and ξ^{new} are, respectively, the n -run designs supported at the original and new sets of points (x_1, \dots, x_n) . Both the temperature T and perturbation size d_T are decreased geometrically.

Bohachevsky et al. (1986) developed a generalized SA method for function optimization and noted that this algorithm has “the ability to migrate through a sequence of local extrema in search of the global solution and to recognize when the global extremum has been located.” Meyer and Nachtsheim (1988) used SA (programmed in ANSI-77 standard FORTRAN) for constructing exact D -optimal designs. They evaluated their methods for both finite and continuous design spaces. Fang and Wiens (2000) used SA for obtaining integer-valued designs. In their paper, they mentioned that, following Haines (1987), they initially chose T in such a way that the acceptance rate is at least 50%. As discussed by Press et al. (1992), they decreased T by a factor of 0.9 after every 100 iterations. Among others, Zhou (2008) and Wilmut and Zhou (2011) used SA for obtaining D -optimal minimax designs.

21.6.3 Particle Swarm Optimization

PSO was proposed about 18 years ago by Kennedy and Eberhart (1995), and it has slowly but surely gained a lot of attention in the past 10 years. Researchers continue to report their successes with the algorithm for solving large-scale, complex optimization problems in several fields, including finance, engineering, biosciences, monitoring power systems, social networking, and behavioral patterns. The rate of success and excitement generated by PSO has led to at least one annual conference solely devoted to PSO for more than a decade and usually sponsored by IEEE. There are several websites that provide in-depth discussions of PSO with codes, tutorials, and both real and illustrative applications. An example of such a website is <http://www.swarmintelligence.org/index.php>. Currently, there are at least three journals devoted to publishing research and applications based on swarm intelligence PSO methods.

The special features of the PSO techniques are that the method itself is remarkably simple to implement and flexible, requires no assumption on the function to be optimized, and requires specification of only a few easy-to-use tuning parameters to values that work well for a large class of problems. This is unlike other algorithms such as the genetic or SA algorithms that can be sensitive to the choice of tuning parameters. For PSO, typically only two parameters seem to matter—the number of particles and the number of iterations—with the rest taking on the default values. A larger number of particles generate more starting designs that more likely cover the search space more adequately and so usually produces a higher quality solution. A probable downside is that it may take longer time to arrive at the optimum because more communication time is required by the larger number of particles to decide where the global optimum is. The user also has to specify the maximum number of iterations allowed for the algorithm, but this usually is inconsequential for small dimensional optimization problems. This is because the search time is typically very short, and so one can try repeatedly and find the optimal design quickly for most problems. Each particle is a potential solution of the optimization problem, and at every iteration, each has a fitness value determined by the design criterion.

There are two key equations in the PSO algorithm that define its search to optimize a user-selected objective function ψ . At the t and $t + 1$ iterations, the movement of particle i is governed by

$$v_i^{t+1} = \omega_i v_i^t + \gamma_1 \beta_1 \odot (p_i - x_i^t) + \gamma_2 \beta_2 \odot (p_g - x_i^t) \quad (21.10)$$

and

$$\mathbf{x}_i^{t+1} = \mathbf{x}_i^t + \mathbf{v}_i^{t+1}. \quad (21.11)$$

Here \mathbf{v}_i^t is the particle velocity at time t and \mathbf{x}_i^t is the current particle position at time t . The inertia weight ω_t adjusts the influence of the former velocity and can be a constant or a decreasing function with values between 0 and 1. For example, Eberhart and Shi (2000) used a linearly decreasing function over the specified time range with initial value 0.9 and end value 0.4. Further, the vector \mathbf{p}_i is the personal best (optimal) position as encountered by the i th particle, and the vector \mathbf{p}_g is the global best (optimal) position as encountered by all particles, up to time t . This means that up to time t , the personal best for particle i is $pbest_i = \psi(\mathbf{p}_i)$ and, for all particles, $gbest = \psi(\mathbf{p}_g)$ is the optimal value. The two random vectors in the PSO algorithm are β_1 and β_2 , and their components are usually taken to be independent random variables from $U(0, 1)$. The constant γ_1 is the cognitive learning factor, and γ_2 is the social learning factor. These two constants determine how each particle moves toward its own personal best position and the overall global best position. The default values for these two constants in the PSO codes are $\gamma_1 = \gamma_2 = 2$, and they seem to work well in practice for nearly all problems that we have investigated so far. Note that in (21.10), the product in the last two terms is the Hadamard product. The pseudocode for the PSO procedure using q particles is given in the succeeding text.

Pseudocode for PSO algorithm is as follows:

1. Initialize particles.
 - a. Initiate positions \mathbf{x}_i and velocities \mathbf{v}_i for $i = 1, \dots, q$.
 - b. Calculate the fitness values $\psi(\mathbf{x}_i)$ for $i = 1, \dots, q$.
 - c. Initialize the personal best positions $\mathbf{p}_i = \mathbf{x}_i$ and the global best position \mathbf{p}_g .
2. Repeat until stopping criteria are satisfied.
 - a. Calculate particle velocity according to Equation (21.10).
 - b. Update particle position according to Equation (21.11).
 - c. Update the fitness values $\psi(\mathbf{x}_i)$.
 - d. Update personal and global best positions \mathbf{p}_i and \mathbf{p}_g .
3. Output $\mathbf{p}_g = \arg \min \psi(\mathbf{p}_i)$ with $gbest = \psi(\mathbf{p}_g)$.

We have set up mirror websites at <http://optimal-design.biostat.ucla.edu/podpack/>, <http://www.math.ntu.edu.tw/~optdesign/>, and <http://www.stat.ncku.edu.tw/optdesign>, where the reader can download our PSO P-codes, run them, and verify some of the results in this chapter. Many of the PSO codes can be readily changed to find another type of optimal design for the same model or for a different model. Typically, the only changes that are required are in the information matrix and the design criterion.

The sites are new and are still under construction as improvements are made. We alert the reader that some of the notation on these sites may be different from that used in this chapter. The sites have instructions for downloading MATLAB[®] P-codes and running the codes for finding various types of optimal designs. On the interface window, we provide default values for two PSO parameters that had successfully found the optimal design before, the number of particles, and the number of iterations; all other parameters are

default values recommended by PSO and are not displayed. The interface window also provides the swarm plot showing how the swarm of initial candidate designs converges or not to the optimal location, along with a plot of the directional derivative of the design criterion for the PSO-generated design to confirm its optimality or not. The ease of use and flexibility of PSO are compelling compared with current methods of finding optimal designs.

As examples, we refer the reader to the download webpage on one of the aforementioned websites, where under the heading Part A, we have codes for finding minimax optimal designs. The first example concerns finding a design to minimize the maximum variance of the fitted response across the design space when errors have a known heteroscedastic structure, and the second example concerns E -optimality where we seek a design that minimizes the maximum eigenvalue of the inverse of the information matrix (Ehnenfeld, 1955). We invite the reader to try out the PSO codes on the website for generating various types of optimal designs for the compartmental model discussed earlier in Section 21.5.1 and compare results in Tables 21.1 and 21.2. Other sample applications of using PSO to design real studies available from the website include finding different locally optimal designs for the simple and quadratic logistic models, D -optimal designs for mixture models, locally D -optimal designs for the four-parameter Hill model used in education and biomedical studies, locally D -optimal designs for an exponential survival model with type I right censoring, and locally D -optimal designs for a double-exponential model used in monitoring tumor regrowth rate.

PSO techniques seem like a very under utilized tool in statistics to date. They seem ideally suited for finding optimal experimental designs. This is because many applications having a design close to the optimum (without knowing the optimum) may suffice for most practical purposes. When we work with approximate designs, the convexity assumption in the design criterion implies that the skeptic can also always check the quality of the PSO-generated design using an equivalence theorem. Intuitively, PSO is attractive because it uses many starting designs (particles) at the initial stage to cover the search space, and so one can expect such an approach is preferable to methods that use only a single starting design.

21.7 Summary

This chapter discusses algorithmic searches for an optimal design for a statistical model described in (21.2). We reviewed a few algorithms commonly used to find an optimal design for the problem and also newer algorithms, such as particle swarm-based algorithms.

Our discussion has focused on a single-objective study. In practice, experiments may have more than one goal or objective, and the implemented design should carefully incorporate all the objectives at the onset. In the past, the practice was to design the study to satisfy the most important objective and hope that the same design also does well in terms of other objectives in the study. Nowadays, multi objective optimization problems can be handled by constructing a multiple-objective optimal design that directly accommodates all experimental goals at the same time. These techniques require that the objectives be first prioritized in order of their importance. By construction, the multiple-objective optimal design captures the varying degrees of importance of the various objectives and

delivers user-specified efficiencies for the objectives according to their importance; see details in Cook and Wong (1994), for example. In particular, it can be shown that many of the algorithms for finding a single-objective optimal design can be directly applied to find multiple-objective optimal designs as well (Cook and Wong 1994; Wong 1999). Chapter 25 constructs some specific multiple-objective optimal designs for event-related fMRI studies.

We have not made a clear distinction between finding optimal exact or approximate designs but note that some algorithms are more flexible than others. There is no algorithm that is universally best for solving all optimization problems. Each algorithm, by construction, has its own unique strengths, weaknesses, and restrictions. For example, both multiplicative algorithms and SDP-based methods require that the search space be discretized, but PSO can work well either in a continuous or a discrete search space. Further, the performance of some algorithms, such as the GA, can be highly dependent on input values of the tuning parameters, while others, such as PSO-based algorithms, are less so. It is therefore important to fully appreciate the properties of the algorithm before implementing it to find the optimal design for the problem at hand. Frequently, for more complicated optimization problems, a hybrid algorithm that combines two or more different algorithms can prove effective because it incorporates the unique strengths from the component algorithms.

We would be remiss not to mention that there are canned software packages in commercial statistical software for finding optimal designs. For example, Atkinson et al. (2007) provided SAS codes for finding a variety of optimal designs based on various types of algorithms. Currently, algorithms are available for generating D -, c -, A -optimal designs and optimal designs found under differentiable criteria. Other statistical packages also focus on such types of optimal designs. Minimax optimal designs are notoriously difficult to find, and we are not aware of any commercial package that handles them. GA and PSO that do not require the objective function to be differentiable would seem more appropriate for such problems. Examples of work in this area are Zhou (2008) and, most recently, Chen et al. (2014) and Qiu et al. (2014).

In summary, algorithms are crucial for finding optimal designs to solve real problems and will be more important as scientists increasingly use more realistic models to reflect the complexities of the underlying process. Consequently, analytical descriptions of optimal designs will be more difficult and most likely impossible to obtain. Numerical methods are therefore necessary, and more effective algorithms should be developed and applied to solve real-world design problems. It therefore behooves the design community to always keep a constant eye of current and new optimization techniques used in the optimization literature and investigate their suitability and efficiency for finding optimal designs for a statistical problem.

Acknowledgments

The research of Dr. Mandal was in part supported by NSF grant DMS-09-05731 and NSA grant H98230-13-1-025. Dr. Wong is especially thankful to his collaborators in Taiwan for introducing him to PSO and their continuing research in this exciting topic. They are Professor Weichung Wang from the Department of Mathematics at National Taiwan University and Professor Ray-Bing Chen from the Department of Statistics at National Cheng Kung University. Thanks also go to several students from National Taiwan

University and UCLA, who all wrote and implemented MATLAB codes to find various types of optimal designs using PSO on the website. The research of Wong reported in this paper was partially supported by the National Institute of General Medical Sciences of the National Institutes of Health under Award Number R01GM107639. The content is solely the responsibility of the author and does not necessarily represent the official views of the National Institutes of Health. Dr. Yu would like to thank Professor Holger Dette for stimulating discussions concerning algorithms as well as minimax designs.

All authors are grateful to the editorial team for taking a lot of time to carefully read earlier versions of this chapter and providing excellent feedback. We thank them very much for their guidance throughout our preparation of this chapter. We also thank Guanghao Qi from the Department of Mathematics at Fudan University for comments on this chapter.

References

- Atashgah, A. B. and A. Seifi, Optimal design of multi-response experiments using semi-definite programming, *Optimization in Engineering*, 10, 75–90, 2009.
- Atkinson, A. C., A. N. Donev, and R. D. Tobias, *Optimum Experimental Designs, with SAS*, Oxford University Press, Oxford, U.K. 2007.
- Bates, S. J., J. Sienze, and D. S. Langley, Formulation of the AudzeEglais uniform Latin hypercube design of experiments, *Advances in Engineering Software*, 34, 493–506, 2003.
- Berger, M. P. F., J. King, and W. K. Wong, Minimax designs for item response theory models, *Psychometrika*, 65, 377–390, 2000.
- Bohachevsky, I. O., M. E. Johnson, and M. L. Stein, Generalized simulated annealing for function optimization, *Technometrics*, 28, 209–217, 1986.
- Böhning, D., A vertex-exchange-method in D -optimal design theory, *Metrika*, 33, 337–347, 1986.
- Broudiscou, A., R. Leardi, and R. Phan-Tan-Luu, Genetic algorithm as a tool for selection of D -optimal design, *Chemometrics and Intelligent Laboratory Systems*, 35, 105–116, 1996.
- Brown, L. D. and W. K. Wong, An algorithmic construction of optimal minimax designs for heteroscedastic linear models, *Journal of Statistical Planning and Inference*, 85, 103–114, 2000.
- Chaloner, K. and K. Larntz, Optimal Bayesian design applied to logistic regression experiments, *Journal of Statistical Planning and Inference*, 21, 191–208, 1989.
- Chaloner, K. and I. Verdinelli, Bayesian experimental design: A review, *Statistical Science*, 10, 273–304, 1995.
- Chang, S., An algorithm to generate near D -optimal designs for multiple response surface models, *IIE Transactions*, 29, 1073–1081, 1997.
- Chen, R. B., S. P. Chang, W. Wang, H. C. Tung, and W. K. Wong, Minimax optimal designs via particle swarm optimization methods, *Statistics and Computing*, 2014. doi:10.1007/s11222-014-9466-0.
- Chen, R. B., W. K. Wong, and K. Y. Li, Optimal minimax designs for estimating response surface over a prespecified region in a heteroscedastic model, *Statistics and Probability Letters*, 78, 1914–1921, 2008.
- Cook, R. D. and C. J. Nachtsheim, A comparison of algorithms for constructing exact D -optimal designs, *Technometrics*, 22(3), 315–324, 1980.
- Cook, R. D. and C. J. Nachtsheim, Model robust, linear-optimal designs, *Technometrics*, 24(1), 49–54, 1982.
- Cook, R. D. and W. K. Wong, On the equivalence of constrained and compound optimal designs, *Journal of American Statistical Association*, 89, 687–692, 1994.
- Crombecq, K. and T. Dhaene, Generating sequential space-filling designs using genetic algorithms and Monte Carlo methods, *Lecture Notes in Computer Science*, Vol. 6457. Springer-Verlag: Berlin, Germany, pp. 80–88, 2010.

- Dean, A. M. and S. M. Lewis (eds.) Screening: Methods for Experimentation in Industry, Drug Discovery and Genetics, Springer-Verlag: New York, 2006.
- Dette, H., A. Pepelyshev, and A. Zhigljavsky, Improving updating rules in multiplicative algorithms for computing D -optimal designs, *Computational Statistics & Data Analysis*, 53, 312–320, 2008.
- Dette, H. and W. J. Studden, *The Theory of Canonical Moments with Applications in Statistics, Probability, and Analysis*, Wiley: New York, 1997.
- Duarte, B. P. M. and W. K. Wong, Finding Bayesian optimal designs for nonlinear models: A semidefinite programming based approach, *International Statistical Review*. 2014a. doi:10.1111/insr.12073.
- Duarte, B. P. M. and W. K. Wong, A semi-infinite programming based algorithm for finding minimax D -optimal designs for nonlinear models, *Statistics and Computing*, 24, 1063–1080, 2014b.
- Eberhart, R. C. and Y. Shi, Comparing inertia weights and constriction factors in particle swarm optimization, *Proceedings of the 2000 Congress on Evolutionary Computation*, 1, 84–88, 2000.
- Ehrenfeld, S., On the efficiencies of experimental designs, *Annals of Mathematical Statistics*, 26, 247–255, 1955.
- Fang, Z. and D. P. Wiens, Integer-valued, minimax robust designs for estimation and extrapolation in heteroscedastic, approximately linear models, *Journal of American Statistical Association*, 95, 807–818, 2000.
- Fedorov, V. V., Theory of optimal experiments, Preprint 7 LSM, Izd-vo Moscow State University, Moscow, Russia, 1969.
- Fedorov, V. V., *Theory of Optimal Experiments*, translated by W. J. Studden and E. M. Klimko, Academic Press: New York, 1972.
- Fellman, J., An empirical study of a class of iterative searches for optimal designs, *Journal of Statistical Planning and Inference*, 21, 85–92, 1989.
- Filová, L., M. Trnovská, and R. Harman, Computing maximin efficient experimental designs using the methods of semidefinite programming, *Metrika*, 75, 709–719, 2012.
- Goldfarb, H. B., C. M. Borror, D. C. Montgomery, and C. M. Anderson-Cook, Using genetic algorithms to generate mixture-process experimental designs involving control and noise variables, *Journal of Quality Technology*, 37, 60–74, 2005.
- Haines, L. M., The application of the annealing algorithm to construction of exact optimal designs for linear regression models, *Technometrics*, 29, 439–447, 1987.
- Hamada, M., H. D. Martz, C. S. Reese, and A. G. Wilson, Finding near-optimal Bayesian designs via genetic algorithms, *The American Statistician*, 55, 175–181, 2001.
- Harman, R., Multiplicative methods for computing D -optimal stratified designs of experiments, *Journal of Statistical Planning and Inference*, 146, 82–94, 2014.
- Harman, R. and L. Pronzato, Improvements on removing nonoptimal support points in D -optimum design algorithms, *Statistics and Probability Letters*, 77, 90–94, 2007.
- Harrington, E. C. Jr., The desirability function, *Industrial Quality Control*, 21, 494–498, 1965.
- Heredia-Langner, A., W. M. Carlyle, D. C. Montgomery, C. M. Borror, and G. C. Runger, Genetic algorithms for the construction of D -optimal designs, *Journal of Quality Technology*, 35, 28–46, 2003.
- Heredia-Langner, A., D. C. Montgomery, W. M. Carlyle, and C. M. Borror, Model-robust optimal designs: A genetic algorithm approach, *Journal of Quality Technology*, 36 (3), 263–279, 2004.
- Holland, J. M. *Adaptation in Natural and Artificial Systems*, University of Michigan Press: Ann Arbor, MI, 1975.
- Johnson, M. E. and C. J. Nachtsheim, Some guidelines for constructing exact D -optimal designs on convex design spaces, *Technometrics*, 25(3), 271–277, 1983.
- Kao, M. H., A. Mandal, N. Lazar, and J. Stufken, Multi-objective optimal experimental designs for event-related fMRI studies, *NeuroImage*, 44, 849–856, 2009.
- Kennedy, J. and R. Eberhart, Particle swarm optimization, *Proceedings of IEEE International Conference on Neural Networks, IV*, Perth, Australia, pp. 1942–1948, 1995.
- Kiefer, J., General equivalence theory for optimum designs (approximate theory), *Annals of Statistics*, 2, 849–879, 1974.

- King, J. and W. K. Wong, Minimax D -optimal designs for the logistic model, *Biometrics*, 56, 1263–1267, 2000.
- Kirkpatrick, S., C. D. Gelatt, and M. P. Vecchi, Optimization by simulated annealing, *Science*, New series, 220, 4598, 671–680, 1983.
- Levy, A. V. and A. Montalvo, The tunnelling algorithm for the global minimization of functions, *SIAM Journal of Scientific and Statistical Computing*, 6, 15–29, 1985.
- Liefvendahl, M. and A. Stocki, A study on algorithms for optimization of Latin hypercubes, *Journal of Statistical Planning and Inference*, 136, 3231–3247, 2006.
- Lin, C. D., C. M. Anderson-Cook, M. S. Hamada, L. M. Moore, and R. R. Sitter, Using genetic algorithms to design experiments: A review, *Quality and Reliability Engineering International*, 31, 155–167, 2015.
- Mandal, A., Johnson, K., Wu, C. F. J., and Bornmeier, D., Identifying promising compounds in drug discovery: Genetic algorithms and some new statistical techniques, *Journal of Chemical Information and Modeling*, 47, 981–988, 2007.
- Mandal, A., P. Ranjan, and C. F. J. Wu, \mathcal{G} -SELC: Optimization by sequential elimination of level combinations using genetic algorithms and Gaussian processes, *Annals of Applied Statistics*, 3, 398–421, 2009.
- Mandal, A., C. F. J. Wu, and K. Johnson, SELC: Sequential elimination of level combinations by means of modified genetic algorithms, *Technometrics*, 48, 273–283, 2006.
- Mandal, S. and B. Torsney, Construction of optimal designs using a clustering approach, *Journal of Statistical Planning and Inference*, 136, 1120–1134, 2006.
- Marseguerra, M., E. Zio, and M. Cipollone, Designing optimal degradation tests via multi-objective genetic algorithms, *Reliability Engineering and System Safety*, 79, 87–94, 2003.
- Metropolis, N., A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, Equation of state calculations by fast computing machines, *Journal of Chemical Physics*, 21, 1087–1092, 1953.
- Meyer, R. K. and C. J. Nachtsheim, Constructing exact D -optimal experimental designs by simulated annealing, *American Journal of Mathematical and Management Sciences*, 8, 329–359, 1988.
- Meyer, R. K. and C. J. Nachtsheim, The coordinate-exchange algorithm for constructing exact optimal experimental designs, *Technometrics*, 37, 60–69, 1995.
- Mitchell, T. J., An algorithm for the construction of D -optimal experimental designs, *Technometrics*, 20, 203–210, 1974.
- Nguyen, N. K. and A. Miller, A review of exchange algorithms for constructing discrete D -optimal designs, *Computational Statistics and Data Analysis*, 14, 489–498, 1992.
- Papp, D., Optimal designs for rational function regression, *Journal of the American Statistical Association*, 107, 400–411, 2012.
- Pázman, A., Foundations of Optimum Experimental Design, Reidel: Dordrecht, the Netherlands, 1986.
- Press, W. H., S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed., Cambridge University Press, Cambridge, U.K., 1992.
- Pukelsheim, F. and B. Torsney, Optimal weights for experimental designs on linearly independent support points, *Annals of Statistics*, 19, 1614–1625, 1991.
- Qiu, J., R. B. Chen, W. Wang, and W. K. Wong, Using animal instincts to design efficient biomedical studies via particle swarm optimization, *Swarm and Evolutionary Computation*, 18, 1–10, 2014.
- Ranjan, P., D. Bingham, and G. Michailidis, Sequential experiment design for contour estimation from complex computer codes, *Technometrics*, 50 (4), 527–541, 2008.
- Reemtsen, R. and J. J. Rückman, *Semi-Infinite Programming*, Kluwer Academic Publishers: Dordrecht, the Netherlands, 1998.
- Reeves, C. L. and C. C. Wright, Genetic algorithms and the design of experiments, In Davis, L. D.; DeJong, K.; Vose, M. D., and Whitley, L. D. (eds.), *Evolutionary Algorithms: IMA Volumes in Mathematics and its Applications*, Vol 111. Springer-Verlag: New York, pp. 207–226, 1999.
- Rodriguez, M., B. Jones, C. M. Borrer, and D. C. Montgomery, Generating and assessing exact G -optimal designs, *Journal of Quality Technology*, 42, 3–20, 2010.

- Rodriguez, M., D. C. Montgomery, and C. M. Borrór, Generating experimental designs involving control and noise variables using genetic algorithms, *Quality and Reliability Engineering International*, 25, 1045–1065, 2009.
- Sexton, C. J., D. K. Anthony, S. M. Lewis, C. P. Please, and A. J. Keane, Design of experiment algorithms for assembled products, *Journal of Quality Technology*, 38, 298–308, 2006.
- Silvey, S. D., *Optimal Design*, Chapman & Hall: London, U.K., 1980.
- Silvey, S. D., D. M. Titterington, and B. Torsney, An algorithm for optimal designs on a finite design space, *Communications in Statistics–Theory and Methods*, 14, 1379–1389, 1978.
- Smith, K., On the standard deviations of adjusted and interpolated values of an observed polynomial function and its constants and the guidance they give towards a proper choice of the distribution of observations, *Biometrika*, 12, 1–85, 1918.
- Syed, M. N., I. Kotsireas, and P. M. Pardalos, D -optimal designs: A mathematical programming approach using cyclotomic cosets, *Informatica*, 22, 577–587, 2011.
- Titterington, D. M., Algorithms for computing D -optimal design on finite design spaces, in *Proceedings of the 1976 Conference on Information Science and Systems*, John Hopkins University: Baltimore, MD, Vol. 3, pp. 213–216, 1976.
- Titterington, D. M., Estimation of correlation coefficients by ellipsoidal trimming, *Applied Statistics*, 27, 227–234, 1978.
- Torsney, B. and R. Martin-Martin, Multiplicative algorithms for computing optimum designs, *Journal of Statistical Planning and Inference*, 139, 3947–3961, 2009.
- Vandenberghe, L. and S. Boyd, Semidefinite programming, *SIAM Review*, 38, 49–95, 1996.
- Welch, W. J., Branch and bound search for experimental designs based on D -optimality and other criteria, *Technometrics*, 24(1), 41–48, 1982.
- Wilmot, M. and J. Zhou, D -optimal minimax design criterion for two-level fractional factorial designs, *Journal of Statistical Planning and Inference*, 141, 576–587, 2011.
- Wong, W. K., A unified approach to the construction of mini-max Designs, *Biometrika*, 79, 611–620, 1992.
- Wong, W. K. and R. D. Cook, Heteroscedastic G -optimal designs, *Journal of Royal Statistical Society, Series B*, 55, 871–880, 1993.
- Wong, W. K., Recent advances in constrained optimal design strategies, *Statistica Neerlandica*, (Invited paper) 53, 257–276, 1999.
- Woods, D. C., Robust designs for binary data: Applications of simulated annealing, *Journal of Statistical Computation and Simulation*, 80, 29–41, 2010.
- Wynn, H. P., Results in the theory and construction of D -optimum experimental designs, *Journal of Royal Statistical Society, Series B*, 34, 133–147, 1972.
- Yang, X. S., *Nature-Inspired Metaheuristic Algorithms*, Luniver Press: Frome, U.K., 2008.
- Yang, X. S., *Nature-Inspired Metaheuristic Algorithms*, 2nd ed., Luniver Press, U.K., 2010.
- Yang, M., S. Biedermann, and E. Tang, On optimal designs for nonlinear models: A general and efficient algorithm, *Journal of the American Statistical Association*, 108, 1411–1420, 2013.
- Yang, J., A. Mandal, and D. Majumdar, Optimal designs for 2^k factorial experiments with binary response, *Statistica Sinica*, 2015. doi:10.5705/ss.2013.265.
- Yu, Y., Monotonic convergence of a general algorithm for computing optimal designs, *Annals of Statistics*, 38, 1593–1606, 2010.
- Yu, Y., D -optimal designs via a cocktail algorithm, *Statistics and Computing*, 21, 475–481, 2011.
- Zhou, J., D -optimal minimax regression designs on discrete design space, *Journal of Statistical Planning and Inference*, 138, 4081–4092, 2008.

Section VII

Design for Contemporary Applications

22

Design for Discrete Choice Experiments

Heiko Grossmann and Rainer Schwabe

CONTENTS

22.1	Introduction.....	787
22.2	Choice Models.....	789
22.2.1	Multinomial Logit Model.....	789
22.2.2	Mixed Logit Model.....	791
22.2.3	Other Extensions.....	792
22.3	Paired Comparisons.....	793
22.3.1	Main Effects Only.....	794
22.3.1.1	Optimal Approximate Designs.....	795
22.3.1.2	Design Constructions for Main Effects Only.....	798
22.3.2	Main Effects and Two-Factor Interactions.....	803
22.3.2.1	Optimal Approximate Designs.....	803
22.3.2.2	Constructions for Main Effects and Two-Factor Interactions.....	804
22.3.3	Partial Profiles.....	807
22.4	Designs for Larger Choice Sets.....	809
22.4.1	Efficiency Gains due to Increased Size of Choice Sets.....	810
22.4.2	Relationship with Block Designs.....	811
22.4.3	Some Local Optimality Results.....	816
22.5	Computational Approaches.....	818
22.5.1	Multinomial Logit Model.....	818
22.5.2	Mixed Logit Model.....	821
22.6	Ranking-Based Methods and Best–Worst Scaling.....	822
22.7	Miscellaneous Problems.....	824
22.8	Concluding Remarks.....	825
	References.....	826

22.1 Introduction

In modern life, people frequently face situations where they have to make choices between a variety of options. Purchasing a smartphone, picking an energy supplier, or deciding upon a cancer treatment are just a few examples where a number of features of the available alternatives have to be considered when making a decision. How people make these decisions is a fascinating topic and contributions to this area have been made from many disciplines including psychology, marketing, econometrics, and statistics. As a consequence, the study of choice behavior is a truly multidisciplinary endeavor.

Choice experiments mimic the real world by using a type of questionnaire or survey that asks respondents to complete a series of choice tasks. Each of these tasks presents a choice set of competing alternatives from which, typically, the most preferred alternative has to be selected. Modes of presentation range from simple paper-and-pencil administration to computerized systems that allow the inclusion of pictures or videos to describe the alternatives. A key feature of these experiments is that the competing alternatives are defined by combinations of the values or levels of a certain number of attributes, which represent the main characteristics of interest. In statistical models, these attributes play a role similar to explanatory variables in a regression model or factors in a multiway analysis of variance. A common goal of the analysis of choice data is then to estimate parameters that reflect the utility of the attribute levels. These estimates may feed further into simulations of market shares or may be used to identify groups of individuals with similar preferences. Generating a design for a choice experiment amounts to combining attribute levels into alternatives and arranging these into choice sets. This chapter mainly focuses on designs that are optimal or efficient under the D -optimality criterion (see Chapter 2).

The origins of the field can be traced back to psychophysical scaling experiments in the late nineteenth century, which used the method of paired comparisons to derive, for example, a functional law of how the perceived weight of objects depends on their actual mass (David 1988). That choices between pairs could be used more widely to measure social values and attitudes was pointed out by Thurstone (1927, 1928). Starting from a simple axiom, Luce (1959) derived a now famous model for choices between two or more alternatives, which for the special case of paired comparisons coincides with the model of Bradley and Terry (1952). Both models contain a separate parameter for each alternative and express the probability of choosing a particular alternative from a choice set as the ratio of the parameter for the chosen alternative over the sum of the parameters corresponding to all the alternatives in the choice set (see Section 22.2.1).

The most popular model for choice data is the econometric version of Luce's choice model known as the multinomial logit or, in short, MNL model. This is obtained by reparameterizing Luce's model so that the original model parameters are replaced with their logarithms, which then enter the MNL model equation for the choice probabilities as exponentials. An appealing feature of the MNL model is that it can be derived from random utility theory (Louviere et al. 2000; Train 2003), that is, from latent random variables representing the utility of the alternatives. Also, the likelihood function can be expressed in closed form, unlike that of some more advanced choice models. A sometimes undesirable feature of the MNL model is, however, the *independence-from-irrelevant-alternatives* (IIA) property, which means that for any two alternatives in a choice set, the ratio of the choice probabilities does not depend on the other alternatives in the set. Notwithstanding, the multinomial logit model continues to play a pivotal role in choice experiments, for many more advanced choice models can be regarded as extensions of the MNL model. Excellent introductions to the general area of choice experiments are available in the monographs by Louviere et al. (2000) and Train (2003).

When discussing the design of choice experiments, some related techniques for measuring preferences that are jointly known under the name of *conjoint analysis* (Green and Srinivasan 1978, 1990) are also worth mentioning. These are similar to choice experiments in that they use experimental designs to compose a questionnaire of alternatives for collecting preference judgments from which utility values for the attribute levels are estimated. The main difference from choice experiments is that the alternatives are typically assessed on a rating scale and that linear models and least squares estimation are used.

Until the mid-1990s, conjoint analysis dominated in marketing applications, but since then, improvements in computational power and software have led to a more widespread use of choice models. The potential usefulness of fractional factorial and other classical designs (see Chapter 7) for applications of conjoint analysis was recognized early on (Green 1974; Green et al. 1976, 1978) and some ideas from this area have also influenced the design of choice experiments.

In a seminal paper, Louviere and Woodworth (1983) discussed the role of designs for choice experiments from a conceptual viewpoint, and an introduction to optimum design theory for a marketing audience can be found in Kuhfeld et al. (1994). Huber and Zwerina (1996) presented a heuristic approach for generating locally D -optimal designs for the MNL model, which stimulated much subsequent research. Summaries of the early work on designs for conjoint and choice experiments have been provided by Großmann et al. (2002) and Louviere et al. (2004). More recent developments are summarized by Rose and Bliemer (2009) and Street and Burgess (2012).

In what follows, we mainly concentrate on so-called generic choice experiments in which the attributes that characterize the choice alternatives are qualitative factors with a finite number of levels. Apart from being popular in marketing applications, this type of choice experiment appears to be typically used in health economics (De Bekker-Grob et al. 2012) and environmental valuation studies (Alriksson and Öberg 2008; Hoyos 2010). Applications in the field of transportation research frequently use other types of choice experiments where, for example, some of the attributes are quantitative continuous variables and different sets of attributes may be used to describe the competing alternatives. We do not cover such alternative-specific and related designs here.

The remainder of the chapter is organized as follows. The next section provides a brief description of the multinomial logit model and some extensions. This is followed by a review of designs for paired comparisons and choice experiments with larger choice sets, with an emphasis on analytic results. After this, computational approaches are considered. We then move on to ranking-based approaches and some more specialized problems before offering some concluding remarks.

22.2 Choice Models

The complexity of choice models has steadily increased over the years, although many of the more sophisticated models can be recognized as generalizations of the MNL model. By contrast, the work on experimental designs for choice models has advanced at a slower pace with most results being available for the MNL model. For this reason, we also focus on the MNL model. Some extensions will be mentioned at the end of this section, and corresponding design results are considered in Sections 22.5 and 22.6.

22.2.1 Multinomial Logit Model

The MNL model describes the probability with which a particular alternative is chosen from a finite set of alternatives. It is also referred to as the conditional logit model. Newcomers to the area will soon realize that in the literature, the terms *alternative*, *option*, *object*, and *profile* are used interchangeably. Likewise, the participants in a choice experiment are

known under various names, such as *respondents*, *individuals*, *persons*, or *subjects*. In what follows, we mainly use the terms *alternatives* for choice options and *respondents* for people making choices, respectively.

Suppose \mathcal{C} is the finite set of all alternatives that are of interest. The features of every alternative in \mathcal{C} are described by a vector of *attributes*, which may be quantitative or qualitative. A major goal of a choice experiment is then to find out to what extent the different attributes influence the preferences for the alternatives. Therefore, when modeling the preferences, every alternative is identified with its corresponding attribute vector.

The designer of a *choice experiment* determines the common size m and the total number N of the *choice sets* for the experiment. Every choice set C_n then contains m alternatives in \mathcal{C} . As already explained, each such alternative is identified with a vector $x_{n,i}$, where $n = 1, \dots, N$ is the number of the choice set and $i = 1, \dots, m$ the number of the alternative within the choice set. Exactly as in a linear model, a vector f of *regression functions* is used to code qualitative attributes or to represent, for example, polynomial effects of quantitative attributes. The coding of qualitative attributes will be considered in somewhat more detail later in the chapter. Associated with f is a p -dimensional *parameter vector* θ representing the effects of the attributes and their levels. Thus, while the length of $x_{n,i}$ is equal to the number of attributes, the vector $f(x_{n,i})$ has p components.

In the MNL model, the probability of choosing the alternative $x_{n,i}$ from the choice set $C_n = \{x_{n,1}, \dots, x_{n,m}\}$ is assumed to be equal to

$$p_{\theta}(i; C_n) = \frac{e^{f'(x_{n,i})\theta}}{\sum_{j=1}^m e^{f'(x_{n,j})\theta}}. \quad (22.1)$$

During the experiment, usually the same choice sets are given to a sample of S individuals. Let \mathbf{y}_s be a vector representing the *responses* of respondent s with components $y_{s,n,i}$ being equal to either one or zero depending on whether or not respondent s has chosen alternative $x_{n,i}$ from the set C_n . Furthermore, let \mathbf{y} be the concatenation of these vectors for the whole sample. Assuming that the choices are independent within and across respondents, the MNL *likelihood function* $L(\theta; \mathbf{y})$ for estimating the model parameters is given by

$$L(\theta; \mathbf{y}) = \prod_{s=1}^S \prod_{n=1}^N \prod_{i=1}^m p_{\theta}(i; C_n)^{y_{s,n,i}}. \quad (22.2)$$

For a *design* d with choice sets C_1, \dots, C_N , the *Fisher information* is obtained in the usual way as $M_{d,\theta} = E(\mathbf{h}_{\theta}(\mathbf{y})\mathbf{h}'_{\theta}(\mathbf{y}))$, where $\mathbf{h}_{\theta}(\mathbf{y}) = \partial \log L(\theta; \mathbf{y}) / \partial \theta$ is the column vector of partial derivatives of the log-likelihood function. Because all respondents receive the same design, for the purpose of comparing different designs on the basis of $M_{d,\theta}$, it is sufficient to consider a single respondent, that is, $S = 1$. Similar to Huber and Zwerina (1996), a convenient representation of the Fisher information matrix of a design d is then

$$M_{d,\theta} = \sum_{n=1}^N \mathbf{X}'_n (\text{diag}(\mathbf{p}_{\theta,n}) - \mathbf{p}_{\theta,n} \mathbf{p}'_{\theta,n}) \mathbf{X}_n, \quad (22.3)$$

where \mathbf{X}_n is an $m \times p$ matrix with rows $f'(x_{n,i})$, $i = 1, \dots, m$, representing the alternatives in C_n and $\mathbf{p}_{\theta,n} = (p_{\theta}(1; C_n), \dots, p_{\theta}(m; C_n))'$ is the vector of the corresponding *choice probabilities*.

Sometimes, we will consider designs that contain a single choice set C , in which case the information matrix in (22.3) will be denoted by $M_{C,\theta}$. By applying this notational convention to the individual choice sets C_1, \dots, C_N of a design d , Equation 22.3 can be rewritten as $M_{d,\theta} = \sum_{n=1}^N M_{C_n,\theta}$.

The MNL model is nonlinear in the unknown model parameters in θ and hence the Fisher information matrix also depends on θ . For the purpose of deriving optimal designs, it is often assumed that $\theta = \mathbf{0}$. This *indifference assumption* means that all alternatives in a choice set C_n are chosen with the same probability $p_0(i; C_n) = 1/m$ and leads to a considerable simplification of the matrix in (22.3) and the design problem. The corresponding information matrix $M_{d,0}$ for $\theta = \mathbf{0}$ is given by

$$M_{d,0} = \frac{1}{m} \sum_{n=1}^N X_n' K_m X_n, \tag{22.4}$$

where $K_m = I_m - \frac{1}{m} J_m$ is a centering matrix. Furthermore, I_m denotes the identity matrix of order m and J_m the $m \times m$ matrix with all elements equal to one.

The important special case where the choice sets are pairs, that is, where $m = 2$, can be recognized as a reparameterization of the *Bradley–Terry model* and also of Luce’s original choice model. More precisely, for a pair of alternatives $x_{n,i}$, $i = 1, 2$, setting $\pi_{n,i} = \exp(f'(x_{n,i})\theta)$ in (22.1) gives the choice probability $p_\theta(i; C_n) = \pi_{n,i} / (\pi_{n,1} + \pi_{n,2})$ where the right-hand side is equivalent to the expressions in Bradley and Terry (1952) and Luce (1959).

22.2.2 Mixed Logit Model

The choices described by the MNL model are based on homogeneous preferences, which is reflected in (22.1) by the dependence of the probabilities on a fixed parameter vector θ . A flexible model that allows for a certain type of preference heterogeneity is the *mixed logit model* (e.g., see Train 2003, Chapter 6) that, like the MNL model, can be motivated by considering models of utility maximizing behavior. In the mixed logit model, the parameters determining the preferences of each individual are assumed to come from a multivariate distribution. Usually, this distribution is assumed to be continuous with a probability density function g and can be thought of as describing the variability of tastes in a population. Often, the parameters are assumed to be normally distributed, but other distributions, such as uniform, triangular, or Rayleigh distributions, have also been used (Train 2003).

The distinction between the situation where the multiple choices made by a respondent are assumed to be independent and the more general situation where a respondent’s choices are correlated leads to two variants of the model that are known as the *cross-sectional* and the *panel mixed logit model*, respectively (Bliemer and Rose 2010). In the cross-sectional formulation, a respondent’s probability of choosing alternative i from choice set C_n is equal to $q(i; C_n) = \int p_\theta(i; C_n) g(\theta) d\theta$, where $p_\theta(i; C_n)$ is the MNL probability (22.1). The likelihood function is then obtained by substituting $q(i; C_n)$ for $p_\theta(i; C_n)$ in (22.2).

The corresponding equations for the panel mixed logit model are more complicated due to the lack of independence of the choices within respondents. The parameters in θ are assumed to be constant for every individual respondent but vary over people. As before, let y_s be the vector of choices made by respondent s from N choice sets C_1, \dots, C_N , where

the choice sets are the same for all respondents. Then the probability $q(\mathbf{y}_s; C_1, \dots, C_N)$ of this particular sequence of responses is the average, as weighted by g , of the choice probabilities in (22.1) over the space of the parameter vectors, that is,

$$q(\mathbf{y}_s; C_1, \dots, C_N) = \int \prod_{n=1}^N \prod_{i=1}^m p_{\theta}(i; C_n)^{y_{s,n,i}} g(\theta) d\theta. \quad (22.5)$$

The likelihood for a sample of S respondents is the product of the expressions in Equation 22.5.

In general, the density function g depends on a set of population or hyperparameters, and finding efficient designs for estimating those parameters is the goal on which most design work for the mixed logit model has concentrated. Usually, the integral in (22.5) cannot be evaluated in closed form. Therefore, when estimating the population parameters, simulation methods are commonly used to approximate the integral and the likelihood function. In principle, the distribution of the θ parameters can also be discrete with g being the probability mass function of a distribution with finite support. The integral in (22.5) then becomes a sum and the model becomes a finite mixture or latent class model (Train 2003).

22.2.3 Other Extensions

Another generalization of the MNL model, which relaxes the independence-from-irrelevant-alternatives assumption, is the *nested logit model* (Train 2003, Chapter 4). In this model, all possible alternatives are grouped into a number of disjoint sets that are referred to as nests and that are such that the IIA property holds for alternatives within each nest, but not across different nests. As for the MNL model, the nested logit choice probabilities can be expressed in closed form, and the model parameters are usually estimated by standard maximum likelihood procedures. Moreover, the probability of choosing any given alternative can be factorized into the probability of choosing the nest that contains the alternative times the conditional probability of selecting the alternative from that nest, where both the probabilities for choices between nests and those for alternatives within nests can be described by separate MNL models.

A common feature of all models considered so far is that the choice task requires the respondents to select a single alternative from every choice set. More information about the underlying preferences may be obtained from complete or partial rankings of the alternatives within each set. For example, a *complete ranking* of the m alternatives in a choice set, may be regarded as the outcome of an iterated choice process where initially the most preferred alternative is selected. After removing this alternative from the choice set the next best alternative is chosen from the remaining alternatives and so forth. This process terminates after $m - 1$ choices have been made. At each stage, the probability of choosing the best from the remaining alternatives can be expressed as in the MNL model and the corresponding model is known as the *rank-ordered logit* (Lancsar et al. 2013) or *exploded logit model* (Chapman and Staelin 1982; Train 2003).

The ranking task is relatively easy when the size m of the choice sets is small but becomes more demanding for larger numbers of alternatives. In general, respondents appear to find it less challenging to place alternatives at the top and bottom of an ordered list than to make reliable judgments about alternatives that are somewhere in the middle of the preference

spectrum. *Best–worst scaling* (Marley and Louviere 2005; Marley and Pihlens 2012) tries to exploit this finding by asking respondents to identify the most and the least preferred alternative within each choice set.

The list of generalizations of, and alternatives to, the MNL model could easily be expanded. Here, we have chosen to mention only those approaches where at least some work on the corresponding experimental designs has been done.

22.3 Paired Comparisons

Experiments with choice sets of size two are closely related to the method of *paired comparisons* (Bradley and Terry 1952; Bradley 1976; Davidson and Farquhar 1976; David 1988), and such *paired comparison experiments* have attracted considerable attention in their own right. Initial contributions to their design were made by Bradley, El-Helbawy, and coworkers (Bradley and El-Helbawy 1976; El-Helbawy and Bradley 1978; El-Helbawy 1984; El-Helbawy and Ahmed 1984; El-Helbawy et al. 1994). This work adopts a modeling approach that is different from the mainstream formulation in Section 22.2.1, although for the purpose of finding optimal or efficient designs, the two approaches are equivalent. Also, the presentation in El-Helbawy and Bradley (1978) has strongly influenced more recent work by Street, Burgess and coworkers (Street et al. 2001; Street and Burgess 2004a, 2007; Burgess and Street 2005).

The results in the aforementioned papers are obtained under the assumption of equal choice probabilities which is equivalent to considering the MNL model (22.1) with $\theta = \mathbf{0}$. In this particular case, the information matrix $M_{d,0}$ in (22.4) for choice sets of size $m = 2$ is proportional to the information matrix in a *linear paired comparison model*

$$y = X\theta + \epsilon \quad (22.6)$$

with standard assumptions of uncorrelated and homoscedastic random errors ϵ , where for a design d with N pairs $(x_{n,1}, x_{n,2})$, the design matrix X has rows $f'(x_{n,1}) - f'(x_{n,2})$ for $n = 1, \dots, N$ and f is defined above (22.1). More precisely, as shown in Großmann et al. (2002), $M_{d,0} = \frac{1}{4}M_d$, where $M_d = X'X = \sum_{n=1}^N (f(x_{n,1}) - f(x_{n,2}))(f'(x_{n,1}) - f'(x_{n,2}))$. Every row of X corresponds to one pair. Moreover, the pairs are ordered so that $(x_{n,1}, x_{n,2})$ is distinguished from $(x_{n,2}, x_{n,1})$. Although both pairs contribute the same information $(f(x_{n,1}) - f(x_{n,2}))(f'(x_{n,1}) - f'(x_{n,2})) = (f(x_{n,2}) - f(x_{n,1}))(f'(x_{n,2}) - f'(x_{n,1}))$ to M_d , distinguishing the order of the alternatives within the pairs adds some symmetry to the design problem that can be exploited for finding optimal designs.

With regard to the usual alphabetic optimality criteria (see Chapter 2), it follows that, under the assumption $\theta = \mathbf{0}$, the efficiency of a design in the MNL model (22.1) for $m = 2$ is the same as the efficiency of the design in the linear paired comparison model. In particular, optimal designs for the MNL model can be derived by considering the linear model, which is the approach taken by Graßhoff et al. (2003, 2004).

The simplification of the information matrix that is achieved by assuming $\theta = \mathbf{0}$ makes it possible to prove analytical optimality results. Also, explicit constructions of optimal or highly efficient designs with practical numbers of choice sets can be derived. Without that assumption, because of the nonlinearity of the MNL model and the implied parameter

dependence of the information matrix, algorithms usually need to be used to generate good designs. We consider algorithmic approaches in Section 22.5.

The remainder of this section considers designs for experiments in which the choice sets are pairs. We return to experiments with larger choice sets in Section 22.4. In what follows, we concentrate on optimality results and design construction methods under the indifference assumption $\theta = 0$. First, we look at experiments where only the main effects of the attributes are to be estimated. Secondly, we consider designs for estimating the main effects and all two-factor interactions. Finally, designs for paired comparison experiments involving incomplete descriptions of the alternatives, known as partial profiles, are discussed.

22.3.1 Main Effects Only

In the simplest type of paired comparison or choice experiment, every attribute is a qualitative *factor*, usually with a small number of levels, and the overall utility of every alternative is equal to the sum of its parts, that is, the sum of so-called *part-worth utilities* of the attribute levels. This setup leads to an expression for $f'(x_{n,i})\theta$ in (22.1) that is similar to the systematic part of a main-effects-only multifactor analysis of variance model. For every alternative $x_{n,i}$, the components of $f(x_{n,i})$ are the suitably coded attribute levels that define $x_{n,i}$ and the parameters in θ are interpreted as the part-worth utilities of the levels. It should be noted however that despite the similarity with the analysis of variance, there is no model parameter corresponding to a grand or overall mean. If such a parameter were to be included, the choice probabilities in expression (22.1) would not change, which shows that the parameter is redundant. This invariance of the choice probabilities is often described by the statement that *only differences in utility matter* (Train 2003, p. 23).

In order to represent the levels of a qualitative attribute, usually, effects coding is recommended (e.g., see Bech and Gyrd-Hansen 2005) and we adopt this coding throughout the chapter. For an attribute with v levels, this coding scheme represents every level by a vector of length $v - 1$. The vector corresponding to level j with $1 \leq j \leq v - 1$ is a unit vector of length $v - 1$, which has a one in position j and contains zeros otherwise. The final level v is represented by the vector $-\mathbf{1}_{v-1}$ of the same length with all elements equal to -1 . For example, for $v = 2$, the level 1 is coded as 1 and the level 2 by -1 where it should be noted that the coded vectors of length $v - 1 = 1$ are just scalars. Similarly, the levels 1, 2, and 3 of an attribute with $v = 3$ levels are represented by the vectors $(1, 0)'$, $(0, 1)'$, and $(-1, -1)'$, respectively. The portion of the parameter vector θ for an attribute with v levels contains $v - 1$ parameters, which correspond to the levels $1, \dots, v - 1$. When effects coding is used, the parameter for level $1 \leq j \leq v - 1$ of the attribute can be interpreted as the utility of j minus the average of the utility values of all v levels of that attribute. Note that the parameter vector does not contain a parameter for the level v of the attribute, but that the utility of the level v minus the average of the utility values for all v levels of the attribute can be recovered by multiplying the sum of the parameters for the levels $1, \dots, v - 1$ with -1 .

For a choice experiment with K attributes in which the k th attribute has v_k levels, $k = 1, \dots, K$, every alternative $x = (x_1, \dots, x_K)$ is a K -tuple whose k th component x_k is one of the levels $1, \dots, v_k$. The effects-coded vector $f(x)$ representing the levels of x in the main effects model is obtained by first applying the coding separately to every level x_k and by secondly stacking the K column vectors corresponding to the coded attribute levels. Note that, equivalently, the vector $f(x)$ can be defined as $f(x) = (f'_1(x_1), \dots, f'_K(x_K))'$ where for

$k = 1, \dots, K$ the vector f_k contains the $v_k - 1$ individual regression functions $f_{k,j} = 1_{\{j\}} - 1_{\{v_k\}}$, $j = 1, \dots, v_k - 1$. Here, for every a , the function $1_{\{a\}}$ is an indicator function with $1_{\{a\}}(z) = 1$ if $z = a$ and $1_{\{a\}}(z) = 0$ if $z \neq a$.

Most design criteria are sensitive to the coding of the attribute levels. A remarkable exception is the D -optimality criterion (e.g., see Chapters 2 and 3), which tries to find designs that maximize the determinant of the information matrix or, equivalently, minimize the determinant of its inverse, because the criterion is not affected by reparameterizations. This invariance property and the fact that the determinant is mathematically more tractable than other functions of the information matrix appear to be the main reasons why most optimality results for choice and paired comparison designs have been derived for the D -criterion.

22.3.1.1 Optimal Approximate Designs

Approximate or continuous designs ξ are defined as discrete probability measures on the design region \mathcal{X} of an experiment and can be represented by t distinct design points x_1, \dots, x_t in \mathcal{X} together with corresponding positive weights w_1, \dots, w_t where $w_i = \xi(x_i)$ for $i = 1, \dots, t$ and $\sum_{i=1}^t w_i = 1$ (see Chapter 2). In the context of choice experiments with choice sets of size $m = 2$, every design point is a pair of alternatives. One approach to proving optimality, adopted by Graßhoff et al. (2004), relies on the *equivalence theorem* of Kiefer and Wolfowitz (1960) for approximate designs. Another approach considers a general class of designs, establishes an upper bound for the determinant of the information matrix, and then tries to characterize those designs for which the upper bound is attained (e.g., see Street et al. 2001; Street and Burgess 2004a; Burgess and Street 2005). Except for some special cases, the optimal designs obtained under both approaches usually contain a prohibitively large number of choice sets. However, although the optimal designs for a given model are not unique, they all have the same information matrix, which can then serve as a benchmark for assessing other smaller designs. Thus, in a subsequent and often challenging step, methods for constructing exact designs with practical numbers of choice sets are devised whose information matrices are either equal to the optimal one or that have a very high efficiency.

Suppose there are K attributes such that the k th attribute, $k = 1, \dots, K$, has v_k levels. The design region \mathcal{X} for the linear paired comparison model (22.6) then consists of all ordered pairs $x = (x_1, x_2)$ of K -tuples x_1 and x_2 whose k th component is one of the levels $1, \dots, v_k$. The information matrix of any approximate design ξ for the paired comparison model (22.6) is given by

$$M_\xi = \int_{\mathcal{X}} (f(x_1) - f(x_2))(f'(x_1) - f'(x_2))\xi(x) dx,$$

where $x = (x_1, x_2)$ denotes a generic pair. If ξ assigns positive weight $w_i = \xi((x_{i,1}, x_{i,2}))$ to exactly t distinct pairs $(x_{i,1}, x_{i,2})$, $i = 1, \dots, t$, then the information matrix can alternatively be written as a weighted sum

$$M_\xi = \sum_{i=1}^t w_i (f(x_{i,1}) - f(x_{i,2}))(f'(x_{i,1}) - f'(x_{i,2})).$$

Graßhoff et al. (2004) showed that when the attribute levels are effects coded the unique information matrix of a D -optimal design ξ^* is a block-diagonal matrix

$$\mathbf{M}_{\xi^*} = \text{diag}(\mathbf{M}_{v_1}, \dots, \mathbf{M}_{v_K}), \quad (22.7)$$

where $\mathbf{M}_{v_k} = \frac{2}{v_k-1}(\mathbf{I}_{v_k-1} + \mathbf{J}_{v_k-1})$. Since the D -criterion is invariant under reparameterizations, the corresponding optimal designs do not depend on the coding or the labeling of the attribute levels. The question of how D -optimal designs with practical numbers of pairs can be constructed will be discussed later in this section.

It follows that the quality of every exact design $d = \{(x_{1,1}, x_{1,2}), \dots, (x_{N,1}, x_{N,2})\}$ with N pairs can be assessed by calculating the D -efficiency

$$\text{eff}_D(d) = \left(\frac{\det(\mathbf{M}_d/N)}{\det(\mathbf{M}_{\xi^*})} \right)^{1/p}, \quad (22.8)$$

where $p = \sum_{k=1}^K (v_k - 1)$ is the number of model parameters and \mathbf{M}_d was defined in connection with model (22.6).

It should be noted that in order to put the comparison of the exact design d and the approximate design ξ^* on an equal footing, the D -efficiency measure compares the determinant of the normalized information matrix $\frac{1}{N}\mathbf{M}_d$ for d with that of \mathbf{M}_{ξ^*} . By using the D -efficiency, it is also possible to compare exact designs with different numbers of choice sets on a common scale. One interpretation of the D -efficiency is as follows. Suppose there exists a D -optimal exact design d^* with N^* pairs for which $\mathbf{M}_{d^*}/N^* = \mathbf{M}_{\xi^*}$. An exact design d with D -efficiency $\text{eff}_D(d) = c$, say, will then require $N = N^*/c$ pairs to estimate the model parameters with the same precision as the optimal design.

The problem of finding optimal designs for the MNL model with choice sets of size $m = 2$ under the assumption $\theta = \mathbf{0}$, where only the main effects of the K attributes with possibly different numbers of levels v_k are to be estimated, was also considered by Burgess and Street (2005) who used the approach of El-Helbawy and Bradley (1978). They derived the optimal information matrix for a class of designs that is less general than the set of all possible designs on the design region \mathcal{X} of all pairs considered here and in Graßhoff et al. (2004). Because it may not be obvious how the results in Graßhoff et al. (2004) and Burgess and Street (2005) are related, we outline how the matrix in (22.7) can be obtained when the latter approach is followed. As a consequence, the optimal designs of Burgess and Street (2005) are also optimal with respect to the fully general class of designs considered by Graßhoff et al. (2004).

For choice sets of size $m = 2$ and K attributes with v_k levels, $k = 1, \dots, K$, Theorem 2 in Burgess and Street (2005) shows that the exact design d^* , whose design points consist of all pairs of alternatives with the property that the two alternatives in a pair have different levels for each attribute, is D -optimal for estimating main effects. The corresponding optimal information matrix, normalized by the number of pairs N , is then equal to

$$\frac{1}{N}\mathbf{A}_{d^*} = \frac{1}{N} \text{diag}(\mathbf{A}_1, \dots, \mathbf{A}_K), \quad (22.9)$$

where $\mathbf{A}_k = \frac{c}{4} \frac{v_k}{v_k-1} \mathbf{I}_{v_k-1}$ for every attribute k (Burgess and Street 2005; Street and Burgess 2007, pp. 187–188) and c is the number of choice sets in the optimal design that contain the alternative whose attribute levels are all equal to 1.

Moreover, it can be shown that $c/N = 2/L$ where $L = \prod_{k=1}^K v_k$. It should be noted that here, we have modified the notation in Burgess and Street (2005) to make it more consistent with the current presentation. For example, where the original paper uses the attribute levels $0, \dots, v_k - 1$, we use the levels $1, \dots, v_k$. The framework in El-Helbawy and Bradley (1978) on which the derivation of the information matrix in Burgess and Street (2005) is based differs from the formulation of the MNL model in Section 22.2.1 and therefore a different letter is used for denoting the information matrix and its diagonal blocks. Moreover, contrary to Graßhoff et al. (2004), Burgess and Street (2005) do not formally distinguish the order of the two alternatives within each pair.

Although for $m = 2$ both Graßhoff et al. (2004) and Burgess and Street (2005) consider essentially the same model, the information matrix M_{ξ^*} in (22.7) for the optimal approximate design ξ^* and the normalized information matrix $\frac{1}{N}A_{d^*}$ in (22.9) for the optimal exact design d^* , which (apart from considering the order of the alternatives in the pairs) are for the same optimal design, are different. The reason for this discrepancy can be understood by looking more closely at the approach that underlies the work of Burgess and Street (2005) and Street and Burgess (2007). Initially, by using the model in El-Helbawy and Bradley (1978), these authors derive the normalized Fisher information matrix Λ , which is a square matrix of order $L = \prod_{k=1}^K v_k$, for all possible alternatives in the full factorial design for the K attributes. They then simplify Λ under an assumption of equal choice probabilities that is equivalent to assuming $\theta = \mathbf{0}$ in (22.1). Next, a $p \times L$ matrix

$$B = \begin{pmatrix} \mathbf{B}_{v_1} \otimes \frac{1}{\sqrt{v_2}} \mathbf{1}'_{v_2} \otimes \dots \otimes \frac{1}{\sqrt{v_K}} \mathbf{1}'_{v_K} \\ \frac{1}{\sqrt{v_1}} \mathbf{1}'_{v_1} \otimes \mathbf{B}_{v_2} \otimes \dots \otimes \frac{1}{\sqrt{v_K}} \mathbf{1}'_{v_K} \\ \vdots \\ \frac{1}{\sqrt{v_1}} \mathbf{1}'_{v_1} \otimes \frac{1}{\sqrt{v_2}} \mathbf{1}'_{v_2} \otimes \dots \otimes \mathbf{B}_{v_K} \end{pmatrix}, \tag{22.10}$$

of orthonormal contrasts for the main effects is introduced, which is partitioned into $(v_k - 1) \times L$ matrices corresponding to the attributes $k = 1, \dots, K$ and where every \mathbf{B}_{v_k} of order $(v_k - 1) \times v_k$ has orthonormal rows. Subsequently, the information matrix for the estimation of main effects is derived as $\mathbf{B}\Lambda\mathbf{B}'$, which further simplifies to $\frac{1}{N}A_{d^*}$ in (22.9). Note that in Equation 22.10 and throughout this chapter, for every natural number a , the all-one column vector of length a is denoted by $\mathbf{1}_a$ and the Kronecker product of vectors and matrices by “ \otimes .”

Although this approach may be technically convenient, the contrasts represented by the matrix \mathbf{B} usually have no natural interpretation for qualitative attributes, which is a consequence of requiring the rows of \mathbf{B} to be orthogonal vectors. For example, in the main-effects-only model, the contrast for each attribute level cannot be interpreted in the same way as the parameters corresponding to the effects-coded attribute levels at the beginning of Section 22.3.1. As was mentioned before, the D -optimality criterion is invariant with respect to reparameterizations but, for some other criteria, different parameterizations lead to different optimal designs and in such cases using a parameterization that has no clear interpretation does not appear to be helpful.

In general, for an arbitrary contrast matrix \mathbf{B} as in Burgess and Street (2005) whose rows are not orthogonal, the corresponding information matrix has to be computed by using generalized inverses as $(\mathbf{B}\Lambda^{-1}\mathbf{B}')^{-}$. The appropriate contrast matrix for estimating the p effects-coded parameters is given by $\mathbf{B} = \mathbf{S}\mathbf{G}$, where \mathbf{G} is obtained from (22.10) by replacing every $(v_k - 1) \times v_k$ matrix \mathbf{B}_{v_k} with the $v_k \times v_k$ centering matrix $\mathbf{K}_{v_k} = \mathbf{I}_{v_k} - \frac{1}{v_k}\mathbf{J}_{v_k}$

and where S is a rectangular block-diagonal matrix with K diagonal blocks of dimension $(v_k - 1) \times v_k$ given by $(\sqrt{v_k/LI_{v_k-1}}, \mathbf{0})$ for $k = 1, \dots, K$. It can then be shown that the information matrix for SG , that is, for the p effects-coded parameters, when using the optimal design is equal to

$$(S(GAG')^{-1}S')^{-1} = M_{\xi^*},$$

where M_{ξ^*} is the matrix in (22.7). Thus, by using the matrix SG , the approach in Burgess and Street (2005) leads to the same optimal information matrix as in the paired comparison setup of Graßhoff et al. (2004).

22.3.1.2 Design Constructions for Main Effects Only

The optimal designs that are used to derive the information matrices in (22.7) or (22.9) are generally not suitable for practical applications since they usually require every respondent to make a prohibitively large number of choices. Therefore, methods have been developed for constructing designs that use more practical numbers of choice sets and that are either optimal or possess a very high efficiency as measured, for example, by eff_D in (22.8). The most far-reaching results are available for situations where all attributes possess the same number of levels, but there is still some scope for further development when the attributes have different numbers of levels.

A simple method (Street and Burgess 2004a) for generating optimal choice designs with choice sets of size $m = 2$ when all attributes have two levels is based on regular fractional factorial designs of resolution III or higher. Each treatment combination of such a fractional factorial design specifies one alternative $x_{n,1}$ in a pair and the second alternative $x_{n,2}$ is obtained as the *foldover* of $x_{n,1}$. That is, for every attribute, the level in $x_{n,2}$ is the level not used in $x_{n,1}$. For example, if there are $K = 4$ attributes and if $x_{n,1} = (1, 2, 2, 1)$, then this alternative is paired with $x_{n,2} = (2, 1, 1, 2)$. The order of the alternatives within the pairs is ignored, and therefore if two rows of the fractional factorial lead to pairs that contain the same alternatives, then only one of those pairs is included in the choice design.

In Burgess and Street (2005), this construction is generalized to asymmetric situations where the K attributes can have different numbers of levels v_1, \dots, v_K and also to larger choice sets. The method starts with the full factorial design F in which the levels of every attribute k are coded as $0, \dots, v_k - 1$. As before, each treatment combination in F defines, without loss of generality, the first alternative in a choice set. For pairs, that is, choice sets of size $m = 2$, the second alternative in each choice set is obtained by adding a so-called *generator* $\mathbf{g} = (g_1, \dots, g_K)$ to the first alternative, where $g_k \in \{0, \dots, v_k - 1\}$ for $k = 1, \dots, K$, and the addition is performed componentwise modulo v_k . In order to construct an optimal design, usually, several of these generators will be needed and every additional generator gives rise to $L = \prod_{k=1}^K v_k$ further choice sets, where, as before, L is the number of treatments in a full factorial. If the generators satisfy the conditions in Theorem 3 of Burgess and Street (2005), then the design is D -optimal. The total number N of pairs that form the optimal design is equal to the size of the full factorial design times the number of generators, and consequently, except for small numbers of attributes and levels, the designs do not appear to be very practical. Street and Burgess (2008) acknowledged that when this method is used "it is something of an art to get the smallest design." In particular, for practitioners, it may not be very clear how appropriate generators can be found.

Several constructions of optimal designs with practical numbers of pairs are presented by Graßhoff et al. (2004). These are particularly useful when all attributes have the same number of levels or when some attributes have two, and the remaining attributes three, levels. If all K attributes have the same number of levels $v_k = v$ for $k = 1, \dots, K$, then one method to generate optimal designs is based on *Hadamard matrices* (see Chapter 9), that is, $\ell \times \ell$ matrices H with all elements equal to ± 1 and $HH' = H'H = \ell I_\ell$. Such matrices are known to exist for at least $\ell \leq 664$ (Kharaghani and Tayfeh-Rezaie 2005), where ℓ is either equal to 1, or 2 or is a multiple of 4. The construction uses K columns h_1, \dots, h_K of a Hadamard matrix of order ℓ where ℓ is the smallest integer greater than or equal to K for which such a matrix exists. These columns are then combined with a $(v - 1)v/2 \times (v - 1)$ matrix X_v , which is the effects-coded design matrix in the one-attribute version of the paired comparison model (22.6) with a single attribute at v levels for the design with all pairs (i, j) where $1 \leq i < j \leq v$. It follows that $\frac{2}{(v-1)v} X'_v X_v = M_v = \frac{2}{v-1} (I_{v-1} + J_{v-1})$. The resulting exact design d^* is D -optimal with design matrix

$$X = (h_1 \otimes X_v, \dots, h_K \otimes X_v) \tag{22.11}$$

and normalized information matrix $\frac{1}{N} M_{d^*} = M_{\xi^*}$, where $N = \ell(v - 1)v/2$ is the number of pairs and M_{ξ^*} is the matrix in (22.7). Since this construction represents the optimal design by its design matrix, for practical application, the coded levels in X need to be replaced with the actual levels. The following example illustrates an equivalent practical procedure for generating the designs, which sidesteps the need of explicitly constructing the matrix X in (22.11).

Example 22.1

Suppose there are $K = 3$ attributes with $v = 4$ levels each. In order to construct a D -optimal paired comparison design for estimating the main effects only, we start by writing down all possible pairs (i, j) of the four levels, where $i < j$, as a vector x^+ . Furthermore, a similar vector x^- is created in which the internal order of the pairs is reversed. This gives the two vectors

$$x^+ = \begin{pmatrix} (1, 2) \\ (1, 3) \\ (1, 4) \\ (2, 3) \\ (2, 4) \\ (3, 4) \end{pmatrix} \quad \text{and} \quad x^- = \begin{pmatrix} (2, 1) \\ (3, 1) \\ (4, 1) \\ (3, 2) \\ (4, 2) \\ (4, 3) \end{pmatrix}.$$

Note that x^+ can be regarded as a paired comparison design for model (22.6) with a single attribute with $v = 4$ levels and that the effects-coded design matrix for x^+ is equal to X_v in (22.11). More precisely, the design matrix for x^+ in the single-attribute model is given by

$$X_v = \begin{pmatrix} 1 & -1 & 0 \\ 1 & 0 & -1 \\ 2 & 1 & 1 \\ 0 & 1 & -1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{pmatrix}.$$

Similarly, when regarded as a single-attribute design for the paired comparison model, x^- possesses the effects-coded design matrix $-X_v$.

For $K = 3$ attributes, the construction requires three columns $h_1, h_2,$ and h_3 of a 4×4 Hadamard matrix, such as $h_1 = (1, 1, -1, -1)'$, $h_2 = (1, -1, 1, -1)'$, and $h_3 = (1, -1, -1, 1)'$. Then, in the matrix (h_1, h_2, h_3) , every 1 is replaced with x^+ and every -1 with x^- which leads to the array of pairs in Table 22.1a. Each row of the array is used to form one pair of alternatives in the optimal design. For every choice set n , the first alternative $x_{n,1}$ consists of the first components of each of the pairs in row n of Table 22.1a, while the second components of the pairs form the second alternative $x_{n,2}$. The $N = 24$ resulting pairs of the optimal design are shown in Table 22.1b. It should be noted that, although the order of the alternatives within the pairs of the optimal design is distinguished, the optimality of the design is not affected by interchanging alternatives within pairs. Thus, in Table 22.1b, every pair $(x_{n,1}, x_{n,2})$ could be replaced with $(x_{n,2}, x_{n,1})$.

Another construction of D -optimal designs (Graßhoff et al. 2004), which does not necessarily require the K attributes to have the same number of levels, combines orthogonal arrays (e.g., see Chapter 9 and Hedayat et al. 1999) with single-attribute designs for the paired comparison model (22.6). The first building block of these optimal designs is an asymmetric orthogonal array of strength two with N rows, K columns, and m_k symbols in

TABLE 22.1

Illustration of Hadamard Construction for Optimal Main-Effects-Only Designs

(a) Array from Substituting 1 and -1 in (h_1, h_2, h_3) with x^+ and x^-			(b) Pairs $(x_{n,1}, x_{n,2})$ of Optimal Design
(1,2)	(1,2)	(1,2)	((1,1,1), (2,2,2))
(1,3)	(1,3)	(1,3)	((1,1,1), (3,3,3))
(1,4)	(1,4)	(1,4)	((1,1,1), (4,4,4))
(2,3)	(2,3)	(2,3)	((2,2,2), (3,3,3))
(2,4)	(2,4)	(2,4)	((2,2,2), (4,4,4))
(3,4)	(3,4)	(3,4)	((3,3,3), (4,4,4))
(1,2)	(2,1)	(2,1)	((1,2,2), (2,1,1))
(1,3)	(3,1)	(3,1)	((1,3,3), (3,1,1))
(1,4)	(4,1)	(4,1)	((1,4,4), (4,1,1))
(2,3)	(3,2)	(3,2)	((2,3,3), (3,2,2))
(2,4)	(4,2)	(4,2)	((2,4,4), (4,2,2))
(3,4)	(4,3)	(4,3)	((3,4,4), (4,3,3))
(2,1)	(1,2)	(2,1)	((2,1,2), (1,2,1))
(3,1)	(1,3)	(3,1)	((3,1,3), (1,3,1))
(4,1)	(1,4)	(4,1)	((4,1,4), (1,4,1))
(3,2)	(2,3)	(3,2)	((3,2,3), (2,3,2))
(4,2)	(2,4)	(4,2)	((4,2,4), (2,4,2))
(4,3)	(3,4)	(4,3)	((4,3,4), (3,4,3))
(2,1)	(2,1)	(1,2)	((2,2,1), (1,1,2))
(3,1)	(3,1)	(1,3)	((3,3,1), (1,1,3))
(4,1)	(4,1)	(1,4)	((4,4,1), (1,1,4))
(3,2)	(3,2)	(2,3)	((3,3,2), (2,2,3))
(4,2)	(4,2)	(2,4)	((4,4,2), (2,2,4))
(4,3)	(4,3)	(3,4)	((4,4,3), (3,3,4))

column k , $k = 1, \dots, K$, which is commonly denoted by $OA(N; m_1, \dots, m_K; 2)$. The property that the orthogonal array has strength two means that in every $N \times 2$ subarray, all pairs of symbols appear equally often as rows. When some attributes have the same number of levels, it is convenient to use the notation $OA(N; m_1^{k_1}, \dots, m_r^{k_r}; 2)$, where m_1, \dots, m_r are distinct numbers of levels of, respectively, k_1, \dots, k_r attributes and $k_1 + \dots + k_r = K$. The second building block consists of designs that are optimal for the one-attribute version of the paired comparison model (22.6). For that model with a single attribute at v_k levels, the design that uses the $v_k(v_k - 1)/2$ pairs (i, j) where $i < j$ is D -optimal. If v_k is odd, then an optimal design of the same size that is *position balanced* in the sense that every level appears the same number of times in the first and second position of the pairs can also be found (Graßhoff et al. 2004), but if v_k is even, such a position balanced optimal design requires $v_k(v_k - 1)$ pairs. For a choice experiment with K attributes with v_1, \dots, v_K levels and choice sets of size $m = 2$, a D -optimal design can then be obtained by replacing the symbols in column k of an $OA(N; m_1, \dots, m_K; 2)$ with the N_k , say, pairs of an optimal design for the model (22.6) with a single attribute at v_k levels, where for all but one attribute the single-attribute design must be position balanced as explained. In order for the construction to be applicable, each m_k must be a multiple of N_k . Therefore, the method is most useful when all attributes have two or three levels, since in this situation, many suitable orthogonal arrays exist. For example, the $OA(36; 2^{11}, 3^{12}; 2)$ can be used to construct optimal designs with $N = 36$ pairs and up to 11 two-level and 12 three-level attributes. Otherwise, orthogonal arrays with the required properties are much less frequently available.

Example 22.2

To illustrate the construction of Graßhoff et al. (2004), we consider $K = 7$ attributes with $v_1 = \dots = v_6 = 3$ and $v_7 = 4$, levels. For $k = 1, \dots, 6$, the design with the $N_k = 3$ pairs $(1, 2)$, $(2, 3)$ and $(3, 1)$ is D -optimal and position balanced. For the last attribute with $v_7 = 4$, the $N_7 = 6$ pairs (i, j) of levels where $i < j$ are also D -optimal in the single-attribute paired comparison model. Clearly, this design is not position balanced and is the one exception allowed by the construction. A suitable orthogonal array is the $OA(18; 3^6 6^1; 2)$ in Table 22.2a. In each of the first six columns of the orthogonal array, the symbols 0, 1, and 2 are replaced with the pairs $(1, 2)$, $(2, 3)$, and $(3, 1)$, respectively. Similarly, in the final column, the symbols 0, 1, 2, 3, 4, and 5 are replaced with the pairs $(1, 2)$, $(1, 3)$, $(1, 4)$, $(2, 3)$, $(2, 4)$, and $(3, 4)$. The resulting array is shown in Table 22.2b, where each row will be used to construct a choice set. For every row, the corresponding choice set in the final design is obtained by collecting the first components of the pairs in Table 22.2b into the first alternative and the second components into the second alternative of the choice set. This gives the D -optimal design displayed in Table 22.2c.

It is worth emphasizing that in order for the method to work, for all but one of the K attributes, a position balanced design as described earlier needs to be used. Otherwise, the resulting design can be very inefficient. For example, in a recent description of the method (Street and Burgess 2012), the situation in Example 22.2 with $K = 7$ attributes and $v_1 = \dots = v_6 = 3$ and $v_7 = 4$ levels was considered, but for the three-level attributes, the symbols 0, 1, and 2 in Table 22.2a were replaced with—in our notation—the imbalanced single-attribute design $(1, 2)$, $(1, 3)$, and $(2, 3)$ leading to a design (Street and Burgess 2012, Table 10.13(b)) with a D -efficiency of 48.075%. This design does however not satisfy the conditions of Theorem 4 in Graßhoff et al. (2004, p. 369). In Table 10.13(c) of Street and

TABLE 22.2

Construction of Optimal Main-Effects-Only Design from Orthogonal Array

(a) $OA(18; 3^6 6^1; 2)$	(b) $OA(18; 3^6 6^1; 2)$ with symbols replaced
0 0 0 0 0 0 0	(1,2) (1,2) (1,2) (1,2) (1,2) (1,2) (1,2)
0 0 1 1 2 2 1	(1,2) (1,2) (2,3) (2,3) (3,1) (3,1) (1,3)
0 1 0 2 2 1 2	(1,2) (2,3) (1,2) (3,1) (3,1) (2,3) (1,4)
0 1 2 0 1 2 3	(1,2) (2,3) (3,1) (1,2) (2,3) (3,1) (2,3)
0 2 1 2 1 0 4	(1,2) (3,1) (2,3) (3,1) (2,3) (1,2) (2,4)
0 2 2 1 0 1 5	(1,2) (3,1) (3,1) (2,3) (1,2) (2,3) (3,4)
1 0 0 2 1 2 5	(2,3) (1,2) (1,2) (3,1) (2,3) (3,1) (3,4)
1 0 2 0 2 1 4	(2,3) (1,2) (3,1) (1,2) (3,1) (2,3) (2,4)
1 1 1 1 1 1 0	(2,3) (2,3) (2,3) (2,3) (2,3) (2,3) (1,2)
1 1 2 2 0 0 1	(2,3) (2,3) (3,1) (3,1) (1,2) (1,2) (1,3)
1 2 0 1 2 0 3	(2,3) (3,1) (1,2) (2,3) (3,1) (1,2) (2,3)
1 2 1 0 0 2 2	(2,3) (3,1) (2,3) (1,2) (1,2) (3,1) (1,4)
2 0 1 2 0 1 3	(3,1) (1,2) (2,3) (3,1) (1,2) (2,3) (2,3)
2 0 2 1 1 0 2	(3,1) (1,2) (3,1) (2,3) (2,3) (1,2) (1,4)
2 1 0 1 0 2 4	(3,1) (2,3) (1,2) (2,3) (1,2) (3,1) (2,4)
2 1 1 0 2 0 5	(3,1) (2,3) (2,3) (1,2) (3,1) (1,2) (3,4)
2 2 0 0 1 1 1	(3,1) (3,1) (1,2) (1,2) (2,3) (2,3) (1,3)
2 2 2 2 2 2 0	(3,1) (3,1) (3,1) (3,1) (3,1) (3,1) (1,2)

(c) Pairs $(x_{n,1}, x_{n,2})$ of optimal design

- ((1,1,1,1,1,1,1), (2,2,2,2,2,2,2))
- ((1,1,2,2,3,3,1), (2,2,3,3,1,1,3))
- ((1,2,1,3,3,2,1), (2,3,2,1,1,3,4))
- ((1,2,3,1,2,3,2), (2,3,1,2,3,1,3))
- ((1,3,2,3,2,1,2), (2,1,3,1,3,2,4))
- ((1,3,3,2,1,2,3), (2,1,1,3,2,3,4))
- ((2,1,1,3,2,3,3), (3,2,2,1,3,1,4))
- ((2,1,3,1,3,2,2), (3,2,1,2,1,3,4))
- ((2,2,2,2,2,2,1), (3,3,3,3,3,3,2))
- ((2,2,3,3,1,1,1), (3,3,1,1,2,2,3))
- ((2,3,1,2,3,1,2), (3,1,2,3,1,2,3))
- ((2,3,2,1,1,3,1), (3,1,3,2,2,1,4))
- ((3,1,2,3,1,2,2), (1,2,3,1,2,3,3))
- ((3,1,3,2,2,1,1), (1,2,1,3,3,2,4))
- ((3,2,1,2,1,3,2), (1,3,2,3,2,1,4))
- ((3,2,2,1,3,1,3), (1,3,3,2,1,2,4))
- ((3,3,1,1,2,2,1), (1,1,2,2,3,3,3))
- ((3,3,3,3,3,3,1), (1,1,1,1,1,1,2))

Burgess (2012), another design is presented that is also based on the orthogonal array in Table 22.2a and that possesses a D -efficiency of 98.10%, whereas the design in Table 22.2c is in fact D -optimal. If the framework of Burgess and Street (2005) with an orthonormal matrix B as in (22.10) and appropriate Λ matrices for the different designs is used, then the information matrix of the design in Table 10.13(c) of Street and Burgess (2012) is not

diagonal, whereas the information matrix of the optimal design in Table 22.2c is. Also, for the latter design, the determinant of the information matrix $B\Lambda B'$ attains the optimal bound reported in Burgess and Street (2005, p. 295).

22.3.2 Main Effects and Two-Factor Interactions

Main effects models and the corresponding designs are not appropriate if interactions of the attributes in a choice experiment can be expected. Frequently, it is assumed that higher-order interactions of three or more attributes are negligible and designs that allow the estimation of all main effects and two-factor (or first-order) interactions are considered. No general results appear to be available for the situation where all attributes can have different numbers of levels. Therefore, we focus on results for experiments in which the number of levels is the same.

22.3.2.1 Optimal Approximate Designs

In what follows, we consider the situation where all K attributes have v levels and the design region \mathcal{X} is the set of all ordered pairs of K -tuples with components in $\{1, \dots, v\}$. In order to represent main effects and two-factor interactions, the vector f of regression functions in the MNL model (22.1) or the paired comparison model (22.6) needs to be defined appropriately. If the attribute levels are effects coded, then for every alternative $x_i = (x_{i,1}, \dots, x_{i,K}) \in \{1, \dots, v\}^K$, the vector is given by

$$f(x_i) = (g'(x_{i,1}), \dots, g'(x_{i,K}), g'(x_{i,1}) \otimes g'(x_{i,2}), \dots, g'(x_{i,1}) \otimes g'(x_{i,K}), \dots, g'(x_{i,K-1}) \otimes g'(x_{i,K}))',$$

where for every $k = 1, \dots, K$, the component $g(x_{i,k})$ represents the effects-coded attribute level $x_{i,k}$. Hence, $g(x_{i,k})$ is a unit vector of length $v-1$ with a one in position $x_{i,k}$ if $1 \leq x_{i,k} < v$ and zero otherwise, and $g(x_{i,k}) = -\mathbf{1}_{v-1}$, if $x_{i,k} = v$. The first K component vectors of $f(x_i)$ are associated with the main effects of the attributes, while the remaining $K(K-1)/2$ Kronecker products represent interaction terms for pairs of attributes. The parameter vector θ can be partitioned similarly and in total, there are $p = (v-1)K + (v-1)^2K(K-1)/2$ parameters.

For every integer $1 \leq t \leq K$, let $\mathcal{X}(t)$ denote the subset of \mathcal{X} that contains all ordered pairs that possess different levels for exactly t of the K attributes. The number t is called the *comparison depth* (e.g., see Graßhoff et al. 2003) and can be recognized as the *Hamming distance* between x_i and x_j for any pair (x_i, x_j) of alternatives in $\mathcal{X}(t)$. Let $\bar{\xi}_t$ be the uniform approximate design on $\mathcal{X}(t)$ that gives the same weight to every pair in this set. The information matrix of every such design is block-diagonal and equal to

$$M_{\bar{\xi}_t} = \begin{pmatrix} h_1(t)\mathbf{I}_K \otimes M_v & \mathbf{0} \\ \mathbf{0} & h_2(t)\mathbf{I}_{K(K-1)/2} \otimes M_v \otimes M_v \end{pmatrix}, \tag{22.12}$$

where $h_1(t) = t/K$, $h_2(t) = t(1 - 1/v - (t-1)/(2(K-1)))/K$ and $M_v = \frac{2}{v-1}(\mathbf{I}_{v-1} + J_{v-1})$.

Furthermore, consider

$$t^* = K - 1 - \left\lceil \frac{K-2}{v} \right\rceil,$$

where $[r]$ denotes the integer part of r so that $[r] \leq r < [r] + 1$, and $w^* = (t^* + 1)/(t^*v + 1)$. It can then be shown (Graßhoff et al. 2003, Theorem 3) that the approximate design $\bar{\xi}_{t^*}$ is D -optimal on the design region \mathcal{X} if $(K - 2)/v$ is not an integer. Otherwise, if $(K - 2)/v$ is an integer, then the approximate design $\bar{\xi}_{t^*,w^*} = w^*\bar{\xi}_{t^*} + (1 - w^*)\bar{\xi}_{t^*+1}$ with information matrix $M_{\bar{\xi}_{t^*,w^*}} = w^*M_{\bar{\xi}_{t^*}} + (1 - w^*)M_{\bar{\xi}_{t^*+1}}$ is D -optimal, where $M_{\bar{\xi}_{t^*}}$ and $M_{\bar{\xi}_{t^*+1}}$ are the matrices in (22.12) for $\bar{\xi}_{t^*}$ and $\bar{\xi}_{t^*+1}$, respectively.

For the special case $v = 2$, where all attributes have two levels, it follows that if K is odd, then the uniform design on the set of all ordered pairs that have different levels for $t^* = (K + 1)/2$ of the attributes is D -optimal with information matrix $2(K + 1)/KI_{K+K(K-1)/2}$. If K is even, then the uniform design on the set of all ordered pairs that differ in $t^* = K/2$ or $t^* + 1 = K/2 + 1$ of the attributes is D -optimal and the information matrix is equal to $2(K + 2)/(K + 1)I_{K+K(K-1)/2}$. These designs for $v = 2$ were obtained before in a regression setting by Van Berkum (1987, pp. 30–31) and later for the multinomial logit model by Street et al. (2001) who use the approach of El-Helbawy and Bradley (1978).

22.3.2.2 Constructions for Main Effects and Two-Factor Interactions

Similar to the situation where only the main effects are to be estimated, the optimal designs in Graßhoff et al. (2003), Van Berkum (1987), and Street et al. (2001) for main effects and two-factor interactions are not practical since, apart from a few special cases, they require respondents to assess too many choice sets. Nevertheless, the designs are useful for the following reasons. First, the corresponding optimal information matrices allow the computation of the efficiency of any proposed design. Secondly, by considering which sorts of pairs are used by the optimal designs, it is possible to get an idea as to which pairs should be used for generating good designs with fewer choice sets.

The available methods for constructing efficient exact designs with more practical numbers of pairs that have been proposed to date are, essentially, limited to the case where all attributes have $v = 2$ levels. Instead of trying to find strictly optimal designs, which often would again lead to prohibitively large numbers of choice sets, these techniques sacrifice some efficiency for the benefit of being able to reduce the size of the designs. In one approach (Street and Burgess 2004a; Street et al. 2005), each treatment combination of the 2^K full factorial or a fractional factorial design of resolution V or higher is used to specify the first alternative in a pair, and the second alternative in each pair is obtained by adding generators in the same way as was described in Section 22.3.1.2. Some examples and limited extensions to situations with different numbers of levels can be found in Street and Burgess (2007).

Another technique (Großmann et al. 2012) uses Hadamard matrices, fractional factorial and incomplete block designs to generate designs that consist entirely of pairs that have different levels for t of the K two-level attributes. In principle, t can be chosen arbitrarily, but since the uniform designs $\bar{\xi}_t$ mentioned earlier are highly efficient if, for odd values of K , t is taken to be $t^* = (K + 1)/2$ and, for even values of K , $t = t^* = K/2$ or $t^* + 1 = K/2 + 1$ is used, at least heuristically, this suggests the use of these values also when constructing exact designs with smaller numbers of pairs. That with this choice of t , highly efficient designs can indeed be found is illustrated by Großmann et al. (2012) for $K = 3, \dots, 8$.

For given values of K and t , the construction of designs for estimating main effects and two-factor interactions in Großmann et al. (2012) requires three building blocks: a Hadamard matrix H of order $a \geq t$, a regular two-level fractional factorial design of

resolution III or higher for $K - t$ attributes (or the 2^{K-t} full factorial if $K - t \leq 2$) represented by a $u \times (K - t)$ matrix F , and a binary incomplete block design for K treatments in b blocks of size t represented by a $t \times b$ matrix B . Elements of F are coded as ± 1 and in B , the integers $1, \dots, K$ are used to indicate the treatments. The construction can then be summarized as follows. First, an $a \times t$ matrix A is obtained by selecting t columns of H . Secondly, two $au \times K$ matrices S and T are formed by combining the rows of A and $-A$ with the rows of F as shown schematically in Figure 22.1. More precisely, S is obtained by using u copies of A and, for each copy, every row of A is concatenated with a fixed row of F . The matrix T is obtained in a similar way by using $-A$. Finally, each of the b columns or blocks of B is used to permute the columns of a copy of S and T to generate au pairs of the design. More specifically, for a particular column g of B with elements $b_{1,g}, \dots, b_{t,g}$, the original columns $1, \dots, t$ of both S and T become the columns $b_{1,g}, \dots, b_{t,g}$ in the permuted matrices, and the remaining $K - t$ columns $t + 1, \dots, K$ of S and T are moved to the positions $c_{1,g} < \dots < c_{K-t,g}$, where $\{c_{1,g}, \dots, c_{K-t,g}\} = \{1, \dots, K\} \setminus \{b_{1,g}, \dots, b_{t,g}\}$. The au pairs for column g of B are obtained by combining every row of the permuted matrix S with the same row of the permuted matrix T . In other words, every choice set consists of a row from one of the permuted copies of S and the corresponding row of T . The procedure is repeated for each of the columns in B . The final design has $N = bau$ pairs. Note that, in the final design, the numbers 1 and -1 represent the effects-coded levels 1 and 2, respectively.

Example 22.3

An example of a design with $N = 48$ pairs that have different levels for $t = 3$ out of $K = 6$ attributes is shown in Table 22.3. For $n = 1, \dots, 48$, the table gives the pairs $(x_{n,1}, x_{n,2})$, where for simplicity of presentation, the values 1 and -1 are indicated by their signs. This design was constructed from a Hadamard matrix of order $a = 4$, a regular half fraction of the 2^3 full factorial, and an incomplete block design with blocks $\{1, 2, 3\}$, $\{1, 4, 5\}$, and $\{2, 4, 6\}$. In the first 16 pairs, the levels of the attributes 1, 2, and 3 in each alternative are determined by three columns of the Hadamard matrix, whereas the levels of the three remaining attributes are the same in both alternatives and correspond to the rows of the fractional factorial design. Likewise, for the pairs 17–32 and 33–48, the levels of the attributes in columns 1, 4, 5 and 2, 4, 6, respectively, depend on the three chosen columns of the Hadamard matrix, while the common levels of the remaining attributes are determined by the fractional factorial design. The final design has a D -efficiency of 91.85% for estimating the main effects and all two-factor interactions.

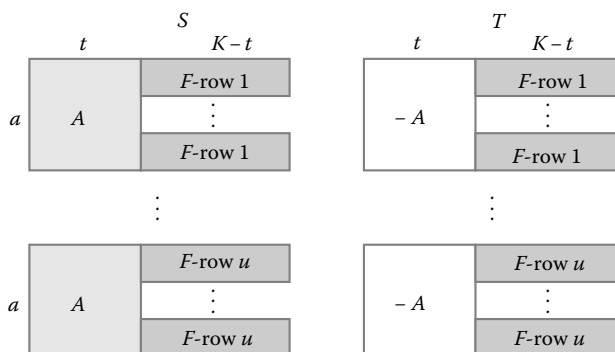


FIGURE 22.1 Matrices S and T in design construction for main effects and interactions.

TABLE 22.3

Design for Estimating Main Effects and Two-Factor Interactions of $K = 6$ Two-Level Attributes Generated by Method in Großmann et al. (2012)

n	Pair $(x_{n,1}, x_{n,2})$	n	Pair $(x_{n,1}, x_{n,2})$
1	((-, -, -, -, -, -), (+, +, +, -, -, -))	25	((-, +, -, -, -, +), (+, +, -, +, +, +))
2	((-, +, +, -, -, -), (+, -, -, -, -, -))	26	((-, +, -, +, +, +), (+, +, -, -, -, +))
3	((+, -, +, -, -, -), (-, +, -, -, -, -))	27	((+, +, -, -, +, +), (-, +, -, +, -, +))
4	((+, +, -, -, -, -), (-, -, +, -, -, -))	28	((+, +, -, +, -, +), (-, +, -, -, +, +))
5	((-, -, -, -, +, +), (+, +, +, -, +, +))	29	((-, +, +, -, -, -), (+, +, +, +, +, -))
6	((-, +, +, -, +, +), (+, -, -, -, +, +))	30	((-, +, +, +, +, -), (+, +, +, -, -, -))
7	((+, -, +, -, +, +), (-, +, -, -, +, +))	31	((+, +, +, -, +, -), (-, +, +, +, -, -))
8	((+, +, -, -, +, +), (-, -, +, -, +, +))	32	((+, +, +, +, -, -), (-, +, +, -, +, -))
9	((-, -, -, +, -, +), (+, +, +, +, -, +))	33	((-, -, -, -, -, -), (-, +, -, +, -, +))
10	((-, +, +, +, -, +), (+, -, -, +, -, +))	34	((-, -, -, +, -, +), (-, +, -, -, -, -))
11	((+, -, +, +, -, +), (-, -, +, -, +, -))	35	((-, +, -, -, -, +), (-, -, -, +, -, -))
12	((+, +, -, +, -, +), (-, -, +, +, -, +))	36	((-, +, -, +, -, -), (-, -, -, -, -, +))
13	((-, -, -, +, +, -), (+, +, +, +, +, -))	37	((-, -, +, -, +, -), (-, +, +, +, +, +))
14	((-, +, +, +, +, -), (+, -, -, +, +, -))	38	((-, -, +, +, +, +), (-, +, +, -, +, -))
15	((+, -, +, +, +, -), (-, +, -, +, +, -))	39	((-, +, +, -, +, +), (-, -, +, +, +, -))
16	((+, +, -, +, +, -), (-, -, +, +, +, -))	40	((-, +, +, +, +, -), (-, -, +, -, +, +))
17	((-, -, -, -, -, -), (+, -, -, +, +, -))	41	((+, -, -, -, +, -), (+, +, -, +, +, +))
18	((-, -, -, +, +, -), (+, -, -, -, -, -))	42	((+, -, -, +, +, +), (+, +, -, -, +, -))
19	((+, -, -, -, +, -), (-, -, -, +, -, -))	43	((+, +, -, -, +, +), (+, -, -, +, +, -))
20	((+, -, -, +, -, -), (-, -, -, -, +, -))	44	((+, +, -, +, +, -), (+, -, -, -, +, +))
21	((-, -, +, -, -, +), (+, -, +, +, +, +))	45	((+, -, +, -, -, -), (+, +, +, +, -, +))
22	((-, -, +, +, +, +), (+, -, +, -, -, +))	46	((+, -, +, +, -, +), (+, +, +, -, -, -))
23	((+, -, +, -, +, +), (-, -, +, +, -, +))	47	((+, +, +, -, -, +), (+, -, +, +, -, -))
24	((+, -, +, +, -, +), (-, -, +, -, +, +))	48	((+, +, +, +, -, -), (+, -, +, -, -, +))

It can be proved (Großmann et al. 2012) that the normalized information matrix in the paired comparison model (22.6) of every design of the aforementioned type is a diagonal matrix that is entirely determined by the incomplete block design \mathbf{B} . More precisely, this information matrix depends only on the number of blocks b , the replication numbers $r_k, k = 1, \dots, K$, of the treatments, and the concurrences $\lambda_{k,\ell}$ of the treatments, where $1 \leq k < \ell \leq K$. The diagonal element of the information matrix corresponding to the main effect of attribute k is equal to $4r_k/b$ and the diagonal element corresponding to the interaction of the attributes k and ℓ is equal to $4(r_k + r_\ell - 2\lambda_{k,\ell})/b$. As a consequence, the efficiency of every design of the aforementioned type can be computed very quickly without actually having to construct the design explicitly. If, for appropriate values of the block size t , we can choose \mathbf{B} as a balanced incomplete block design, so that the replication numbers are equal and also the concurrences are equal, then the corresponding choice design will usually possess a high D -efficiency (Großmann et al. 2012). Moreover, for given values of K and t and after choosing a Hadamard matrix \mathbf{H} of order a and a fractional factorial design F with u treatment combinations, the best design of the aforementioned type with $N = bau$ choice sets can be found by performing an exhaustive computer search over all binary incomplete block designs \mathbf{B} for K treatments in b blocks of size t , which again can be done very quickly. A comparison (Großmann et al. 2012) shows that for the same number of

pairs, the aforementioned designs are equally or more efficient than the designs proposed by Street and Burgess (2004a).

22.3.3 Partial Profiles

The designs in Tables 22.2c and Table 22.3 illustrate a potential issue in choice experiments that use very detailed descriptions of the alternatives, as reflected by a large number of attributes. On the one hand, increasing the number of attributes adds to realism but, on the other hand, empirical studies in psychology and other areas have demonstrated that when confronted with large amounts of information, respondents may resort to simplifying strategies in their preference judgments, such as ignoring part of the information or basing decisions only on subsets of the attributes. We do not attempt to review the voluminous literature here, but note only that such strategies may lead to violations of the assumption of the model (22.1) or (22.6) that the overall utility of every alternative $x_{n,i}$ is, essentially, a linear function $f'(x_{n,i})\theta$ of the parameters. There does not appear to be a consensus about the extent of this so-called information overload problem but, in agreement with common sense, empirical studies of task complexity show (e.g., see Caussade et al. 2005) that the number of attributes is one of the dimensions that influence complexity.

An early suggestion by Green (1974) for handling many attributes that was made in the context of conjoint analysis research proposes to describe each alternative by only a subset of the attributes in a study. A variant of this idea still uses all attributes but restricts the number of attributes for which the alternatives in every choice set are allowed to have different levels. Such incomplete or restricted descriptions of the alternatives are known as *partial profiles* (e.g., see Chrzan 2010). It is hoped that, when these are used, respondents are more likely to apply a decision process that conforms with the linearity assumption in the choice model about how the effects of the attributes are combined.

The problem of finding optimal designs for model (22.6) when the design region is restricted to pairs of partial profiles and when only the main effects of the attributes are to be estimated was considered by Graßhoff et al. (2004) for attributes with equal numbers of levels and by Großmann et al. (2006) for the general situation of unequal numbers of levels. For K attributes with v_1, \dots, v_K levels, respectively, the design region $\mathcal{X}(t)$ is the set of all ordered pairs $(x_{n,1}, x_{n,2})$ that have different levels for exactly t of the attributes. The levels of the other $K - t$ attributes are the same in the ordered pair, but not necessarily the same from pair to pair (see Example 22.4).

In the terminology of Section 22.3.2, the number t is the comparison depth. However, in the discussion of main-effects-only models, t is more frequently referred to as the *profile strength* (e.g., see Graßhoff et al. 2004), mainly because it is envisaged that for every pair in $\mathcal{X}(t)$, only the t attributes that have different levels will be shown to the respondents. In what follows, for reasons of consistency and in accordance with Graßhoff et al. (2003), we reserve the term *comparison depth* for the number t of attributes for which the two alternatives in a pair have different levels, whereas the term *profile strength* is used to designate the number q of attributes that are used when presenting pairs to respondents. We assume that $t \leq q \leq K$ and that the t attributes that differentiate the two alternatives in every pair are always among the q attributes that are shown. For the main-effects-only models considered here, it is then worthwhile to note that for a given comparison depth t , the same designs are D -optimal regardless of the value of q .

If all attributes have the same number of levels $v_k = v$, say, for $k = 1, \dots, K$, then the uniform approximate design $\bar{\xi}$ on $\mathcal{X}(t)$ is D -optimal with a block-diagonal information matrix. Under effects coding, this matrix is given by $M_{\bar{\xi}} = \frac{t}{K} I_K \otimes M_v$, where $M_v = \frac{2}{v-1} (I_{v-1} + J_{v-1})$.

Graßhoff et al. (2004) present a method for constructing D -optimal exact designs with practical numbers of choice sets that combines certain incomplete block designs and Hadamard matrices. These designs require $N = v(v - 1)\ell K / (2 \gcd(t, K))$ pairs of partial profiles, where ℓ is the smallest integer greater than or equal to the comparison depth t for which a Hadamard matrix of order ℓ exists and $\gcd(t, K)$ denotes the greatest common divisor of t and K .

In the general case of attributes with different numbers of levels, there exist D -optimal approximate designs for estimating main effects that are convex combinations or mixtures of uniform designs on certain subsets of the design region (Großmann et al. 2006). Under effects coding, the information matrix of such an optimal design for estimating main effects is block diagonal with a block for each attribute, that for $k = 1, \dots, K$ is equal to $c_k \mathbf{M}_{v_k}$, where c_k is a constant that depends on the design and $\mathbf{M}_{v_k} = \frac{2}{v_k - 1} (\mathbf{I}_{v_k - 1} + \mathbf{J}_{v_k - 1})$. It can be shown (Großmann et al. 2006) that if $c_k = (v_k - 1)t/p$ for every k , where $p = \sum_{k=1}^K (v_k - 1)$ is the total number of model parameters, then the corresponding design is D -optimal with a constant variance function that takes the value p over the whole design region. It is worth pointing out, however, that there also exist D -optimal designs whose variance function is not constant (Großmann et al. 2009) and for which the constants c_k , $k = 1, \dots, K$, are different from $(v_k - 1)t/p$.

Which designs are optimal depends very much on the number of attributes K , the comparison depth t , and the numbers of levels v_1, \dots, v_K . In applications, often some of the attributes have the same number of levels and by grouping the attributes accordingly, characterizations of optimal designs can be obtained more easily. For the situation in which there are K_1 attributes with u_1 levels and $K_2 = K - K_1$ attributes with u_2 levels, where $K \geq 2$ and $u_1 < u_2$, optimal approximate and exact designs for comparison depth $t < K$ are presented in Großmann et al. (2009). These designs use at most two different types of pairs in $\mathcal{X}(t)$ where, by definition, the alternatives in a pair of type (n_1, n_2) with $n_1 + n_2 = t$ have different levels for n_1 attributes with u_1 levels and n_2 attributes with u_2 levels.

Example 22.4

Table 22.4 shows an exact D -optimal design with $N = 28$ pairs for $K = 5$ attributes and $t = 3$, which was constructed by using the methods of Großmann et al. (2009), where the $K_1 = 3$ attributes in the first group have $u_1 = 2$ levels and the $K_2 = 2$ attributes in the second group have $u_2 = 3$ levels. In the table, attribute levels that are held constant in the two alternatives of a pair are represented by asterisks. Since in the main effects model the design matrix and, hence, the information matrix and the design efficiency are not affected by the choice of the constant levels, asterisks in the two alternatives of a pair that represent a level of the same attribute can be replaced with any of the possible levels of that attribute. For example, the first pair in Table 22.4 is $((2, 2, 2, *, *), (1, 1, 1, *, *))$. In both alternatives the asterisks indicating the constant level of the fourth attribute could be replaced with 1, whereas the asterisks corresponding to the constant level of the fifth attribute could be replaced with 3 to give $((2, 2, 2, 1, 3), (1, 1, 1, 1, 3))$. Alternatively, replacing the constant levels of the fourth and fifth attribute with 3 and 2, respectively, would be another possibility that would give the pair $((2, 2, 2, 3, 2), (1, 1, 1, 3, 2))$. Moreover, in different pairs, asterisks representing levels of the same attribute can be replaced with different levels of the attribute. For example, whereas in the first pair in Table 22.4 level 1 could be used for the fourth and level 3 for the fifth attribute, in the second pair of the table, levels 3 and 2 could be used, respectively, to replace the asterisks representing the constant levels of the fourth and fifth attribute, and so the first two pairs in the table would become $((2, 2, 2, 1, 3), (1, 1, 1, 1, 3))$ and $((2, 1, 1, 3, 2), (1, 2, 2, 3, 2))$.

TABLE 22.4

Optimal Partial Profile Design Generated by Method in Großmann et al. (2009)

n	Pair $(x_{n,1}, x_{n,2})$	n	Pair $(x_{n,1}, x_{n,2})$
1	$((2, 2, 2, *, *), (1, 1, 1, *, *))$	15	$((*, 2, *, 3, 1), (*, 1, *, 1, 3))$
2	$((2, 1, 1, *, *), (1, 2, 2, *, *))$	16	$((*, *, 2, 3, 2), (*, *, 1, 2, 3))$
3	$((1, 2, 1, *, *), (2, 1, 2, *, *))$	17	$((1, *, *, 1, 2), (2, *, *, 2, 1))$
4	$((1, 1, 2, *, *), (2, 2, 1, *, *))$	18	$((*, 1, *, 1, 3), (*, 2, *, 3, 1))$
5	$((2, *, *, 2, 2), (1, *, *, 1, 1))$	19	$((*, *, 1, 2, 3), (*, *, 2, 3, 2))$
6	$((*, 2, *, 3, 3), (*, 1, *, 1, 1))$	20	$((2, *, *, 1, 1), (1, *, *, 2, 2))$
7	$((*, *, 2, 3, 3), (*, *, 1, 2, 2))$	21	$((*, 2, *, 1, 1), (*, 1, *, 3, 3))$
8	$((1, *, *, 2, 1), (2, *, *, 1, 2))$	22	$((*, *, 2, 2, 2), (*, *, 1, 3, 3))$
9	$((*, 1, *, 3, 1), (*, 2, *, 1, 3))$	23	$((2, *, *, 1, 2), (1, *, *, 2, 1))$
10	$((*, *, 1, 3, 2), (*, *, 2, 2, 3))$	24	$((*, 2, *, 1, 3), (*, 1, *, 3, 1))$
11	$((1, *, *, 2, 2), (2, *, *, 1, 1))$	25	$((*, *, 2, 2, 3), (*, *, 1, 3, 2))$
12	$((*, 1, *, 3, 3), (*, 2, *, 1, 1))$	26	$((1, *, *, 1, 1), (2, *, *, 2, 2))$
13	$((*, *, 1, 3, 3), (*, *, 2, 2, 2))$	27	$((*, 1, *, 1, 1), (*, 2, *, 3, 3))$
14	$((2, *, *, 2, 1), (1, *, *, 1, 2))$	28	$((*, *, 1, 2, 2), (*, *, 2, 3, 3))$

In practice, often the attributes with constant levels would be omitted from the alternatives when pairs are presented to respondents. For example, in applications of the design in Table 22.4, typically only the three attributes with nonconstant levels would be used to characterize the alternatives in every pair. Thus, when making their choices, respondents only have to consider three of the five attributes which reduces the amount of information to be processed considerably. As was mentioned before, omitting the attributes with constant levels does not change the efficiency of the design for estimating the main effects.

Optimal designs for experiments where the attributes can be arranged into three groups with different numbers of levels are presented by Großmann et al. (2014), where due to the interplay between the group sizes, the comparison depth and the numbers of levels, more than 40 different designs have to be distinguished. The paper reports D -optimal approximate designs, but does not give any constructions of exact designs with practical numbers of pairs. This remains a challenging problem.

Deriving optimal designs for pairs of partial profiles when all two-factor interactions in addition to the main effects are to be estimated becomes even more challenging. In this setting, it is necessary to distinguish explicitly between the profile strength, that is, the number of attributes that are presented, and the comparison depth, that is, the number of attributes for which the alternatives in a pair have different levels. Optimal approximate designs and the corresponding information matrices can be found in Graßhoff et al. (2003) for the case where all attributes have the same number of levels. The construction of relatively small exact designs that possess good efficiency properties appears to be an open problem.

22.4 Designs for Larger Choice Sets

The use of choice sets with three or more alternatives is often motivated by reference to the similarity between the corresponding choice experiments and choice problems in the

real world. More specifically, experiments with larger choice sets are regarded as being more realistic, since, for example, purchasing a product usually requires choosing one of potentially many alternatives. In addition, there appears to be a widespread belief that larger choice sets provide more information than can be obtained from paired comparisons. This section considers this claim in some detail before looking at how optimal designs for the MNL model with choice sets of size $m \geq 3$ can be derived. Here, the emphasis is again on analytic results which, for the most part, can only be obtained under the assumption of indifference, that is, for $\theta = 0$. Computational approaches for the case where $\theta \neq 0$ will be described in Section 22.5.

22.4.1 Efficiency Gains due to Increased Size of Choice Sets

According to formula (22.3), the asymptotic information for a single choice set with m alternatives is given by

$$M_{C,\theta} = X'(\text{diag}(p_\theta) - p_\theta p'_\theta)X,$$

where we have suppressed the index n specifying the choice set C . It is not hard to see that $M_{C,\theta}$ can also be expressed as

$$M_{C,\theta} = \sum_{1 \leq i < j \leq m} p_\theta(i; C) p_\theta(j; C) (f(x_i) - f(x_j))(f'(x_i) - f'(x_j)),$$

which is a weighted sum

$$M_{C,\theta} = \sum_{1 \leq i < j \leq m} \left(\frac{e^{f'(x_i)\theta} + e^{f'(x_j)\theta}}{\sum_{k=1}^m e^{f'(x_k)\theta}} \right)^2 M_{\{x_i, x_j\}, \theta}$$

of the information matrices

$$M_{\{x_i, x_j\}, \theta} = p_\theta(i; \{x_i, x_j\}) p_\theta(j; \{x_i, x_j\}) (f(x_i) - f(x_j))(f'(x_i) - f'(x_j)) \tag{22.13}$$

for the individual pairs (x_i, x_j) in a paired comparison experiment that can be formed from the alternatives within the choice set C . In the case of indifference, that is, for $\theta = 0$, this representation of $M_{C,\theta}$ simplifies to

$$M_{C,0} = \frac{4}{m^2} \sum_{1 \leq i < j \leq m} M_{\{x_i, x_j\}, 0}. \tag{22.14}$$

By the linearity of the information matrix, this decomposition carries over to any design ξ on the set of choice sets, so that

$$M_{\xi,0} = \frac{4}{m^2} \sum_{1 \leq i < j \leq m} M_{\xi_{ij},0},$$

where ξ_{ij} is the two-dimensional marginal of ξ for the i th and j th alternative. Note that ξ_{ij} uses choice sets of size two, which are obtained by selecting the i th and j th alternative from

every choice set of the design ξ . Hence, for any homogeneous criterion function Φ , defined on the set of nonnegative definite symmetric matrices, which is to be maximized, we obtain an upper bound

$$\Phi(M_{\xi^*, \mathbf{0}}) \leq 2 \frac{m-1}{m} \Phi(M_{\xi_{(12)}^*, \mathbf{0}})$$

for any optimal design ξ^* on the choice sets of size m with respect to an optimal paired comparison design $\xi_{(12)}^*$, say, for choice sets of size 2. This immediately implies that the relative efficiency of ξ^* with respect to $\xi_{(12)}^*$ satisfies the inequality

$$\text{eff}_{\Phi}(\xi^*; \xi_{(12)}^*) = \frac{\Phi(M_{\xi^*, \mathbf{0}})}{\Phi(M_{\xi_{(12)}^*, \mathbf{0}})} \leq 2 \frac{m-1}{m}. \quad (22.15)$$

In particular, even if the size m gets large, this relative efficiency cannot exceed the value of 2. It is worth noting that the bound in (22.15) is not only valid for main effects models but for any vector f of regression functions in the MNL model (22.1).

In main-effects-only models, in general, this upper bound can only be achieved, if all two-dimensional marginal designs ξ_{ij}^* , $1 \leq i < j \leq m$, of ξ^* are optimal for the paired comparison model (22.6). This can and will only be the case if, for each attribute, the number of levels is at least equal to the size m of the choice sets. For smaller numbers of levels, $\text{eff}_{\Phi}(\xi^*; \xi_{(12)}^*)$ may drop below 1, which means that sometimes using choice sets of size two may be preferable to using sets of size m , as will be seen in Section 22.4.2. The homogeneous function Φ that corresponds to the D -criterion is defined by $\Phi(M) = (\det(M))^{1/p}$ for every nonnegative definite symmetric $p \times p$ matrix M , where p is the number of model parameters.

Equation 22.15 can be rearranged to give a lower bound for the efficiency of the optimal paired comparison design $\xi_{(12)}^*$ relative to the optimal choice design ξ^* with choice sets of size $m > 2$. It follows that for $m = 3, \dots, 6$, the lower bounds for the efficiency of the optimal paired comparison design are equal to 75%, 66.66%, 62.5%, and 60%, respectively. These bounds are in agreement with results for the D -criterion reported in Burgess and Street (2006), where the relative D -efficiency of the paired comparison design with respect to designs with larger choice sets can be obtained by dividing the value for $m = 2$ in their tables by the corresponding values for larger choice set size m . In the cases where for main-effects-only models the aforementioned lower bounds are attained, it can be seen that the number of levels of each attribute is greater than or equal to m as expected.

By using experiments with choice sets of size $m = 3$ or $m = 4$, sometimes, substantial efficiency gains may be achieved, although whether this is possible depends on the number of attribute levels. Increasing the size of the choice sets further leads only to diminishing improvements. Also, when the size of the choice sets for an empirical application is determined, the potential statistical efficiency gains for $m > 2$ need to be balanced with the increased information processing requirements that may affect the quality of the responses. Thus, although studies with very large choice sets have been reported, it appears that choice sets of size $m = 2, 3$, or 4 are probably most useful for applications.

22.4.2 Relationship with Block Designs

Under the indifference assumption $\theta = \mathbf{0}$, there exists an intimate relationship between the MNL model and models for blocked experiments in which both the attributes and the

blocks have fixed effects. To this end, every choice set is regarded as a block. Consider first the linear model

$$y = X\theta + Z\gamma + \epsilon, \tag{22.16}$$

where X is obtained by stacking the $m \times p$ matrices X_1, \dots, X_N corresponding to N choice sets C_1, \dots, C_N of size m . Moreover, $Z = I_N \otimes \mathbf{1}_m$ is an $mN \times N$ matrix of zeros and ones that for every row of X indicates the block (choice set) to which the alternative represented by that row belongs. The parameter vectors θ and γ represent the effects of the attributes and the effects of the blocks, respectively. Provided θ is estimable under the given design, the information matrix for estimating θ in this model is equal to (e.g., see Goos 2002, p. 15)

$$\tilde{M} = X'X - X'Z(Z'Z)^{-1}Z'X = \sum_{n=1}^N X'_n K_m X_n, \tag{22.17}$$

where K_m denotes again a centering matrix. Under indifference, the information matrix for the MNL model in Equation 22.4 is equal to $\frac{1}{m}\tilde{M}$, where $1/m$ is the constant probability of choosing an alternative. For optimality criteria that are based on the information matrix, this relationship shows that blocked designs that are optimal for estimating θ in the linear model with fixed block effects are also optimal for the MNL choice model under the indifference assumption.

When there is only a single attribute at v levels, a design for a choice experiment with N sets of size m can be regarded as a classical block design d with v treatments in N blocks of size m . We consider this situation in some detail and use results for this setting later in the section to generate choice designs with choice sets of size m that are optimal for estimating main effects in experiments with K attributes.

Assuming that blocks and treatments have fixed effects, the model for v treatments in N blocks of size m is

$$y = \mathbf{1}_{mN}\mu + \tilde{X}\tilde{\theta} + Z\gamma + \epsilon, \tag{22.18}$$

where \tilde{X} is now an $mN \times v$ matrix of zeros and ones indicating which treatments are used on the units corresponding to the rows of the matrix and the parameter vector $\tilde{\theta}$ of length v has a component for every treatment or attribute level. The matrix Z and the vector γ are the same as in (22.16) and the parameter μ represents a general mean. The information matrix for d is based on the intrablock analysis (John and Williams 1995) and usually expressed in the form

$$C_d = \text{diag}(r_1, \dots, r_v) - \frac{1}{m}NN', \tag{22.19}$$

where r_1, \dots, r_v are the replication numbers of the levels and the $v \times N$ matrix N is the incidence matrix (e.g., see Shah and Sinha 1989). It should be noted that, in accordance with Kiefer (1958, p. 686), here, the elements of N are not restricted to be equal to zero or one but can be nonnegative integers. Thus, incomplete block designs with or without replication in the same block as well as complete block designs with possibly unequal replication of the

treatments in the same block are covered. The matrix C_d can be written as \tilde{M} in the first equation of (22.17) where only X has to be replaced with \tilde{X} .

Following Kiefer (1958, p. 689), in this setting with v treatments and N blocks of size m , a design is called a *balanced block design* (BBD) if (1) all elements n_{ij} of N are equal to m/v if this ratio is an integer and are either of the two integers closest to m/v otherwise, (2) all v treatments have the same replication $r_i = r$, and (3) all $\lambda_{ij} = \sum_{k=1}^N n_{ik}n_{jk}$ are equal to, say, λ for $i \neq j$. Also see Chapter 3. For a BBD d^* , the matrix C_{d^*} is completely symmetric and hence has only two distinct eigenvalues (Rao 1973, p. 67). The eigenvalue zero has multiplicity one with corresponding eigenvector $\mathbf{1}_v$ and the eigenvalue $r + (\lambda - \rho)/m$ has multiplicity $v - 1$, where ρ is the element on the diagonal of NN' .

Now consider the problem of estimating a vector of $v - 1$ orthonormal contrasts of $\tilde{\theta}$, in the block design model (22.18), that is, of a vector $P\tilde{\theta}$ where $PP' = I_{v-1}$ and $P\mathbf{1}_v = \mathbf{0}$. In order for $P\tilde{\theta}$ to be estimable, it is necessary and sufficient that C_d has rank $v - 1$. A design d with this property is called *connected*. The covariance matrix of a connected design for estimating $P\tilde{\theta}$ is $(PC_dP')^{-1}$ (Kiefer 1958) and the corresponding information matrix is just the inverse. If a BBD d^* exists, then, among other things, this is D -optimal for estimating $P\tilde{\theta}$ in the class of all connected designs (Kiefer 1958, Theorem 3.1). In other words, $\det(PC_{d^*}P') \geq \det(PC_dP')$ for every connected design d . The columns of P' form an orthonormal basis of the eigenspace for eigenvalue $r + (\lambda - \rho)/m$ and so $PC_{d^*}P' = (r + (\lambda - \rho)/m)I_{v-1}$ is a multiple of the identity matrix.

For every $(v - 1) \times v$ matrix L whose rows are not necessarily orthogonal but which span the same vector space as the rows of P , there exists a nonsingular matrix Q of order $v - 1$ such that $L = QP$. The covariance matrix of a block design d for estimating $L\tilde{\theta}$ in the block design model (22.18) is then equal to $Q(PC_dP')^{-1}Q'$ and the corresponding information matrix is $(Q^{-1})'(PC_dP')Q^{-1}$. It follows that the same designs are D -optimal for estimating $P\tilde{\theta}$ and $L\tilde{\theta}$. Such functions $L\tilde{\theta}$ are of interest since the components of $L\tilde{\theta}$ may have a clearer interpretation than those of $P\tilde{\theta}$. For example, the matrix $L = (I_{v-1}, \mathbf{0}) - \frac{1}{v}\mathbf{1}_{v-1}\mathbf{1}'_v$ transforms $\tilde{\theta}$ into a vector of effects-coded parameters. For this choice of L and every P with orthonormal rows, the matrix Q that satisfies $L = QP$ is given by $Q = (I_{v-1}, \mathbf{0})P'$ with inverse $Q^{-1} = P(I_{v-1}, \mathbf{0})'(I_{v-1} + J_{v-1})$.

By means of L , the classical block design setting can be linked to the model (22.16). More precisely, there exists a design matrix $X = \tilde{X}\tilde{L}$ such that estimating θ in (22.16) is equivalent to estimating $L\tilde{\theta}$ in (22.18). For example, if L is the matrix given earlier that accomplishes the effects coding, then $\tilde{L} = (I_{v-1}, -\mathbf{1}_{v-1})'$ is the corresponding matrix for transforming the design matrix. By letting $X = \tilde{X}\tilde{L}$, it can be seen that for every connected design d , the information matrix in (22.17) is equal to $(Q^{-1})'(PC_dP')Q^{-1}$, that is, the information matrix for the estimation of $L\tilde{\theta}$ in the model (22.18). This relationship implies that BBDs are also D -optimal for choice experiments with a single attribute at v levels and N choice sets of size m .

In particular, for a single qualitative attribute with v levels, *balanced incomplete block designs* (BIBDs) with N blocks of size m , are optimal if v is greater than the size m of the choice sets. Note that the three parameters v , m , and N determine the other parameters of the BIBD, that is, the replication r and the index λ . If $m = cv$ for a positive integer c , then a design with N complete blocks is optimal, where all treatments appear c times in every block. Finally, if $m > v$ is not a multiple of v so that $m = cv + s$, where c and s are positive integers and $s < v$, a general BBD is optimal. This design may be regarded as being

obtained by combining blockwise a complete block design where in each of N blocks all treatments are replicated c times with a BIBD for v treatments in N blocks of size s . For example, for a binary attribute with $v = 2$ levels, a complete block design with equal replication of both levels within each block is optimal, if m is even. If m is odd, then in half of the choice sets of an optimal design, the first and second level occur with frequencies $(m + 1)/2$ and $(m - 1)/2$, respectively, while in the remaining sets these frequencies are reversed.

In all of the aforementioned three cases, the normalized information matrix $\frac{1}{N}PC_{d^*}P'$ for estimating $P\theta$ in model (22.18) is a multiple c^*I_{v-1} of the identity matrix and the corresponding matrix when effects coding is used is $c^*(I_{v-1} + J_{v-1})$. Table 22.5 gives values of the common multiplier c^* , where the first two cases can be recognized as special cases of the last row for general BBDs by setting, in the expression $m = cv + s$ for the size of the choice sets, c or s to 0. When considering the normalized information matrix for the MNL choice model, the constants c^* have to be divided further by m .

From Table 22.5, it can be seen that, if the size m of the choice sets is not a multiple of the number of levels v , then the optimal design for size m carries less information than an optimal design where the size of the choice sets is a multiple of v . More precisely, the relative D -efficiency in the MNL model of a design that is optimal for choice sets of size $m > v$ with respect to a design that is optimal for choice sets of size cv , where c is a positive integer, is equal to $1 - \frac{s(v-s)}{m^2(v-1)}$. For example, for $v = 2$ levels, the relative efficiency of the optimal design with choice sets of size $m = 3$ drops down to 89% compared to designs that are optimal for choice sets of even size. Similarly, the relative D -efficiency in the MNL model of a design that is optimal for choice sets of size $m < v$ with respect to an optimal design for choice sets of size cv is $\frac{mv-v}{mv-m}$.

As indicated at the start of the section, the optimality results for a single attribute can serve as building blocks for constructing optimal designs with larger numbers of attributes. To this end, every choice set is treated as a block and the corresponding responses are regarded as a multivariate observation in a linear model with additive effects. By generalizing results of Schwabe (1996b, Section 5.2) to multivariate responses and using similar ideas as in Schwabe (1996a), optimal designs can be generated as *product-type designs*. If there are K attributes in total, then such a design is obtained by firstly choosing for every attribute k with v_k levels a block design that is optimal in the single-attribute model and where for all but at most one attribute, the single-attribute designs satisfy an additional orthogonality condition to be given below. Secondly, the choice sets in the model with K attributes are generated as a *full factorial* of entire blocks in the sense that every block from the design for the first attribute is combined with every block from the design for the second attribute, every such combination of blocks with every block from the design for the third attribute, and so forth.

More precisely, if for every attribute $k = 1, \dots, K$, the matrices $X_1^{(k)}, \dots, X_{N_k}^{(k)}$ denote the design matrices in the choice model with a single attribute that are obtained from an

TABLE 22.5
Summary of Optimal BBDs

Case	Optimal Design	Multiplier c^*
$m < v$	BIBD	$\frac{m-1}{v-1}$
$m = cv, c \in \mathbb{N}$	Complete block design	$\frac{m}{v}$
$m = cv + s$, where $c, s \in \mathbb{N}$ and $s < v$	General BBD	$\frac{1}{v} \left(m - \frac{s(v-s)}{m(v-1)} \right)$

optimal block design with v_k treatments in N_k blocks of size m , then the optimal product-type design for the choice model with K attributes has $N = \prod_{k=1}^K N_k$ choice sets, where for every choice set, the design matrix is of the form $(\mathbf{X}_{n_1}^{(1)}, \dots, \mathbf{X}_{n_K}^{(K)})$ and each $n_k \in \{1, \dots, N_k\}$. For all but one attribute, the single-attribute designs need to satisfy the orthogonality condition that the sum of $\mathbf{X}_1^{(k)}, \dots, \mathbf{X}_{N_k}^{(k)}$ is a zero matrix.

Example 22.5

A D -optimal product-type design with choice sets of size $m = 3$ for estimating the main effects of $K = 3$ qualitative attributes with $v_1 = 2, v_2 = 3,$ and $v_3 = 4$ levels, respectively, is shown in Table 22.6. For every choice set C_n and every attribute k , the block of the single-attribute design that was used to construct the choice set consists of the k th component of each of the three alternatives (in that order) that make up the choice set. For example, the three blocks underlying the first choice set in Table 22.6 are 1,2,2 for the first attribute, 3,1,2 for the second attribute, and 4,3,2 for the third attribute, respectively.

The normalized information matrix $\frac{1}{N} \mathbf{M}_{d^*, \mathbf{0}}$ of an optimal product-type design d^* for estimating main effects in the MNL model under indifference is block diagonal. Under effects-coding, the block on the diagonal for attribute k with v_k levels is equal to $\frac{c_k^*}{m} (\mathbf{I}_{v_k-1} + \mathbf{J}_{v_k-1})$, where c_k^* is the multiplier from Table 22.5 for v_k levels and m is the size of the choice sets. Hence, when comparing optimal designs for different values of the size of the choice sets, the efficiency results carry over from the marginal models with only a single attribute to the model with several attributes: maximal information can only be attained if the size m of the choice sets is a multiple of all numbers of levels, that is, if the least common multiple of the numbers of levels is a divisor of m . If all attributes have the same number of levels, then the optimal product-type design that employs pairs has the same efficiency as each of its marginal components relative to optimal designs with choice sets of size $m > 2$. For binary attributes, these efficiencies are equal to 1 if m is even and to $1 - 1/m^2$ if m is odd.

TABLE 22.6

Optimal Design for $m = 3$ and Three Attributes with Levels $v_1 = 2, v_2 = 3,$ and $v_3 = 4$

n	Choice Set C_n			n	Choice Set C_n		
1	(1, 3, 4)	(2, 1, 3)	(2, 2, 2)	13	(2, 3, 4)	(1, 1, 3)	(1, 2, 2)
2	(1, 3, 2)	(2, 1, 1)	(2, 2, 3)	14	(2, 3, 2)	(1, 1, 1)	(1, 2, 3)
3	(1, 3, 1)	(2, 1, 2)	(2, 2, 4)	15	(2, 3, 1)	(1, 1, 2)	(1, 2, 4)
4	(1, 3, 1)	(2, 1, 4)	(2, 2, 3)	16	(2, 3, 1)	(1, 1, 4)	(1, 2, 3)
5	(1, 1, 4)	(2, 2, 3)	(2, 3, 2)	17	(2, 1, 4)	(1, 2, 3)	(1, 3, 2)
6	(1, 1, 2)	(2, 2, 1)	(2, 3, 3)	18	(2, 1, 2)	(1, 2, 1)	(1, 3, 3)
7	(1, 1, 1)	(2, 2, 2)	(2, 3, 4)	19	(2, 1, 1)	(1, 2, 2)	(1, 3, 4)
8	(1, 1, 1)	(2, 2, 4)	(2, 3, 3)	20	(2, 1, 1)	(1, 2, 4)	(1, 3, 3)
9	(1, 2, 4)	(2, 3, 3)	(2, 1, 2)	21	(2, 2, 4)	(1, 3, 3)	(1, 1, 2)
10	(1, 2, 2)	(2, 3, 1)	(2, 1, 3)	22	(2, 2, 2)	(1, 3, 1)	(1, 1, 3)
11	(1, 2, 1)	(2, 3, 2)	(2, 1, 4)	23	(2, 2, 1)	(1, 3, 2)	(1, 1, 4)
12	(1, 2, 1)	(2, 3, 4)	(2, 1, 3)	24	(2, 2, 1)	(1, 3, 4)	(1, 1, 3)

The three types of BBDs in Table 22.5 can be characterized in terms of the number of pairs of units in each block that receive a different treatment. More precisely, if $\{b_1, \dots, b_m\}$ is a generic block, then this number is

$$\delta_m(v) = \sum_{i=1}^{m-1} \sum_{j=i+1}^m \delta(b_i, b_j)$$

where $\delta(b_i, b_j) = 1$ if $b_i \neq b_j$ and $\delta(b_i, b_j) = 0$ otherwise. If $m < v$, then $\delta_m(v) = m(m - 1)/2$ since no treatment is replicated in a block. For $m = cv$, within each block, every treatment has replication $c = m/v$ which implies that $\delta_m(v) = m^2(v - 1)/(2v)$. Finally, if $m = cv + s$ for positive integers c and s where $s < v$, it follows that $\delta_m(v) = (m^2(v - 1) - s(v - s))/(2v)$.

For K attributes with levels v_1, \dots, v_K , the optimal product-type designs can be characterized by extending the concept of the comparison depth for pairs of alternatives (see Section 22.3.2) to larger choice sets. For any two alternatives x_i and x_j in a choice set, let the *comparison depth* between x_i and x_j be defined as the number of attributes for which x_i and x_j have different levels (e.g., see Graßhoff et al. 2013). Moreover, for a choice set C of size m , define the *total comparison depth* of C as the sum of the comparison depths for all pairs of alternatives x_i and x_j in C where $i < j$. For an optimal product-type design, the total comparison depth is the same for all choice sets and equal to

$$\delta = \sum_{k=1}^K \delta_m(v_k).$$

An optimal design then consists of choice sets with maximum total comparison depth δ .

When all K attributes are binary, the maximum comparison depth is $\delta = m^2K/4$ if m is even and $\delta = (m^2 - 1)K/4$ if m is odd. Optimal designs can then be obtained from a single choice set with maximal comparison depth by permuting the levels of all attributes independently. This gives a kind of *full factorial* design, which is optimal and from which smaller optimal designs can be obtained as *fractional factorials* in the usual way. For attributes with different numbers of levels, orthogonal arrays may be used instead of full factorials in order to generate optimal designs with smaller numbers of choice sets.

The values of $\delta_m(v)$ given earlier coincide with the upper bounds for sums of differences reported by Burgess and Street (2005, Theorem 1). Likewise, the requirement that the choice sets of optimal designs need to have maximum total comparison depth is in agreement with their Theorem 2. For the special case of binary attributes, the approach presented here immediately gives the optimality results for main effects in Burgess and Street (2003). The cited papers also contain some additional results for estimating main effects and two-factor interactions. Corresponding construction methods are presented by Street et al. (2005) and further examples can be found in Street and Burgess (2007). It should be noted however that in general, these methods require a very large number of choice sets (Großmann et al. 2007). Some constructions that can be used for generating optimal designs with smaller numbers of choice sets have been presented by Mukerjee et al. (2002) and Das and Dey (2004).

22.4.3 Some Local Optimality Results

Analytic results for the general case in which the parameter vector θ in the choice model (22.1) is not assumed to be equal to a zero vector are hard to come by because of the

nonlinearity of the model. Typically, numerical optimization techniques or search algorithms have to be invoked in order to obtain designs that are locally optimal for a given $\theta \neq 0$.

A notable exception to this general rule is the MNL model proposed by Kanninen (2002), in which initially all K attributes are assumed to be continuous. All but one of these attributes take values in a bounded interval. The settings of these $K - 1$ attributes are chosen from the boundary values of the intervals. The remaining attribute, which the author refers to as a *manipulator*, is not bounded and can be adjusted continuously. In practice, this means that the manipulator must be allowed to take any value in a sufficiently large interval. In this situation, optimal designs can be constructed by selecting a particular full or fractional factorial design for the bounded attributes and by adjusting the manipulator to an optimal level that depends on the parameter θ . Moreover, the designs remain optimal for the model in which each of the bounded attributes is replaced by a qualitative two-level factor.

An appealing feature of Kanninen's (2002) model is that optimal settings for the $K - 1$ bounded attributes can be chosen without having to worry about the values of the parameters in the vector θ . Only the settings of the unbounded manipulator variable depend on the parameters in θ . For choice sets of size $m = 2$, the settings of the $K - 1$ bounded attributes of a locally D -optimal design for estimating the vector $\theta \neq 0$ in Kanninen's model with the additional manipulator variable can be chosen as those of any design that, under the indifference assumption that the parameters are equal to zero, is optimal for the model that contains only the $K - 1$ bounded attributes (Graßhoff et al. 2007). For larger choice sets of size $m \geq 3$, the settings for the $K - 1$ bounded attributes in Kanninen's model have to be selected carefully from those that are optimal under indifference in the choice model (22.1) with $K - 1$ attributes at two levels. In particular, for optimal choice sets of size $m = 3$, the alternatives have to be selected in such a way that, with respect to the binary attributes, within every choice set two of the alternatives are identical while the third alternative is the *foldover* of the previous two, that is, all its levels are different from those of the other alternatives in the choice set (Graßhoff et al. 2013). From a practical point of view, these optimal designs do not seem to be very appealing, since the choice sets will necessarily contain alternatives that differ only in the values of the manipulator variable. If the manipulator represents, for example, a price variable, then every choice set will contain at least one alternative that is dominated in the sense that it is clearly less attractive than the other alternatives. Therefore, other models that account for dependencies between similar alternatives may be more appropriate. Alternatively, so-called Pareto optimal designs (Raghavarao and Wiley 2006) that impose restrictions on the choice sets that are permitted to be used in a choice experiment may be considered.

The discussion of the relative efficiency of designs that use choice sets of different sizes in Section 22.4.2 implies that, under the indifference assumption, choice sets of larger size are not advantageous for binary attributes, since for such attributes, $m = 2$ is the optimal size of the choice sets. The opposite is true for the model of Kanninen (2002). For example, if, in addition to the manipulator variable with values in a large interval, the model contains only one further binary attribute, then the efficiency of the optimal design with choice sets of size 3 attains remarkable 150% compared to choice sets of size 2. This means that 50% more pairs of alternatives have to be presented than choice sets of size 3 in order to estimate the model parameters with the same precision.

The efficiency gain due to larger choice sets is even more striking for the case where only a single quantitative attribute is present, so that $f'(x)\theta = \theta x$. As in every situation with a one-dimensional parameter, the optimal design requires only one experimental setting, for

which the modulus of the regression function attains its maximum. For pairs, the locally optimal choice set for a given parameter $\theta \neq 0$ can be chosen as (x_1^*, x_2^*) with alternatives satisfying $x_1^* - x_2^* = 2.40/\theta$, which results in choice probabilities of 0.92 and 0.08, respectively, and the information for this optimal pair of alternatives equals 0.44. The optimal choice set (x_1^*, x_2^*, x_3^*) of size $m = 3$ is given by $x_1^* - x_2^* = 2.65/\theta$ and $x_2^* = x_3^*$ with choice probabilities of 0.88 for x_1^* and 0.06 for x_2^* and x_3^* , respectively. The corresponding information can be calculated as 0.76, which provides an efficiency of 173% for the optimal choice set of size 3 relative to the optimal pair of alternatives.

The previous examples give some indication that, for nonzero parameter values, the use of larger choice sets may be preferable. However, it has to be taken into account that a respondent's task may become harder when choices have to be made from a larger number of alternatives. This may cause the values of the parameters to vary with the size of the choice sets. In particular, these values may tend to become smaller as m increases, which may diminish the attributes' discriminatory power.

22.5 Computational Approaches

In addition to obtaining analytic results for the multinomial logit choice model under the indifference assumption, there has been considerable interest in algorithmic approaches to generating optimal or efficient designs for the general case of nonzero parameter vectors θ . In this work, locally optimal designs for a fixed $\theta \neq \mathbf{0}$ play a relatively minor role. Instead, the parameter vectors associated with individual respondents are frequently regarded as realizations of a multivariate distribution, which is usually assumed to be normal. This distribution may be interpreted as reflecting prior knowledge as well as uncertainty about the model parameters or it may be used to model heterogeneous preferences in a population. Likewise, for most generalizations of the MNL model, such as the mixed logit model, efficient designs can only be found algorithmically.

The corresponding design algorithms usually optimize variants of the D -optimality criterion. For example, in the context of the MNL model for every given design d , the value of the determinant of the inverse of the information matrix $M_{d,\theta}$ in (22.3) is averaged over the distribution of the θ values, and designs that minimize this average are found. The resulting designs are often referred to as *Bayesian* or *semi-Bayesian designs*, although typically, a standard maximum likelihood approach is adopted for estimating the parameters. Many of the corresponding ideas were originally presented in the context of the MNL model but have subsequently been generalized to the mixed logit model.

When designs are generated algorithmically, the term *optimal design* usually needs to be interpreted as the best design the algorithm is able to find. Bearing this in mind seems to be worthwhile, for example, when drawing conclusions from simulation studies that compare different types of designs.

22.5.1 Multinomial Logit Model

Kuhfeld et al. (1994) were among the first to popularize the idea of using efficient designs in choice experiments. Based on the later challenged assumption that "an efficient design for a linear model is a good design for the multinomial logit (MNL) model," they illustrated how the OPTEX procedure in the SAS software can be used for a variety of problems.

Shortly thereafter, Huber and Zwerina (1996) presented the concept of locally optimal designs for arbitrary vectors $\theta \neq \mathbf{0}$ and a corresponding design algorithm to a wider audience in marketing. The algorithm examines certain types of changes—known as swapping and relabeling—to the attribute levels of an initial design and performs those operations that lead to the largest improvement of the D -criterion. This approach is similar in vein to standard exchange or interchange algorithms (see Chapter 21) but differs in the details of how the initial design is modified. More precisely, the changes to the attribute levels are motivated by the heuristic principle of *utility balance* that aims at making the MNL choice probabilities of the alternatives within every choice set as similar as possible. Other possibilities for generating locally optimal designs include the %ChoiceEff macro for SAS (Kuhfeld 2010, pp. 806–955) that uses a modified Fedorov algorithm and the JMP 10 software (JMP 1989–2012).

For choice sets of size $m = 2$ and some simple situations with one or two attributes, Graßhoff and Schwabe (2008) determined analytically the locally D -optimal designs for all possible values of the parameters. The results show that the structure of a locally optimal design depends strongly on the parameters for which the design is optimized. Moreover, designs that are locally optimal for some choice of the parameters can be very inefficient for other values. For this reason, the authors recommend maximin efficient designs as a robust alternative to locally optimal designs (see Chapter 20).

The work of Sándor and Wedel (2001) extends the ideas of Huber and Zwerina (1996) to a Bayesian framework by assuming that the vector θ has a distribution that may be interpreted as reflecting uncertainty about the parameters. By using the swapping and relabeling operations of Huber and Zwerina (1996) plus a new cycling operation, the authors generate D_B -optimal designs d that minimize

$$\int_{\mathbb{R}^p} \det(\mathbf{M}_{d,\theta})^{-1/p} \pi(\theta) d\theta, \quad (22.20)$$

where π is the probability density function of the distribution of θ and p is the number of model parameters.

The matrix $\mathbf{M}_{d,\theta}$ is given in (22.3). Typically, the strong assumption is made that, at the design stage, π is completely known. In practice, the integral (22.20), which is often called the D_B -error, is approximated by Monte Carlo methods, that is, by averaging values of $\det(\mathbf{M}_{d,\theta_r})^{-1/p}$ for a large number of random draws θ_r from the distribution represented by π . Based on simulation results and an empirical application, the authors conclude that the D_B -optimal designs are more efficient and have higher predictive validity than the locally optimal designs of Huber and Zwerina (1996). Some discussion on how the distribution represented by the density π should be chosen is available in Kessels et al. (2008b). A further generalization of the Bayesian approach presented by Sándor and Wedel (2005) considers the construction of so-called heterogeneous designs that can be thought of as a collection of a small number of different designs, exactly one of which is to be given to every individual respondent. This idea is related to, but not the same as, the idea of regarding the respondents as blocks of observations (Kessels et al. 2008a).

Kessels et al. (2006) adapt the Bayesian approach of Sándor and Wedel (2001) to other optimality criteria and generate the designs by means of a modified Fedorov algorithm. A faster alternative to that algorithm is presented by Kessels et al. (2009). A comparison of seven numerical methods for approximating the multidimensional integral in (22.20) can be found in Yu et al. (2010). The results show that the traditional Monte Carlo approximation of the integral can be improved by using other techniques that need fewer random

draws to achieve the same precision. Alternative Bayesian design criteria that are based on generalizations of, and alternatives to, the Fisher information matrix are compared by Yu et al. (2012).

Algorithms for choice experiments involving partial profiles have been presented by Kessels et al. (2011a). The authors generalize ideas of Green (1974) and construct the designs in two stages. First, the attributes that are held constant within the choice sets are determined and second, the levels for the nonconstant attributes are chosen. Bayesian optimal designs are found by maximizing $\int_{\mathbb{R}^p} \log(\det(\mathbf{M}_{d,\theta}))\pi(\theta)d\theta$, where $\mathbf{M}_{d,\theta}$ is the information matrix in the MNL model and π is the probability density of a multivariate normal distribution. The so-called unrestricted version (see Kessels et al. 2011a) of the algorithm has been implemented in the JMP 10 software (JMP 1989–2012) and can also be used to generate partial profile designs under the indifference assumption $\theta = \mathbf{0}$. In the latter case, a limited comparison of designs with choice sets of size $m = 2$ that were generated for $\theta = \mathbf{0}$ by the unrestricted JMP 10 algorithm with the optimal designs of Großmann et al. (2009) shows that, in several of the cases considered, the designs produced by the software possess an efficiency of less than 80% (Großmann 2013).

For some simple situations with a small number of attributes and/or levels and a few choice sets per respondent, Bayesian optimal designs have been compared with designs obtained under the indifference assumption $\theta = \mathbf{0}$ (Kessels et al. 2008b, 2011b), which in this context are referred to as *utility neutral designs*. The results show that Bayesian designs that were generated by using a normal *prior distribution* for the parameter vector θ are robust in the sense that, for the majority of vectors in the parameter space considered, these designs are more efficient than the utility neutral designs when the two types of designs are compared in terms of local non-Bayesian D -efficiency. For only a small portion of the parameter space is this ordering reversed. In a discussion of Kessels et al. (2011b), it was pointed out by Holling and Schwabe (2011) that other Bayesian design criteria may be more appropriate. They compared locally optimal designs for the means of the priors with the Bayesian designs of Kessels et al. and, depending on the variance of the underlying normal distribution of the parameters, found only small-to-moderate differences between the two types of designs.

Moreover, for a model with six two-level attributes, Kessels et al. (2011b) compared the precision of estimates obtained from a Bayesian and a utility neutral design for eight choice sets of size 2. For both designs, 100 samples of responses from 200 respondents were simulated and parameters were estimated for each of the 100 data sets. Data were simulated from a population in which the parameters for three of the attributes were equal to zero, whereas the parameter for each of the remaining three attributes was equal to -0.8 . These specifications mean that respondents make their choices by focusing on only three attributes and these three attributes are equally important. For this choice of parameters, two of the choice sets of the utility neutral design contained an alternative that had a probability of being chosen, which was greater than 0.99. In the simulation study this led to five outliers in the reported estimates with the effect that for each of the nonzero parameters, the average of the one hundred estimates was biased. When trying to replicate these results, Holling and Schwabe (2011) found that outlying estimates were caused by data sets for which the parameters were not estimable. After excluding such estimates the averages of the parameters were close to the true values and hence did not indicate any bias.

Whether the conclusions from simulation studies such as Kessels et al. (2008b, 2011b) remain valid in more representative design comparisons and for less extreme choices of the parameter vector seems to await further clarification. Also, it would seem to be worthwhile to include optimal analytic designs that are constructed under the indifference assumption

in such studies. Probably, an even more important question is whether the promise of optimal designs, whether Bayesian or utility neutral, to give more precise estimates can be corroborated in empirical studies (Großmann et al. 2005; Bliemer and Rose 2011).

22.5.2 Mixed Logit Model

Sándor and Wedel (2002) compared locally optimal designs for the cross-sectional mixed logit model from Section 22.2.2 with locally optimal designs for the MNL model. Designs for the mixed logit model were generated by using a modified version of the algorithm of Sándor and Wedel (2001) that minimizes the determinant of the inverse of the information matrix in Equation (5) of Sándor and Wedel (2002). Based on simulation results, the authors concluded that the optimal mixed logit designs give more precise parameter estimates and have higher predictive validity.

The work of Yu et al. (2009) extends the ideas of Sándor and Wedel (2001) to a so-called semi-Bayesian approach for the cross-sectional mixed logit model. Following Sándor and Wedel (2002), it is assumed that the parameter vectors θ have a multivariate normal distribution whose mean vector and covariance matrix are to be estimated efficiently. Basically, this means that instead of minimizing the integral in (22.20), a similar integral given in Equation (8) of Yu et al. (2009) is minimized. For computational reasons, only a prior distribution for the mean vector of the normal distribution is assumed, whereas the covariance matrix is treated as known. The authors propose a design algorithm that combines methods for facilitating the approximate computation of integrals, such as Halton draws (Train 2003) and systematic sampling ideas of Kessels et al. (2009), with a coordinate exchange procedure. The resulting semi-Bayesian D -optimal designs are compared with locally optimal designs for the mixed logit and the MNL model, with Bayesian designs for the MNL model, as well as with so-called nearly orthogonal designs and with fully Bayesian mixed logit designs.

For the panel mixed logit model in Section 22.2.2, only little design work has been done so far. Bliemer and Rose (2010) compared locally optimal designs for the cross-sectional and the panel mixed logit model with locally optimal designs for the MNL model in three case studies. They found that efficient designs for the cross-sectional mixed logit model are usually not efficient for estimating the parameters of the panel mixed logit model and vice versa. Interestingly, locally optimal designs for the MNL model are relatively efficient for estimating the panel version of the mixed logit model. Furthermore, the authors showed that the misspecification of the prior parameters that were used at the design stage could substantially reduce the efficiency of the designs.

An adaptive Bayesian approach for sequentially generating designs for the panel mixed logit model that are customized to the preferences of individual respondents has been presented by Yu et al. (2011). The approach generalizes the idea of heterogeneous designs in Sándor and Wedel (2005) and represents an alternative to the polyhedral method of Toubia et al. (2004, 2007). Simulation results for the polyhedral approach indicate that using a respondent's earlier responses to adapt the subsequent choice sets to that person's preferences seems to work well when the components of the respondent's parameter vector θ are relatively large and, hence, the response accuracy is high (Toubia et al. 2004), but less well when the response accuracy is low. The method of Yu et al. (2011) appears to be superior in this respect, although a direct comparison of the sequential Bayesian and the polyhedral approach has not yet been reported. Customization of designs within the panel mixed logit model is also considered by Crabbe and Vandebroek (2012),

who model the parameter vector of every individual respondent as a function of several covariates.

22.6 Ranking-Based Methods and Best–Worst Scaling

Selecting the most preferred alternative from a choice set does not give any information about the preference ordering of the remaining alternatives in the set. More information can be obtained by asking the respondents to rank order all or, at least, the most attractive alternatives in a choice set. We illustrate this approach for the situation where the top two alternatives have to be selected from a larger choice set of size $m \geq 3$ and ranked according to preference. For $m = 3$, this task is equivalent to a complete ranking of the three alternatives in the choice set.

The *rank-ordered logit model* mentioned in Section 22.2.3 regards every ranking as the outcome of a series of multinomial choices from a shrinking choice set. More precisely, given a choice set C , first the most preferred alternative is assumed to be chosen as in the multinomial logit model with probability

$$p_{\theta}(i; C) = \frac{e^{f'(x_i)\theta}}{\sum_{k=1}^m e^{f'(x_k)\theta}}.$$

The second best alternative is then chosen from the remaining alternatives with probability

$$p_{\theta}(j; C_{-i}) = \frac{e^{f'(x_j)\theta}}{\sum_{k=1, k \neq i}^m e^{f'(x_k)\theta}},$$

according to a multinomial logit model for the reduced choice set C_{-i} of size $m - 1$ obtained by removing x_i from C , but with the same parameters θ as in the first selection step.

Hence, the probability $p_{\theta}((i, j); C)$ of ranking alternative x_i highest and x_j second within the choice set C can be factorized as

$$p_{\theta}((i, j); C) = p_{\theta}(i; C) \times p_{\theta}(j; C_{-i}).$$

Denoting the design matrices corresponding to C and C_{-i} by X and X_{-i} , respectively, the information matrix $M_{C, \theta}^{(\text{rank})}$ for a choice set C in this experiment can be readily computed to be equal to

$$M_{C, \theta}^{(\text{rank})} = X'(\text{diag}(p_{\theta}) - p_{\theta}p'_{\theta})X + \sum_{i=1}^m p_{\theta}(i; C)X'_{-i}(\text{diag}(p_{\theta, -i}) - p_{\theta, -i}p'_{\theta, -i})X_{-i},$$

where $p_{\theta} = (p_{\theta}(1; C), \dots, p_{\theta}(m; C))'$ and $p_{\theta, -i}$ is the corresponding vector of length $m - 1$ that contains the choice probabilities $p_{\theta}(j; C_{-i})$ for $j \neq i$.

The first term on the right-hand side of the expression for $M_{C, \theta}^{(\text{rank})}$ is the information $M_{C, \theta}$ for a standard multinomial choice experiment given in (22.3), and the second

term represents the additional information that can be attributed to the ranking task (Vermeulen et al. 2011).

Equivalently, the information matrix can be expressed as a weighted sum of the information contained in the comparisons of the pairs of alternatives within the choice set, so that

$$\begin{aligned} M_{C,\theta}^{(\text{rank})} &= M_{C,\theta} + \sum_{i < j} (1 - p_{\theta}(i; C) - p_{\theta}(j; C)) p_{\theta}(i; C_{-j}) p_{\theta}(j; C_{-i}) M_{\{x_i, x_j\}, \theta} \\ &= M_{C,\theta} + \sum_{i < j} \frac{(e^{f'(x_i)\theta} + e^{f'(x_j)\theta})^2 \cdot \sum_{i \neq k \neq j} e^{f'(x_k)\theta}}{\sum_{k=1}^m e^{f'(x_k)\theta} \cdot \sum_{k \neq i} e^{f'(x_k)\theta} \cdot \sum_{k \neq j} e^{f'(x_k)\theta}} M_{\{x_i, x_j\}, \theta}, \end{aligned}$$

where $M_{\{x_i, x_j\}, \theta}$ is the matrix in (22.13). Under the indifference assumption $\theta = \mathbf{0}$ and by using (22.14), this expression can be simplified further to

$$M_{C,0}^{(\text{rank})} = \left(\frac{4}{m^2} + \frac{4(m-2)}{m(m-1)^2} \right) \sum_{i < j} M_{\{x_i, x_j\}, 0} = \left(2 - \frac{1}{(m-1)^2} \right) M_{C,0},$$

which shows that, compared to a standard multinomial choice experiment, the ranking task provides 75% more information for choice sets of size $m = 3$ and about twice as much information if the size m of the choice set is large.

Vermeulen et al. (2011) adapt the Bayesian approach of Sándor and Wedel (2001) to the rank-ordered logit model. Designs are generated by using a modified version of the algorithm in Kessels et al. (2009). The authors present simulation results for complete and partial rankings for a model with five attributes, with two at two levels and the other three at three levels, and a design with nine choice sets of size four. For the case of complete rankings, Bayesian D -optimal designs for the rank-ordered logit model lead to only slight improvements in terms of prediction and estimation accuracy over Bayesian D -optimal designs for the multinomial logit model that minimize (22.20). For partial rankings of two out of the four alternatives in each set, this ordering of the designs is reversed. The authors note that similar results were obtained for some additional models and scenarios.

Similar to the rankings-based approach, *best-worst scaling* aims to extract more information from every individual choice set. To this end, respondents are asked to choose the most and the least preferred alternative from every choice set. A model for this situation that is closely related to the MNL model is known as the *maximum difference* or *maxdiff model* (Marley and Louviere 2005; Marley and Pihlens 2012). Vermeulen et al. (2010) use a coordinate exchange algorithm for generating Bayesian D -optimal designs for the maxdiff model and perform a simulation study that essentially parallels the setup in Vermeulen et al. (2011) and leads to similar results in that the differences between the Bayesian optimal designs for the maxdiff and the MNL model are relatively small. For another variant of best-worst scaling, known as *attribute-level best-worst choice experiments*, Street and Knox (2012) show that, under the indifference assumption $\theta = \mathbf{0}$, fractional factorial designs of resolution III and V perform equally well as the full factorial for estimating main effects or for estimating main effects and two-factor interactions, respectively.

22.7 Miscellaneous Problems

This section briefly reviews some contributions that have addressed more specialized topics. In particular, we consider model uncertainty at the design stage, designs for situations where the order of presentation of the alternatives may matter, and choice experiments that take into account that sometimes a choice set may not contain a single “best” alternative.

Most of the current work on the design of choice experiments is based on the a priori assumption that either only the main effects of the attributes or the main effects and all two-factor interactions are to be estimated. Little work has considered how choice experiments should be designed when there is uncertainty at the design stage about whether or not some or all interaction effects should be included in the choice model. For this type of model uncertainty, Yu et al. (2008) propose a Bayesian composite design criterion and a corresponding algorithm for generating designs that perform well irrespective of whether or not the interaction effects are present. A complementary approach presented by Raghavarao and Wiley (2006) outlines a sequential strategy, where initially choice sets are generated that allow the estimation of main effects and the performance of a lack-of-fit test to find out whether a main-effects-only model is appropriate. If the test is not significant, then conclusions are based on the main effects. Otherwise, additional choice sets are presented from which two-factor interactions can be estimated and another lack-of-fit test is performed in order to check if interactions involving three attributes need to be included. Depending on the outcome of the test, more data are collected so that, if necessary, final inferences can be based on the main effects and the interactions of up to three attributes.

In some applications, the alternatives in a choice set need to be considered sequentially, for example, when the alternatives need to be tasted or tried out. In such cases, the order in which the alternatives are presented may have an effect on the responses. Davidson and Beaver (1977) proposed a generalization of the Bradley–Terry model for paired comparisons that contains additional parameters representing potential order effects. For a simple variant of the model where the order effect is assumed to be constant across all pairs, Goos and Großmann (2011) show that under the indifference assumption $\theta = \mathbf{0}$, special versions of designs of Graßhoff et al. (2004) remain optimal for estimating main effects. They also list practical designs with up to 40 pairs for the case where the attributes have two or three levels. Bush et al. (2012) extend the model of Davidson and Beaver (1977) to experiments with choice sets of size $m > 2$ and adapt the results of Burgess and Street (2005), essentially confirming that under indifference, and subject to some additional restrictions, D -optimal designs for the MNL model without order effects remain optimal when position effects are taken into account. Note that for models that incorporate order or position effects, the order of the alternatives within the choice sets is important.

The choice experiments we have considered employ a forced-choice format, where exactly one alternative must be selected from every choice set. In order to make the choice task more realistic, in practice, sometimes a variant is used, where every set contains a so-called no-choice option. By selecting this alternative, respondents are able to circumvent making a definite choice when, for example, none of the alternatives in a choice set offers a sufficient amount of utility. Another type of choice task uses choice sets, all of which contain a common *base option* or *reference alternative*.

For the MNL model under indifference, Street and Burgess (2004b) show that, for the purpose of estimating main effects in a setting with a reference alternative and choice

sets of size $m = 2$, all choice designs of the same size that are based on fractional factorial designs of resolution III are equally efficient regardless of the particular fractional factorial chosen and also irrespective of which combination of attribute levels from the fractional factorial is used as the reference alternative. Moreover, for experiments with a no-choice option and choice sets of arbitrary size, they prove that optimal designs for estimating main effects from forced choices remain optimal if a no-choice option is adjoined to every choice set. Experiments with a base option are also discussed by Sándor and Wedel (2002) in their work on locally optimal designs for the cross-sectional mixed logit model.

Vermeulen et al. (2008) consider variants of the MNL and the nested logit model with no-choice options. They adapt the Bayesian approach of Sándor and Wedel (2001) to generate designs that are optimal for estimating the parameters for the attribute levels but not those for the no-choice option. Bayesian optimal designs for the MNL model with and without a no-choice option and for the nested logit model are then compared on simulated data from an extended no-choice MNL and the nested logit model. From the numerical results, the authors conclude that the finding of Street and Burgess (2004b) that optimal designs for the MNL model without a no-choice option can also be used for models that include a no-choice option can be generalized to Bayesian optimal designs. Although Vermeulen et al. (2008) seem to indicate that a similar statement also holds for the nested logit model, in another paper, the authors recommend taking the presence of a no-choice option into account at the design stage whenever prior information about the respondents' preferences is available (Goos et al. 2010).

Another possibility to account for the fact that a choice set may not contain a single "best" alternative is to allow tied responses by giving respondents the opportunity to designate some alternatives as being equally attractive. Bush et al. (2010) consider extensions of the Bradley–Terry and the MNL model for this situation and show that, under the indifference assumption, optimal designs for the forced-choice versions of the models remain optimal when ties are permitted.

22.8 Concluding Remarks

Over the past decade, the design of choice experiments has attracted a lot of interest as can be seen from the volume of publications that have been published since the reviews of Großmann et al. (2002) and Louviere et al. (2004). Much progress has been made both for analytic designs under the indifference assumption as well as with regard to computational approaches and, in particular, Bayesian designs. Despite these advances, it appears that the area would benefit from some attempts to consolidate what has been achieved.

For practitioners, it is probably already difficult to identify from the variety of existing approaches the method that suits their needs best. For example, it is frequently claimed that Bayesian designs represent the state of the art and that this approach is superior to the use of designs that are generated under the indifference assumption. While this is probably true in many situations, often, the supporting simulation studies compare the Bayesian designs with suboptimal orthogonal or utility neutral designs instead of using designs that have been proved to be optimal under indifference. Typically, these studies consider designs with relatively small numbers of choice sets, and it would be interesting to investigate if similar results are obtained for designs with, say, around 40 choice sets, which is not an uncommon number in practical applications (Louviere et al. 2005). Moreover, several of the

simulation studies reported to date have used similar specifications of the prior distribution for generating Bayesian optimal designs, and, for some, one may wonder whether the assumed covariance structure is consistent with the modeling assumptions, in particular, with the coding of the attribute levels. Although the practice of using similar prior specifications facilitates comparisons across different studies, it may also limit the generalizability of the results.

The results of Graßhoff and Schwabe (2008) show that, in principle, any design that is locally optimal for some values of the model parameters can be very inefficient for other values in the parameter space and this applies, in particular, to designs that are generated under the indifference assumption. Anticipating when this happens and quantifying the efficiency loss appears to be difficult. As suggested by these authors, the use of maximin efficient designs appears to be another worthwhile approach for dealing with parameter uncertainty. Corresponding results for a related logistic model are presented by Graßhoff et al. (2012).

An alternative reading of the simulation results for Bayesian designs may suggest that the gains in efficiency or estimation accuracy are often not enormous when compared to utility neutral designs. Given that in order to generate Bayesian designs specialized software is needed, practitioners will need to decide whether the benefits that may be obtained from Bayesian designs do justify the additional efforts. Often, the required software is not publicly available. Simulation studies also sidestep the crucial question of how the required prior information can be reliably obtained in practice. What appears to be a more important lesson from studies reported by Vermeulen et al. (2008), Bliemer et al. (2009) and Bliemer and Rose (2010) is that the performance of efficient designs strongly depends on whether or not the model that is used for optimizing the design corresponds to the mechanism that generates the data.

In order to clarify some of these issues and to be able to distill practical recommendations perhaps a way forward could be to create a repository that would provide easy access to tables of optimal analytic designs and to source code of algorithms for generating, for example, Bayesian optimal designs. Such a resource would enable researchers to replicate and extend simulation studies from which, hopefully, a more complete picture of what works best and when would emerge.

References

- Alriksson, S. and Öberg, T. (2008), Conjoint analysis for environmental evaluation: A review of methods and applications, *Environmental Science and Pollution Research*, 15, 244–257.
- Bech, M. and Gyrd-Hansen, D. (2005), Effects coding in discrete choice experiments, *Health Economics*, 14, 1079–1083.
- Bliemer, M. C. J. and Rose, J. M. (2010), Construction of experimental designs for mixed logit models allowing for correlation across choice observations, *Transportation Research Part B*, 44, 720–734.
- Bliemer, M. C. J. and Rose, J. M. (2011), Experimental design influences on stated choice outputs: An empirical study in air travel choice, *Transportation Research Part A*, 45, 63–79.
- Bliemer, M. C. J., Rose, J. M., and Hensher, D. A. (2009), Efficient stated choice experiments for estimating nested logit models, *Transportation Research Part B*, 43, 19–35.
- Bradley, R. A. (1976), Science, statistics, and paired comparisons, *Biometrics*, 32, 213–232.

- Bradley, R. A. and El-Helbawy, A. T. (1976), Treatment contrasts in paired comparisons: Basic procedures with application to factorials, *Biometrika*, 63, 255–262.
- Bradley, R. A. and Terry, M. E. (1952), Rank analysis of incomplete block designs: I. The method of paired comparisons, *Biometrika*, 39, 324–345.
- Burgess, L. and Street, D. J. (2003), Optimal designs for 2^k choice experiments, *Communications in Statistics—Theory and Methods*, 32, 2185–2206.
- Burgess, L. and Street, D. J. (2005), Optimal designs for choice experiments with asymmetric attributes, *Journal of Statistical Planning and Inference*, 134, 288–301.
- Burgess, L. and Street, D. J. (2006), The optimal size of choice sets in choice experiments, *Statistics*, 40, 507–515.
- Bush, S., Burgess, L., and Street, D. (2010), Optimal designs for stated choice experiments that incorporate ties, *Journal of Statistical Planning and Inference*, 140, 1712–1718.
- Bush, S., Street, D. J., and Burgess, L. (2012), Optimal designs for stated choice experiments that incorporate position effects, *Communications in Statistics—Theory and Methods*, 41, 1771–1795.
- Caussade, S., Ortúzar, J. de D., Rizzi, L. I., and Hensher, D. A. (2005), Assessing the influence of design dimensions on stated choice experiments, *Transportation Research, Part B*, 39, 621–640.
- Chapman, R. G. and Staelin, R. (1982), Exploiting rank ordered choice set data within the stochastic utility model, *Journal of Marketing Research*, 19, 288–301.
- Chrzan, K. (2010), Using partial profile choice experiments to handle large numbers of attributes, *International Journal of Market Research*, 52, 827–840.
- Crabbe, M. and Vandebroek, M. (2012), Improving the efficiency of individualized designs for the mixed logit choice model by including covariates, *Computational Statistics and Data Analysis*, 56, 2059–2072.
- Das, A. and Dey, A. (2004), Optimal main effect plans with non-orthogonal blocks, *Sankhya*, 66, 378–384.
- David, H. A. (1988), *The Method of Paired Comparisons*, 2nd edn., London, U.K. Griffin.
- Davidson, R. R. and Beaver, R. J. (1977), On extending the Bradley-Terry model to incorporate within-pair order effects, *Biometrics*, 33, 693–702.
- Davidson, R. R. and Farquhar, P. H. (1976), A bibliography on the method of paired comparisons, *Biometrics*, 32, 241–252.
- De Bekker-Grob, E. W., Ryan, M., and Gerard, K. (2012), Discrete choice experiments in health economics: A review of the literature, *Health Economics*, 21, 145–172.
- El-Helbawy, A. T. (1984), Asymptotic relative efficiency of designs for factorial paired comparison experiments, *Journal of Statistical Planning and Inference*, 10, 105–113.
- El-Helbawy, A. T. and Ahmed, E. A. (1984), Optimal-design results for 2^n -factorial paired comparison experiments, *Communications in Statistics—Theory and Methods*, 13, 2827–2845.
- El-Helbawy, A. T., Ahmed, E. A., and Alharbey, A. H. (1994), Optimal designs for asymmetrical factorial paired comparison experiments, *Communications in Statistics—Simulation and Computation*, 23, 663–681.
- El-Helbawy, A. T. and Bradley, R. A. (1978), Treatment contrasts in paired comparisons: Large-sample results, applications, and some optimal designs, *Journal of the American Statistical Association*, 73, 831–839.
- Goos, P. (2002), *The Optimal Design of Blocked and Split-Plot Experiments*, New York: Springer.
- Goos, P. and Großmann, H. (2011), Optimal design of factorial paired comparison experiments in the presence of within-pair order effects, *Food Quality and Preference*, 22, 198–204.
- Goos, P., Vermeulen, B., and Vandebroek, M. (2010), D -optimal conjoint choice designs with no-choice options for a nested logit model, *Journal of Statistical Planning and Inference*, 140, 851–861.
- Graßhoff, U., Großmann, H., Holling, H., and Schwabe, R. (2003), Optimal paired comparison designs for first-order interactions, *Statistics*, 37, 373–386.
- Graßhoff, U., Großmann, H., Holling, H., and Schwabe, R. (2004), Optimal designs for main effects in linear paired comparison models, *Journal of Statistical Planning and Inference*, 126, 361–376.

- Graßhoff, U., Großmann, H., Holling, H., and Schwabe, R. (2007), Design optimality in multi-factor generalized linear models in the presence of an unrestricted quantitative factor, *Journal of Statistical Planning and Inference*, 137, 3882–3893.
- Graßhoff, U., Großmann, H., Holling, H., and Schwabe, R. (2013), Optimal design for discrete choice experiments, *Journal of Statistical Planning and Inference*, 143, 167–175.
- Graßhoff, U., Holling, H., and Schwabe, R. (2012), Optimal designs for the Rasch model, *Psychometrika*, 77, 710–723.
- Graßhoff, U. and Schwabe, R. (2008), Optimal design for the Bradley-Terry paired comparison model, *Statistical Methods and Applications*, 17, 275–289.
- Green, P. E. (1974), On the design of choice experiments involving multifactor alternatives, *Journal of Consumer Research*, 1, 61–68.
- Green, P. E., Carroll, J. D., and Carmone, F. J. (1976), Superordinate factorial designs in the analysis of consumer judgments, *Journal of Business Research*, 4, 281–295.
- Green, P. E., Carroll, J. D., and Carmone, F. J. (1978), Some new types of fractional factorial designs for marketing experiments, in *Research in Marketing*, ed. Sheth, J. N., Greenwich, CT: JAI Press, vol. 1, pp. 99–122.
- Green, P. E. and Srinivasan, V. (1978), Conjoint analysis in consumer research: Issues and outlook, *Journal of Consumer Research*, 5, 103–123.
- Green, P. E. and Srinivasan, V. (1990), Conjoint analysis in marketing: New developments with implications for research and practice, *Journal of Marketing*, 54, 3–19.
- Großmann, H. (2013), Differences between analytic and algorithmic choice designs for pairs of partial profiles, in *mODa 10—Advances in Model-Oriented Design and Analysis*, eds. Uciński, D., Atkinson, A. C., and Patan, M., Cham, Switzerland: Springer, pp. 125–133.
- Großmann, H., Graßhoff, U., and Schwabe, R. (2009), Approximate and exact optimal designs for paired comparisons of partial profiles when there are two groups of factors, *Journal of Statistical Planning and Inference*, 139, 1171–1179.
- Großmann, H., Graßhoff, U., and Schwabe, R. (2014), A catalogue of designs for partial profiles in paired comparison experiments with three groups of factors, *Statistics*, 48, 1268–1281.
- Großmann, H., Holling, H., Brocke, M., Graßhoff, U., and Schwabe, R. (2005), On the empirical relevance of optimal designs for the measurement of preferences, in *Applied Optimal Designs*, eds. Berger, M. and Wong, W. K., Chichester, U.K.: Wiley, pp. 45–65.
- Großmann, H., Holling, H., Graßhoff, U., and Schwabe, R. (2006), Optimal designs for asymmetric linear paired comparisons with a profile strength constraint, *Metrika*, 64, 109–119.
- Großmann, H., Holling, H., Graßhoff, U., and Schwabe, R. (2007), A comparison of efficient designs for choices between two options, in *mODa 8—Advances in Model-Oriented Design and Analysis*, eds. López-Fidalgo, J., Rodríguez-Díaz, J. M., and Torsney, B., Heidelberg, Germany: Physica-Verlag, pp. 83–90.
- Großmann, H., Holling, H., and Schwabe, R. (2002), Advances in optimum experimental design for conjoint analysis and discrete choice models, in *Econometric Models in Marketing. Advances in Econometrics*, eds. Franses, P. H. and Montgomery, A. L., Amsterdam, the Netherlands: JAI Press, vol. 16, pp. 93–117.
- Großmann, H., Schwabe, R., and Gilmour, S. G. (2012), Designs for first-order interactions in paired comparison experiments with two-level factors, *Journal of Statistical Planning and Inference*, 142, 2395–2401.
- Hedayat, A. S., Sloane, N. J. A., and Stufken, J. (1999), *Orthogonal Arrays: Theory and Applications*, New York: Springer.
- Holling, H. and Schwabe, R. (2011), ‘The usefulness of Bayesian optimal designs for discrete choice experiments’ by Roselinde Kessels, Bradley Jones, Peter Goos and Martina Vandebroek, *Applied Stochastic Models in Business and Industry*, 27, 189–192.
- Hoyos, D. (2010), The state of the art of environmental valuation with discrete choice experiments, *Ecological Economics*, 69, 1595–1603.
- Huber, J. and Zwerina, K. (1996), The importance of utility balance in efficient choice designs, *Journal of Marketing Research*, 33, 307–317.

- JMP (1989–2012), Version 10, Cary, NC: SAS Institute Inc.
- John, J. A. and Williams, E. R. (1995), *Cyclic and Computer Generated Designs*, 2nd edn., London, U.K. Chapman & Hall.
- Kanninen, B. J. (2002), Optimal design for multinomial choice experiments, *Journal of Marketing Research*, 39, 214–227.
- Kessels, R., Goos, P., and Vandebroek, M. (2006), A comparison of criteria to design efficient choice experiments, *Journal of Marketing Research*, 43, 409–419.
- Kessels, R., Goos, P., and Vandebroek, M. (2008a), Optimal designs for conjoint experiments, *Computational Statistics and Data Analysis*, 52, 2369–2387.
- Kessels, R., Jones, B., and Goos, P. (2011a), Bayesian optimal designs for discrete choice experiments with partial profiles, *Journal of Choice Modelling*, 4, 52–74.
- Kessels, R., Jones, B., Goos, P., and Vandebroek, M. (2008b), Recommendations on the use of Bayesian optimal designs for choice experiments, *Quality and Reliability Engineering International*, 24, 737–744.
- Kessels, R., Jones, B., Goos, P., and Vandebroek, M. (2009), An efficient algorithm for constructing Bayesian optimal choice designs, *Journal of Business and Economic Statistics*, 27, 279–291.
- Kessels, R., Jones, B., Goos, P., and Vandebroek, M. (2011b), The usefulness of Bayesian optimal designs for discrete choice experiments, *Applied Stochastic Models in Business and Industry*, 27, 173–188.
- Kharaghani, H. and Tayfeh-Rezaie, B. (2005), A Hadamard matrix of order 428, *Journal of Combinatorial Designs*, 13, 435–440.
- Kiefer, J. (1958), On the nonrandomized optimality and randomized nonoptimality of symmetrical designs, *Annals of Mathematical Statistics*, 29, 675–699.
- Kiefer, J. and Wolfowitz, J. (1960), The equivalence of two extremum problems, *Canadian Journal of Mathematics*, 12, 363–366.
- Kuhfeld, W. F. (2010), Marketing research methods in SAS: Experimental design, choice, conjoint and graphical techniques. SAS 9.2 edition, Technical Report MR-2010, SAS Institute Inc., Cary, NC.
- Kuhfeld, W. F., Tobias, R. D., and Garratt, M. (1994), Efficient experimental design with marketing research applications, *Journal of Marketing Research*, 31, 545–557.
- Lancsar, E., Louviere, J., Donaldson, C., Currie, G., and Burgess, L. (2013), Best worst discrete choice experiments in health: Methods and an application, *Social Science & Medicine*, 76, 74–82.
- Louviere, J. J., Eagle, T. C., and Cohen, S. H. (2005), Conjoint analysis: Methods, myths and much more, Technical Report CenSoC Working Paper No. 05-001, Centre for the Study of Choice, Faculty of Business, University of Technology, Sydney, New South Wales, Australia.
- Louviere, J. J., Hensher, D. A., and Swait, J. D. (2000), *Stated Choice Methods: Analysis and Application*, Cambridge, U.K.: Cambridge University Press.
- Louviere, J. J., Street, D. J., and Burgess, L. (2004), A 20+ years' retrospective on choice experiments, in *Market Research and Modeling: Progress and Prospects*, eds. Wind, Y. and Green, P. E., Boston, MA: Kluwer Academic Publishers, pp. 201–214.
- Louviere, J. J. and Woodworth, G. (1983), Design and analysis of simulated consumer choice or allocation experiments: An approach based on aggregate data, *Journal of Marketing Research*, 20, 350–367.
- Luce, R. D. (1959), *Individual Choice Behavior: A Theoretical Analysis*, New York: Wiley.
- Marley, A. A. J. and Louviere, J. J. (2005), Some probabilistic models of best, worst, and best-worst choices, *Journal of Mathematical Psychology*, 49, 464–480.
- Marley, A. A. J. and Pihlens, D. (2012), Models of best-worst choice and ranking among multiattribute options (profiles), *Journal of Mathematical Psychology*, 56, 24–34.
- Mukerjee, R., Dey, A., and Chatterjee, K. (2002), Optimal main effect plans with non-orthogonal blocking, *Biometrika*, 89, 225–229.
- Raghavarao, D. and Wiley, J. B. (2006), Design strategies for sequential choice experiments involving economic alternatives, *Journal of Statistical Planning and Inference*, 136, 3287–3306.

- Rao, C. R. (1973), *Linear Statistical Inference and Its Applications*, 2nd edn., New York: Wiley.
- Rose, J. M. and Bliemer, M. C. J. (2009), Constructing efficient stated choice experimental designs, *Transport Reviews*, 29, 587–617.
- Sándor, Z. and Wedel, M. (2001), Designing conjoint choice experiments using managers' prior beliefs, *Journal of Marketing Research*, 38, 430–444.
- Sándor Z. and Wedel, M. (2002), Profile construction in experimental choice designs for mixed logit models, *Marketing Science*, 21, 455–475.
- Sándor Z. and Wedel, M. (2005), Heterogeneous conjoint choice designs, *Journal of Marketing Research*, 42, 210–218.
- Schwabe, R. (1996a), Optimal design of experiments for multivariate response in two-factor linear models, in *Multidimensional Statistical Analysis and Theory of Random Matrices: Proceedings of the Sixth Eugene Lukacs Symposium*, eds. Gupta, A. K. and Girko, V. L., Utrecht, the Netherlands: VSP, pp. 235–242.
- Schwabe, R. (1996b), *Optimum Designs for Multi-Factor Models*, New York: Springer.
- Shah, K. R. and Sinha, B. K. (1989), *Theory of Optimal Designs*, New York: Springer.
- Street, D. J., Bunch, D. S., and Moore, B. J. (2001), Optimal designs for 2^k paired comparison experiments, *Communications in Statistics—Theory and Methods*, 30, 2149–2171.
- Street, D. J. and Burgess, L. (2004a), Optimal and near-optimal pairs for the estimation of effects in 2-level choice experiments, *Journal of Statistical Planning and Inference*, 118, 185–199.
- Street, D. J. and Burgess, L. (2004b), Optimal stated preference choice experiments when all choice sets contain a specific option, *Statistical Methodology*, 1, 37–45.
- Street, D. J. and Burgess, L. (2007), *The Construction of Optimal Stated Choice Experiments: Theory and Methods*, Hoboken, NJ: Wiley.
- Street, D. J. and Burgess, L. (2008), Some open combinatorial problems in the design of stated choice experiments, *Discrete Mathematics*, 308, 2781–2788.
- Street, D. J. and Burgess, L. (2012), Designs for choice experiments for the multinomial logit model, in *Design and Analysis of Experiments: Special Designs and Applications*, ed. Hinkelmann, K., New York: Wiley, pp. 331–378.
- Street, D. J., Burgess, L., and Louviere, J. J. (2005), Quick and easy choice sets: Constructing optimal and nearly optimal stated choice experiments, *International Journal of Research in Marketing*, 22, 459–470.
- Street, D. J. and Knox, S. A. (2012), Designing for attribute-level best-worst choice experiments, *Journal of Statistical Theory and Practice*, 6, 363–375.
- Thurstone, L. L. (1927), The method of paired comparisons for social values, *Journal of Abnormal and Social Psychology*, 21, 384–400.
- Thurstone, L. L. (1928), Attitudes can be measured, *American Journal of Sociology*, 33, 529–554.
- Toubia, O., Hauser, J., and Garcia, R. (2007), Probabilistic polyhedral methods for adaptive choice-based conjoint analysis: Theory and application, *Marketing Science*, 26, 596–610.
- Toubia, O., Hauser, J. R., and Simester, D. I. (2004), Polyhedral methods for adaptive choice-based conjoint analysis, *Journal of Marketing Research*, 41, 116–131.
- Train, K. E. (2003), *Discrete Choice Methods with Simulation*, Cambridge, U.K.: Cambridge University Press.
- Van Berkum, E. E. M (1987), *Optimal Paired Comparisons Designs for Factorial Experiments*, vol. 31 of CWI Tracts, Amsterdam, the Netherlands: Centrum voor Wiskunde en Informatica.
- Vermeulen, B., Goos, P., and Vandebroek, M. (2008), Models and optimal designs for conjoint choice experiments including a no-choice option, *International Journal of Research in Marketing*, 25, 94–103.
- Vermeulen, B., Goos, P., and Vandebroek, M. (2010), Obtaining more information from conjoint experiments by best-worst choices, *Computational Statistics and Data Analysis*, 54, 1426–1433.
- Vermeulen, B., Goos, P., and Vandebroek, M. (2011), Rank-order choice-based conjoint choice experiments: Efficiency and design, *Journal of Statistical Planning and Inference*, 141, 2519–2531.
- Yu, J., Goos, P., and Vandebroek, M. (2008), Model-robust design of conjoint choice experiments, *Communications in Statistics—Simulation and Computation*, 37, 1603–1621.

- Yu, J., Goos, P., and Vandebroek, M. (2009), Efficient conjoint choice designs in the presence of respondent heterogeneity, *Marketing Science*, 28, 122–135.
- Yu, J., Goos, P., and Vandebroek, M. (2010), Comparing different sampling schemes for approximating the integrals involved in the efficient design of stated choice experiments, *Transportation Research Part B*, 44, 1268–1289.
- Yu, J., Goos, P., and Vandebroek, M. (2011), Individually adapted sequential Bayesian conjoint-choice designs in the presence of consumer heterogeneity, *International Journal of Research in Marketing*, 28, 378–388.
- Yu, J., Goos, P., and Vandebroek, M. (2012), A comparison of different Bayesian design criteria for setting up stated preference studies, *Transportation Research Part B*, 46, 789–807.

23

Plate Designs in High-Throughput Screening Experiments for Drug Discovery

Xianggui Qu (Harvey) and Stanley Young

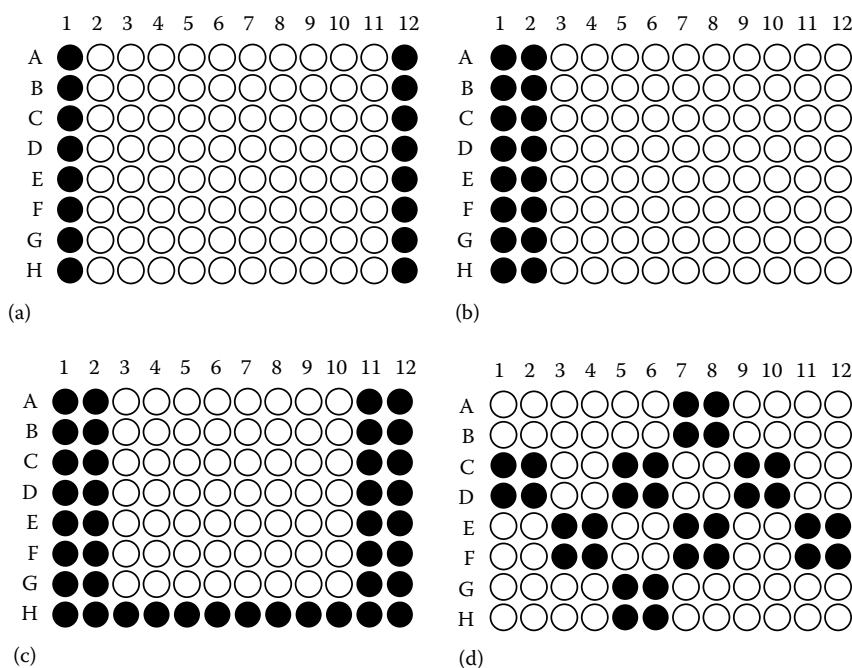
CONTENTS

23.1 Introduction.....	833
23.2 Saturated Row–Column Designs.....	838
23.3 Design and Analysis of HTS Experiments with One Replicated Treatment.....	840
23.3.1 Designs with One Replicated (Control) Treatment.....	840
23.3.2 Analysis of Row–Column Designs with One Replicated Treatment.....	842
23.4 Designs of HTS Experiments with Multiple Replicated Treatments.....	849
23.4.1 (<i>M, S</i>)-Optimal Designs.....	849
23.4.2 Designs with a Cyclic Structure.....	852
23.5 Conclusions.....	854
References.....	854

23.1 Introduction

High-throughput screening (HTS) is a large-scale process that screens hundreds of thousands to millions of compounds in order to identify biologically active compounds as candidates for further validation and confirmation. HTS started and developed as an early discovery platform in the pharmaceutical industry in the late 1980s to mid-1990s. The goal of HTS is to generate chemical structures that will lead to drug discovery through testing chemicals for their biological activity against target molecules (An and Tolliday 2009). Devlin (1997) and Hüser et al. (2006) provide more details on HTS in drug discovery. Recently, advances in HTS technologies in drug discovery have been used in other research fields. For example, approaches developed for small-molecule screening have been applied to, and improved, identification and validation of specific gene functions in RNA interference (Zhang 2011).

A key piece of HTS equipment is the microplate: a small container that features a grid of small, open divots called wells. Figure 23.1 presents diagrams of 96 (8×12)-well microplates commonly used in HTS practice where solid and empty circles are wells. Other HTS microplates such as 384 (16×24)-well, 1536 (32×48)-well, and 3456 (48×72)-well plates are also used in HTS experiments. Even though the plating format and the number of compounds per plate vary, primary HTS operations typically measure a single observation from most compounds incubated in the wells of the rectangular microplate.

**FIGURE 23.1**

Current HTS designs, where solid circles contain controls and empty circles have compounds to be screened. (a) <http://chembank.broadinstitute.org/assays/view-plate.htm?id=1029542>, (b) id = 1029548, (c) id = 1003379, and (d) = 2001521.

There are many statistical challenges in data preprocessing and active-compound identification in primary HTS. Active compounds are usually called hits in HTS. Malo et al. (2006) reviewed statistical issues in the current HTS practice. One challenge is to correct various systematic errors in HTS data. Systematic errors in HTS are usually caused by factors such as unintended differences in compound concentrations due to aging, reagent evaporation, or cell decay; errors in liquid handling or malfunction of pipettes; variations in incubation time and time drift in measuring different wells or different plates; and reader effects (Heuer et al. 2005).

Unlike random errors that cause measurement discrepancies and result in random variation in hit identification, systematic errors contaminate HTS data and bias (over- or underestimate) hit selection systematically in almost every HTS experiment (Kevorkov and Makarenkov 2005). Often, systematic errors are revealed as row and column effects on microplates in HTS practice. That is, observations from the same compound may vary systematically as well as randomly from well to well on the microplate. This is the issue discussed in this chapter.

Figure 23.2 shows the heat map of data from an HTS experiment. In this experiment, one assay at a level of five ng concentration was incubated into all 384 wells. The incubated plate was then read by an FL×800 Fluorescence Microplate Reader (BioTek Instruments Inc.). The fluorescence intensity levels of assay in all wells were recorded. The heat map shows that levels of intensity are different for wells in different rows and columns.

Table 23.1 shows the analysis of variance (ANOVA) of the data, and it is clear that both row and column effects are statistically significant. Brideau et al. (2003) studied more than




FIGURE 23.2
Heat map of an HTS experiment.

TABLE 23.1

ANOVA of the HTS Experiment

	Degree of Freedom	Sum of Squares	Mean Sum of Squares	F Value	P Value
Row	15	288.2396	19.2160	29.9782	<0.0001
Column	23	504.8646	21.9506	34.2443	<0.0001
Residuals	345	221.1354	0.6410		

1000 384-well microplates and discovered that activity levels for wells in row one were, on average, 14% lower than those in row 16.

Due to the tremendous impact of row and column effects on hit selection, various remedies have been proposed to eliminate row and column effects in HTS. For instance, Lundholt et al. (2003) experimentally used a preincubation technique to reduce row and column effects in cell-based assays. However, experimental remedies usually increase the cost of material inventory and time in HTS. On the other hand, analytic methods are more cost-effective.

There are two types of analytic methods. One applies normalization techniques such as the median polish (Tukey 1977) method followed by the B score (Brideau et al. 2003) method to HTS data directly without considering the design of the experiment. Tukey's median polish method estimates row and column effects through an iterative procedure and eliminates row and column effects by subtracting them from their corresponding observations. The difference between the observation from each well and the sum of the estimated general mean and corresponding row and column effects is called the residual and is used for hit selection. The B score method uses B scores to select candidates. B scores are the ratios of residuals to their median absolute deviation multiplied by 1.4826.

The other normalization method uses experimental design to separate row and column effects from treatment effects and then eliminates row and column effects (see Section 3.5) by using estimated treatment effects for hit selection. Our focus is of the latter type. One important goal of various normalization techniques is to identify row and column

effects so that they can be effectively adjusted in the downstream analysis of hit identification, which results in fewer false positives and false negatives. Brideau et al. (2003) provided more details on statistical normalization techniques.

Figure 23.1 displays four commonly used designs of HTS experimentation in ChemBank database (<http://chembank.broadinstitute.org>) where solid circles contain controls and empty circles have compounds to be screened. ChemBank is a public database created by the Broad Institute's Chemical Biology Program. It includes freely available data of more than 2500 HTS experiments. For example, the HTS experiment studied in Section 23.3 has a design of type (c) where wells in the microplate represented by empty circles of the diagram contain the targeted enzyme, dihydroorotate dehydrogenase, and wells of solid circles contain dimethyl sulfoxide solvent without the enzyme. As is shown later, designs (a), (b), and (d) are unable to de-alias treatment effects from the contamination of both the row and column effects simultaneously.

Various patterns of row and column effects have been discovered in HTS. In this chapter, we focus on designs and analysis that are appropriate for models in which row and column effects are additive, that is, where row–column interactions are not needed.

Arranging compounds or treatments on microplates to eliminate row and column effects is the same as constructing row–column designs to eliminate additive two-way heterogeneity. Latin squares, Youden squares, and their generalizations are classic row–column designs with various optimal properties (Kiefer 1975; Cheng 1981; and Section 3.5.3). However, these classic designs usually have strict combinatorial constraints on the numbers of rows, columns, and treatments. For example, the numbers of rows, columns, as well as treatments have to be equal in Latin square designs and the numbers of columns and treatments are the same in Youden squares. Because of such constraints and because HTS experiments in drug discovery and microarray experiments have very large numbers of treatments (Hüser et al. 2006; Kerr 2003), classic designs are not applicable, and more general types of row–column designs are needed.

Since scientists in HTS are primarily interested in the selection and identification of superior treatments for further improvement as opposed to precise estimation or prediction of their effects, HTS experiments are often characterized by a shortage of experimental materials to evaluate hundreds of thousands of compounds (called treatments, hereafter). Therefore, multiple replications of all treatments in the primary screening stage are neither feasible nor cost-effective. Moreover, the need for replication depends on the variability of biological assays. The demand for replication in HTS experiments with high-quality assays is much less than that with low-quality assays. On the other hand, if the strategy is to use initial screening data for modeling or optimizing responses, then, as is shown in most textbooks of design of experiments, multiple replications of a small number of compounds are better than a few or no replication of a large number of treatments because multiple replications result in smaller variances of effect estimates.

Two computer packages CycDesign (Whitaker et al. 2008) and Gendex DOE Toolkit (Nguyen 1997) can generate row–column designs of any dimensions with all treatments equally replicated and occurring at least twice. Such designs are not popular in the early stages of the HTS selection processes (Lin and Poushinsky 1983) because they can study, at most, $N/2$ treatments, where N is the number of wells. Row–column designs in which no constraints are imposed on numbers of rows and columns are discussed in the following, and the majority of treatments in these designs are not replicated.

There are many criteria for selecting row–column designs. In this chapter, we discuss D -, A -, E -, and (M, S) optimality. We base the first three on “canonical efficiency factors,”

which we define as the nonzero eigenvalues of the information matrix C_d , defined by (23.2) in Section 23.2, but with respect to a diagonal matrix R_d whose diagonal elements are the replicate numbers of treatments in the design (see Section 3.1 for details), that is, $R_d^{-1/2} C_d R_d^{-1/2}$. A design whose geometric mean of canonical efficiency factors is at least as large as that of any other design is D -optimal. An A -optimal design has the largest harmonic mean of canonical efficiency factors, and a design whose smallest canonical efficiency factor is at least as large as that of any other design is E -optimal (Jacroux 1990).

For equally replicated designs, an A -optimal design has the smallest average variance over all pairwise differences of treatments (John and Williams 1995, p. 30). As described in Chapter 3, the A -criterion is popular for assessing block designs (see also John and Williams 1995, pp. 31–32). However, unequally replicated A -optimal row–column designs are computationally impractical to find when the number of experimental units is not a multiple of the number of treatments and the number of treatments is large. Sonnemann (1985) studied A -optimal row–column designs for two treatments. Morgan and Parvu (2007) solved the A -optimality problem for three treatments using information matrix C_d (23.2), as discussed in Chapter 3. To the authors' knowledge, Sonnemann (1985) and Morgan and Parvu (2007) provide the only published work on unequally replicated A -optimal row–column designs.

The (M, S) -optimality criterion is used in the following. As explained in Section 3.3.1 of Chapter 3, a design is said to be (M, S) -optimal (or have the (M, S) -property) within a class of designs if its information matrix, C_d , has the maximum trace and the square of the information matrix has the minimum trace over all designs with the maximum trace. The rationale of using the (M, S) -criterion is as follows. First, it is desirable in HTS experimentation to estimate all paired comparisons with the same precision, and it is known, for equally replicated row–column designs, that all paired comparisons can be estimated with the same precision if and only if all nonzero eigenvalues of the information matrix are equal (Raghavarao 1971, p. 52). Let λ_i , for $i = 1, 2, \dots, n$, be the nonzero eigenvalues of the information matrix. Since $\sum_{i=1}^n (\lambda_i - \frac{1}{n} \sum_{i=1}^n \lambda_i)^2 = \sum_{i=1}^n \lambda_i^2 - (\sum_{i=1}^n \lambda_i)^2/n$, $\sum_{i=1}^n \lambda_i^2$ is the trace of the squared information matrix, and $\sum_{i=1}^n \lambda_i$ is the trace of the information matrix, (M, S) -optimal designs have the least variable eigenvalues among those for which the average is largest.

Second, (M, S) -optimal designs are easier to find than A - and E -optimal designs. The (M, S) -criterion has huge computational advantages over other criteria. Such advantages have been popularly adopted in searching A - and E -optimal row–column designs by computers. For example, Nguyen (1997) constructed optimal or near-optimal row–column designs with up to 100 treatments by permuting or interchanging the treatments within the blocks of an optimal or near-optimal incomplete block design used as the column component of an row–column design, and, using this technique, he searched for the A - and E -optimal designs among (M, S) -optimal designs. Although (M, S) -optimal designs are not necessarily A - and E -optimal, Cheng (1978) shows that, at least for regular graph designs, they do tend to have good A - and E -criterion performance. In Section 23.4.2, we show that our (M, S) -optimal designs perform well under the A - and E -criteria compared with two other classes of designs that we discuss.

The discussion in the following text is arranged as follows. Section 23.2 introduces a class of row–column designs and some preliminaries. Designs with one replicated treatment are discussed in Section 23.3. Section 23.4 studies designs with multiple replicated treatments and concluding remarks are given in Section 23.5.

23.2 Saturated Row–Column Designs

In an HTS experiment, wells in the microplate are experimental units, and they are grouped by two blocking factors with one factor representing the rows of the microplate and the other representing the columns. The third factor of the experiment is the treatment such as test compound or biological assay. An appropriate row–column design for the experiment increases the accuracy of estimating treatment comparisons by enabling row and column effects to be eliminated in the analysis.

In this section, a class of row–column designs (Qu et al. 2010) is introduced. Consider allocating $v = (b - 1)(k - 1) + 1$ treatments, T_1, T_2, \dots, T_v , to the bk experimental units in $b(\geq 3)$ rows and $k(\geq 3)$ columns. The class of designs, say, $\mathcal{D}(b, k)$, is constructed as follows:

- a. Treatment $T_{(i-1)(k-1)+j}$ is arranged in the (i, j) th cell of the design for $i = 1, 2, \dots, b - 1, j = 1, 2, \dots, k - 1$.
- b. Treatment T_v is placed in the (b, k) th cell of the design.
- c. For $1 \leq i \leq b - 1$, η_i in column k is any treatment in $\{T_1, T_2, \dots, T_v\}$ that has not appeared in the i th row.
- d. For $1 \leq j \leq k - 1$, θ_j in row b is any treatment in $\{T_1, T_2, \dots, T_v\}$ that has not appeared in the j th column.

The treatment arrangement of designs in Table 23.2 provides great flexibility for HTS experimentation where only treatments in the b th row and k th column are restricted and replicated and all others are nonrepeated. A variant of this type of designs has been used in Hsiao et al. (2012). There are no constraints on b and k in designs of class $\mathcal{D}(b, k)$ other than $v = (b - 1)(k - 1) + 1$. Note that there is no loss of generality if we restrict $b \leq k$, since the design can always be rotated by 90° to achieve this.

The following discussion is devoted to notations of row–column designs. More details can be found in Section 3.5.3 or John and Williams (1995, Chapter 5). Let $y_{ij(\ell)}$ be the observation from treatment T_ℓ in the (i, j) th cell, where $\ell \in \{1, 2, \dots, v\}$, $i = 1, 2, \dots, b$, and $j = 1, 2, \dots, k$. The following additive fixed-effect model is considered in this chapter,

TABLE 23.2
Class of Row–Column Designs

Row\Column	1	2	3	...	$k - 1$	k
1	T_1	T_2	T_3	...	T_{k-1}	η_1
2	T_k	T_{k+1}	T_{k+2}	...	$T_{2(k-1)}$	η_2
3	T_{2k-1}	T_{2k}	T_{2k+1}	...	$T_{3(k-1)}$	η_3
⋮	⋮	⋮	⋮	⋮	⋮	⋮
$b - 1$	$T_{(b-2)k-(b-3)}$	$T_{(b-2)k-(b-4)}$	$T_{(b-2)k-(b-5)}$...	T_{v-1}	η_{b-1}
b	θ_1	θ_2	θ_3	...	θ_{k-1}	T_v

$$y_{ij(\ell)} = \mu + \alpha_i + \beta_j + \tau_\ell + \epsilon_{ij}, \tag{23.1}$$

where μ is the general mean, α_i is the effect of the i th row, β_j is the effect of the j th column, τ_ℓ is the effect of the treatment T_ℓ that appears in the (i, j) th cell, $\ell = 1, 2, \dots, v$, and ϵ_{ij} 's are independently and identically distributed random errors with mean zero and standard deviation σ .

Since treatment, row, and column effects in an HTS experiment form a three-way classification, (23.1) is the simplest linear model that separates treatment effects from random errors as well as from row and column effects. However, least squares analysis based on this model has not been popularly used for hit selection in the HTS literature with the appropriate row–column designs. Currently, most selection procedures are ad hoc. Few experimenters attempt to de-alias treatment effects from the contamination of row and column effects as well as from random errors using row–column designs. The popular B score method (Brideau et al. 2003) eliminates row and column effects by estimating such effects through Tukey’s median polish procedure and subtracting them from corresponding observations. Such elimination does reduce the influence of row and column effects but cannot separate treatment effects from random errors. In fact, the residuals used in the B score procedure are sums of treatment effects and random errors in model (23.1). As is shown later, model (23.1) cannot be applied to designs (a), (b), and (d) of Figure 23.1.

The information matrix C_d of a row–column design d (Section 3.5.3) has the following form under model (23.1):

$$C_d = \mathbf{R}_d - \frac{1}{k}\mathbf{N}_d\mathbf{N}'_d - \frac{1}{b}\mathbf{M}_d\mathbf{M}'_d + \frac{1}{bk}\mathbf{r}_d\mathbf{r}'_d, \tag{23.2}$$

or, equivalently,

$$bkC_d = bk\mathbf{R}_d - b\mathbf{N}_d\mathbf{N}'_d - k\mathbf{M}_d\mathbf{M}'_d + \mathbf{r}_d\mathbf{r}'_d, \tag{23.3}$$

where $\mathbf{r}'_d = (r_1, \dots, r_v)$ is the treatment replicate vector, $\mathbf{R}_d = \text{diag}(r_1, r_2, \dots, r_v)$ is a diagonal matrix with entries r_1, \dots, r_v , $\mathbf{N}_d = (n_{ij})$ of order $v \times b$ denotes the treatment-row incidence matrix, that is, n_{ij} is the number of times treatment T_i appears in the j th row, and $\mathbf{M}_d = (m_{ij})$ of order $v \times k$ is the treatment-column incidence matrix; that is, m_{ij} is the number of times treatment T_i occurs in the j th column. If all entries in \mathbf{N}_d and \mathbf{M}_d are either 0 or 1, design d is said to be binary.

It is known that a row–column design is treatment-connected (i.e., all paired comparisons of treatments are estimable) if and only if the rank of C_d is $v - 1$. All treatment effects in model (23.1) are estimable in a treatment-connected row–column design in spite of the existence of row and column effects. A row–column design is called row-connected if the rank of $C_d^{(r)} = \mathbf{R}_d - \frac{1}{k}\mathbf{N}_d\mathbf{N}'_d$ is $v - 1$. In a row-connected design, treatment effects may be confounded with column effects and correction can be made for row effects. Similarly, a row–column design is called column-connected if the rank of $C_d^{(c)} = \mathbf{R}_d - \frac{1}{b}\mathbf{M}_d\mathbf{M}'_d$ is $v - 1$. In a column-connected design, treatment effects may be confounded with row effects, and correction can be made for column effects. A treatment-connected row–column design is row- and column-connected. However, there are designs that are both row- and column-connected but not treatment-connected. In fact, design (d) in Figure 23.1 is such a design.

If all solid circles in designs in Figure 23.1 are filled with one treatment, and all open circles represent different, unreplicated treatments, there are $v = 81$ distinct treatments in design (a), and a straightforward QR decomposition shows that $\text{rank}(C_d) = 70$, which is smaller than $v - 1 = 80$. Therefore, design (a) is not treatment-connected and not all paired comparisons of treatments are estimable. In fact, design (a) is row-connected. Similarly, design (b) is not treatment-connected but it is row-connected. For design (c) with $v = 57$ distinct treatments, $\text{rank}(C_d) = 56$, and therefore, the design is treatment-connected. Design (d), with $v = 65$ distinct treatments, is row- and column-connected but not treatment-connected because $\text{rank}(C_d^{(r)}) = 64$, $\text{rank}(C_d^{(c)}) = 64$, and $\text{rank}(C_d) = 63$. As is pointed out next, the maximum number of treatments that can be arranged in an 8×12 treatment-connected design is $(b-1)(k-1)+1 = 78$. Zhang (2008) proposed a series of plate designs by arranging control treatments in certain patterns to offset row-column effects. The designs labeled C1 to C5 in the article are not treatment-connected. Therefore, one cannot estimate all the parameters in model (23.1) using those designs.

Since there are bk observations in the design of Table 23.2 for each b and k , the total number of degrees of freedom is $bk - 1$. According to Equation 23.1, $b - 1$ and $k - 1$ degrees of freedom are used to estimate row and column effects, respectively. There are $(b - 1)(k - 1) = v - 1$ degrees of freedom left for treatment effects. Thus, the design of Table 23.2 is "saturated" because $v = bk - b - k + 2$ is the maximum number of treatments that can be arranged in a row-column layout to eliminate nonnegligible two-way heterogeneity. There are no degrees of freedom left to estimate the variance of the random error σ^2 (see Section 23.3.2).

23.3 Design and Analysis of HTS Experiments with One Replicated Treatment

We discuss design and analysis with one replicated (control) treatment on each microplate in Sections 23.3.1 and 23.3.2, and we are equally interested in comparisons of nonreplicated treatments with the replicated treatment and comparisons of nonreplicated treatments with other nonreplicated treatments in the following discussion.

23.3.1 Designs with One Replicated (Control) Treatment

In this section, row-column designs with one (control) replicated treatment are discussed. For $b \geq 3$ and $k \geq 3$, consider the row-column layout in Table 23.3 where cells in the first $b - 1$ rows and $k - 1$ columns are labeled T_1, T_2, \dots, T_{v-1} with one label per cell and the $v - 1$ treatments are assigned at random to the labels T_1, T_2, \dots, T_{v-1} . The control treatment, labeled T_v , is assigned to cells in the b th row and k th column, where $v = (b - 1)(k - 1) + 1 = bk - b - k + 2$.

Note that the design in Table 23.3 is a special case of the design in Table 23.2 with $\eta_i = \theta_j = T_v$ for $i = 1, 2, \dots, b - 1$ and $j = 1, 2, \dots, k - 1$. In practice, the row-column design in Table 23.3 is usually randomized before use by randomizing the orders of rows, columns, and treatments, respectively. If the k th column of the design in Table 23.3 is discarded, Dey et al. (1995) showed that such a design minimizes the maximum variances of the pairwise

TABLE 23.3
Row-Column Designs with One Control

Row\Column	1	2	3	...	$k - 1$	k
1	T_1	T_2	T_3	...	T_{k-1}	T_v
2	T_k	T_{k+1}	T_{k+2}	...	$T_{2(k-1)}$	T_v
3	T_{2k-1}	T_{2k}	T_{2k+1}	...	$T_{3(k-1)}$	T_v
⋮	⋮	⋮	⋮	⋮	⋮	⋮
$b - 1$	$T_{(b-2)k-(b-3)}$	$T_{(b-2)k-(b-4)}$	$T_{(b-2)k-(b-5)}$...	T_{v-1}	T_v
b	T_v	T_v	T_v	...	T_v	T_v

differences of the best linear unbiased estimators among all incomplete block designs with v treatments and $k - 1$ blocks of size b .

Properties of the design in Table 23.3 have been studied in the literature. Ogunyemi et al. (2007) proved that the $b \times k$ design in Table 23.3 has smaller trace of squared information matrix than any other $b \times k$ binary designs in Table 23.2. It can be shown (Theorem 23.1) that a $b \times k$ design in Table 23.3 has the minimum trace of the information matrix among all $b \times k$ designs in Table 23.2. Nevertheless, it is shown in Section 23.4.1 that these nonbinary designs perform nearly as well as the (M, S) -optimal designs under A - and D -efficiency (defined in Section 23.1) and are sometimes better under E -efficiency.

Theorem 23.1 For any design d in the class $\mathcal{D}(b, k)$ of designs in Table 23.2,

$$\text{trace}(C_d) \geq (bk - b - k + 2) \left(1 - \frac{1}{b} - \frac{1}{k} + \frac{1}{bk} \right). \tag{23.4}$$

The lower bound of $\text{trace}(C_d)$ in inequality (23.4) is attained when d is a $b \times k$ design in Table 23.3.

Least squares estimates of treatment effects in the design in Table 23.3 can be obtained from the reduced normal equations after eliminating row and column effects in model (23.1), that is, $C_d \hat{\tau} = q$, where C_d is the information matrix and q is the vector of adjusted treatment totals, that is,

$$q = t - \frac{1}{b} N_d b - \frac{1}{k} M_d c + \frac{G}{bk} r_d,$$

where r_d is the replication vector, $t = (y_{\cdot(1)}, y_{\cdot(2)}, \dots, y_{\cdot(v)})'$, $b = (y_{1(\cdot)}, y_{2(\cdot)}, \dots, y_{b(\cdot)})'$, and $c = (y_{1(\cdot)}, y_{2(\cdot)}, \dots, y_{k(\cdot)})'$ are vectors of treatment, row, and column totals, respectively, and G is the total sum of observations. It is known that $\text{Var}(q) = \sigma^2 C_d$ (John and Williams 1995, p. 89) and a solution to the reduced normal equation is given by $\hat{\tau} = C_d^- q$, where C_d^- is a generalized inverse of C_d .

The information matrix of the design in Table 23.3 is

$$C_d = \begin{pmatrix} (\mathbf{I}_{b-1} - \frac{1}{b} \mathbf{J}_{b-1}) \otimes (\mathbf{I}_{k-1} - \frac{1}{k} \mathbf{J}_{k-1}) & -\frac{1}{bk} \mathbf{1}_{b-1} \otimes \mathbf{1}_{k-1} \\ -\frac{1}{bk} \mathbf{1}'_{b-1} \otimes \mathbf{1}'_{k-1} & \frac{(b-1)(k-1)}{bk} \end{pmatrix},$$

where \otimes is the Kronecker product of matrices, $J_n = \mathbf{1}_n \mathbf{1}'_n$, and $\mathbf{1}_n$ is the $n \times 1$ vector of 1's for $n = b - 1$ and $k - 1$. Since $\text{rank}(C_d) = (b - 1)(k - 1) = v - 1$, the design is treatment-connected, and its treatment-connectedness guarantees that all treatment contrasts are estimable regardless of row and column effects. Therefore, all treatments within one microplate are comparable.

Note that one generalized inverse of C_d is

$$\begin{pmatrix} (\mathbf{I}_{b-1} + J_{b-1}) \otimes (\mathbf{I}_{k-1} + J_{k-1}) & \mathbf{0}_{(b-1)(k-1)} \\ \mathbf{0}'_{(b-1)(k-1)} & 0 \end{pmatrix}.$$

It is straightforward to show that the least squares estimate of $\tau_h - \tau_v$ using $C_d^- \mathbf{q}$ is

$$\widehat{\tau}_h - \widehat{\tau}_v = y_{ij(h)} - y_{ik(v)} - y_{bj(v)} + y_{bk(v)},$$

if treatment T_h is in the i th row and j th column.

It follows that the least squares estimate of every contrast of treatment versus control, that is, $\tau_h - \tau_v$, has the same variance. In fact, $\text{Var}(\widehat{\tau}_h - \widehat{\tau}_v) = 4\sigma^2$ for $h = 1, 2, \dots, v - 1$. If T_i and T_j are treatments in the same row or column ($i \neq v, j \neq v$ and $i \neq j$), the variance of the least squares estimate of their effect difference ($\tau_i - \tau_j$) is

$$\text{Var}(\widehat{\tau}_i - \widehat{\tau}_j) = \text{Var}(\widehat{\tau}_i - \widehat{\tau}_v) + \text{Var}(\widehat{\tau}_j - \widehat{\tau}_v) - 2\text{Cov}(\widehat{\tau}_i - \widehat{\tau}_v, \widehat{\tau}_j - \widehat{\tau}_v) = 4\sigma^2.$$

If T_i and T_j are in different rows and columns, then $\text{Var}(\widehat{\tau}_i - \widehat{\tau}_j) = 6\sigma^2$. Therefore, treatments in the same row or column are compared more precisely than those in different rows or columns. These formulas provide a convenient way of finding least squares estimates of treatment differences and their corresponding variances when large microplates such as those with 1536(32×48)- or 3456(48×72)-wells are used in HTS experiments.

23.3.2 Analysis of Row–Column Designs with One Replicated Treatment

The analysis of the design in Table 23.3 is complicated by the fact that there are no degrees of freedom left to estimate the process error and the design is nonorthogonal. Though many methods have been proposed for the analysis of orthogonal, saturated designs (see Hamada and Balakrishnan 1998, Kinaterer et al. 2000b for details), only a few articles in the literature deal with nonorthogonal, saturated designs. Two methods proposed, respectively, by Kunert (1997) and Wang and Voss (2001) will be discussed and compared in this section because effect sparsity (see Chapter 7), one of the key feature of HTS, is used in both methods to estimate the random error. The two methods will also be compared in this section with other normalization techniques such as the B score method.

Kinaterer et al. (2000a) provided a variation of Kunert's method to control the type I error of individual tests. While type I error is controlled, Kinaterer's method does not take the effect sparsity in HTS into consideration. Moreover, since the main target of primary HTS is to discover hits for future research, it is usually more important to control type II or false-negative errors than type I or false-positive errors. Therefore, Kinaterer's method will not be used in the following discussion.

False-positive treatments add cost to the downstream research of HTS but will eventually be filtered out, while false-negative treatments are targets that are missed and that

could have been invaluable discoveries. Therefore, we recommend that the type I error in an individual test not be adjusted for multiple comparisons. We believe that emphasizing the reduction of the type II error while maintaining type I errors as low as possible is the most beneficial approach when using statistical tests for scientific discovery in HTS.

The data used in the following discussion are from an HTS experiment on enzyme inhibition of the dihydroorotate dehydrogenase assay and are available from <http://chembank.broad-institute.org/assays/view-plate.htm?id=1003379>. The experiment was done by Dr. Derek Martyn at the Broad Institute.

There are two replicates: Plates 1021.0006.1335.A and 1021.0006.1335.B. The two replicates of data make it possible to analyze the data from each plate individually and to compare the consistency of results from various analysis methods. In the experiment, three hundred treatments with different chemical structures are arranged in columns 3 to 22 and rows 1 to 15 of two 384(16 × 24)-well microplates. Following a 30 min incubation, the absorbance value (between 0 and 1) was measured by EnVision readers (PerkinElmer Inc.). The ANOVA table of data from plate A is given in Table 23.4. It is observed that row and column effects are statistically significant. Note that the statistical significance of the row and column effects in this experiment may be due to the tiny random error. Whether such row and column effects are practically significant has to be determined by scientists in the HTS practice.

To apply the Kunert and the Wang and Voss (WV, hereafter) methods to the design with one replicated treatment, all wells except those in the 16th row or 24th column will be regarded as treatments for the purpose of illustration. Thus, although wells in columns 1, 2, 23, and 24 or in row 16 contain the control, the 45 wells in rows 1 to 15 of columns 1, 2, and 23 will be regarded here as distinct treatments and be used to evaluate false-positive rates. If an analysis method is effective, treatments in these 45 wells are unlikely to be selected. So in our example there are 345 treatments and one control.

Kunert's method is carried out as follows:

Step 1: Use the baseline constraints where α_{16} , β_{24} , and τ_{346} are set to be zero for row, column, and treatment effects, respectively; calculate the least squares estimates of contrasts $\alpha_i - \alpha_{16}$, $\beta_j - \beta_{24}$, and $\tau_\ell - \tau_{346}$ in model (23.1) for $1 \leq i \leq 15$, $1 \leq j \leq 23$, and $1 \leq \ell \leq 345$; and select the order of the absolute sizes of these estimates (the largest estimate comes first) as the order of contrasts to calculate the adaptive standard error (ASE) in Step 3.

Step 2: Rearrange the model matrix according to the order of estimates in Step 1. Let $x_0 = \mathbf{1}_{384}$ be the 384-dimensional vector of ones and x_1, \dots, x_{383} be vectors in the

TABLE 23.4

ANOVA of Dihydroorotate Dehydrogenase Data

Source	DF	Sum of Squares	Mean Square	F Value	P Value
Model	338	8.7374	0.0259	129.5	<0.0001
Row	15	0.0297	0.0020	10.0	<0.0001
Column	23	4.6799	0.2035	1017.5	<0.0001
Treatment	300	0.9801	0.0033	16.5	<0.0001
Error	45	0.0084	0.0002		
Total	383	8.7458			

model matrix corresponding to the ordered estimates, respectively. For $1 \leq i \leq 383$, let $\mathbf{U}_{<i} = [x_0, x_1, \dots, x_{i-1}]$ be the model matrix of the first $i - 1$ effects, $\text{pr}(\mathbf{U}_{<i}) = \mathbf{U}_{<i}(\mathbf{U}'_{<i}\mathbf{U}_{<i})^{-1}\mathbf{U}'_{<i}$ be the projection matrix to its column space, and $\mathbf{W}_{<i} = \text{pr}^\perp(\mathbf{U}_{<i}) = \mathbf{I}_{384} - \text{pr}(\mathbf{U}_{<i})$. Define the i th contrast estimate as $s_i = \mathbf{x}'_i\mathbf{W}_{<i}\mathbf{y}/\sqrt{\mathbf{x}'_i\mathbf{W}_{<i}\mathbf{x}_i}$.

Step 3: Divide the i th ordered contrast estimate, say, $|s_{(i)}|$, by $\text{ASE}/\sqrt{\mathbf{x}'_i\mathbf{W}_{<i}\mathbf{x}_i}$, where

$$\mathbf{W}_{<i} = \text{pr}^\perp(x_0, x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_{383}), \quad \text{ASE} = \sqrt{\frac{1.08}{383 - m} \sum_{|s_i| \leq 2.56\hat{\sigma}_M} s_i^2}$$

m is the number of $|s_i|$ that are larger than $2.56\hat{\sigma}_M$, and $\hat{\sigma}_M = 3\text{median}\{|s_1|, \dots, |s_{383}|\}/2$. Treatment effects whose corresponding $|s_{(i)}|\sqrt{\mathbf{x}'_i\mathbf{W}_{<i}\mathbf{x}_i}/\text{ASE}$ ratios have p -values less than 0.05 under the student t distribution with $0.69n$ degrees of freedom are declared to be significantly different from the control.

Kunert’s method assumes that treatment effects can be ordered so that active ones come early and a data-driven ordering is adopted in Step 1 for the relative importance of treatment effects. Unlike Kunert’s method where least squares estimates of treatment effects and adaptive standard error are used, the WV method uses adjusted least squares estimates of effects. More specifically, the WV method is applied as follows:

Step 1: Specify the number of negligible contrasts among row, column, and treatment contrasts, say, ν , where one anticipates having at least ν negligible contrasts under the full model of Equation 23.1, and identify those that yield the ν smallest normalized estimates as the contrasts that are negligible.

Step 2: Let I be the index set of those ν parameters with the smallest normalized estimates. Drop the parameters corresponding to contrasts with the smallest contrast estimates from the full model and fit a reduced model to obtain adjusted estimates of contrasts. That is, for each retained parameter θ , calculate its adjusted estimate $\tilde{\theta} = \hat{\theta} - \mathbf{c}'_I\mathbf{V}_I^{-1}\hat{\boldsymbol{\gamma}}_I$ and the standard error of the adjusted estimate

$$s(\tilde{\theta}) = \sqrt{(a - \mathbf{c}'_I\mathbf{V}_I^{-1}\mathbf{c}_I)(\hat{\boldsymbol{\gamma}}'_I\mathbf{V}_I^{-1}\hat{\boldsymbol{\gamma}}_I/\nu)}$$

where $\hat{\theta}$, $\hat{\boldsymbol{\gamma}}_I$ are least squares estimates of θ and negligible parameters $\boldsymbol{\gamma}_I$, $a\sigma^2$ is the variance of $\hat{\theta}$, $\sigma^2\mathbf{V}_I$ is the variance matrix of $\hat{\boldsymbol{\gamma}}_I$, and $\sigma^2\mathbf{c}_I$ is the covariance between $\hat{\theta}$ and $\hat{\boldsymbol{\gamma}}_I$. Step 3: For each θ , the null distribution of statistic $\tilde{\theta}/s(\tilde{\theta})$ is obtained from 10,000 simulations, and the p -value of testing $H_0 : \theta = 0$ is calculated from the simulated null distribution.

Step 4: Treatment effects whose p -values in Step 3 are less than 0.05 are declared to be significantly different from the control.

In practice, data from two plates are usually combined for analysis. Since, here, we are interested in the consistency of results from various methods, selection procedures are applied separately to data from each plate. Significant effects at the 5% level commonly selected from both plates using the WV and the Kunert methods are listed in Table 23.5. For example, 46 and 53 treatments are selected from plates A and B by the WV method

TABLE 23.5

Hits Selected by Various Methods

Method	Wells of Hits
WV ($v = 327$) (46 and 53)	A04 A07 A10 A11 A12 A15 A19 A23
	B04 B09 B11 B12 B13 B15 B16
	C05 C06 C07 C11 C12 C13 C14 C15 C16 C19
	D03 D04 D07 D08 D17 D18 D19 D21
	E05 E07 E09 E10 E12 E16 E17 E18 E19 E21
	F3
Kunert (148 and 76)	A01 A02 A03 A04 A05 A06 A07 A08 A09 A10 A11 A12 A13 A14 A15 A16
	A17 A18 A19 A20 A21 A22
	B04 B06 B09 B10 B11 B12 B13 B15 B16
	C04 C05 C06 C07 C08 C10 C11 C12 C13 C14 C15 C16 C17 C19 C20 C21
	D03 D04 D07 D08 D12 D14 D17 D18 D19 D21
	E05 E07 E09 E10 E12 E16 E17 E18 E19 E21
	F03
	I03 I07 I11
O13	
B score (79 and 80)	A04 A07 A10 A11 A15 A19 A20 A23 A24
	B04 B06 B09 B11 B12 B13 B15 B16 B19
	C01 C02 C03 C07 C09 C12 C13 C14 C15 C16 C17 C18 C19 C22 C23 C24
	D02 D03 D04 D07 D08 D10 D11 D13 D15 D17 D19 D19 D20 D21 D22
	E01 E02 E05 E10 E11 E12 E16 E17 E18 E19 E21
	F03
	K04
	L24
N05	
Z score (27 and 21)	B24 D23 D24 E23 E24 F24 G23 H23 I23 J24 K24
	L23 L24 M23 M24 N23 N24 O23 O24 P23 P24
LASSO (97 and 96)	A01 A02 A04 A06 A07 A10 A11 A14 A15 A16 A18 A19 A22 A23
	B01 B02 B04 B05 B07 B09 B11 B12 B13 B15 B16 B19
	C01 C02 C03 C05 C06 C07 C09 C11 C12 C13 C14 C15 C17 C18 C19 C22 C23
	D01 D02 D03 D04 D05 D07 D08 D10 D11 D13 D15 D17 D18 D19 D20 D21 D22
	E01 E02 E04 E05 E10 E11 E12 E13 E14 E15 E16 E17 E18 E19 E21
F03 K04 L23 M23 N05	
Bradu–Hawkins (9 and 11)	B09 B11 B12 C07 C14 D21 E17 F03

Notes: Letters stand for the rows and numbers for the columns of the microplate. Note that a false positive occurs when any treatment in rows 1, 2, 23, 24 or row P is selected.

based on 10,000 simulations for the null distribution and $v = 327$ (or assume 5% of 345 treatments are active), respectively. The 44 treatments in Table 23.5 are selected from both plates *A* and *B* where *A04* stands for the compound in row *A* and column four. Kunert's method chooses 148 and 76 significant treatments from plates *A* and *B* of which 72 treatments are from both plates. All treatments selected by the WV method except *A23* (a false

positive) are also chosen by Kunert's method. Kunert's method selects more treatments than the WV method.

Selections from the B score and the Z score methods with a popularly used threshold of three (Brideau et al. 2003) are also listed in Table 23.5. The B score method selects 79 and 80 treatments from plates A and B of which 64 treatments are in common. Z scores are standardized observations (centered by the mean and scaled by the standard deviation) without any adjustment of row and column effects. The Z score method picks 27 and 21 from plates A and B of which 21 treatments are in common. Note that both the B score and the Z score methods select hits from all 384 wells on the microplates, and thus do not use comparisons with the control, while the WV and the Kunert methods select hits from 345 wells because the 39 wells in the 16th row and 24th column have been used as the control reference in effect estimation. It is observed that the treatments selected by the methods of WV, Kunert, and B score overlap substantially.

Table 23.5 also lists treatments selected by the least absolute shrinkage and selection operator (hereafter, LASSO; Tibshirani 1996) and the Bradu–Hawkins method (Bradu and Hawkins 1982). The LASSO is a method for high-dimensional estimation in linear regression. It minimizes the residual sum of squares subject to the sum of the absolute value of the coefficients being less than a constant. LASSO is a penalized likelihood approach and is known for its accuracy in prediction and variable selection coupled with its computational feasibility. It is designed for cases like HTS experiments where signals are sparse. Bühlmann and van de Geer (2011) has more details on LASSO. To obtain more stable cross-validation results, the LASSO selection is done by averaging 10 outcomes of fivefold cross validation.

The Bradu–Hawkins method uses the median of tetrads associated with a cell to identify multiple hits in a single step. A tetrad associated with cell (i, j) is defined as $y_{ij} - y_{sj} - y_{it} + y_{st}$ for $i \neq s$ and $j \neq t$, where y_{ij} is the observation from the (i, j) th cell and so on. The half-normal plot provides an indicator of the actual number of hits. Half-normal plots of plates A and B are given in Figure 23.3. For example, since there is a break at the ordered-absolute-median-tetrad value of 0.2 in the half-normal plot of plate A , all treatments whose absolute-median-tetrad values larger than 0.2 are empirically identified as hits in plate A , where 4, 21 stands for the treatment in the fourth row and 21st column or $D21$. While the LASSO method seems to consistently select more treatments than other methods, the Bradu–Hawkins method picks much smaller numbers of hits from both plates. Note that the number of hits selected by the Bradu–Hawkins method depends on the subjective threshold on the half-normal plot. If there is no obvious breaking points on the half-normal plot, the Bradu–Hawkins method may select more hits with a low threshold.

If plate B is regarded as a confirmation of plate A , the positive confirmation rate (PCR) in Zhang et al. (2005) can be used to compare the screening efficiency where

$$PCR = \frac{\text{Number of treatments selected in both plates } A \text{ and } B}{\text{Number of treatments selected in plate } A}.$$

The positive confirmation rates of the WV, the Kunert, the B score, the Z score, the LASSO, and the Bradu–Hawkins methods are 95.65%, 48.65%, 81.01%, 77.78%, 82.47%, and 88.89%, respectively. Therefore, the WV method is the most consistent one followed by the Bradu–Hawkins, the LASSO, the B score, the Z score, and the Kunert methods.

All treatments selected by the Z score method are control treatments and, therefore, are false positives. Simulation studies in Qu (2011a) show that the nonnormality of random errors may be responsible for the high false-positive rate of the Z score method. The Z




FIGURE 23.3
Half-normal selection of plates *A* (a) and *B* (b).

score method is based on the normal distribution of the observations in HTS, and it is very sensitive to any deviations from normality. Moreover, the Z score method does not adjust for row and column effects.

The B score method is robust to outliers since it uses the median absolute deviation as its denominator. However, the B score method selects more treatments in this example than the Kunert and the WV methods and therefore has higher chance to collect more false positives. One reason is that the B score method uses the residuals rather than estimates of effects. The difference between residuals and estimates of effects is clearly shown in Figure 23.4. Although the two graphs in the first row of Figure 23.4 seem to suggest that there may be a shift difference between residuals and effect estimates, different shapes of the two graphs in the second row of Figure 23.4 show a subtle difference between residuals and effect estimates. The correlation coefficient between estimated effects in plates *A* and *B* (0.9244) is higher than that between residuals (0.9096), which shows estimated effects as the estimates of true activity levels of compounds are more consistent than residuals. More false positives in the B score selection could also be due to the error distribution. Simulation studies in Qu (2011a) show that the B score method chooses more false positives when the error distribution is skewed.

The ordering of effect estimates in Kunert's method contributes to the large difference between two numbers of treatments selected from plates *A* and *B* (148 and 76). As pointed out by Kunert (1997), active effects have a large chance to be detected if they come early in the ordering. The data-driven ordering used by Kunert's method varies from data set to data set, and it does not guarantee that the probability of an active effect entering an early position is larger than that of a late position. The adjusted estimates and

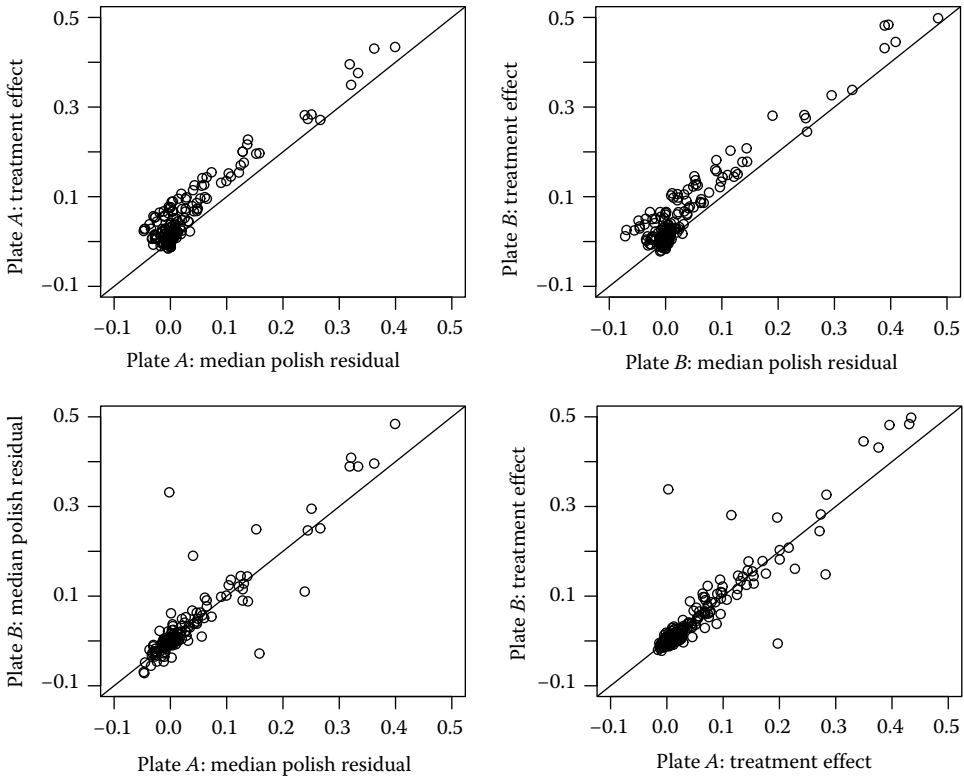


FIGURE 23.4
Estimates of effects versus residuals.

nonparametric determination of significant effects make the WV method more consistent than others.

Given the sparsity of hits in HTS culture (say, less than 5%), selection rates of 13%(44/345) of the WV method, 17%(64/384) of the B score method, 21%(80/384) of the LASSO method, and 21%(72/345) of Kunert’s method indicate that each method may select false positives. Ten false positives (A23, A24, C01, C02, C23, C24, D02, E01, E02, L24) have been selected by the B score method and 14 (A01, A02, A23, B01, B02, C01, C02, C23, D01, D02, E01, E02, L23, and M23) by the LASSO method because these treatments are from control wells. Two false positives (A01, A02) were selected by Kunert’s method, and only one (A23) is picked by the WV method. The Bradu–Hawkins method has no such false positives. The Bradu–Hawkins and WV methods are the best in terms of picking up the fewest treatments from control wells. However, since no experimental confirmation is available for the noncontrol wells, it is unclear which method has the fewest false positives, and the information on false negatives is not known either.

Since different methods use different standards to select hits, they may not have the same type I error rate. Methods with a high type I error rate may yield more hits than those with a low type I error rate. Therefore, the number of hits selected by a specific method should not be used as the only standard in comparing with other methods of selection. In our opinion, the WV method is the best one for HTS practice.

23.4 Designs of HTS Experiments with Multiple Replicated Treatments

Control treatments play an important role in HTS experiments. Replicated measurements from those control treatments on microplates are often used to calculate quality indices of screening. For example, the ratio of the average response from the controls in one microplate to that of another can be used to assess plate-to-plate variability. A ratio of close to one shows low variation from plate to plate.

Many HTS experiments need positive, negative, and other controls due to scientific specifications. For example, plate design C2 in Zhang (2008) has two positive, one negative, and one other controls residing in various wells of the microplate. Controls are usually repeated unequally in the microplate. The negative control in plate design C2 in Zhang (2008) is repeated 24 times, and the first positive control is repeated 12 times.

In an HTS experiment with positive and negative controls, a clear distinction between the two types of controls is an indication of high-quality screening. The Z-factor, $[|\bar{y}_P - \bar{y}_N| - 3(s_P + s_N)]/|\bar{y}_P - \bar{y}_N|$, proposed by Zhang et al. (1999) has been a popular quality index, where \bar{y}_P , \bar{y}_N , s_P , and s_N are the sample means and standard deviations of responses from positive and negative controls, respectively. A Z-factor value of one is an ideal level of screening quality, and a Z-factor value between 0.5 and 1 is an indication of excellent quality of screening.

How to construct optimal, treatment-connected designs with multiple replicated treatments is challenging. (M, S) -optimal designs are discussed in Section 23.4.1 and designs with a cyclic structure are presented in Section 23.4.2. Any replicated treatment in these designs can be used as a control treatment in HTS practice. For simplicity, we do not differentiate replicated treatments from other treatments in the following discussion.

23.4.1 (M, S) -Optimal Designs

For $b \geq 4$, $k = b + s$, and $0 \leq s \leq b - 1$, Qu (2010) constructed treatment-connected row-column designs that are (M, S) -optimal in $\mathcal{D}(b, k)$ by the following four steps where $1, 2, \dots, v = bk - b - k + 2$ are treatment labels:

1. Treatment $(i - 1)(k - 1) + j$ is in row i and column j for $i = 1, 2, \dots, b - 1$, $j = 1, 2, \dots$, and $k - 1$, respectively.
2. Treatment v is in row b and column k .
3. Treatments in rows 1 to $b - 1$ of column k are $k + 1, 2k + 1, \dots, (b - 2)k + 1$, and 1, respectively.
4. Treatments $k + 1, 2k + 1, \dots, (b - s - 1)k + 1$ are in columns 1 to $b - s - 1$ of row b ; those in columns $b - s$ to $b - 1$ of row b are treatments $(b - s - 1)(k - 2) + 2(b - 1)$, $(b - s)(k - 2) + 2(b - 1), \dots, (b - 2)(k - 2) + 2(b - 1)$; those in columns b to $k - 1$ of row b are treatments 1, $(b - 2)k + 1, \dots$, and $(b - s)k + 1$, respectively.

The design constructed by these four steps has $b - 1$ treatments replicated three times and $k - b$ treatment replicated twice. In fact, an (M, S) -optimal design for $b \geq 4$, $k = b + s$ is a design in Table 23.2 with $\eta_1 = k + 1, \eta_2 = 2k + 1, \dots, \eta_{b-1} = (b - 2)k + 1$ and $\theta_1 = k + 1, \theta_2 = 2k + 1, \dots, \theta_{b-s-1} = (b - s - 1)k + 1, \theta_{b-s} = (b - s - 1)(k - 2) + 2(b - 1), \theta_{b-s+1} = (b - s)(k - 2) + 2(b - 1), \dots, \theta_{b-1} = (b - 2)(k - 2) + 2(b - 1), \theta_b = 1, \theta_{b+1} = (b - 2)k + 1, \dots, \theta_{k-1} = (b - s)k + 1$.

TABLE 23.6

(M, S)-Optimal Design in the Class $\mathcal{D}(8, 12)$ of Designs in Table 23.2 ($b = 8, k = 12$, and $s = 4$)

Row\Column	1	2	3	4	5	6	7	8	9	10	11	12
A	1	2	3	4	5	6	7	8	9	10	11	13
B	12	13	14	15	16	17	18	19	20	21	22	25
C	23	24	25	26	27	28	29	30	31	32	33	37
D	34	35	36	37	38	39	40	41	42	43	44	49
E	45	46	47	48	49	50	51	52	53	54	55	61
F	56	57	58	59	60	61	62	63	64	65	66	73
G	67	68	69	70	71	72	73	74	75	76	77	1
H	13	25	37	44	54	64	74	1	73	61	49	78

For example, Table 23.6 presents a treatment-connected, (M, S)-optimal design in the class $\mathcal{D}(8, 12)$ of designs in Table 23.2 for 96(8×12)-well microplates. Table 23.7 provides a treatment-connected, (M, S)-optimal design in the class $\mathcal{D}(16, 24)$ of designs in Table 23.2 for 384(16×24)-well microplates. (M, S)-optimal row-column designs in the class $\mathcal{D}(b, k)$ of designs in Table 23.2 for popularly used microplates of other dimensions such as 32×48 and 48×72 plates can be constructed accordingly.

For $b \geq 4, k > b + s$, and $0 \leq s \leq b - 1, k$ can be expressed as $k = t(b - 1) + s + 1$ where $t \geq 2$. Treatment-connected, (M, S)-optimal designs in $\mathcal{D}(b, k)$ can be obtained from designs in Table 23.2 with specified η_i and θ_j for $i = 1, 2, \dots, b - 1$ and $j = 1, 2, \dots, k - 1$. For simplicity and clarity, steps of construction are given as follows, where $1, 2, \dots, v - 1$ and $v = bk - b - k + 2$ are treatment labels:

1. Treatment $(i - 1)(k - 1) + j$ is in row i and column j for $i = 1, 2, \dots, b - 1, j = 1, 2, \dots, k - 1$, respectively.
2. Treatment v is in row b and column k .
3. Treatments in rows 1 to $b - 1$ of column k are $k + 1, 2k + 1, \dots, (b - 2)k + 1$, and 1, respectively.
4. If t is even, treatments in columns 1 to $(t - 1)(b - 1)$ of row b are $2(b - 1), (k - 1) + (b - 1) + (b - 2), \dots, (b - 3)(k - 1) + (b - 1) + 2, (b - 2)(k - 1) + (b - 1) + 1; (b - 2)(k - 1) + 3(b - 1), (b - 3)(k - 1) + 3(b - 1) - 1, \dots, (k - 1) + 2(b - 1) + 2, 2(b - 1) + 1; \dots; t(b - 1), (k - 1) + t(b - 1) - 1, \dots, (b - 3)(k - 1) + (t - 1)(b - 1) + 2$, and $(b - 2)(k - 1) + (t - 1)(b - 1) + 1$; those in columns $(t - 1)(b - 1) + 1$ to $t(b - 1) - s$ of row b are $1, (b - 2)k + 1, \dots, (s + 1)k + 1$; those in columns $t(b - 1) - s + 1$ to $t(b - 1)$ of row b are $(s - 1)(k - 1) + t(b - 1) + s, (s - 2)(k - 1) + t(b - 1) + s - 1, \dots, (k - 1) + t(b - 1) + 2$, and $t(b - 1) + 1$; and those in columns $t(b - 1) + 1$ to $k - 1$ of row b are $k + 1, 2k + 1, \dots, sk + 1$, respectively.
5. If t is odd, treatments in columns 1 to $(t - 1)(b - 1)$ of row b are $2(b - 1), (k - 1) + 2(b - 1) - 1, \dots, (b - 3)(k - 1) + (b - 1) + 2, (b - 2)(k - 1) + (b - 1) + 1; (b - 2)(k - 1) + 3(b - 1), (b - 3)(k - 1) + 3(b - 1) - 1, \dots, (k - 1) + 2(b - 1) + 2, 2(b - 1) + 1; \dots; (b - 2)(k - 1) + t(b - 1), (b - 3)(k - 1) + t(b - 1) - 1, \dots, (k - 1) + (t - 1)(b - 1) + 2$, and $(t - 1)(b - 1) + 1$; those in columns $(t - 1)(b - 1) + 1$ to $t(b - 1) - s$ of row b are $k + 1, 2k + 1, \dots, (b - 1 - s)k + 1$;

TABLE 23.7

(M, S) -Optimal Design in the Class $\mathcal{D}(16, 24)$ of Designs in Table 23.2 ($b = 16, k = 24,$ and $s = 8$)

Row\Column	1	2	3	4	5	6	7	8	9	10	11	12
A	1	2	3	4	5	6	7	8	9	10	11	12
B	24	25	26	27	28	29	30	31	32	33	34	35
C	47	48	49	50	51	52	53	54	55	56	57	58
D	70	71	72	73	74	75	76	77	78	79	80	81
E	93	94	95	96	97	98	99	100	101	102	103	104
F	116	117	118	119	120	121	122	123	124	125	126	127
G	139	140	141	142	143	144	145	146	147	148	149	150
H	162	163	164	165	166	167	168	169	170	171	172	173
I	185	186	187	188	189	190	191	192	193	194	195	196
J	208	209	210	211	212	213	214	215	216	217	218	219
K	231	232	233	234	235	236	237	238	239	240	241	242
L	254	255	256	257	258	259	260	261	262	263	264	265
M	277	278	279	280	281	282	283	284	285	286	287	288
N	300	301	302	303	304	305	306	307	308	309	310	311
O	323	324	325	326	327	328	329	330	331	332	333	334
P	25	49	73	97	121	145	169	184	206	228	250	272
Row\Column	13	14	15	16	17	18	19	20	21	22	23	24
A	13	14	15	16	17	18	19	20	21	22	23	25
B	36	37	38	39	40	41	42	43	44	45	46	49
C	59	60	61	62	63	64	65	66	67	68	69	73
D	82	83	84	85	86	87	88	89	90	91	92	97
E	105	106	107	108	109	110	111	112	113	114	115	121
F	128	129	130	131	132	133	134	135	136	137	138	145
G	151	152	153	154	155	156	157	158	159	160	161	169
H	174	175	176	177	178	179	180	181	182	183	184	193
I	197	198	199	200	201	202	203	204	205	206	207	217
J	220	221	222	223	224	225	226	227	228	229	230	241
K	243	244	245	246	247	248	249	250	251	252	253	265
L	266	267	268	269	270	271	272	273	274	275	276	289
M	289	290	291	292	293	294	295	296	297	298	299	313
N	312	313	314	315	316	317	318	319	320	321	322	337
O	335	336	337	338	339	340	341	342	343	344	345	1
P	294	316	338	1	337	313	289	265	241	217	193	346

those in columns $t(b - 1) - s + 1$ to $t(b - 1)$ of row b are $(b - s - 1)(k - 1) + t(b - 1) + s,$
 $(b - s)(k - 1) + t(b - 1) + s - 1, \dots, (b - 2)(k - 1) + t(b - 1) + 1;$ and those in columns
 $t(b - 1) + 1$ to $k - 1$ of row b are $1, (b - 2)k + 1, \dots, (b - s)k + 1,$ respectively.

For example, Table 23.8 presents a treatment-connected, (M, S) -optimal design in the class $\mathcal{D}(8, 16)$ of designs in Table 23.2 although no microplates of such dimensions have been used so far in HTS experiments.

TABLE 23.8

(M, S)-Optimal Design in the Class $\mathcal{D}(8, 16)$ of Designs in Table 23.2 ($b = 8, k = 16, t = 2$, and $s = 1$)

Row\Column	1	2	3	4	5	6	7	8
A	1	2	3	4	5	6	7	8
B	16	17	18	19	20	21	22	23
C	31	32	33	34	35	36	37	38
D	46	47	48	49	50	51	52	53
E	61	62	63	64	65	66	67	68
F	76	77	78	79	80	81	82	83
G	91	92	93	94	95	96	97	98
H	14	28	42	56	70	84	98	1
Row\Column	9	10	11	12	13	14	15	16
A	9	10	11	12	13	14	15	17
B	24	25	26	27	28	29	30	33
C	39	40	41	42	43	44	45	49
D	54	55	56	57	58	59	60	65
E	69	70	71	72	73	74	75	81
F	84	85	86	87	88	89	90	97
G	99	100	101	102	103	104	105	1
H	97	81	65	49	33	15	17	106

23.4.2 Designs with a Cyclic Structure

In HTS practice, the (M, S)-optimal design in Section 23.4.1 may be difficult for experimental setup because of the complicated locations of repeated treatments. Special attention or programming is needed for the experimenter or the robotic distributor to bring repeated treatments to their specific locations in the microplate. Plate designs with a simple allocation of repeated treatments will definitely simplify the distribution process. In this section, we construct a type of row-column designs that is treatment-connected but has a simpler structure. Consider allocating $v = bk - b - k + 2$ treatments, T_1, T_2, \dots, T_v , to the bk experimental units in $b(\geq 2)$ rows and $k(\geq 2)$ columns in Table 23.9. Such designs were proposed in Qu (2011b) where treatments $T_1, T_2, \dots, T_k, T_{2k-1}, T_{3k-2}, \dots, T_{(b-2)(k-1)+1}$, and T_v are replicated twice and others are nonreplicated.

TABLE 23.9

Row-Column Designs of $b \times k$ with a Cyclic Structure

Row\Column	1	2	...	$k - 1$	k
1	T_1	T_2	...	T_{k-1}	T_k
2	T_k	T_{k+1}	...	$T_{2(k-1)}$	T_{2k-1}
3	T_{2k-1}	T_{2k}	...	$T_{3(k-1)}$	T_{3k-2}
⋮	⋮	⋮	⋮	⋮	⋮
$b - 1$	$T_{(b-2)(k-1)+1}$	$T_{(b-2)(k-1)+2}$...	$T_{(b-1)(k-1)}$	T_v
b	T_v	T_1	...	T_{k-2}	T_{k-1}




FIGURE 23.5 Efficiencies of row-column designs with eight rows.

Figure 23.5 displays *A*, *D*, and *E* efficiencies of row-column designs of 8 rows and 8–20 columns, where “Cyclic design, *A*” stands for the *A*-efficiency of the design in Table 23.9, “*(M, S)*-optimal, *D*” is the *D* efficiency of the *(M, S)*-optimal design in Section 23.4.1, and “Nonbinary, *E*” is the *E* efficiency of the design studied in Table 23.3. As described in Section 23.1, the *A* efficiency of a design is the harmonic mean of its canonical efficiency factors or the nonzero eigenvalues of matrix $R_d^{-1/2} C_d R_d^{-1/2}$ where C_d and R_d are defined in (23.2) and (23.3), the *D* efficiency is the geometric mean of those nonzero eigenvalues, and the *E* efficiency is the minimum nonzero eigenvalue of $R_d^{-1/2} C_d R_d^{-1/2}$. It is observed that *(M, S)*-optimal designs have the highest *A* and *D* efficiencies among all three types of designs but have lower *E* efficiencies than nonbinary designs when $18 \leq k \leq 20$. Nonbinary designs have higher *A* and *E* efficiencies than cyclic designs. For $12 \leq k \leq 20$, cyclic designs are the worst in terms of *A*, *D*, and *E* efficiencies.

If we treat controls and treatments equally, it is desirable to estimate all elementary contrasts with the same precision. A design is said to be balanced if all elementary contrasts are estimated with the same precision. A necessary and sufficient condition for a treatment-connected design to be balanced is that all nonzero eigenvalues of its information matrix are equal (Rao 1958).

None of the designs in Table 23.9, the *(M, S)*-optimal, and the nonbinary designs with one control are equireplicated and balanced. To see how unbalanced these designs are, Qu (2011b) introduced a balance index for row-column designs, that is, $\Delta_d(b, k) = b^2 k^2 \text{trace}(C_d^2) - (b k \text{trace}(C_d))^2 / (v - 1)$. For a row-column design *d*, $\Delta_d(b, k)$ measures its deviation from balance. $\Delta_d(b, k) = 0$ is a necessary and sufficient condition for a row-column design *d* to be balanced. Therefore, it is reasonable to say that a row-column design with a small value of $\Delta_d(b, k)$ is more balanced than that with a large value of $\Delta_d(b, k)$. Qu (2011b) showed that cyclic designs in Table 23.9 are more balanced than the

(M, S) -optimal designs in Section 23.4.1 but less balanced than the nonbinary designs in Table 23.3.

Kshirsagar (1958) showed that balanced designs are the most efficient among the class of equireplicated designs of the same size. Recall that the cyclic designs are less A -, D -, and E -efficient than the (M, S) -optimal and the nonbinary designs for $12 \leq k \leq 20$ in Figure 23.5. It seems that Kshirsagar's result does not hold for unequally replicated row-column designs. The relationship between balance and efficiency is complex for unequally replicated row-column designs because balance is based on information matrix C_d but A -, D -, and E efficiencies are based on $R_d^{-1/2} C_d R_d^{-1/2}$.

23.5 Conclusions

We have discussed saturated row-column designs for HTS experiments. These designs use the minimum number of experimental units to compare the maximum number of treatments in a row-column layout. Unlike most row-column designs in the literature where many combinatorial restrictions have been put on numbers of rows and columns, designs considered here can be constructed for any dimensions. This is extremely important to HTS practice because microplates of various dimensions have been manufactured and used in HTS experimentation. All row-column designs constructed are treatment-connected so that treatment effects can be separated from the contamination of row and column effects. All row-column designs constructed in Sections 23.4.1 and 23.4.2 are binary, that is, no treatments are repeated in any row or column.

As is pointed out by Morgan and Parvu (2007), finding A - and E -optimal row-column designs with a large number of nonreplicated treatments is computationally impractical at the current stage. When it is not feasible to find A - and E -optimal designs, the designs we proposed can be good surrogates in HTS experimentation.

This chapter has shown that unequally replicated row-column designs play an important role in the HTS experimentation for drug discovery. The (M, S) -optimal designs constructed in this chapter are based on the information matrix C_d . In addition to finding A - and E -optimal designs, finding (M, S) -optimal designs based on $R_d^{-1/2} C_d R_d^{-1/2}$ is an interesting and challenging open question for unequally replicated row-column designs.

References

- An, W.F. and Tolliday, N.J. (2009), Cell-based assays for high-throughput screening, *Methods in Molecular Biology*, 486, 1–12.
- Bradu, D. and Hawkins, D.M. (1982), Location of multiple outliers in two-way tables, using tetrads, *Technometrics*, 24(2), 103–108.
- Brideau, C., Gunter, B., Pikounis, B., and Liaw, A. (2003), Improved statistical methods for hit selection in high-throughput screening, *Journal of Biomolecular Screening*, 8, 634–647.
- Bühlmann, P. and van de Geer, S. (2011), *Statistics for High-Dimensional Data: Methods, Theory and Applications*, Springer-Verlag, Berlin, Germany.
- Cheng, C.S. (1978), A note on (M, S) -optimality, *Communications in Statistics-Theory and Methods*, 7, 1327–1338.

- Cheng, C.S. (1981), Optimality and construction of pseudo-Youden designs, *Annals of Statistics*, 9, 201–205.
- Devlin, J.P. (1997), *High Throughput Screening: The Discovery of Bioactive Substances*, Marcel Dekker, New York.
- Dey, A., Shah, K.R., and Das, A. (1995), Optimal block designs with minimal and nearly minimal number of units, *Statistica Sinica*, 5, 547–558.
- Hamada, M. and Balakrishnan, N. (1998), Analyzing unreplicated fractional factorials: A review with some new proposals, *Statistica Sinica*, 8, 1–41.
- Heuer, C., Haenel, T., and Prause, B. (2005), A novel approach for quality control and correction of HTS data based on artificial intelligence, Pharmaceutical discovery & development report, PharmaVentures, Oxford, U.K.
- Hsiao, A.Y., Tung, Y.-C., Qu, X., Patel, L.R., Pienta, K.J., and Takayama, S. (2012), 384 hanging drop arrays give excellent Z-factors and allow versatile formation of co-culture spheroids, *Biotechnology and Bioengineering*, 109, 1293–1304.
- Hüser J., Mannhold, R., Kubinyi, H., and Folkers, G. (2006), *High-Throughput Screening in Drug Discovery*, Methods and principles in medicinal chemistry, Vol. 35, Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, Germany.
- Jacroux, M. (1990), Some E-optimal row-column designs having unequally replicated treatments, *Journal of Statistical Planning and Inference*, 26, 65–81.
- John, J. A. and Williams, E. R. (1995), *Cyclic and Computer Generated Designs*, 2nd ed. Monographs on statistics & applied probability, Chapman & Hall/CRC, London.
- Kerr, M. K. (2003), Design considerations for efficient and effective microarray studies, *Biometrics*, 59, 822–828.
- Kevorkov, D. and Makarenkov, V. (2005), Statistical analysis of systematic errors in HTS, *Journal of Biomolecular Screening*, 10, 557–567.
- Kiefer, J. (1975), Construction and optimality of generalized Youden designs, in *A Survey of Statistical Designs and Linear Models*, Ed. J.N. Srivastava, pp. 333–353, North-Holland, Amsterdam, the Netherlands.
- Kinader, K.K.J., Voss, D.T., and Wang, W. (2000a), “Exact confidence intervals in the analysis of nonorthogonal saturated factorial designs,” *American Journal of Mathematical and Management Sciences*, 20, 71–84.
- Kinader, K.K.J., Voss, D.T., and Wang, W. (2000b), Analysis of saturated and super-saturated factorial designs: A review, in *Proceeding of the Indian International Statistical Association 1998 International Conference*, Ed. N. Balakrishnan, pp. 325–347, Gordon and Breach, Newark, NJ.
- Kunert, J. (1997), On the use of the factor-sparsity assumption to get an estimate of the variance in saturated designs. *Technometrics*, 39, 81–90.
- Kshirsagar, A.M. (1958), A note on incomplete block designs, *The Annals of Mathematical Statistics*, 29, 907–910.
- Lin, C.S. and Poushinsky, G. (1983), A modified augmented design for an early stage of plant selection involving a large number of test lines without replication, *Biometrics*, 39, 553–561.
- Lundholt, B.K., Scudder, K. M., and Pagliaro, L. (2003), A simple technique for reducing edge effect in cell-based assays, *Journal of Biomolecular Screening*, 8, 566–569.
- Malo, N. Hanley, J.A. Cerquozzi, S., Pelletier, J., and Nadon, R. (2006), Statistical practice in high-throughput screening data analysis, *Nature Biotechnology*, 24, 167–175.
- Morgan, J.P. and Parvu, V. (2007), Optimal row-column design for three treatments, *Journal of Statistical Planning and Inference*, 137, 1474–1487.
- Nguyen, N.-K. (1997), Construction of optimal row-column designs by computer, *Computing Science and Statistics*, 28, 471–475.
- Ogunyemi, T., Qu, X., and Rao, M.B. (2007), On a nonbinary s-optimal design over a class of minimally connected binary row-column designs, *Journal of Applied Statistical Science*, 15, 345–354.
- Qu, X. (2010), Optimal row-column designs in high-throughput screening experiments, *Technometrics*, 52, 409–420.

- Qu, X. (2011a), Design and analysis of high-throughput screening experiments, *Journal of Systems Science and Complexity*, 24, 1–14.
- Qu, X. (2011b), Row-column designs with minimal units, *Journal of Statistical Planning and Inference*, 141, 3193–3200.
- Qu, X., Ogunyemi, T., and Kushler, R. (2010), A class of saturated row-column designs, *Journal of Statistical Planning and Inference*, 140, 781–790.
- Raghavarao, D. (1971), *Constructions and Combinatorial Problems in Design of Experiments*, John Wiley & Sons, Inc., New York.
- Rao, V.R. (1958), A note on balanced designs, *The Annals of Mathematical Statistics*, 29, 290–294.
- Sonnemann, E. (1985), U-optimum row-column designs for the comparison of two treatments, *Metrika*, 32, 57–63.
- Tibshirani, R. (1996), Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society, Series B*, 58, 267–288.
- Tukey, J.W. (1977), *Exploratory Data Analysis*, Addison-Wesley, Cambridge, MA.
- Wang, W. and Voss, D.T. (2001), On the analysis of nonorthogonal saturated designs using effect sparsity, *Statistics and Applications*, 3, 177–192.
- Whitaker, D., Williams, E.R., and John, J.A. (2008), CycDesignN: A package for computer generation of experimental designs, Version 5.1, <http://www.vsni.co.uk/software/cycdesign/>.
- Zhang, J.H., Chung, T.D.Y., and Oldenburg, K.R. (1999), A simple statistical parameter for use in evaluation and validation of high throughput screening assays, *Journal of Biomolecular Screening*, 4, 67–73.
- Zhang, J.H., Wu, X., and Sills, M.A. (2005), Probing the primary screening efficiency by multiple replicate testing: A quantitative analysis of hit confirmation and false screening results of a biochemical assay, *Journal of Biomolecular Screening*, 10(7), 695–704.
- Zhang, X.D. (2008), Novel analytic criteria and effective plate designs for quality control in genome-scale RNAi screens, *Journal of Biomolecular Screening*, 13, 363–377.
- Zhang, X.D. (2011), *Optimal High-Throughput Screening, Practical Experimental Design and Data Analysis for Genome-Scale RNAi Research*, Cambridge University Press, Cambridge, U.K.

24

Up-and-Down Designs for Dose-Finding

Nancy Flournoy and Assaf P. Oron

CONTENTS

24.1	Introduction.....	858
24.1.1	Background and Motivation.....	858
24.1.2	Basic Terminology.....	859
24.1.3	Transition Probability Matrix.....	859
24.1.4	Up-and-Down Design's Balance Point.....	860
24.2	Tour through First-Order Up-and-Down Designs.....	862
24.2.1	Classical Up-and-Down Rule.....	862
24.2.2	Biased Coin Up-and-Down Designs.....	864
24.2.2.1	Derman's Biased Coin Design.....	864
24.2.2.2	Durham and Flournoy's Biased Coin Design.....	864
24.2.3	Asymptotic Dose-Allocation Distribution.....	865
24.2.3.1	Calculating the Asymptotic Dose-Allocation Distribution for the Biased Coin and Classical Up-and-Down Designs.....	866
24.2.3.2	Unimodality of the Asymptotic Dose-Allocation Distribution and Its Relationship to the Balance Point.....	866
24.2.3.3	Sharpness of the Asymptotic Dose-Allocation Distribution.....	868
24.2.4	Up-and-Down Convergence.....	869
24.2.4.1	Convergence to Stationary Behavior.....	869
24.2.4.2	Convergence of the Empirical Allocation Frequencies.....	870
24.2.5	Estimating the Quantile of Interest.....	871
24.2.5.1	Reversal Averaging Estimators.....	871
24.2.5.2	Isotonic Regression-Based Estimators.....	873
24.3	Higher-Order Up-and-Down Designs and Other Extensions.....	873
24.3.1	Group Up-and-Down Designs.....	874
24.3.2	k -in-a-Row ("Geometric") Design.....	875
24.3.3	Biased Coin Design Extensions.....	877
24.3.4	Comparing Up-and-Down Designs.....	878
24.3.4.1	Convergence Rates.....	878
24.3.4.2	Peakedness of the Asymptotic Dose-Allocation Distribution.....	878
24.3.5	Up-and-Down Designs for Joint Studies of Toxicity and Efficacy.....	879
24.4	Estimation Following an Up-and-Down Design.....	880
24.4.1	Dose Averaging Estimators.....	880
24.4.1.1	Case Against Reversal-Only Estimators.....	881
24.4.1.2	Alternative Solutions for Averaging Estimators.....	882
24.4.2	Regression Estimators.....	882

24.5 Other Approaches to Dose-Finding.....	883
24.5.1 Long-Memory Dose-Finding Designs.....	885
24.5.2 Advantages and Drawbacks of Long-Memory Designs.....	886
24.6 Summary and Practical Recommendations.....	889
24.7 Historical Notes.....	890
References.....	891

24.1 Introduction

24.1.1 Background and Motivation

Up-and-down (UD) dose-finding designs are used extensively in a variety of scientific and engineering fields. The first technical documents describing UD designs date to the 1940s and involve military explosive testing to find the optimal height from which to drop a bomb (Anderson et al. 1946; Dixon and Mood 1948) and hearing-threshold determination (von Békésy 1947). In both of these applications, UD designs are still the method of choice. Other common UD applications include failure-threshold determination in electrical and material engineering (Lagoda and Sonsino 2004) and finding the median effective dose (ED_{50}) in anesthesiology (Pace and Stylianou 2007). As these examples suggest, *dose-finding* is a general term to describe experiments whose outcomes are binary (“yes/no”), conducted to find the treatment that would trigger a “yes” response at a prespecified frequency. UD procedures comprise a family of *sequential* dose-finding designs, meaning that treatments are ordered in time and each treatment (except the first one) is determined by previous outcomes rather than being predetermined before the experiment’s start.

In numerous applications, UD designs are a standard method (e.g., JSME 1981; ASTM 1991; OECD 1998; NIEHS 2001). Prominent among these are animal toxicity studies, which attempt to estimate the median lethal dose (LD_{50}) of various toxins while sacrificing as few animals as possible. Toxicity is also the outcome of interest in first-in-human (Phase I) clinical trials, which are dose-finding experiments attracting considerable methodological attention and innovation. The most popular Phase I design in current use is superficially related to UD and is often confounded with it. More generally, UD plays an important role in the statistical debate regarding Phase I design choices.

One obstacle to the proliferation of UD usage and best practices has been the lack of standard references, such as exists for statistical methods of comparable popularity. In this chapter, we attempt to provide a condensed version of such a standard reference, presenting an unprecedented compilation of UD methodological knowledge. The presentation is closer to a textbook than a review article. This includes the definition of numerous terms we deem crucial to the understanding and proper use of UD designs.

The remainder of the Introduction will present the basic terminology and define some essential concepts. Section 24.2 explores these concepts and demonstrates some basic UD properties, using two of the simplest UD designs. Section 24.3 presents UD design variants whose properties are more complicated, ways to compare UD designs, and some results from such comparisons. Section 24.4 discusses UD estimation methods. Section 24.5 provides a quick overview of other popular approaches to dose-finding and a basic

comparison between their properties and those of UD designs. The chapter ends with a brief summary, some general design recommendations, and historical notes.

24.1.2 Basic Terminology

One must begin with the definition of UD itself, because the term has been used rather loosely by different authors. We prefer the consensus interpretation of the term among modern UD methods researchers: UD designs are dose-finding designs which, under simple assumptions, generate dose-assignment sequences that are *Markov chains* over a discrete set of doses, \mathcal{X} . This chapter does not cover methods that operate on a continuous dose space.

Without loss of generality, we refer to treatment magnitudes as *doses* and to experimental outcomes as *toxicities*. Let $Y_i = 1$ if the i th subject (a plant, animal, or human) of a dose-finding experiment exhibits toxicity, and 0 otherwise ($i = 1, \dots, n$). Given that the dose to which the i th subject was exposed is $X_i = x$, the probability of toxicity is $F(x) = P\{Y_i = 1 \mid X_i = x\}$. Despite the experiment using only a discrete set of doses, the dose-magnitude variable itself, x , is assumed to be continuous, and the toxicity rate is assumed to increase continuously with increasing x . The experiment's goal is to estimate the exact dose x (on a continuous scale) that produces a fixed target toxicity rate $\Gamma = P\{Y = 1 \mid X = x\}$, $\Gamma \in (0, 1)$. This can be expressed as estimation of the quantile $F^{-1}(\Gamma)$ of a cumulative distribution function that models the dose-toxicity curve $F(x)$. The density function $f(x)$ associated with $F(x)$ is interpretable as the distribution of *toxicity thresholds* of the population under study. We focus on $\Gamma \leq 0.5$, the range of target rates used for most toxicity studies.

The general defining characteristics of a UD design follow:

1. Doses, labeled X_1, \dots, X_n , are administered to a sequence of subjects $1, \dots, n$.
2. The doses, X_1, \dots, X_n are restricted to a fixed set of M possible dose levels, which is denoted by

$$\mathcal{X} = \{d_1, \dots, d_M : d_1 < \dots < d_M\}.$$

3. Suppose that $X_i = d_m$, then $X_{i+1} \in \{d_{m-1}, d_m, d_{m+1}\}$, $i \in \{1, \dots, n\}$ according to simple constant rules based on recent toxicity responses, hence the name *up and down*. This means that the Markov chains generated by UD designs are also *random walks*.

24.1.3 Transition Probability Matrix

Given that a subject receives dose d_m , denote the probability that the next subject receives dose d_{m-1} , d_m , or d_{m+1} by $p_{m,m-1}$, p_{mm} or $p_{m,m+1}$, respectively. These *transition probabilities* obey the constraints $p_{m,m-1} + p_{mm} + p_{m,m+1} = 1$ and the boundary conditions $p_{1,0} = p_{M,M+1} = 0$. A specific set of UD rules enables the symbolic calculation of these probabilities, usually as a function of $F(x)$. Assume for now that transition probabilities are fixed in time, depending only upon the current allocation and its outcome, that is, upon (X_i, Y_i) and through them upon $F(x)$ (and possibly on a set of fixed parameters). The probabilities are then best represented via a tridiagonal transition probability matrix (TPM) P :

$$\mathbf{P} = \begin{pmatrix} p_{11} & p_{12} & 0 & \cdots & \cdots & 0 \\ p_{21} & p_{22} & p_{23} & 0 & \ddots & \vdots \\ 0 & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 \\ \vdots & \ddots & 0 & p_{M-1,M-2} & p_{M-1,M-1} & p_{M-1,M} \\ 0 & \cdots & \cdots & 0 & p_{M,M-1} & p_{MM} \end{pmatrix}. \tag{24.1}$$

Transition probabilities between any pair of doses d_j and d_m , administered either to consecutive subjects or to subjects separated by any number $l \geq 1$ of intervening subjects, are denoted as follows:

$$\begin{aligned}
 p_{jm} &= p_{jm}(1) = P\{X_{i+1} = d_m | X_i = d_j\}. \\
 p_{jm}(l) &= P\{X_{i+l} = d_m | X_i = d_j\}; \quad (j, m) \in \{1, \dots, M\}.
 \end{aligned}$$

These probabilities can be calculated via matrix multiplication: $p_{jm}(l) = \left(\mathbf{P}^l\right)_{jm}$, where \mathbf{P}^l denotes the matrix \mathbf{P} multiplied by itself l times. For a given j, m pair, if $p_{jm}(l) > 0$ and $p_{mj}(l) > 0$ for some $l > 0$, then dose levels d_j and d_m are said to *communicate*. If all levels communicate with each other, the matrix \mathbf{P} is called *regular*. Under realistic conditions, the TPMs of typical UD designs are regular. This condition guarantees the existence of positive stationary dose-allocation frequencies (see Section 24.2.3).

24.1.4 Up-and-Down Design’s Balance Point

UD designs generate Markov chains with a central tendency, meaning that dose assignments tend to meander back and forth around some dose that can be calculated from the design parameters (Durham and Flournoy 1994; Hughes 1995). This dose was recently termed the *balance point* by Oron and Hoff (2009). While the term and some related methodological considerations are recent, the intuition giving rise to designs with balance points is as old as the UD design itself.

Continuous *up* and *down* transition functions $p(x)$ and $q(x)$ are now defined such that $\{p_{m;m+1}\}$ and $\{p_{m;m-1}\}$, respectively, are points on these curves.

Definition 24.1 Consider a UD with an “up” transition probability $p_{m,m+1}$ that might vary over time, that is, with the accruing sample size $n \geq 1$.^{*} If, in the interior of \mathcal{X} (specifically, $1 \leq m < M$),

$$\lim_{n \rightarrow \infty} (p_{m,m+1}) = g_{\text{up}}(d_m),$$

where $g_{\text{up}}(d_m)$ is an algebraic function of d_m “only,” then the extension of $g_{\text{up}}(d_m)$ onto its natural continuous domain $\mathcal{D}_{\text{up}} \supseteq [d_1, d_M]$ will be called the “the ‘up’ transition function” $p(x)$,

^{*} A terminology note: for most designs discussed in this chapter, the random-walk transition probabilities are constant in time, and therefore, the use of “ $\lim_{n \rightarrow \infty}$ ” in Definition 24.1 is redundant. However, Section 24.3.2 presents a widely used UD method whose dose-transition probabilities only stabilize in the $n \rightarrow \infty$ limit. Therefore, we use a limit notation.

$x \in [d_1, d_M]$. The “down” function $q(x)$ is similarly defined from the “down” transition probability $p_{m,m-1}$, $1 < m \leq M$.

In words, $p(x)$ and $q(x)$ are defined on a continuous domain that includes \mathcal{X} , while $\{p_{m,m+1}\}$ and $\{p_{m,m-1}\}$ are defined only on \mathcal{X} itself.

As will be seen later, for practically all UD designs, the transition probabilities are simple functions of the dose-toxicity rate function $F(x)$ evaluated on \mathcal{X} . Therefore, their continuous extension as $p(x)$ and $q(x)$ is straightforward and enables a useful definition for the balance point:

Definition 24.2 Consider a UD design with “up” and “down” functions $p(x)$ and $q(x)$, respectively. If $p(x)$ is monotone decreasing and $q(x)$ monotone increasing in x , then the UD balance point is the dose x^* such that

$$x^* = \arg_x \{p(x) = q(x)\}. \quad (24.2)$$

The toxicity rate at x^* will be denoted F^* , that is, $F(x^*) = F^*$.

The dual monotonicity condition in Definition 24.2—in words, as the dose increases, the “up” probability goes down while the “down” probability goes up—will be referred to as *the Durham–Flournoy conditions*, after the first study to specify them in generic form (Durham and Flournoy 1994, 1995). As shown in Theorem 24.1, designs that meet these conditions generate a central tendency around their balance point. In UD experiments, x^* usually is not one of the doses in \mathcal{X} , but falls between two of the designs’ doses.

One would like the center of the dose-allocation distribution to be close to the experiment’s designated target quantile $F^{-1}(\Gamma)$. This can be accomplished by judicious selection of design parameters, so that $\Gamma \approx F^* \equiv F(x^*)$. An illustration of $p(x), q(x)$ and the balance point is shown in Figure 24.1 for the classical UD design and for a BCD design, both described in Section 24.2.2.

In dose-finding, in general and in UD designs in particular, it is easy to focus on the central tendency and on asymptotic properties and forget the magnitude of sampling variability that might be observed in practice. While UD designs do guarantee an x^* -centered random walk, it is still a random walk. Since most dose-finding applications use small samples, the relative magnitude of sampling variability is often rather substantial. In order to keep the reader cognizant of this variability and to illustrate the concept of a random walk on \mathcal{X} , Figure 24.2 shows the experimental trajectories generated by three $n = 30$ random draws from the same toxicity threshold distribution, on an $M = 8$ dose set, for each of the two designs described in Section 24.2. The plots follow the graphical conventions common to dose-finding studies, with filled circles representing toxicities and empty ones nontoxicities. Balance points are denoted by dashed horizontal lines. The squares denote *reversal points* in the experimental trajectories, whose properties are explored later.

We now describe a few first-order designs, called thus because they generate first-order Markov chains. For these chains, the probability of transiting to an adjacent dose in \mathcal{X} depends upon the dose-allocation history, only through the current dose level and the current experimental outcome.




FIGURE 24.1

The “up” and “down” transition probability functions $p(x)$ (solid line) and $q(x)$ (dashes), for classical UD (a) and BCD with $F^* = \Gamma = 0.3$ (b), using hypothetical logistically distributed toxicity threshold scenario and $M = 10$ levels. The crosshairs denote $x = x^*$ and $y = F^*$.

24.2 Tour through First-Order Up-and-Down Designs

24.2.1 Classical Up-and-Down Rule

The original (hereafter, “classical”) UD design follows this rule: given $X_i = d_m$,

$$X_{i+1} = \begin{cases} d_{m+1} & \text{if } Y_i = 0; \\ d_{m-1} & \text{if } Y_i = 1, \end{cases}$$

for $m = 2, \dots, M - 1$. In words, the experiment moves *up* one level following a nontoxicity and *down* one level following a toxicity, hence the name *up and down*. On the boundaries of \mathcal{X} , replace d_0 with d_1 whenever the former is mandated and similarly replace d_{M+1} with d_M (i.e., the experiment stays on the boundary rather than venture outside it, which is impossible by design). Subsequently, boundary conditions are omitted since they always follow this basic adjustment.

The classical UD remains the most commonly used UD design. In the interior of \mathcal{X} , it leads to the following transition probabilities:

$$\begin{aligned} p_{m,m+1} &= P\{Y_i = 0 | X_i = d_m\} = 1 - F(d_m). \\ p_{m,m-1} &= P\{Y_i = 1 | X_i = d_m\} = F(d_m). \end{aligned} \tag{24.3}$$

The balance point (24.2) is the dose x^* whose toxicity rate F^* obeys

$$1 - F^* = F^* \longrightarrow F^* = 0.5.$$

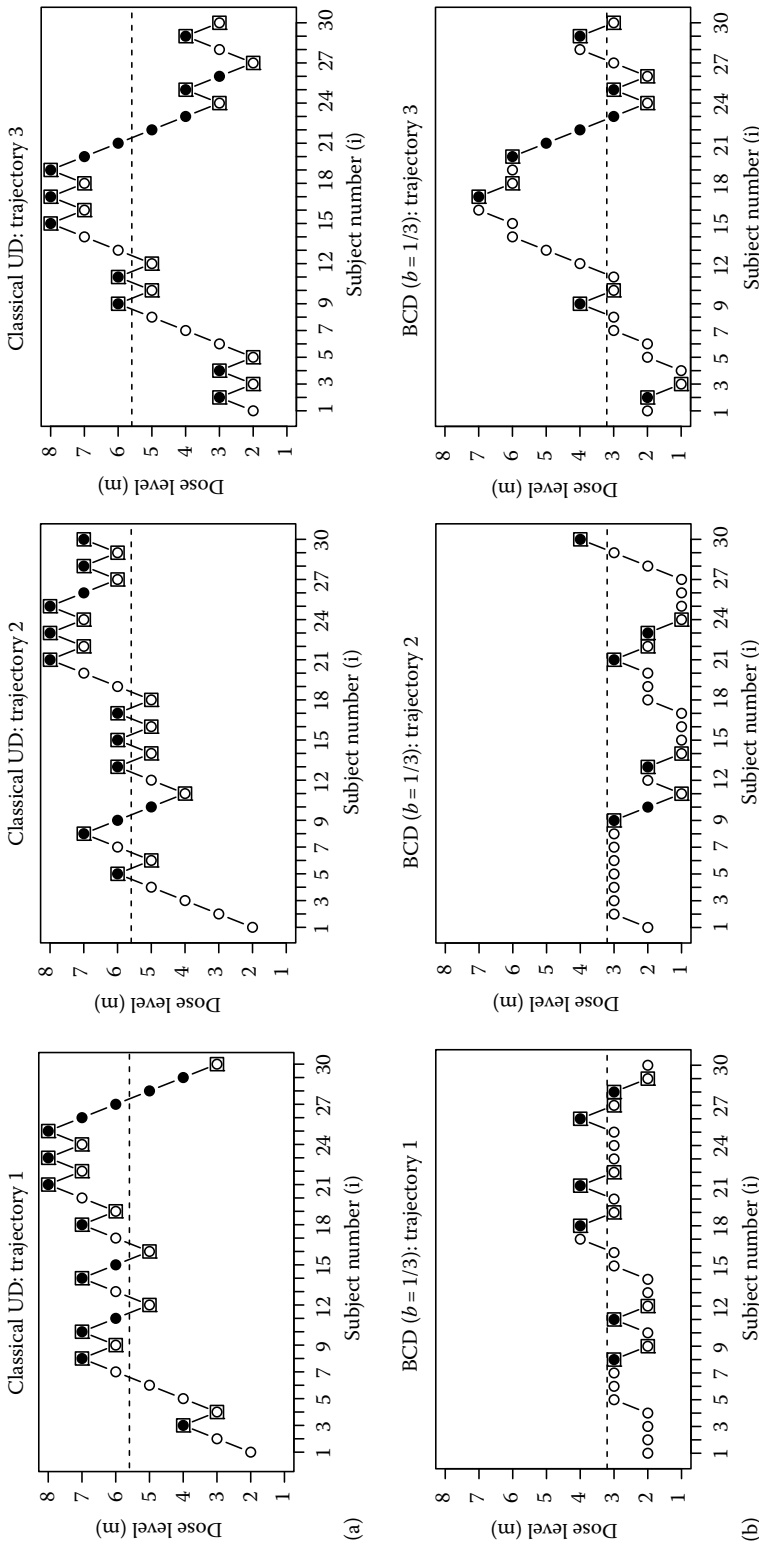


FIGURE 24.2 Randomly generated 'experimental' trajectories of classical UD for median estimation (a) and BCD with $b = 1/3$ for estimating the 25th percentile (b), under a gamma-distributed toxicity threshold density f and $M = 8$ levels. All runs begin at $X_1 = d_2$. Toxicity responses are marked with filled circles and nontoxicities with empty circles. The balance point x^* is marked by a horizontal dashed line. The squares denote *reversal* points.

As expected from its symmetric transition rules, the classical UD design is best suited for estimating the median toxicity threshold.

24.2.2 Biased Coin Up-and-Down Designs

24.2.2.1 Derman’s Biased Coin Design

Derman (1957) was the first to develop a UD extension for nonmedian percentiles. His design requires a random toss of a metaphoric *coin*, with $b = P\{\text{heads}\} \in [0.5, 1]$. Given $X_i = d_m$,

$$X_{i+1} = \begin{cases} d_{m+1} & \text{if } Y_i = 0, \text{ or} \\ d_{m+1} & \text{if } Y_i = 1 \text{ and tails;} \\ d_{m-1} & \text{if } Y_i = 1 \text{ and heads.} \end{cases}$$

The balance point x^* is given by

$$\begin{aligned} p(x^*) &= q(x^*); \\ 1 - F(x^*) + (1 - b)F(x^*) &= bF(x^*); \\ F^* = F(x^*) &= \frac{1}{2b} \in [0.5, 1]. \end{aligned} \tag{24.4}$$

For example, to produce a central tendency around the dose $F^{-1}(2/3)$, set $b = 0.75$. When $b = 1$, the balance point is the median, and indeed this value makes Derman’s design identical to the classical UD design. A mirror-image version of the design (i.e., de-escalate after a toxicity and toss a *coin* after a nontoxicity) can be used to estimate percentiles below the median.

24.2.2.2 Durham and Flournoy’s Biased Coin Design

Derman’s design has a practically unappealing property: if the coin points *tails*, the experiment will move in the direction opposite to that indicated by the most recent response (i.e., the dose will escalate despite having just observed a toxic response). This is especially undesirable in clinical trials, where such transitions might actually halt the experiment due to safety concerns (e.g., Neuenschwander et al. 2008). Not surprisingly, Phase I researchers prefer designs that preclude such occurrences, which Cheung (2005) called *incoherent* dose transitions.*

Durham and Flournoy (1994) suggested a design that uses a biased coin that centers the dose-allocation distribution around any percentile without incoherent transitions. Given $X_i = d_m$, the next dose allocation will be

$$X_{i+1} = \begin{cases} d_{m+1} & \text{if } Y_i = 0 \text{ and heads;} \\ d_{m-1} & \text{if } Y_i = 1; \\ d_m & \text{if } Y_i = 0 \text{ and tails.} \end{cases} \tag{24.5}$$

* Not to be confused with the usage of *coherence* in Bayesian theory.

The *heads* probability b can take any value in $[0, 1]$. The balance point is given by

$$\begin{aligned}
 b(1 - F^*) &= F^*, \\
 F^* &= \frac{b}{1 + b} \in [0, 0.5].
 \end{aligned}
 \tag{24.6}$$

Given a desired toxicity rate Γ , the BCD balance point can be made identical to $F^{-1}(\Gamma)$ by setting the *heads* probability to $b = \Gamma / (1 - \Gamma)$. For example, for $\Gamma = 0.3$ (commonly used in Phase I cancer trials), set $b = 3/7$. Once again, setting $b = 1$ makes this design identical to the classical UD, and inverting the rules produces above-median balance points. Hereafter, we refer to this design simply as the BCD.

The conditioning of dose escalation upon the outcome of a random draw, in addition to a nontoxicity, keeps experimental trajectories at lower doses, on average, than the classical UD design does. In Figure 24.2b, each BCD run with $b = 1/3$ (i.e., $F^* = 0.25$) uses exactly the same toxicity thresholds as the classical UD trajectory shown (i.e., $F^* = 0.50$) immediately above it in the top row (all drawn randomly from a Gamma distribution). While the classical UD design runs escalate rather easily to the upper half of the $M = 8$ dose range, the BCD runs are confined mostly to the bottom half.

24.2.3 Asymptotic Dose-Allocation Distribution

Because UD sequences of assigned doses x_1, \dots, x_n possess the properties of a Markov chain, as $n \rightarrow \infty$, the relative proportions of subjects assigned to each dose in \mathcal{X} approach the *asymptotic* or *stationary allocation distribution* $\boldsymbol{\pi} = (\pi_1, \dots, \pi_M)'$; $\sum_m \pi_m = 1$. If \mathbf{P} is regular, then $\pi_m > 0 \forall m$. The distribution $\boldsymbol{\pi}$ can be found by solving the M *balance equations*, which are greatly simplified by the tridiagonal structure of \mathbf{P} for first-order UD designs (24.1):

$$\pi_m = \pi_{m-1}p_{m-1,m} + \pi_m p_{mm} + \pi_{m+1}p_{m+1,m}, \quad m = 2, \dots, M - 1.
 \tag{24.7}$$

The solution is a recursive formula for calculating $\boldsymbol{\pi}$ from the transition probabilities:

$$\pi_{m+1} = \lambda_m \pi_m, \quad \lambda_m = \frac{p_{m,m+1}}{p_{m+1,m}}, \quad m = 1, \dots, M - 1,
 \tag{24.8}$$

where λ_m , the *adjacent-level ratio*, is the ratio between the single-step probabilities of escalating from d_m and de-escalating to d_m . To ensure that $\sum_m \pi_m = 1$, the base-level frequency π_1 is set by

$$\pi_1^{-1} = 1 + \sum_{m=1}^{M-1} \prod_{j=1}^m \lambda_j.$$

24.2.3.1 Calculating the Asymptotic Dose-Allocation Distribution for the Biased Coin and Classical Up-and-Down Designs

For the BCD, calculation of the adjacent-level ratio λ_m is straightforward:

$$\lambda_m = \frac{p_{m,m+1}}{p_{m+1,m}} = \frac{b(1 - F(d_m))}{F(d_{m+1})}. \tag{24.9}$$

The numerator is monotone decreasing in m , while the denominator is increasing, and therefore, λ_m as a whole is monotone decreasing. The doses for which $\{\lambda_m\}$ straddle unity bound the balance point x^* .

Interestingly, the classical UD presents a somewhat nonstandard asymptotic allocation distribution. Since $p_{mm} = 0$ except on the boundaries, classical UD Markov chains are *quasiperiodical*. For example, suppose the experiment begins at d_5 , and $M = 10$; then for subjects with an odd index i , the allocated dose level m will be odd and vice versa—unless and until the trajectory hits upon a boundary and remains there for two consecutive subjects. If the design had no boundaries, then the chain would be purely periodical with a period of 2, yielding two noncommunicating stationary distributions $\pi^{(odd)}, \pi^{(even)}$. With finite \mathcal{X} , the classical UD design has a single π with the monotone decreasing adjacent-level ratio $\lambda_m = (1 - F(d_m))/F(d_{m+1})$.

24.2.3.2 Unimodality of the Asymptotic Dose-Allocation Distribution and Its Relationship to the Balance Point

If λ_m is monotone decreasing in m , as was just shown for the classical UD design and BCD, then it immediately follows that π is *unimodal*:

- If $\lambda_1 \leq 1$ (or $\lambda_M \geq 1$), then $\pi_1 \geq \pi_2 \geq \dots \geq \pi_M$ (or vice versa), and π has a single mode on the lower (upper) boundary.
- If $\lambda_1 \geq 1 \geq \lambda_M$, then the π_m initially increase as m increases; then once $\lambda_m \leq 1$, they monotonically decrease—thereby creating a single mode.

Unimodality of π is very useful if the mode’s location is close to x^* and if—as stipulated earlier—the design is chosen so that $F^* \approx \Gamma$, because then the UD information collection rate about $F(x)$ eventually peaks around the percentile of interest (see, e.g., Figure 24.3).* The following theorem guarantees this property and is a cornerstone of UD design theory:

Theorem 24.1 Consider a UD design generating a Markov chain, with an asymptotic allocation distribution $\pi = (\pi_1, \dots, \pi_M)$ over \mathcal{X} , with $p(x)$ ($q(x)$) strictly monotone decreasing (increasing) in x , and a balance point x^* as defined in (24.2). If there exist two adjacent levels d_{m^*} and d_{m^*+1} such that $x^* \in [d_{m^*}, d_{m^*+1}]$, then π ’s mode is either at d_{m^*} or at d_{m^*+1} . Otherwise, the mode is on the boundary closest to x^* .

* As long as $F(x)$ is strictly increasing, and transition probabilities are strictly monotone functions of $F(x)$ (as is the case for both classical UD and BCD), then scenarios when $\lambda_m = 1$ will be all but nonexistent. If such a case does occur, the maximum value in π will be shared by two adjacent levels, called a *modal set*, which together can be still seen as forming a single mode. In other words, qualitatively, the design will still behave as do designs with single-dose mode.




FIGURE 24.3 The asymptotic dose-allocation distribution π for the classical UD (a) and BCD with $\Gamma = 0.3$ (b), using the same \mathcal{X} and F as in Figure 24.1.

Proof: First, assume that $x^* \in [d_{m^*}, d_{m^*+1}]$. Under the specified monotonicity conditions, the following inequalities always hold for any $2 \leq m < M$:

$$\lambda_{m-1} = \frac{p(x)|_{x=d_{m-1}}}{q(x)|_{x=d_m}} > \frac{p(x)|_{x=d_m}}{q(x)|_{x=d_m}} > \frac{p(x)|_{x=d_m}}{q(x)|_{x=d_{m+1}}} = \lambda_m. \tag{24.10}$$

The ratio in the middle of this inequality is monotone decreasing in the dose x and becomes exactly 1 at $x = x^*$. Therefore, for all $m < m^*$, $\lambda_m > 1$ and for all $m > m^*$, $\lambda_m < 1$. This means that the mode of π has to be at either d_{m^*} or d_{m^*+1} .

On the lower boundary of \mathcal{X} , if $x^* < d_1$, then $\lambda_1 = p_{12}/p_{21} < p(x^*)/q(x^*) = 1$, and therefore, the mode is at d_1 . The upper-boundary case is analogous. \square

Versions of Theorem 24.1 in varying degrees of generality have appeared in the literature from Derman (1957) onward. The aforementioned formulation used is adapted from Oron and Hoff (2009). In plain terms, the mode of π is guaranteed to be at one of the two dose levels straddling the balance point x^* .

Whether the asymptotic mode is in fact on the single closest level to x^* depends upon the design and upon how *closest* is defined. Defining *closeness* on the toxicity frequency (F) scale, rather than on the dose (x) scale, is generally more useful. For example, the following property was proven for the classical UD:

Theorem 24.2 (Oron and Hoff 2009) *For the classical UD design, π 's mode is at the level whose toxicity rate is closest to 0.5.*

Proof: Let $x^* = F^{-1}(0.5) \in [d_{m^*}, d_{m^*+1}]$. Write $F(d_{m^*}) = 0.5 - \Delta p_1$ and $F(d_{m^*+1}) = 0.5 + \Delta p_2$, where $\Delta p_1 \geq 0$ and $\Delta p_2 \geq 0$. Then the adjacent-level ratio (24.8) at level m^* can be written as

$$\lambda_{m^*} = \frac{p_{m^*, m^*+1}}{p_{m^*+1, m^*}} = \frac{0.5 + \Delta p_1}{0.5 + \Delta p_2}. \tag{24.11}$$

Because $\lambda_{m^*} \geq 1$ if and only if $\Delta p_1 \geq \Delta p_2$, the result immediately follows from this inequality. □

For the BCD, the mode is not necessarily the dose closest to the target. Let $x^* = F^{-1}(\Gamma)$, $\Gamma < 0.5$, and cast the BCD adjacent-level ratio (24.9) into a form analogous to (24.11), to obtain

$$\lambda_{m^*} = \frac{p_{m^*, m^*+1}}{p_{m^*+1, m^*}} = \frac{\Gamma(1 - \Gamma + \Delta p_1)}{(1 - \Gamma)(\Gamma + \Delta p_2)}.$$

This adjacent-level ratio is greater than 1 if and only if $\Delta p_1 / \Delta p_2 > (1 - \Gamma) / \Gamma$. For example, if $\Gamma = 1/4$ and $x^* \in [d_2, d_3]$, then the asymptotic mode will be at d_2 unless $F(d_3)$ is three times closer to $1/4$ than $F(d_2)$. This means that the BCD's dose-allocation behavior is somewhat more conservative (i.e., less amenable to dose escalation) than indicated by its nominal target toxicity rate.

24.2.3.3 Sharpness of the Asymptotic Dose-Allocation Distribution

The unimodality of π helps concentrate information collection around x^* . It is also of practical interest to know how steeply π decreases to either side of its mode, in order to know the degree of this concentration and the frequency of excursions away from x^* . Assume, without loss of generality, that the mode is at a single interior level d_{m^*} . Due to monotonicity conditions, we know that $\lambda_{m^*-1} > 1$ and $\lambda_{m^*} < 1$. If the adjacent-level ratios remained constant at these values on either side of the mode, then the rate of decrease in π values away from the mode would be geometric—the discrete analogue of exponentially decreasing tails. However, due to their monotonicity, the λ 's become *more* extreme as one moves away from the mode. Hence, the drop in asymptotic allocation frequencies away from π 's mode is faster than geometric for UD designs.

Durham and Flournoy (1994) obtained a valuable result (shown in Theorem 24.3) regarding the form of the BCD's π and the steepness of this decrease away from the mode. This result requires the following definition:

Definition 24.3 Let Z be a random variable that is defined on a set of discrete real-valued points \mathcal{X} . Let $\tilde{\Phi}(z)$ denote $P\{Z = z\}$. Then Z is said to have a discrete normal distribution if

$$\tilde{\Phi}(z_j) = \frac{e^{-\frac{1}{2}z_j^2}}{\sum_{i=1}^M e^{-\frac{1}{2}z_i^2}}, \quad z_j \in \mathcal{X}.$$

Theorem 24.3 Assume that subjects' toxicity thresholds are logistically distributed, with

$$1 - F(x) = (1 + \exp(\mu + \sigma x))^{-1}, \quad \sigma > 0,$$

and that the doses are equally spaced: $\mathcal{X} = \{d_1, d_1 + \Delta, d_1 + 2\Delta, \dots, d_1 + (M - 1)\Delta\}$. Then the BCD's asymptotic allocation distribution is a mixture of two discrete normal distributions over \mathcal{X} , for which the mixing parameter is the target toxicity rate Γ , that is,

$$\pi_m = (1 - \Gamma)\tilde{\Phi}\left(\frac{d_m - (x^* - 0.5\Delta)}{\sqrt{\Delta/\sigma}}\right) + \Gamma\tilde{\Phi}\left(\frac{d_m - (x^* + 0.5\Delta)}{\sqrt{\Delta/\sigma}}\right), \quad d_m \in \mathcal{X}. \quad (24.12)$$

See Durham and Flournoy (1994) for the complete proof. The modes of the two discrete normal distributions are separated by one dose spacing unit Δ , so they have considerable overlap. The first (second) component of (24.12) is the asymptotic allocation distribution for subjects with nontoxic (toxic) responses. If $\Gamma = 0.5$, that is, the design is the classical UD procedure, the mixture is equally weighted, and, as shown in more generality in Theorem 24.2, the asymptotic mode will be on the dose whose toxicity rate is closest to 0.5. Conversely, as $\Gamma \rightarrow 0$, the first component will dominate, shifting the asymptotic mode as far as $\Delta/2$ downward.

Despite the specific parametric assumptions on F , (24.12) provides a useful guideline about the stationary frequency of BCD excursions, because many common response function models are very similar except on the extreme tails and because (as will be shown later) several common UD designs produce π s at least as sharp as those of the BCD. Under the assumptions of Theorem 24.3, roughly 97.5% of the allocations will be within $x^* \pm 2\sqrt{\Delta/\sigma}$. This approximation can be used to choose Δ so as to control for highly toxic events, given approximate prior knowledge of the dispersion parameter σ (Durham et al. 1997).

24.2.4 Up-and-Down Convergence

24.2.4.1 Convergence to Stationary Behavior

The presentation of asymptotic UD properties has thus far sidestepped an important practical aspect. Dose-finding experiments are by their nature small-to-medium-sample affairs ($n < 100$ and more often than not $n < 30$), so asymptotic behavior would generally not dominate the magnitude of observed variability—as we reminded the reader at the close of Section 24.1. So why bother about asymptotic behavior at all? Indeed, in recent years, dose-finding analyses have rarely examined asymptotic behavior, relying almost exclusively upon numerical simulations instead.

As Oron et al. (2011) argue, this recent view is misguided. An obvious practical reason is that simulations are, generally speaking, haphazard rather than systematic samples from the space of all relevant dose-toxicity scenarios and are therefore prone to misuse and abuse—as anyone who has read enough simulation-based dose-finding articles can attest. But an even stronger, *conceptual* reason is that all sequential dose-finding designs rely upon some mechanism for concentrating dose allocations near the target. How good is the mechanism?

A straightforward test is to allow the design to run indefinitely and characterize its behavior. A design with a good concentrating mechanism must demonstrate reasonably desirable behavior as $n \rightarrow \infty$. Conversely, if desirable asymptotic behavior *cannot* be guaranteed or if the type of convergence or conditions for convergence make asymptotic behavior all but impossible for dose-finding applications, then the design should probably not be trusted even for small-to-medium n . Hence we turn our attention now to the importance of asymptotics (discussed earlier) and convergence.

With sequential procedures such as UD designs, the term *convergence* carries many different meanings. Results have been published about the convergence of UD dose-allocation frequencies (Durham and Flournoy 1995; Durham et al. 1995) and about the convergence of observed toxicity rates (Flournoy et al. 1995; Oron et al. 2011). A third type of convergence is the convergence of *the dose-allocation process itself*, which in the case of UD designs is convergence to stationary sampling from π .

For Markov chains, convergence of the dose-allocation process to stationary behavior erases, or *mixes*, the arbitrary initial conditions. A deterministically chosen starting dose can be represented by an allocation distribution ρ , which is 0 everywhere and 1 at $d_m = x_1$. In the case of a completely random starting point, ρ has a discrete-uniform distribution over \mathcal{X} . The allocation distribution after the first subject can be written in matrix form as $\rho'P$ (see (24.1)). The probability distribution of the $n+1$ -th assignment, given the starting conditions only, is $\rho'P^n$ (cf., Diaconis and Stroock 1991). All rows of P^n become arbitrarily close to π as $n \rightarrow \infty$, meaning that from some n onward, the marginal allocation distribution of subsequent doses becomes essentially stationary and independent of the initial conditions ρ . Specifically,

$$\max_m |P_m^n - \pi'| \propto \beta_2^n, \quad m = 1, \dots, M, \tag{24.13}$$

where $|\cdot|$ denotes Euclidean distance, P_m^n is the m th row of P^n , and β_2 is the eigenvalue of P with the second-largest absolute value.*

Since it is known that $|\beta_2| < 1$, the allocation-distribution convergence rate is geometric in n —far faster than the usual root- n rate of the laws of large numbers and many other processes.

The practical implication of geometric convergence is that while individual UD allocation sequences might vary (see Figure 24.2), the impact of initial conditions is erased at a rate fast enough to bear fruit even during a small-sample experiment. Oron and Hoff (2009) translated the geometric rate to terms relevant for experimentalists. They progressed the allocation distribution from the initial conditions $X_1 = d_1$ and observed how many subjects it would take the expectation $E_{(n)}(X)$ of the allocation distribution at subject n to reach 99% of the way from d_1 to the stationary mean $E_{(\infty)}(X) = \sum_m \pi_m d_m$. Several UD designs targeting $\Gamma = 0.3$ were examined under several dose-toxicity scenarios each. With $M = 5$ levels, it took 10–20 subjects to arrive 99% of the way to $E_{(\infty)}(X)$, and with $M = 10$ about twice as long.

24.2.4.2 Convergence of the Empirical Allocation Frequencies

The convergence of *empirical* dose allocations at the individual dose levels is also of great practical importance. For example, the quality of the isotonic regression estimator of the

* The largest absolute eigenvalue of P is 1, associated with the stationary behavior itself.

target percentile $F^{-1}(\Gamma)$ (see the following and Section 24.4) depends upon this distribution being clustered around the target. We know that if $F^* \approx \Gamma$, then asymptotically, this is the case. But is this true for real-life sample sizes?

Let $W(n) = (W_1(n), \dots, W_M(n))$ denote the empirical allocation distribution, that is, the proportion of subjects allocated to each level when n subjects have been allocated in total. Then for UD designs, $W_m(n) \rightarrow \pi_m$ as $n \rightarrow \infty$ almost surely, by the law of large numbers for regular Markov chains. The asymptotic distribution of W is determined by the central limit theorem (CLT) for regular Markov chains:

$$\sqrt{n}((W_1(n), \dots, W_M(n)) - \boldsymbol{\pi}) \xrightarrow{n \rightarrow \infty} \mathcal{N}\left(0, \text{diag}\left[\sigma_1^2(\mathbf{P}), \dots, \sigma_M^2(\mathbf{P})\right]\right). \quad (24.14)$$

The rate of convergence is of order $O(1/\sqrt{n})$ —the usual CLT rate. This is one reason why the empirical mode of allocated doses, examined in some recent studies as an estimator for $F^{-1}(\Gamma)$, performs rather poorly in that role despite its asymptotic consistency (Giovagnoli and Pintacuda 1998). At $n = 30$ or even 60, the variability in allocated proportions is too high to reliably equate the empirical mode with the stationary mode.

Note: from a practical perspective, choosing a design with smaller asymptotic variances $(\sigma_1^2(\mathbf{P}), \dots, \sigma_M^2(\mathbf{P}))$ is equivalent to increasing the scaling constant of the convergence rate, that is, accelerating the convergence. Both will help $W(n)$ achieve a prespecified precision threshold with fewer subjects. However, the overall root- n convergence rate cannot be changed, and it is far slower than the convergence of the dose-allocation process itself to stationary behavior.

24.2.5 Estimating the Quantile of Interest

More often than not, estimation of UD experiments has been approached inconsistently by both applied and methodological researchers. Many researchers use outdated and questionable estimators of $F^{-1}(\Gamma)$, just because these estimators happen to be prevalent in their field. Many method developers illogically produce UD estimates in the same way as estimates from an entirely different class of dose-finding designs (e.g., Zacks 2009). To help alleviate the confusion, Section 24.4 is dedicated mostly to UD estimation. Right now, a brief description is provided for two estimators of $F^{-1}(\Gamma)$: the one most commonly used for the classical UD design and a more recent estimator considered as standard for the BCD and recommended for all UD designs.

24.2.5.1 Reversal Averaging Estimators

The most common classical UD estimator for the median-toxicity dose uses *reversal points*—those subjects in the experiment whose toxicity outcome is different from the outcome of the immediately previous subject:

$$\mathcal{R} = \{i : Y_i \neq Y_{i-1}, \quad i = 2, \dots, n\}, \quad (24.15)$$

where Y_i is subject i 's binary toxicity outcome. Reversal points are surrounded by squares in Figure 24.2. The reversal estimator of $F^{-1}(\Gamma)$ is a simple average of the doses at which reversals are observed:

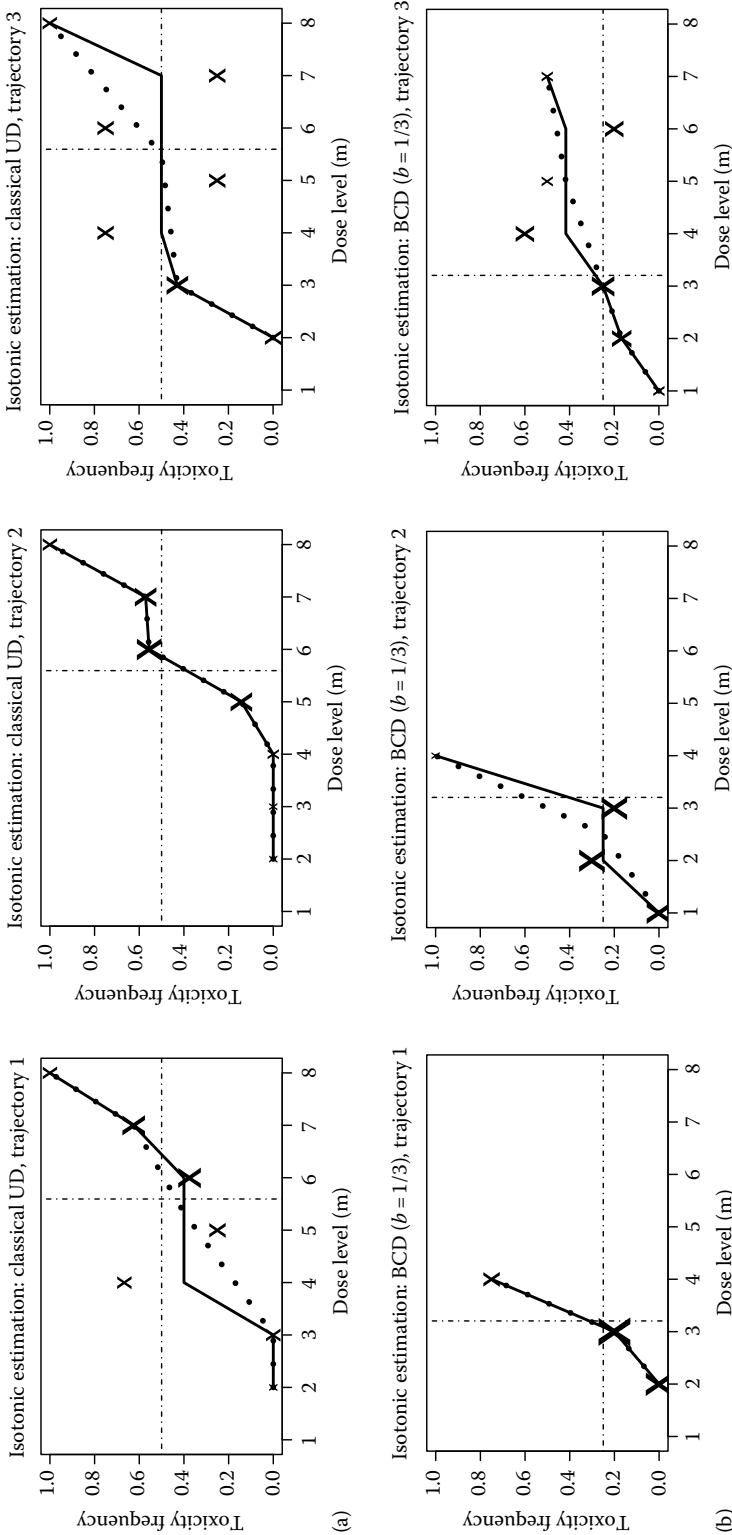


FIGURE 24.4 Dose-response plots display cumulative toxicity rates \hat{F} for each of the randomly generated scenarios of Figure 24.2 and for both the classical UD design ($\Gamma = 0.50$, a) and the BCD with $b = 1/3$ ($\Gamma = 0.25$, b): “x” symbols mark the proportion of subjects receiving dose d_{thr} ; isotonic regression estimates are given by the solid lines; and CIR estimates are given by the dotted lines. The respective point estimates of $F^{-1}(\Gamma)$ are the x values where the solid lines cross the dash-dot horizontal $y = \Gamma$ lines. The true quantile is marked with a dash-dot vertical line.

$$\bar{X}^{(\text{rev.})} = \frac{\sum_i X_i I[i \in \mathcal{R}]}{\sum_i I[i \in \mathcal{R}]}, \quad (24.16)$$

where $I[\cdot]$ is the indicator function. Sometimes the first few reversals are excluded in order to mitigate the starting-dose bias. Using reversals rather than all assigned doses is done mostly for historical reasons (Wetherill et al. 1966). As explained more fully in Section 24.4, we recommend discontinuing the averaging of only reversal doses.

24.2.5.2 Isotonic Regression–Based Estimators

Stylianou and Flournoy (2002) adapted isotonic regression (Barlow et al. 1972) for estimating $F^{-1}(\Gamma)$ following UD experiments. For dose-toxicity applications, isotonic regression begins by calculating the empirical toxicity frequencies at each dose:

$$\hat{F}_m = \frac{\sum_{i=1}^n Y_i I[X_i = d_m]}{\sum_{i=1}^n I[X_i = d_m]}, \quad m = 1, \dots, M. \quad (24.17)$$

The somewhat awkward formulation in (24.17) is equivalent to simply tabulating how many observations were made at each dose and what proportion of these are toxicities. For example, if four subjects have thus far been given dose d_m with two toxicities, then $\hat{F}_m = 0.5$.

In the case of a monotonicity violation among the $\{\hat{F}_m\}$, isotonic regression recursively replaces adjacent pairs of violating values by their weighted average, until all violations disappear and the resulting estimate conforms to the assumption of an increasing dose-toxicity relationship. Stylianou and Flournoy's estimate for F in $[d_1, d_M]$ linearly interpolates between the isotonic regression pointwise results. The dose where \hat{F} crosses the horizontal line $y = \Gamma$ becomes the UD point estimate of $F^{-1}(\Gamma)$. This estimator is a special case of *regression estimators*: estimators that fit some curve $y = \hat{G}(x)$ through the points $\{(d_m, \hat{F}_m)\}$. Dose–response plots and isotonic regression estimators for the six simulated runs of Figure 24.2 are shown in Figure 24.4, along with a minor modification called *centered isotonic regression* (CIR) (Oron 2007, Section 3.3).

24.3 Higher-Order Up-and-Down Designs and Other Extensions

Via first-order designs, we have now completed a brief tour of the main UD design properties and estimation approaches. The first-order BCD (with the classical UD design as a special case) is simple and intuitive and allows for the estimation of any percentile. However, there are cases when neither the classical UD design nor BCD are adequate, such as when

- Subjects are treated in cohorts rather than individually, in order to save time and resources

- In some applications, such as oncology, researchers are reluctant to use randomization in a toxicity study, even for selecting between two adjacent doses (as required for the BCD)
- Other UD designs may yield better estimation efficiency

This section surveys additional, somewhat more sophisticated UD designs that are either in common use or have been studied in some detail. When relevant, they are compared with the BCD and to each other. The section ends with a description of an innovative UD design for dose-finding in which the target dose involves joint toxicity and efficacy probabilities.

24.3.1 Group Up-and-Down Designs

Group UD rules were first introduced in the 1960s for median estimation, by Wetherill (1963) and Tsutakawa (1967a,b), and more recently studied in greater depth by Gezmu and Flournoy (2006), Ivanova (2006), and Baldi Antognini et al. (2008) for the estimation of any percentile. Subjects are treated in cohorts of a fixed size s . Let l and u be two integers such that $0 \leq l < u \leq s$. With a Group UD design, Y_i is the number of toxic responses in the i th cohort, rather than a binary outcome. Given that the i th cohort is treated at $X_i = d_m$ on the interior of \mathcal{X} , the $i + 1$ -th cohort is assigned to

$$X_{i+1} = \begin{cases} d_{m+1} & \text{if } Y_i \leq l; \\ d_{m-1} & \text{if } Y_i \geq u; \\ d_m & \text{if } l < Y_i < u. \end{cases}$$

Conditional upon X_i and $F(x)$, Y_i has a binomial distribution with parameters s and $F(d_m)$, with the “up” and “down” probabilities being the binomial distribution’s tails and the “stay” probability its center (it is zero if $u = l + 1$). This design rule is abbreviated as $GUD_{(s,l,u)}$, or GUD in general.

Nominally, a GUD is an s -order design, since the s most recent observations are needed to determine the next allocation. However, with cohorts viewed as single entities (with each entity containing s subjects), the Group UD generates a first-order Markov chain having a tridiagonal TPM as in (24.1). In this *condensed form*, it is straightforward to show that the Durham–Flournoy monotonicity conditions on $p(x)$ and $q(x)$ are met, and therefore, Theorem 24.1 holds and π is unimodal. Unlike the classical UD design and the BCD, except for special cases, GUD balance points can be calculated only numerically, by finding the dose x^* with toxicity rate F^* such that

$$\sum_{r=u}^s \binom{s}{r} (F^*)^r (1 - F^*)^{s-r} = \sum_{t=0}^l \binom{s}{t} (F^*)^t (1 - F^*)^{s-t}. \tag{24.18}$$

Any numerical root-finding algorithm can be used to solve for F^* in (24.18). Gezmu and Flournoy (2006) and Ivanova et al. (2007) present tables with (s, l, u) combinations and the corresponding values of F^* satisfying (24.18). For example, GUD designs that might be useful for $\Gamma = 0.3$ include the following combinations:

s	2	3	4	5		6			
(l, u)	(0,1)	(0,2)	(0,2)	(0,3)	(1,2)	(0,3)	(1,2)	(0,4)	(1,3)
F^*	0.293	0.347	0.267	0.302	0.314	0.253	0.264	0.326	0.341

Notable special cases of group UD designs, whose F^* can be calculated by hand, are symmetric ones ($u + l = s$) for which x^* is the median (the classical UD design again is a special case of this family viz., $GUD_{(1,0,1)}$). In addition, the family $GUD_{(k,0,1)}$ is related to the design we describe next.

24.3.2 k -in-a-Row (“Geometric”) Design

The design called here the “ k -in-a-row design” is presented in some detail, due to its widespread use in sensory studies (Tretwein 1995) and to a now-resolved controversy regarding its properties. The rule, whose development and introduction to sensory studies date back to Wetherill’s work in the 1960s (Wetherill 1963; Wetherill and Levitt 1966), is extremely simple:

$$X_{i+1} = \begin{cases} d_{m+1} & \text{if } Y_{i-k+1} = \dots = Y_i = 0, \text{ all observed at } d_m; \\ d_{m-1} & \text{if } Y_i = 1; \\ d_m & \text{otherwise,} \end{cases} \tag{24.19}$$

where $k \geq 1$ is an integer constant. In words, under this rule, every dose escalation requires k consecutive nontoxicities at the current level, while de-escalation only requires a single toxicity. Using straightforward reasoning, Wetherill concluded that this design’s balance point at dose x^* is given by

$$F^* = F(x^*) = 1 - \left(\frac{1}{2}\right)^{1/k}. \tag{24.20}$$

The $k = 1$ case is simply the classical UD design, but $k = 2, 3, 4$ would be associated with $F^* \approx 0.293, 0.206$ and 0.159 , respectively.*

While Wetherill correctly identified F^* , the characterization of k -in-a-row’s overall Markov chain behavior was provided only a generation later by Gezmu (1996) and completed by Oron and Hoff (2009). Recent work also gave the design the following names: first, *geometric UD design* (for reasons that will be apparent soon) and then *k -in-a-row design* (Ivanova et al. 2003), the name used here. Meanwhile, in sensory studies, where it is most commonly used, the design is known as *forced-choice fixed-staircase methods* (Tretwein 1995), but this name is used for related non-Markovian designs as well.

The k -in-a-row design generates a k th order Markov chain because knowledge of the last k responses might be needed. Like any k -order chain, it can be represented as a first-order chain with Mk states, useful for writing out a TPM.† Equivalently, it can be modeled

* The method used in sensory studies is actually the mirror image of (24.19), with k successive responses required for a de-escalation and only one nonresponse for escalation, yielding $F^* \approx 0.707, 0.794, 0.841, \dots$

† It should be noted that for the highest level d_M , internal states are redundant: no escalation is possible, and therefore, dose-transition behavior is identical and depends only on the most recent outcome. So the most concise representation of a k -in-a-row TPM would have $(M - 1)k + 1$ states. However, for symmetry and aesthetics, we use Mk states.

as a Markov chain with M levels, each having k internal states labeled 0 to $k - 1$ (Weiss 1994; Oron and Hoff 2009). The internal state serves as a counter of the number of immediately recent consecutive nontoxicities observed at level m . This description is closer to the physical dose-allocation process, because subjects at different internal states of the same level m are all assigned the same dose d_m . Either way, the TPM is $Mk \times Mk$. Equation 24.21 shows the expanded k -in-a-row TPM with $k = 2$ and $M = 6$, using the abbreviations $F_m = F(d_m), \bar{F}_m = 1 - F_m, m = 1, \dots, M$. Dose levels are labeled.

$$P = \begin{matrix} & \begin{matrix} m=1 & m=2 & m=3 & m=4 & m=5 & m=6 \end{matrix} \\ \begin{matrix} m=1 \\ m=2 \\ m=3 \\ m=4 \\ m=5 \\ m=6 \end{matrix} & \begin{bmatrix} F_1 & \bar{F}_1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ F_1 & 0 & \bar{F}_1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ F_2 & 0 & 0 & \bar{F}_2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ F_2 & 0 & 0 & 0 & \bar{F}_2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & F_3 & 0 & 0 & \bar{F}_3 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & F_3 & 0 & 0 & 0 & \bar{F}_3 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & F_4 & 0 & 0 & \bar{F}_4 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & F_4 & 0 & 0 & 0 & \bar{F}_4 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & F_5 & 0 & 0 & \bar{F}_5 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & F_5 & 0 & 0 & 0 & \bar{F}_5 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & F_6 & 0 & 0 & \bar{F}_6 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & F_6 & 0 & \bar{F}_6 & 0 \end{bmatrix} \end{matrix} \tag{24.21}$$

The matrix is regular, assuring that there is an asymptotic distribution $\pi^{(states)}$, with positive values for all Mk states. This is not the asymptotic *dose-allocation* distribution π , since it counts separately different internal states of the same level m . The latter can be found by summing $\pi^{(states)}$ over the internal states of each level.

Gezmu (1996) found $\pi^{(states)}$ by noting that nonbase (>0) internal states can only be accessed from the internal state immediately below in the same level. The balance equation for such states is, therefore,

$$\pi_{m(j)}^{(states)} = (1 - F_m) \pi_{m(j-1)}^{(states)}, \quad j = 1, \dots, k - 1, \tag{24.22}$$

where $m(j)$ denotes the internal state j within dose level m . This is the formula of a decreasing geometric sequence with a quotient of $(1 - F_m)$, hence the name Gezmu suggested for the design (*geometric UD design*).

Now, the balance equation between d_m and d_{m+1} is

$$\pi_{m(k-1)}^{(states)} (1 - F_m) = \pi_{m+1} F_{m+1}. \tag{24.23}$$

In words, “up” transitions are possible only from the top internal state, while “down” transitions are mandated upon any toxicity regardless of internal state. After applying standard geometric-series formulae to express $\pi_{m(k-1)}^{(states)}$ as a function of π_m , one obtains the stationary adjacent-level ratio of the dose-allocation distribution π :

$$\lambda_m = \frac{F_m (1 - F_m)^k}{F_{m+1} [1 - (1 - F_m)^k]}. \tag{24.24}$$

This ratio is monotone decreasing in m , conforming to the Durham–Flournoy conditions (see Section 24.1.4). The k -in-a-row balance point can now be calculated; indeed, it is the one found by Wetherill (24.20). It is also identical to the balance point of $\text{GUD}_{(k,0,1)}$. The same unimodality and mode-location properties, with respect to x^* shown for π of the classical UD design and the BCD, hold for k -in-a-row as well (Oron and Hoff 2009).*

24.3.3 Biased Coin Design Extensions

The versatile framework of the BCD lends itself to further extensions. Baldi Antognini et al. (2008) generalized the Group UD design for any user-specified toxicity rate Γ by adding a biased coin to the dose allocation procedure. Flournoy (1998) proposed the use of an urn model to integrate parallel BCD runs, helping shorten an experiment’s duration. This strategy is also useful when clinical trials are performed at different locations.

Taking advantage of the flexibility in constructing biased coin designs, Bortot and Giovagnoli (2005) also examined a generalized first-order approach in order to study their properties. Let there be two coins having $P\{\text{heads}\}$ equal to b_1 and b_2 . Then

$$X_{i+1} = \begin{cases} d_{m+1} & \text{if } Y_i = 0 \text{ and coin 1 yields heads;} \\ d_{m-1} & \text{if } Y_i = 1 \text{ and coin 2 yields heads;} \\ d_m & \text{if } Y_i = 0 \text{ and coin 1 yields tails} \\ & \text{or } Y_i = 1 \text{ and coin 2 yields tails.} \end{cases} \quad (24.25)$$

The balance point maintains

$$F^*b_2 = (1 - F^*)b_1; \quad F^* = \frac{b_1}{(b_1 + b_2)}.$$

Setting $b_2 = 1$ yields Durham and Flournoy’s BCD. A family of generalized BCDs for the same Γ can be created by setting

$$\frac{b_2}{b_1} = \frac{1 - \Gamma}{\Gamma}, \quad b_2 \in [0, 1],$$

so it can be indexed by a single coin-toss parameter. Although all members of the family share the same balance point, they differ in other properties, as described in the following section.

Bortot and Giovagnoli (2005) also developed a second-order BCD, again using two coins. Given $X_i = d_m$,

$$X_{i+1} = \begin{cases} d_{m+1} & \text{if } Y_i = 0, Y_{i-1} = 0 \text{ and coin 1 yields heads} \\ & \text{or } Y_i = 0, Y_{i-1} = 1 \text{ and coin 2 yields heads;} \\ d_{m-1} & \text{if } Y_i = 1; \\ d_m & \text{otherwise.} \end{cases}$$

*The derivation of (24.24) assumes stationarity. For finite n , the marginal transition probabilities are very cumbersome to calculate; it is for the k -in-a-row design that the $n \rightarrow \infty$ limit terminology is needed in Definition 24.1.

Bortot and Giovagnoli derived π and showed that it is unimodal by the Durham–Flournoy conditions (see Section 24.1.4) and if $(1-\Gamma)b_1 + \Gamma b_2 = \Gamma/(1-\Gamma)$. However, their simulations indicate only a modest estimation performance improvement over the first-order BCD.

24.3.4 Comparing Up-and-Down Designs

With several UD variants to choose from for estimating a given quantile, some ways to compare UD designs are needed. An obvious comparison criterion is estimation performance—for example, mean squared error (MSE) of the quantile estimates. Since estimates at the various dose levels are dependent, in general, this practical metric can only be gauged via simulation studies. A more didactic approach that can often save simulation time is to consider individual properties and find *theoretical* relationships between designs for these properties. Some properties suggested and examined by researchers are

- Convergence rates of the dose-allocation process and of empirical dose-allocation frequencies
- How tight (or *peaked*, see definition as follows) is the asymptotic distribution π
- How well is π centered around x^*

24.3.4.1 Convergence Rates

Bortot and Giovagnoli (2005) showed that for the family of generalized BCDs of the form (24.25) having the same Γ , the asymptotic allocation variance is a decreasing function of b_2 , and, therefore, Durham and Flournoy's BCD ($b_2 = 1$) is the minimum-asymptotic-variance member of this family. Moreover, the geometric rate of convergence to stationary behavior increases with increasing b_2 , up to a certain $b_2^{(\max)}$, which in numerical simulations tended to be greater than 1. In summary, the simplest member of the family (i.e., the BCD) appears to have the best convergence properties.

Oron and Hoff (2009) found the following result comparing the convergence of k -in-a-row with the analogous group UD design.

Theorem 24.4 *Consider a k -in-a-row and a $GUD_{(k,0,1)}$ design with the same k, \mathcal{X} , and dose-toxicity rate function F . Then the k -in-a-row design converges more quickly to its stationary behavior.*

The proof appears in Oron and Hoff (2009). In simulations, k -in-a-row also held a pervasive and even larger advantage over a BCD with the same x^* . BCD designs took 10% – 70% more subjects than k -in-a-row to reach the 99% stationary benchmark defined in Section 24.2.4 (Oron and Hoff 2009). However, this numerical advantage over BCD was not theoretically proven.

24.3.4.2 Peakedness of the Asymptotic Dose-Allocation Distribution

Giovagnoli and Pintacuda (1998) introduced the concept of *peakedness* as an easy-to-analyze measure of precision for discrete unimodal distributions of the type generated by UD designs.

Definition 24.4 For two UD designs D_1, D_2 operating over the same \mathcal{X} and having the same balance point x^* , D_1 's asymptotic distribution is more "peaked" if it increases more quickly to the left of x^* and decreases more quickly to the right of x^* . That is, if one considers $\lambda_m^{(D_1)}, \lambda_m^{(D_2)}$, the adjacent-level ratio of π under designs D_1, D_2 , respectively, then

$$\begin{cases} \lambda_m^{(D_1)} \geq \lambda_m^{(D_2)} \geq 1 \text{ for } m = 1, \dots, m^* - 1; \\ \lambda_m^{(D_1)} \leq \lambda_m^{(D_2)} \leq 1 \text{ for } m = m^* + 1, \dots, M - 1, \end{cases} \quad (24.26)$$

where m^* denotes the dose level immediately below x^* .

As the stationary allocation distribution π becomes more *peaked*, excursions away from x^* are of shorter duration and estimation efficiency of $F^{-1}(\Gamma)$ generally improves. Therefore, all other things being equal, a design with a more *peaked* π is preferable.

Consider the family of symmetric Group UD designs $\text{GUD}_{(s,l,s-l)}$, a family designed to estimate the median (see Section 24.3.1). Oron (2007) proved that designs of this family with smaller l (i.e., ones with a higher probability of staying at the same dose) have a more peaked π than designs with larger l and the same s . On the other hand, it appears that designs with larger l converge faster, although a proof has not been found.

For comparing the *peakedness* of UD design families, Oron and Hoff (2009) prove the following results.

Theorem 24.5 (i) *k-in-a-row designs generate stationary allocation distributions π that are more peaked than those of BCD designs over the same \mathcal{X} and with the same x^* .* (ii) *For a k-in-a-row design and a $\text{GUD}_{(k,0,1)}$ with the same k , neither design's π can be said to be consistently more peaked than the other.*

24.3.5 Up-and-Down Designs for Joint Studies of Toxicity and Efficacy

Most dose-finding experiments focus upon toxicity only, or efficacy only. For these goals, the assumption of a monotone increasing dose-toxicity function $F(x)$ or a monotonically increasing dose-efficacy function $R(x)$ is usually adequate. One approach for experiments involving both toxicity and efficacy is to dichotomize observations into two binary responses per subject, one each for toxicity and efficacy. They can be modeled via a pair of monotone-increasing functions: the dose-toxicity rate function $F(x)$ and the dose-efficacy rate function $R(x)$. However, attempts to optimize a single function of toxicity and efficacy are becoming increasingly popular. For example, one could use a single unimodal utility function $U(x)$, the simplest form of which is $U(x) = aR(x) - F(x)$, with $a > 0$ a constant quantifying researchers' assessment of the relative overall benefit of efficacy as compared with the harm of toxicity.

Kpamegan and Flournoy (2001, 2008) developed a size 2 cohort UD design for locating the maximum of $U(x)$. For each pair, doses are split with one subject receiving a dose exactly Δ units larger than the other. Outcomes are dichotomized to *success* and *failure*, with success defined as a positive, toxicity-free efficacy response. Each cohort produces a pair of such binary outcomes— (F, F) , (S, F) , (F, S) , or (S, S) . The transition rule moves the dosage in the direction of observed success in the case of two opposite responses and repeats the same dose in the case of two identical responses. Kpamegan and Flournoy (2008)

proved that the *left* (smaller) and *right* dose allocations of each pair form Markov chains, with unimodal asymptotic distributions identical except for an offset of Δ . The asymptotic properties of unimodality and closeness of the mode to the maximum of $U(x)$ were also proven.

One caveat to this design is that, if the starting point has an extremely low or high success probability, information accrues very slowly and the experiment does not progress very well toward the optimum. Therefore, it is preferable to begin the experiment where a mixture of successes and failures is expected within a few early allocations, and not where they are highly unlikely as at a very low or very high dose. Alternatively, or additionally, the rate of convergence improves if the dose is increased or decreased according to the flip of a coin when pairs of responses are either both successes or both failures; however, the resulting stationary distribution will be less peaked.

24.4 Estimation Following an Up-and-Down Design

A crucial property of UD designs that is often missed by the casual observer is that the estimation procedure may be completely separate from the design. Because the design is ancillary (Rosenberger et al. 1997) to the outcome process, a researcher can design and run an UD experiment and then choose essentially any estimator that is a function of the dose and response sequences, to produce a point estimate of $F^{-1}(\Gamma)$. We view this property favorably, since the task of information collection around the target's vicinity (accomplished by the design) is different from the task of estimation, once information collection is completed. However, as hinted in Section 24.2, while UD sampling and convergence properties have been fairly thoroughly studied and documented, the same cannot be said of estimation methods. As a result, patterns of estimator use are uneven and quite often rely upon each field's tradition.

For example, numerous environmental toxicity studies (Lichtman 1998; Sunderam et al. 2004; Sweeney et al. 2010) estimate LD_{50} using reference tables prepared in the 1960's by Dixon (1965) for classical UD experiments with extremely small samples ($n < 10$). Dixon assumed that the toxicity threshold distribution is normal, with its standard deviation exactly equal to the UD dose spacing. At that time, UD Markov properties were only beginning to be understood. These tables from Dixon should be considered obsolete, but they are still being used occasionally.

Another example of a misinformed and very poor estimator is the use of the UD design's last allocation X_n or the last allocation recommendation X_{n+1} . Methodologists comparing UD designs with procedures for which the last allocation is a reasonable estimate (see Section 24.5 for such designs) sometimes assume that this estimator is sensible for UD designs as well (O'Quigley and Chevret 1991; Zacks 2009). But using an arbitrary allocation from a random walk as an estimator of anything makes no sense. Needless to say, the UD design does not perform well in such misguided comparisons.

24.4.1 Dose Averaging Estimators

In the overwhelming majority of UD experiments, $F^{-1}(\Gamma)$ is estimated using some average of a subset of the dose-allocation chain $\{X_n\}$. This seems a plausible approach: as shown earlier, convergence to asymptotic behavior is fast, π is centered fairly tightly around x^* , and

if dose spacing is uniform, then the distribution is also approximately symmetric (boundaries permitting). Hence, once initial conditions wear off, the mean of $\{X_n\}$ should be fairly close to x^* , and being the mean, it is also a relatively precise statistic. Dixon and Mood's (1948) original estimator (which is no longer used) was also a type of weighted average.

However, averaging estimators suffer from numerous limitations:

- The most commonly used estimators incorporate limited knowledge, at best, of the UD design's Markov properties—in particular, starting-point-induced biases and insight regarding what constitutes a plausibly stable sample from π .
- As the balance point approaches a design boundary (d_1 or d_M), π gradually loses its symmetry (compare Figure 24.3a vs. b). If x^* is less than two levels away from the boundary, the average will be substantially biased in the opposite direction.
- The BCD and k -in-a-row designs have π s that are not centered on x^* as well as those of the classical UD design. Rather, they are centered as far as half a dose spacing below x^* , resulting in a downward averaging-estimator bias for these designs (see Section 24.2.3).
- Furthermore, designs such as k -in-a-row and the GUD offer a limited set of balance points that are often slightly different from the quantile of interest $F^{-1}(\Gamma)$. For example, with $\Gamma = 0.3$, the natural k -in-a-row choice is $k = 2$, whose balance point is approximately $x^* = F^{-1}(0.293)$. This introduces yet another bias to averaging estimators, since these estimators' reference point is x^* rather than $F^{-1}(\Gamma)$.

Modern developers and users of averaging estimators must be cognizant of these limitations and use estimation methods that mitigate them. First and foremost, averaging estimators are predicated upon a uniformly spaced \mathcal{X} . If log-doses are uniformly spaced, then it is the logarithm of doses that should be averaged, rather than the doses themselves (García-Perez 1998; Oron 2007). If dose spacing is not uniform on any reasonable scale, researchers are strongly advised not to use an averaging estimator. In the following discussion, we assume uniform dose spacing.

24.4.1.1 Case Against Reversal-Only Estimators

Estimators that average only the allocations at reversal points originated in the 1960s with Wetherill et al. (1966), were slightly modified by Choi (1971), and have proliferated widely since then. They are probably the single most commonly used UD estimator. Recall that reversal points are points in the experiment when the current response differs from the immediately preceding one (Figure 24.2).

A classical UD experiment's trajectory can be fully reconstructed from the reversal points. It can be shown that for the classical UD design, the reversal estimator $\bar{X}^{(\text{rev.})}$ (24.16) is equivalent to a *weighted* average of all assigned doses between the first and last reversal included in the average. The weights are inversely proportional to the length of the between-reversal stretches, except for the first and last reversal whose weights are larger. Oron (2007) showed that under stationarity, using only reversals is equivalent to sampling from a distribution whose variance is larger than that of π . Moreover, all allocations after the last reversal (including X_{n+1} , the would-be next allocation at the experiment's end which is already determined by X_n and Y_n (Brownlee et al. 1953)) are excluded from reversal estimators—reducing overall precision and giving up the data points least affected by starting-dose bias. In summary, reversal-only estimators are not recommended.

24.4.1.2 Alternative Solutions for Averaging Estimators

Averaging estimators can still be viable when used with care. The recommended averaging estimator should include *all* allocated doses from X_c (c being some cutoff point, $1 < c < n/2$) through X_{n+1} . The first and third reversal are reasonable cutoff points. Another option is choosing c retroactively as the earliest point in the trajectory, at which there is empirical evidence that the starting-point bias had been eliminated (Oron 2007, Section 3.3).

As mentioned earlier, when the balance point x^* is close to a boundary of \mathcal{X} , the approximate symmetry of π is destroyed, and dose averages become biased in the direction away from the boundary. A simple mitigating fix, which should work reasonably well as long as the empirical mode is not on the boundary itself, is to *impute* the dose-allocation chain by adding a virtual dose of d_0 or d_{M+1} each time the boundary condition had been triggered and use the imputed trajectory in lieu of the original one. This adjustment has yet to be thoroughly studied (Oron 2007, Section 3.3).

Confidence intervals are sometimes reported together with the reversal estimator (Capogna et al. 2001; Camorcia et al. 2004). They assume asymptotic normality of the reversal average and independence of every pair of reversal points. The second assumption is not supported by recent theory. Oron (2007) explored alternative options for estimating the effective sample size for averaging estimators. The most robust option was the number of times the experimental trajectory had crossed the average or visited the dose levels closest to it. According to Markov chain theory, the subchains between consecutive visits to a specific level are conditionally independent of each other (Tsutakawa 1967b). In numerical runs, the resulting standard error formula using this effective-sample-size estimate was found to be conservative in estimating the variance, but, due to biases associated with averaging estimators, the empirical interval coverage was usually close to the specified confidence level.

24.4.2 Regression Estimators

Parametric regression (usually logistic or probit) is occasionally encountered as a way to estimate $F^{-1}(\Gamma)$ following a UD experiment. Given the small sample size and the limited number of dose levels visited during a typical UD experiment, using parametric regression to estimate the dose-toxicity rate function F is, generally speaking, inefficient. This is because estimates of $F^{-1}(\Gamma)$ typically will be a function of a slope parameter, which requires allocations far from $F^{-1}(\Gamma)$ in order to be estimated efficiently (Ford et al. 1992). UD procedures, by design concentrating allocations around $F^{-1}(\Gamma)$, will typically result in relatively poor estimates of a slope parameter, even with moderately large sample sizes.

The nonparametric isotonic regression dose-finding estimate, introduced by Stylianou and Flournoy (2002), is a simple piecewise linear interpolation between point estimates calculated via the pool-adjacent-violators algorithm (PAVA, Barlow et al. 1972). The point estimates themselves are identical to the empirical frequencies $\hat{F} = (\hat{F}_1, \dots, \hat{F}_M)$ (see (24.17)), as long as the latter do not violate monotonicity. In case of a violation, violating values are recursively replaced by an average weighted by the number of observations, until all violations disappear. In the binary dose-toxicity case, the weighted average is equivalent to *pooling* all observations from the violating doses together and calculating the overall toxicity frequency, hence the name *pool-adjacent-violators algorithm*. The procedure generates the characteristic flat stretches, clearly visible in the solid lines of Figure 24.4.

Isotonic regression is also the nonparametric maximum likelihood estimator for F on \mathcal{X} , under a monotonicity assumption. However, if one makes the plausible assumption that F is *strictly* increasing, then the estimate violates it, ostensibly producing biases in opposite directions at each flat stretch.* Quantile estimation requires *inverse estimation* from the isotonic estimate of the regression function, and a flat stretch near $y = \Gamma$ would push some quantile estimates sideways (see again Figure 24.4), increasing their MSE.

Oron (2007) proposed a simple modification of PAVA that alleviates this problem while retaining the method's simplicity. Each flat stretch is replaced by a single point with the same \hat{F} value, located at an x value that is a weighted average calculated using the same weights used for \hat{F} . With linear interpolation between the isotonic estimates, this modification is called CIR; see the dotted lines in Figure 24.4. In numerical simulations, it nearly always produces lower MSEs than ordinary isotonic regression. R code for CIR is available online (<https://www.github.com/assaforon/cir>).

Estimators based on isotonic regression are the most robust all-around choice currently available for any limited-sample dose-finding application, including UD designs. Unlike averaging estimators, they do not suffer from peculiarities related to dose spacing, boundaries, or the starting point. As stated earlier, without large sample sizes, isotonic estimates outperform parametric estimates because data are concentrated near x^* .

One drawback of isotonic estimators is the current lack of an adequate interval estimate. Wright (1984) and Mukerjee (1993) study percentile estimators under the condition that the dose space becomes dense. This assumption is clearly invalid for UD designs with fixed doses. Chao and Fuh (2001) introduced the bootstrap to UD interval estimation. Their motivating example was a classical-UD explosive testing experiment, and the estimator used was from a parametric regression model. Bootstrap samples are generated by running virtual UD experiments, with toxicity probabilities determined via the regression estimator that uses the observed empirical toxicity rates \hat{F} . Stylianou et al. (2003) implemented the method in a nonparametric context via isotonic regression, but only for forward estimation (i.e., estimation of F values at specified doses, rather than estimates of the quantile $F^{-1}(\Gamma)$). Oron (2007) explored the bootstrap for UD quantile estimation and found that interval coverage was generally insufficient (i.e., the intervals were too narrow). We are exploring some ideas for improving coverage; working solutions are available in the aforementioned online source.

24.5 Other Approaches to Dose-Finding

UD procedures are not the only approach to dose-finding on a discrete dose set. As a contrast to UD designs, we now focus on a family of designs at the center of heated debate among dose-finding methodologists. The field in question is Phase I cancer clinical trials (hereafter, "Phase I"). Sample sizes for Phase I experiments are small: as a rule, $n < 30$. First, we describe the outdated method that novel Phase I designs seek to replace, which is often (Rogatko et al. 2007; Zacks 2009) mistaken for a UD design.

* It is one of those cases when the MLE is on the boundary of the constrained parameter region. Strict monotonicity is equivalent to excluding the boundaries themselves from the allowed region.

3+3 Rule (Precise Origins Unknown)

1. Begin at a safe dose: $X_1 = d_1$ or d_2 , and treat three subjects at a time.
2. After the first cohort at any given dose, $X_i = d_m$,

$$X_{i+1} = \begin{cases} d_{m+1} & \text{if 0 of 3 are toxic;} \\ d_{m-1} & \text{if at least 2 of 3 responses are toxic;} \\ d_m & \text{if exactly 1 of 3 is toxic.} \end{cases}$$

3. After a *second* cohort at any dose d_m , go to

$$\begin{cases} d_{m+1} & \text{if 0 – 1 are toxic, cumulatively out of all 6 responses;} \\ d_{m-1} & \text{if } \geq 2 \text{ of the 6 responses are toxic.} \end{cases}$$

4. If a third cohort is mandated at any level, stop.
5. Once the experiment has de-escalated, it cannot escalate back, that is, after a transition from d_m to d_{m-1} , the experiment cannot return to d_m anymore. If such a reescalation is mandated, the experiment stops. If only three patients have been treated at the stopping dose (d_{m-1} in the example above), add three more patients at that dose.
6. After final stopping, the maximum-tolerated-dose (MTD) estimate is the highest dose for which $\hat{F}_m < 1/3$.*

The “3+3” rule has been generalized as the “A+B rule” (Lin and Shih 2001), but the “3+3” rule is by far the most commonly used.

The “3+3” rule is a *toxicity-controlling escalation protocol* rather than an experimental design in the research sense. However, the initial cohort at each level follows rules similar to $GUD_{(3,0,2)}$, and after the second cohort, the rule is reminiscent of $GUD_{(6,1,2)}$ except for the nontrivial difference that under “3+3” the six observations are not necessarily the most recent and three of them had already participated in one transition decision.

Several methodological studies have found that “3+3”s statistical properties are poor and that the recommended MTD is strongly biased downward compared with its ostensible target, which is usually presented as $F^{-1}(0.3)$ or $F^{-1}(1/3)$ (Reiner et al. 1999; Lin and Shih 2001). The “3+3” rule is also disproportionately cumbersome and arcane when contrasted with its simplistic approach to the dose-toxicity relationship. Yet, it continues to be the off-the-shelf method of choice for most Phase I trials.

Interestingly, while statisticians are generally critical of this design, the “3+3” paradigm of *selecting* a discrete dose from \mathcal{X} as the best dose (i.e., the one to be used in subsequent studies of efficacy), in contrast to the universal UD practice of *estimating* $F^{-1}(\Gamma)$ on a continuous dose scale, has been rather influential upon recent methodological approaches to Phase I designs. The effect of this influence is described immediately as follows.

* There exists a variant allowing the final estimate to have $\hat{F}_m = 1/3$.

24.5.1 Long-Memory Dose-Finding Designs

The overwhelming majority of recently developed Phase I designs belong to a class that Oron and Hoff (2013) call *long-memory Phase I designs* (LMP1s). LMP1s calculate a regression estimate after each cohort, using results to select the MTD according to some design-specific optimization criterion and then assigning this selection (\widehat{MTD}) to the next cohort. In contrast with UD designs, LMP1s potentially use *all* previous observations in each dose-allocation decision, rather than only the most recent ones. LMP1s might vary considerably in their external properties, but most share these two core principles:

1. Likelihood-based estimation (in the generic sense that includes Bayesian methods).
2. Assigning \widehat{MTD} to the next cohort.

We now describe the allocation rule for the most commonly used LMP1, the one-parameter continual reassessment method (CRM) (O’Quigley et al. 1990).

Continual Reassessment Method

1. Choose a *parametric curve family* $\mathcal{G}(\theta, \boldsymbol{\phi})$ to describe the dose–response relationship, where θ is a single parameter with a prior distribution and $\boldsymbol{\phi}$ is a vector of prior parameters. A single curve belonging to this family is denoted $G(x | \theta, \boldsymbol{\phi}) \in \mathcal{G}$.
2. The experiment can start at any agreed-upon dose. Cohort size can vary *ad lib*, even within the same experiment.
3. Given $X_i = d_m$, calculate the posterior estimate

$$\hat{G}_i = G(\hat{\theta} | \mathbf{x}, \hat{F}, \boldsymbol{\phi}),$$

using all available data (where $\mathbf{x} = (x_1, \dots, x_i)$ is the observed trajectory and \hat{F} is the vector of observed toxicity frequencies at the doses visited so far).

4. Find the dose whose \hat{G}_i value is closest to the target toxicity frequency Γ ,

$$\widehat{MTD}_i = \operatorname{argmin}_{\mathcal{X}} \left| \hat{G}_i - \Gamma \right|$$

5. For the next cohort, set $X_{i+1} = \widehat{MTD}_i$, unless the modification suggested by Goodman et al. (1995), preventing multidose jumps, is in effect—in which case $X_{i+1} = \min(\widehat{MTD}_i, d_{m+1})$.
6. At the end of the experiment, the recommended MTD is \widehat{MTD}_n .

Despite using estimation to guide dose allocation, the final product of CRM experiments is not a direct statistical estimate of $F^{-1}(\Gamma)$, but rather a dose selected out of \mathcal{X} . This insistence upon dose selection might be the “3+3” mindset’s lingering effect.

The most common probability structure used in CRM experiments is the *power* model, in which a *skeleton* of prior toxicity assessments at the dose levels is raised to the same power:

$$G(d_m) = \phi_m^\theta, \quad \phi_1 < \phi_2 < \cdots < \phi_M, \quad \phi_m \in (0, 1) \forall m, \theta > 0. \quad (24.27)$$

The only data-estimable parameter is θ . It is accompanied by $M + 2$ fixed, nonestimable prior parameters (the vector $\boldsymbol{\phi}$ and two parameters for the prior distribution of θ itself, which is most commonly specified as log-normal). Between its introduction in 1990 and 2006, the CRM was used in only about a dozen out of over a thousand published Phase I studies (Rogatko et al. 2007). Since then, the number of experiments utilizing CRM or closely related LMP1s has substantially increased. Over the same time period, hundreds of *methodological* articles discussing CRM and closely related designs have been published, making CRM one of the most widely discussed and promoted recent statistical designs.

Another well-known LMP1 is the two-parameter escalation with overdose control (EWOC, Babb et al. 1998). Since 2008, Novartis has been widely implementing its own version of two-parameter LMP1s in Phase I trials; their designs are based on EWOC and on Neuenschwander et al. (2008). A nonparametric LMP1, called the Cumulative Cohort Design (CCD) was developed by Ivanova et al. (2007), and its convergence behavior was proven by Oron et al. (2011). There are many dozens of other LMP1s, most of which have yet to be used in practice.

24.5.2 Advantages and Drawbacks of Long-Memory Designs

Unlike for UD designs, when using an LMP1 design, researchers do not have to commit to a fixed cohort size. Choosing a nonuniformly spaced \mathcal{X} also does not carry as many consequences for estimation as with some estimators used with UD designs. But the main selling point of LMP1s has been their claimed asymptotic behavior—together with the misconception that this behavior takes effect early, with small samples.

Since LMP1s allocate the \widehat{MTD} to future patients, and since they potentially use all available data for this estimation, as $n \rightarrow \infty$, they promise to allocate subjects *exclusively* to the MTD. LMP1 method developers contrast this ostensibly razor-sharp asymptotic distribution with UD design's indefinite random walk, obviously recommending the former.

The problem with LMP1s lies in the gap between these promises and reality (Fedorov et al. (2011), recently called LMP1s *best-intention designs*). Theoretical work in the past few years has shown that conditions under which the sequence of doses allocated using the CRM actually converges to an MTD-only distribution are fairly restrictive and cannot be expanded any further (Lee and Cheung 2009; Oron et al. 2011; Azriel 2012). Given a general form of F , the most that can be guaranteed with CRM is asymptotic exclusive allocation to a dose whose F value is within some *tolerance interval* in the vicinity of Γ^* . The interval depends upon the skeleton $\boldsymbol{\phi}$ and is difficult to calculate by hand. Lee and Cheung (2009) provide software that back calculates the prior skeleton to match a user-specified tolerance interval. However, the nonparametric LMP1 mentioned earlier can guarantee the same type of convergence by using the interval width directly, with no need for a reverse-engineered model (Ivanova et al. 2007; Oron et al. 2011).

* If the interval is too narrow to contain any dose level, the asymptotic behavior slowly oscillates back and forth between the two levels straddling it.

More generally, Azriel et al. (2011) proved that no *best-intention* LMP1, parametric or nonparametric, can converge almost surely to MTD-only allocations without substantive restrictions on F . Many LMP1s, including the Novartis designs currently implemented in actual trials, have never had their asymptotic behavior proven.

Further complicating matters is the commonly misunderstood issue of *small-sample* LMP1 behavior. When compared in simulation on an equal footing, UD and LMP1 designs achieve rather similar success rates in selecting the true MTD from doses in \mathcal{X} (Oron and Hoff 2013; Durham et al. 1997). On average, LMP1s assign substantially more subjects to the MTD during the simulated experiment itself. Advocates for LMP1s have focused on this statistic, hereafter called n^* . For example, Rogatko et al. (2007) argued that the use of any non-LMP1 method "...implies needless loss of treatment efficacy and, possibly, lives," because LMP1s assign more patients to the true MTD. Simulated means of n^* now regularly appear in new methodological Phase I articles, presumably as an important quality metric.

This metric is fundamentally debatable: Phase I trials recruit volunteers under informed consent. A trial's main goal is to study a treatment's safety, not to ensure participants will get the optimal treatment combination. But the focus on n^* is questionable for *statistical* reasons as well.

The statistic for n^* reported in simulation studies is really an ensemble average of n^* , taken over a large number of simulated experiments under identical conditions. In real life, Phase I trials are one-of-a-kind affairs. What about n^* 's run-to-run variability? Oron and Hoff (2013) examined the entire *ensemble distribution* of n^* . Typical results are shown in Figure 24.5, under three different forms of F . Each simulated experiment had 16 cohorts of size 2. Distributions of n^* from CRM are on the left, a UD design (specifically, GUD_(2,0,1)) in the center, and a CCD on the right. Both CRM and CCD targeted $\Gamma = 0.3$ and were calibrated at settings recommended by their developers. All histograms use the same scale, and the first cohort was excluded.

With LMP1s, the mean values of n^* (marked by a solid vertical line in each histogram) are indeed larger than with UD designs. With UD designs, n^* values in different runs are tightly clustered around their theoretical expectation. By contrast, LMP1s display far greater variability. With gamma-distributed toxicity thresholds (middle row), the most common CRM n^* value—by far—was *zero* allocations to the true MTD during the entire experiment, except for the opening cohort that was also the true MTD in this case. The nonparametric CCD displays similar behavior to CRM, indicating that this is a universal LMP1 trait, almost unaffected by specific design choice.

The importance of Figure 24.5's message should not be understated. Phase I researchers spend years planning and running a single experiment and would benefit very little from optimistic reports about high average n^* values, when the underlying individual-experiment distribution is (at best) nearly uniform between zero and n . The focus on LMP1's high ensemble-average n^* values has been misguided and misleading. LMP1 designs cannot reliably promise researchers an advantage over UD designs in this respect, because the increase in ensemble-average values comes at the price of greatly increasing the risk of individual experiments with very low n^* .

A thorough inspection of the phenomenon illustrated in Figure 24.5 and its causes is beyond the scope of this chapter, but it can be easily confirmed and reproduced by anyone running a dose-finding simulation study. In any case, contrary to prevalent opinion in many circles, when all aspects are considered, UD designs do in fact hold their own compared with the most sophisticated, cutting-edge LMP1s, even in Phase I cancer applications to which the latter are tailored. Some hybrid designs attempting to combine the advantages of UD and LMP1 designs are in development (Oron 2007, Chapter 5).

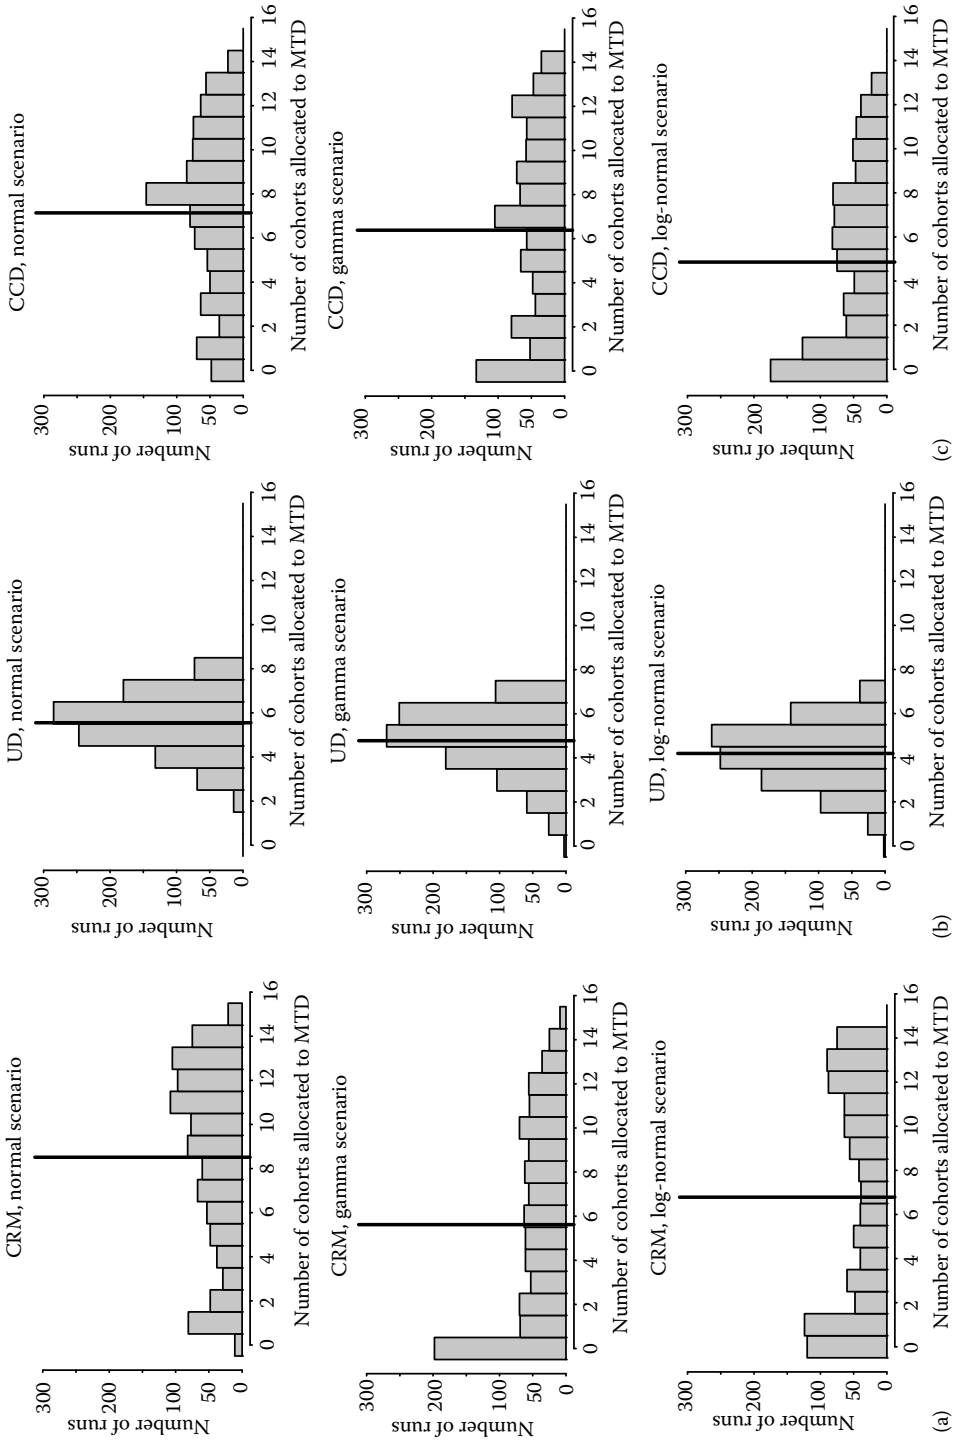


FIGURE 24.5 Simulation ensemble distributions of n^* , the number of cohorts assigned the true MTD during the experiment (excluding the first cohort), under three forms of the dose-toxicity rate function F (top to bottom), and the CRM (a), UD (b) and CCD (c) designs, as described in the text. Ensemble size was 1000 runs.

24.6 Summary and Practical Recommendations

UD procedures constitute a family of simple designs, empowering researchers with limited statistical training to use them independently. However, beneath the simplicity lie the power and elegance of UD's Markov chain behavior, whose fundamentals were explored here. With this additional knowledge, researchers can bolster their experiment's design against the known pitfalls. As this chapter shows, there are still many open challenges, especially with regard to estimation of the target quantile. Researchers are encouraged to consult with a statistician cognizant of modern UD methodology, even when their experiment seems to follow a well-trodden path. When choosing a specific UD design, it should be noted that the scientifically mandated toxicity target rate Γ may not necessarily equal F^* , the expected toxicity frequency of the balance point x^* . Certainly, this is the case with k -in-a-row and GUD, whose F^* values are usually irrational.

The issue of appropriate sample sizes for UD designs is taken very lightly in some fields, where $n \approx 10$ or even less are considered acceptable. One should keep in mind that despite the geometric-rate erasure of initial conditions, UD estimation precision is still largely governed by the same root- n convergence rates of nonsequential estimation (Oron 2007, Section 1.2). Moreover, the available data are a small sample of binary responses scattered over several dose levels (for regression estimators) or a small sample of correlated dose allocations (for averaging estimators). While some experiments might appear very well-behaved, that impression might be misleading unless the underlying variance of toxicity thresholds is far smaller than anticipated. Ultimately, increasing the sample size is the only sure method for substantially reducing estimation error.

In many fields, researchers use the number of reversals as stopping criterion. García-Perez (1998) surveyed all experiments determining vision thresholds published in the 1990s. The most common design used by researchers was the k -in-a-row UD design, and most experiments stopped after a fixed number of reversals, typically 6–12, equivalent to roughly 20–50 observations. Many environmental toxicity studies also use reversals for stopping, albeit far earlier than in vision research (2–4 reversals using the classical UD design, equivalent to roughly 8–15 subjects). Due to the properties of reversals outlined earlier, a reversal-based stopping rule will generally stop earlier precisely for those experiments that would actually most benefit from continuing longer, because their trajectory was more erratic and generated more frequent reversals. Other fields consistently use a fixed sample size. For example, in anesthesiology, $n = 30$ seems to be a common choice for classical UD experiments (Capogna et al. 2001; Camorcia et al. 2004). Pace and Stylianou (2007) recommend a sample size of $n \geq 20$ for the classical UD design.

Another practical challenge, which should be considered in conjunction with sample size, is choosing the number of dose levels M . In order to develop nicely peaked UD allocation distributions that greatly improve estimation precision and help reduce the extent of excursions (Figure 24.3), ideally at least 3–4 levels are needed on each side of the balance point x^* . This translates to at least 8–10 levels. However, increasing M also increases the number of subjects needed to erase the starting-dose effect—and there is no guarantee that the experiment will take a beeline from the starting point x_1 directly to the vicinity of x^* without a detour on the way (see again Figure 24.2). In essence, this is the well-known bias–variance tradeoff in a somewhat unusual reincarnation.

As to starting dose, when regulatory authorities allow beginning anywhere in \mathcal{X} rather than at its lower boundary, we would recommend a starting dose that minimizes the

expected length of this initial tour—that is, a level lying midway between the boundaries, in terms of the number of subjects required to reach each boundary under the transition rules. When starting in the middle is not possible or not advisable and when sample size is constrained ($n \leq 30$), M will probably need to be smaller, in order to allow the experiment to comfortably reach the top boundary of \mathcal{X} if so needed.

In summary, the most generic sample-size and dose-spacing recommendation we can provide is this: use $M \geq 8$ dose levels and choose n large enough to allow at least half the experiment to meander around the balance point x^* , assuming the worst-case distance from the starting point. If this is technically infeasible, decrease M .

From Wetherill (1963) onward, researchers have examined designs that begin with coarse spacing, which is reduced gradually as the experiment progresses. This brings UD procedures closer to the paradigm of stochastic approximation (Robbins and Monro 1951), which allocates gradually decreasing dose increments on a continuous dose scale. However, numerical studies have repeatedly shown little benefit from these elaborate downscaling schemes; simply put, it is difficult to evade the UD design's bias-variance tradeoff in this manner (García-Perez 1998). In the same vein, even though UD design's adaptive nature might invite researchers to implement adaptive stopping rules, we advise against the practice. With an experiment already constrained by low information content (binary observations) and a small sample size, adaptive stopping rules introduce an unnecessary source of additional variability.

When the experiment's goal is dose-selection from \mathcal{X} rather than estimation of $F^{-1}(\Gamma)$ (see Section 24.5), the tradeoff between n and M is somewhat different. Here, with too many dose levels, the experiment might never provide enough data to tell their toxicity rates apart. A simple rule of thumb is that in order to distinguish between the toxicity rates of adjacent doses, n must be large enough, and M small enough, so that for each dose interval around x^* , there will be at least 3–4 subjects, preferably 5 or more, whose toxicity thresholds fall in the interval. Thus, with $n = 20$ and $\Gamma \approx 0.3$, one should probably design no more than 5–7 dose levels, and with $n = 10$ no more than $M = 4$. These rough guidelines assume a starting point in the middle of \mathcal{X} . If toxicity concerns dictate a start near the bottom of \mathcal{X} , n should be increased by at least the minimal number of subjects required to travel halfway up \mathcal{X} according to the design's rules. To attain the same expected sampling density around the median, the classical UD design can make do with perhaps 20%–40% fewer subjects than the $\Gamma \approx 0.3$ rules of thumb, while using the same M . Conversely, very extreme quantiles ($\Gamma \leq 0.1$) would require considerably more subjects.

24.7 Historical Notes

The most widely cited *origin paper* for the classical UD design is Dixon and Mood (1948). The article was part of an adaptation of a memorandum submitted to the Applied Mathematics Panel by the Statistical Research Group, Princeton University. The Statistical Research Group operated under a contract with the Office of Statistical Research and Development, which was directed by the Applied Mathematics Panel of the National Defense Research Committee. Their work was motivated by experiments at the Bruceston naval explosive testing site in Pennsylvania during World War II. As often happens in research, Nobel laureate von Békésy (1947) came up with a similar design for hearing-threshold determination, almost simultaneously. An internal navy report of a lesser known UD design that is more

similar to k -in-a-row (Ted Anderson 1993, personal communication) predates Dixon and Mood by 2 years (Anderson et al. 1946). While that report is widely cited, we have been unable to locate the original manuscript.

The first study known to have written about UD design's Markov properties is Derman (1957), the same article that presented the first biased coin UD design. Wetherill et al. (1966) introduced the still hugely popular reversal estimator a few years later, mostly based on the narrow edge this estimator had over averages of all treatments, in numerical simulations using logistically distributed thresholds. The time being the early 1960s, Wetherill had to travel across the Atlantic to the United States and be awarded precious computer time there, in order to perform the simulations—a feat that was a main topic in the Royal Statistical Society discussion of the Wetherill (1963) article. From that point onward, dose-finding research has been wedded—sometimes blissfully, at other times less so—to numerical simulation studies.

Other important contributions to UD theory and design in the 1950s and 1960s were made by Brownlee et al. (1953) and Tsutakawa (1967a,b). Then the methodological UD design trail gradually ran cold, until its relative reawakening in recent decades as described throughout this chapter.

References

- Anderson, T. W., McCarthy, P. J., and Tukey, J. W. (1946), 'Staircase' method of sensitivity testing, Naval Ordnance Report 65-46, Statistical Research Group, Princeton University, Princeton, NJ.
- ASTM (1991), *Standard Test Method for Estimating Acute Oral Toxicity in Rats*. American Society for Testing and Materials, Philadelphia, PA, 1163–1190.
- Azriel, D. (2012), A note on the robustness of the continual reassessment method. *Stat. Probab. Lett.*, 82, 902–906.
- Azriel, D., Mandel, M., and Rinott, Y. (2011), The treatment versus experimentation dilemma in dose finding studies, *J. Stat. Plann. Inference*, 141, 2759–2768.
- Babb, J., Rogatko, A., Rogatko, A., and Zacks, S. (1998), Cancer phase I clinical trials: Efficient dose escalation with overdose control, *Stat. Med.*, 17, 1103–1120.
- Baldi Antognini, A., Bortot, P., and Giovagnoli, A. (2008), Randomized group up and down experiments, *Ann. Inst. Stat. Math.*, 60, 45–59.
- Barlow, R. E., Bartholomew, D. J., Bremner, J. M., and Brunk, H. D. (1972), *Statistical Inference under Order Restrictions*, John Wiley & Sons, New York.
- Bortot, P. and Giovagnoli, A. (2005), Up-and-down experiments of first and second order, *J. Stat. Plann. Inference*, 134, 236–253.
- Brownlee, K. A., Hodges, Jr., J., and Rosenblatt, M. (1953), The up-and-down method with small samples, *J. Am. Stat. Assoc.*, 48, 262–277.
- Camorcica, M., Capogna, G., Lyons, G., and Columb, M. (2004), Epidural test dose with levobupivacaine and ropivacaine determination of ED50 motor block after spinal administration, *Br. J. Anaesth.*, 92, 850–853.
- Capogna, G., Parpaglionni, R., Lyons, G., Columb, M., and Celleno, D. (2001), Minimum dose of epidural sufentanil for first-stage labor analgesia, *Anesthesiology*, 94, 740–744.
- Chao, M. T. and Fuh, C. D. (2001), Bootstrap methods for the up and down test on pyrotechnics sensitivity analysis, *Stat. Sinica*, 11, 1–21.
- Cheung, Y. K. (2005), Coherence principles in dose-finding studies, *Biometrika*, 92, 863–873.
- Choi, S. (1971), An investigation of Wetherill's method of estimation for the up-and-down experiment, *Biometrics*, 27, 961–970.

- Derman, C. (1957), Non-parametric up-and-down experimentation, *Ann. Math. Stat.*, 28, 795–798.
- Diaconis, P. and Stroock, D. (1991), Geometric bounds for eigenvalues of Markov chains, *Ann. Appl. Probab.*, 1, 36–61.
- Dixon, W. J. (1965), The up-and-down method for small samples, *J. Am. Stat. Assoc.*, 60, 967–978.
- Dixon, W. J. and Mood, A. (1948), A method for obtaining and analyzing sensitivity data, *J. Am. Stat. Assoc.*, 13, 109–126.
- Durham, S. D. and Flournoy, N. (1994), Random walks for quantile estimation, in *Statistical Decision Theory and Related Topics, V*, eds. Gupta, S. S. and Berger, J. O., Springer-Verlag Inc., Hayward, CA, pp. 467–476.
- Durham, S. D. and Flournoy, N. (1995), Up-and-down designs I: Stationary treatment distributions, in *Adaptive Designs*, eds. Flournoy, N. and Rosenberger, W. F., Institute of Mathematical Statistics, Hayward, CA, pp. 139–157.
- Durham, S. D., Flournoy, N., and Montazer-Haghighi, A. A. (1995), Up-and-down designs II: Exact treatment moments, in *Adaptive Designs*, eds. Flournoy, N. and Rosenberger, W. F., Institute of Mathematical Statistics, pp. 158–178.
- Durham, S. D., Flournoy, N., and Rosenberger, W. F. (1997), A random walk rule for phase I clinical trials, *Biometrics*, 53, 745–760.
- Fedorov, V., Flournoy, N., Wu, Y., and Zhang, R. (2011), Best intention designs in dose-finding studies, Technical Report, Isaac Newton Institute for the Mathematical Sciences, Cambridge, U.K., <http://www.newton.ac.uk/preprints/NI11065.pdf>. Accessed February 2, 2015.
- Flournoy, N. (1998), A biased coin urn design for estimating a target quantile from a sequence of binary responses: Asymptotic results, in *Proceedings of the 20th International Conference on Information Technology Interfaces*, eds. Kalpić, D. and Dobrić, V. H., SRCE University Computing Centre, University of Zagreb, Zagreb, Croatia, pp. 13–18.
- Flournoy, N., Durham, S. D., and Rosenberger, W. F. (1995), Toxicity in sequential dose-response experiments, *Seq. Anal.*, 14, 217–227.
- Ford, I., Torsney, B., and Wu, C. F. J. (1992), The use of a canonical form in the construction of locally optimal designs for non-linear problems, *J. R. Stat. Soc., Ser. B*, 54, 569–583.
- García-Perez, M. A. (1998), Forced-choice staircases with fixed step sizes: Asymptotic and small-sample properties, *Vision Res.*, 38, 1861–1881.
- Gezmu, M. (1996), The geometric up-and-down design for allocating dosage levels, PhD thesis, American University, Washington, DC.
- Gezmu, M. and Flournoy, N. (2006), Group up-and-down designs for dose-finding, *J. Stat. Plann. Inference*, 136, 1749–1764.
- Giovagnoli, A. and Pintacuda, N. (1998), Properties of frequency distributions induced by general ‘up-and-down’ methods for estimating quantiles, *J. Stat. Plann. Inference*, 74, 51–63.
- Goodman, S. N., Zahurak, M. L., and Piantadosi, S. (1995), Some practical improvements in the continual reassessment method for Phase I studies, *Stat. Med.*, 14, 1149–1161.
- Hughes, B. D. (1995), *Random Walks and Random Environments*. Vol. 1: Random walks, Oxford Science Publications/The Clarendon Press Oxford University Press, New York.
- Ivanova, A. (2006), Escalation, group and $A + B$ designs for dose-finding trials, *Stat. Med.*, 25, 3668–3678.
- Ivanova, A., Flournoy, N., and Chung, Y. (2007), Cumulative cohort design for dose-finding, *J. Stat. Plann. Inference*, 137, 2316–2317.
- Ivanova, A., Montazer-Haghighi, A., Mohanty, S. G., and Durham, S. D. (2003), Improved up-and-down designs for phase I trials, *Stat. Med.*, 22, 69–82.
- JSME. (1981), *Standard Method of Statistical Fatigue Testing*, JSME S 002. Japan Society of Mechanical Engineers, Tokyo, Japan.
- Kpamegan, E. and Flournoy, N. (2001), An optimizing up-and-down design, in *Optimum Design 2000*, eds. Atkinson, A., Bogacka, B., and Zhigljavsky, A., Kluwer Academic Publishers, Dordrecht, the Netherlands, pp. 211–224.
- Kpamegan, E. and Flournoy, N. (2008), Up-and-down designs for selecting the dose with maximum success probability, *Seq. Des.*, 27, 78–96.

- Lagoda, T. and Sonsino, C. (2004), Comparison of different methods for presenting variable amplitude loading fatigue results, *Mat. Wiss. Werk.*, 35, 13–20.
- Lee, S. M. and Cheung, Y. K. (2009), Model calibration in the continual reassessment method, *Clin. Trials*, 6, 227–238.
- Lichtman, A. (1998), The up-down method substantially reduces the number of animals required to determine antinociceptive ED50 values, *J. Pharm. Toxicol Methods*, 40, 81–85.
- Lin, Y. and Shih, W. J. (2001), Statistical properties of the traditional algorithm-based designs for phase I cancer clinical trials, *Biostatistics (Oxford)*, 2, 203–215.
- Mukerjee, H. (1993), An improved monotone conditional quantile estimator, *Ann. Stat.*, 21, 924–942.
- Neuenschwander, B., Branson, M., and Gsponer, T. (2008), Critical aspects of the Bayesian approach to Phase I cancer trials, *Stat. Med.*, 27, 2420–2439.
- NIEHS. (2001), The revised up-and-down procedure: A test method for determining the acute Oral toxicity of chemicals, Technical Report 2-4501, Washington, DC.
- OECD (1998), *The Revised Up-and-Down Procedure: A Test Method for Determining the Acute Oral Toxicity of Chemicals*, Organization for Economic Co-Operation and Development, Paris, France.
- O’Quigley, J. and Chevret, S. (1991), Methods for dose finding studies in cancer clinical trials: A review and results of a Monte Carlo study, *Stat. Med.*, 10, 1647–1664.
- O’Quigley, J., Pepe, M., and Fisher, L. (1990), Continual reassessment method: A practical design for Phase 1 clinical trials in cancer, *Biometrics*, 46, 33–48.
- Oron, A. P. (2007), Up-and-down and the percentile-finding problem, PhD thesis, University of Washington, Washington, DC, <http://arxiv.org/abs/0808.3004>.
- Oron, A. P., Azriel, D., and Hoff, P. D. (2011), Dose-finding designs: The role of convergence properties. *Int. J. Biostat.*, 7, 39.
- Oron, A. P. and Hoff, P. D. (2009), The k-in-a-row up-and-down design, revisited, *Stat. Med.*, 28, 1805–1820.
- Oron, A. P. and Ho, P. D. (2013), Small-sample behavior of novel phase I cancer trial designs, *Clin. Trials*, 10, 63–80.
- Pace, N. L. and Stylianou, M. P. (2007), Advances in and limitations of up-and-down methodology: A précis of clinical use, study design, and dose estimation in anesthesia research, *Anesthesiology*, 107, 144–152.
- Reiner, E., Paoletti, X., and O’Quigley, J. (1999), Operating characteristics of the standard phase I clinical trial design, *Comput. Stat. Data Anal.*, 30, 303–315.
- Robbins, H. and Monro, S. (1951), A stochastic approximation method, *Ann. Math. Stat.*, 22, 400–407.
- Rogatko, A., Schoeneck, D., Jonas, W., Tighiouart, M., Khuri, F. R., and Porter, A. (2007), Translation of innovative designs into Phase I trials, *J. Clin. Oncol.*, 25, 4982–4986.
- Rosenberger, W. F., Flournoy, N., and Durham, S. D. (1997), Asymptotic normality of maximum likelihood estimators from multiparameter response-driven designs, *J. Stat. Plann. Inference*, 60, 69–76.
- Stylianou, M. and Flournoy, N. (2002), Dose finding using the biased coin up-and-down design and isotonic regression, *Biometrics*, 58, 171–177.
- Stylianou, M., Proschan, M., and Flournoy, N. (2003), Estimating the probability of toxicity at the target dose following an up-and-down design, *Stat. Med.*, 22, 535–543.
- Sunderam, R. M., Patra, R., Julli, M. et al. (2004), Use of the up-and-down acute toxicity test procedure to generate LC50 data for fish, *Bull. Environ. Contam. Toxicol.*, 72, 873–880.
- Sweeney, D., Cui, X., Solomon, S. et al. (2010), Anthrax lethal and edema toxins produce different patterns of cardiovascular and renal dysfunction and synergistically decrease survival in canines, *J. Infect. Dis.*, 202, 1885–1896.
- Treutwein, B. (1995), Minireview: Adaptive psychophysical procedures, *Vision Res.*, 35, 2503–2522.
- Tsutakawa, R. K. (1967a), Asymptotic properties of the block up-and-down method in bio-assay, *Ann. Math. Stat.*, 38, 1822–1828.
- Tsutakawa, R. K. (1967b), Random walk design in bio-assay, *J. Am. Stat. Assoc.*, 62, 842–856.
- von Békésy, G. (1947), A new audiometer, *Acta Otolaryngol.*, 35, 411–422.

- Weiss, G. H. (1994), *Aspects and Applications of the Random Walk (Random Materials and Processes)*, North-Holland Publishing Company, Amsterdam, the Netherlands.
- Wetherill, G. B. (1963), Sequential estimation of quantal response curves, *J. R. Stat. Soc. B*, 25, 1–48.
- Wetherill, G. B., Chen, H., and Vasudeva, R. B. (1966), Sequential estimation of quantal response curves: A new method of estimation, *Biometrika*, 53, 439–454.
- Wetherill, G. B. and Levitt, H. (1966), Sequential estimation of on a psychometric function, *Br. J. Math. Stat. Psychol.*, 18, 1–10.
- Wright, F. T. (1984), The asymptotic behavior of monotone percentile regression estimates, *Can. J. Stat.*, 12, 229–236.
- Zacks, S. (2009), *Stage-Wise Adaptive Designs*, John Wiley & Sons, Inc., Hoboken, NJ.

Jason Ming-Hung Kao and John Stufken

CONTENTS

25.1 Introduction.....	895
25.2 Basic Concepts and Terminology.....	897
25.2.1 fMRI Designs.....	897
25.2.2 Hemodynamic Response Function and fMRI Time Series.....	898
25.3 Designs for Linear Models.....	900
25.3.1 Models and Design Selection Criteria.....	901
25.3.1.1 Linear Model for Detection.....	901
25.3.1.2 Linear Model for Estimation.....	903
25.3.1.3 Design Selection Criteria.....	903
25.3.2 Efficient fMRI Designs.....	905
25.3.3 Genetic Algorithms for Finding Efficient fMRI Designs.....	907
25.3.4 Illustrative Examples.....	908
25.4 Designs for Nonlinear Models.....	910
25.4.1 Illustrative Example for the Nonlinear Model.....	914
25.5 Additional Considerations.....	915
25.5.1 Nonadditive HRFs.....	916
25.5.2 Uncertain Correlation Coefficient.....	918
25.5.3 Multiple Scanning Sessions.....	919
25.6 Summary and Discussion.....	920
Acknowledgments.....	922
References.....	922

25.1 Introduction

Functional magnetic resonance imaging (fMRI) is a brain-mapping technique widely used in research fields such as cognitive neuroscience, medical science, psychology, and education for studying functions of the brain. It helps with understanding the inner workings of our brains and holds great promise for, among others, identification of Alzheimer's disease, pre-neurosurgical planning, and postneurosurgical evaluations. The clinical importance of fMRI can also be seen in a special issue on clinical applications of fMRI in *Neuropsychology Review*, Vol. 7, No. 2, 2007.

A primary objective of an fMRI experiment is typically to investigate how brains react to mental stimuli. In such an experiment, each subject is exposed to a sequence of stimuli (e.g., pictures or sounds) interlaced with periods of rest or visual fixation, to which we

refer as the control. While the subject is performing the mental tasks in response to the stimuli inside a magnetic resonance (MR) scanner, the scanner scans the subject's brain at regular time intervals to collect hundreds of thousands of time series. Each brain voxel (a 3D image unit, or volumetric pixel) gives rise to one such time series. For example, the subject's brain may be mapped into an array of $64 \times 64 \times 30$ voxels, each of size about $3 \times 3 \times 5 \text{ mm}^3$ (cf. Lazar 2008, Section 2.1.1), leading to $64 \times 64 \times 30$ time series from the subject. The size and complexity of fMRI data have spurred much statistical research in this area on data collection and methods for a better analysis and interpretation of the data. The high demand for statistics in fMRI is evident in a survey article by Lindquist (2008); recent books by Lazar (2008) and Poldrack et al. (2011); and papers in special issues on neuroimaging data analysis by several statistics journals, such as *Statistica Sinica*, Vol. 18, No. 4, 2008; *Annals of Applied Statistics*, Vol. 5, No. 2B, 2011; and *Journal of Statistical Software*, Vol. 44, 2011. These references also highlight that there remain many statistical challenges in this area.

We will focus on designing experiments for event-related (ER) fMRI studies. ER-fMRI takes advantage of ultrafast imaging techniques to allow studying transient brain activity due to stimuli of short duration (Josephs and Henson 1999; Culham 2006). Mental stimuli as brief as several milliseconds (e.g., 34 ms) can give rise to effects that are detectable by high-speed MR scanners (see also Rosen et al. 1998; Blamire 2012). This is an important advance. It allows us to move away from traditional fMRI studies where long-duration (e.g., 1 min) stimuli are employed. ER-fMRI designs with brief stimuli can now be considered for studying subtle brain activity. With a brief stimulus duration, the number of stimuli presented to a subject can be increased significantly, which increases the information available from fMRI studies (Josephs and Henson 1999; Huettel 2012) and presents many challenging research questions related to designing ER-fMRI experiments.

An ER-fMRI experiment may involve one or more types of stimuli. In its simplest form, a design for an ER-fMRI experiment presents the types of stimuli given to a subject, their presentation order, and the times at which these stimuli are given. Each stimulus type appears repeatedly in a design, and a stimulus can last from several milliseconds to a few seconds. An ER-fMRI design can be written as a finite sequence of finite numbers such as $d = \{10112002\dots 0\}$ with $q = 1, 2, \dots, Q$ representing a q th-type stimulus and 0 denoting the control when no stimulus is given; see Section 25.2 for a detailed description. The primary design problem is selecting the 'best' sequence of stimuli based on a specific optimality criterion. Since an ER-fMRI design sequence can easily consist of hundreds of stimuli and controls, there is a huge number of candidate designs, making the problem of selecting the best design very challenging. Additionally, while it is crucial to target designs yielding high statistical efficiencies, practical and psychological constraints can also play a role when selecting a design. For example, a design that repeats the same pattern, such as $\{111000\}$, over and over may induce unwanted confounding with psychological effects such as anticipation or habituation. Furthermore, experimental settings such as time between MR scans, time between stimulus onsets, duration of each stimulus, duration of the experiment, and the underlying statistical models for data analysis must also be taken into account at the design stage. These factors combine to make selecting an optimal design very challenging.

Nonetheless, there are some important recent advances in the optimal design problem for ER-fMRI experiments. These include the pioneering works by Dale (1999) and Friston et al. (1999); theoretical results by Liu et al. (2001) and Liu and Frank (2004); and computational approaches by Wager and Nichols (2003); Kao et al. (2009a); Kao et al. (2013); and Kao et al. (2012). These studies provide guidance and powerful computational tools

for obtaining good ER-fMRI designs. However, there still are many open questions. As noted by Lindquist (2008), “While the area of experimental design is a natural domain for statisticians to conduct research, it has so far been largely unexplored by members of the field.”

This chapter provides an introduction to this exciting research area and surveys existing methods for ER-fMRI design selection. The focus is primarily on state-of-the-art computational approaches for obtaining high-quality ER-fMRI designs. We provide background information and introduce terminology and notation in Section 25.2. In Section 25.3, we present the very popular general linear model approach for analyzing ER-fMRI data, followed by a discussion of available results for selecting designs under this approach. In Section 25.4, we introduce a nonlinear model that accommodates consideration of disparities between the responses from different brain voxels. We also present an efficient method for obtaining optimal designs under this approach. Additional considerations for fMRI design problems can be found in Section 25.5, followed by brief remarks and open research questions in Section 25.6. For brevity, we will omit the term “event-related” hereinafter, but our focus throughout this chapter is always on ER-fMRI experiments.

25.2 Basic Concepts and Terminology

This section begins with a description of fMRI experiments and designs. Such experiments are conducted in order to better understand the underlying brain activity. The hemodynamic response function (HRF), which we also introduce in this section, plays an important role in the models that are used to accomplish this. Some of these models will be introduced in the subsequent sections. This chapter focuses on selecting optimal and efficient designs for fMRI experiments. A reader who is interested in statistical analysis of fMRI data may want to consult Lazar (2008) and Lindquist (2008).

Since the notation can at times be intimidating, we present a short list of selected notation that the reader can easily refer to:

d : an fMRI design (or a sequence of mental stimuli of Q types), which is written as a finite sequence of elements $0, 1, \dots, Q$ for indicating the onset times and order of the stimuli with 0 representing *no stimulus onset*.

L : the number of elements of an fMRI design.

Q : total number of stimulus types.

τ_{ISI} : the interstimulus interval, or the time between time points where a stimulus can possibly occur.

τ_{TR} : the time to repetition, or time between two consecutive observations in an fMRI time series.

ΔT : the minimal difference between two time points at which an HRF height can contribute to an observed response; it is computed as the greatest real value making both $(\tau_{ISI}/\Delta T)$ and $(\tau_{TR}/\Delta T)$ integers.

25.2.1 fMRI Designs

An fMRI design d determines the onset times and order of mental stimuli of one or more types to be presented to an experimental subject in an fMRI experiment. It can be written as




FIGURE 25.1

The design $d = \{110202\dots0\}$ in Example 25.1 with 1 s stimuli of pictures of happy and sad faces ($Q = 2$). The control is a visual fixation and the interstimulus interval τ_{ISI} is 3 s.

a sequence of finite numbers, such as $d = \{110202\dots0\}$. The j th element of d corresponds to the j th time point when a stimulus can possibly occur, $j = 1, \dots, L$, where L is the number of elements of d . These time points are τ_{ISI} seconds apart, where τ_{ISI} (e.g., 3 s) is a pre-specified time called interstimulus interval. A positive integer q in the j th position of d thus indicates an onset of the q th-type stimulus at time $(j - 1)\tau_{ISI}$, $q = 1, 2, \dots, Q$. The presentation duration of each stimulus is short (e.g., 1 s), compared with τ_{ISI} . These stimuli are immediately followed by the control (e.g., a period of rest or visual fixation) until the onset of the next stimulus. We might have no stimulus onset at some time points. These occasions are indicated by 0 in d .

Example 25.1

Figure 25.1 presents the first 18 s of a design, $d = \{110202\dots0\}$, with two types ($Q = 2$) of 1 s pictures, a happy or sad face. The interstimulus interval τ_{ISI} is set to 3 s. With a 1 appearing in the first and second position of d , the onset of the first two pictures of happy faces (first stimulus type) is at 0 s and 3 s, respectively. Note that time 0 is typically synchronized with the first valid MR scan. The design also indicates that the onset of sad faces is at 9 s, 15 s, and so on. Each picture lasts for 1 s, and the remaining time until the next onset consists of visual fixation. This applies also for the entire interval (e.g., 4–9 s and 10–15 s) for which the design specifies the use of a control. Visual fixation means here that the experimental subject will fixate on a central cross hair, and for design d this happens from 1–3 s, 4–9 s, 10–15 s, etc. □

25.2.2 Hemodynamic Response Function and fMRI Time Series

A stimulus may activate some voxels of a subject’s brain, whereas other voxels will remain inactive. Such brain activity will be reflected in the signals collected by the MR scanner. These signals are not direct measurements of the underlying brain neuronal activity, but are linked to the change in the ratio of oxy- to deoxyhemoglobin in the cerebral blood vessels. Specifically, for a brain voxel that responds to a stimulus, the stimulus evokes neuronal activity followed by an influx of oxygenated cerebral blood. This leads to a decrease in the concentration of deoxyhemoglobin, which is paramagnetic and tends to diminish MR signals, so that elevated MR signals can be observed. Without additional stimuli, the MR signals would fall back to baseline. Before settling at the baseline, the MR signal intensity can drop below it, resulting in a brief undershoot (see the solid curves in Figure 25.2). Such an undershoot is often attributed to the enlarged blood volume that allows for more deoxygenated blood; see, for example, Song et al. (2006) and Raichle (2006) for more details.

The fluctuation of MR signals provoked by a single stimulus is typically described by a smooth function over time called the HRF. The HRF, which is of primary interest to




FIGURE 25.2

The accumulated HRF (broken curve) is formed by the three HRFs (solid curves), which are evoked by the three stimuli (shaded bars) occurring at 3 s, 6, and 12 s. The interstimulus interval is $\tau_{ISI} = 3$ s, and the time to repetition is $\tau_{TR} = 2$ s, resulting in the discretization interval $\Delta T = 1$ s. The small squares on the last HRF curve show the HRF heights that can possibly contribute to MR measurements.

neuroscientists, may last for many seconds, for example, 32 s, and may look like one of the solid curves in Figure 25.2. When stimuli occur in close succession, the corresponding HRFs overlap, and the signals measured by the MR scanner every τ_{TR} seconds, for example, every 2 s, are assumed to be the result of an accumulation of the separate HRFs. Here, τ_{TR} denotes the time to repetition, which is the prespecified time between consecutive MR scans of the same voxel.

Example 25.2

Consider a design $d = \{011010 \dots 1\}$ with $\tau_{ISI} = 3$ s. The first three onsets of the stimulus occur thus at 3 s, 6 s, and 12 s and are presented as the three shaded bars in Figure 25.2. If they activate a voxel, they give rise to three HRFs (shown as solid curves), which accumulate to become the broken curve in the figure. Here, we consider a common linear time invariant (LTI) system in which stimuli of the same type evoke the same HRF throughout the experiment and the heights of overlapping HRFs sum linearly; see also Section 25.3 and Lindquist (2008). The accumulated HRFs along with nuisance signals (drift or trend) and noise model the fMRI time series, which is in this example sampled every $\tau_{TR} = 2$ s. For a scan at 14 s, say, the accumulated value that is observed is affected by the height at 11 s from the stimulus with an onset at 3 s, by the height at 8 s from the stimulus with an onset at 6 s, and by the height at 2 s from the stimulus with an onset at 12 s. Taking the subsequent stimuli and scanning times into account, with this combination of τ_{ISI} and τ_{TR} , the HRF heights that can possibly contribute to an observed MR signal occur every $\Delta T = 1$ s. These HRF heights are presented as the small squares on the third HRF curve in the figure. In general, the discretization interval ΔT is the greatest time making both $m_{ISI} = (\tau_{ISI}/\Delta T)$ and $m_{TR} = (\tau_{TR}/\Delta T)$ integers. \square

Once the fMRI time series, one from each voxel, have been obtained, they are analyzed to help in understanding the underlying unobserved neuronal activity (see also

Lindquist et al. 2009). There are two aspects that are typically of primary interest. For the first, researchers are interested in inferences for the amplitudes (or maximum heights) of the HRFs. This helps to identify voxels that are activated by the stimuli and allows us to understand the strengths of activations. This objective is referred to as *detection of activation* or simply *detection*. Another common interest pertains to studying the evolution of the HRF over time, for example, by estimating the HRF heights (the small squares on the third solid curve in Figure 25.2). This helps to elicit subtler effects of the stimuli and is referred to as the *estimation of the HRF* or simply *estimation*.

A major difficulty for both detection and estimation is that the fMRI time series are usually very noisy and involve nuisance signals due to machine instability and physiological effects such as heartbeat and respiration; see also Smith et al. (1999) and Lund et al. (2006). It is thus crucial to select designs that can help to achieve the highest signal-to-noise ratio at minimal cost. In Section 25.3, we present a very popular general linear models framework for analyzing fMRI data. Within this common framework, we present optimality criteria for evaluating the performance of competing fMRI designs. We consider both widely used fMRI designs and designs obtained via a powerful computational approach.

25.3 Designs for Linear Models

The most popular approach for analyzing fMRI time series is probably the voxel-wise general linear model approach (Worsley and Friston 1995; Dale 1999); see also Section 6.2.1 of Lindquist (2008) and Loh et al. (2008). With this approach, the multiple time series collected from the multiple brain voxels are analyzed separately. The LTI system described in Example 25.2 is often assumed. This system is thought to be reasonable for most studies (Miezin et al. 2000; Lindquist et al. 2009) although violations can occur; see also Wager et al. (2005) and Grinband et al. (2008). With these assumptions, if $h_q(t)$ represents the HRF for the stimulus of the q th type and $x_q(t)$ is a stimulus function indicating the time periods where a q th-type stimulus occurs, then the accumulated HRF $h(t)$ at time t is (Josephs et al. 1997)

$$h(t) = \sum_{q=1}^Q \int_0^t x_q(t - \tau) h_q(\tau) d\tau. \tag{25.1}$$

With brief stimuli, $x_q(t)$ is normally set to $\sum_{j=1}^{J_q(t)} \delta(t - \tau_{q,j})$, where $0 \leq \tau_{q,1} < \dots < \tau_{q,J_q(t)} \leq t$ are the onset times of the q th-type stimuli, $J_q(t)$ is the number of q th-type stimuli before time t , and $\delta(t)$ is the Dirac delta function, which is infinity at the origin and zero elsewhere and satisfies $\int_{-\infty}^{\infty} \delta(t) dt = 1$, and $\int_{-\infty}^{\infty} h_q(t) \delta(t - \tau) dt = h_q(\tau)$. The MR signals $y(t)$ are then modeled as (see also Josephs et al. 1997)

$$y(t) = h(t) + s(t) + \epsilon(t) = \sum_{q=1}^Q \sum_{j=1}^{J_q(t)} h_q(t - \tau_{q,j}) + s(t) + \epsilon(t), \tag{25.2}$$

where $s(t)$ represents the drift or trend of the response and $\epsilon(t)$ is noise.

In the following sections, we discuss the two most popular models of the form in (25.2), one for detection and the other for estimation. These models differ in the assumptions that are made for $h_q(t)$. Criteria for selecting designs based on these two models are then presented.

25.3.1 Models and Design Selection Criteria

25.3.1.1 Linear Model for Detection

Detection, that is, the identification of the brain voxels that are activated by the stimuli, is one of the primary interests to neuroscientists. For this purpose, the HRF $h_q(t)$ in (25.2) is commonly approximated by the product of an assumed shape of the HRF and an unknown amplitude. Specifically, $h_q(t) = \theta_q b(t)$, where $b(t)$ is an assumed common HRF shape and θ_q an unknown amplitude parameter (or coefficient) for the HRF of the q th-type stimuli, $q = 1, \dots, Q$. Using notation similar to that for (25.2), the model for detection can thus be written as

$$y(t) = \sum_{q=1}^Q \sum_{j=1}^{J_q(t)} \theta_q b(t - \tau_{q,j}) + s(t) + \epsilon(t).$$

Since the MR signals $y(t)$ are collected every τ_{TR} s, the model can be rewritten in the following matrix form:

$$\mathbf{y} = \sum_{q=1}^Q \mathbf{X}_q \mathbf{h}^* \theta_q + \mathbf{S} \boldsymbol{\gamma} + \boldsymbol{\epsilon}. \quad (25.3)$$

Here, \mathbf{y} is a $T \times 1$ vector with the j th element, $y_j = y((j-1)\tau_{TR})$ being the signal collected at the j th MR scan of the voxel; \mathbf{h}^* corresponds to the common shape $b(t)$ for the HRFs of the Q types of stimuli and is obtained by evaluating $b(t)$ at k regular time points (the choice of k is discussed in the next paragraph); \mathbf{X}_q is the $T \times k$, 0–1 design matrix for the q th-type stimuli that depends on design d , and thus, on the $\tau_{q,j}$'s and $J_q(t)$ described after (25.1); $\mathbf{S} \boldsymbol{\gamma}$ is the nuisance term corresponding to $s(t)$ in (25.2) with \mathbf{S} being a known matrix and $\boldsymbol{\gamma}$ an unknown parameter vector; and the $T \times 1$ vector $\boldsymbol{\epsilon}$ represents noise. Some of these terms are further discussed in the succeeding text.

With interstimulus interval τ_{ISI} and time to repetition τ_{TR} , the heights of the HRF that can possibly contribute to the observed MR signals occur at 0 s and every ΔT s following a stimulus onset (Figure 25.2). To take these HRF heights into account, the elements of \mathbf{h}^* in Model (25.3) are set to the heights of $b(t)$ evaluated at $t = (j-1)\Delta T$, $j = 1, 2, \dots$. There are many choices for $b(t)$ (e.g., Lu et al. 2006). One popular choice is a particular *double-gamma function*, which is also part of the Statistical Parametric Mapping (SPM) software package (The Wellcome Trust Centre for Neuroimaging 2003), a widely used software package written in MATLAB[®] for fMRI data analysis. This double-gamma function is a linear combination of two gamma probability density functions and can be written as

$$g(t) = \frac{t^5 e^{-t}}{5!} - \frac{1}{6} \frac{t^{15} e^{-t}}{15!}. \quad (25.4)$$

The first density, with a mean and variance of 6, describes the leading wave of the HRF. The second density, with a mean and variance of 16, approximates the undershoot of the HRF. Following Wager et al. (2005), we use a normalized version of $g(t)$, and set $b(t) = g(t)/\max_s g(s)$, so that the maximum of $b(t)$ is 1. Since $g(t)$ is nearly zero after $t = 32$ s, the number of elements of \mathbf{h}^* is set to $k = 1 + \lfloor 32/\Delta T \rfloor$, where $\lfloor a \rfloor$ is the integer floor of a .

The 1's in the j th row of the $T \times k$ design matrix \mathbf{X}_q correspond to the HRF heights that contribute to the j th response y_j . For illustrative purposes, Example 25.3 provides the design matrix for the design in Example 25.1. The method for constructing \mathbf{X}_q can be easily extended to other cases; see also Kao et al. (2012).

Example 25.3

As in Example 25.1, suppose the design is $d = \{110202\dots 0\}$, $\tau_{ISI} = 3$ s, and $\tau_{TR} = 2$ s. To construct \mathbf{X}_1 and \mathbf{X}_2 , we first generate two indicator vectors, δ_1 and δ_2 , each of the same length as the design. The j th element of δ_q is $\delta_{q,j} = \mathbb{1}_{\{d_j=q\}}$, where d_j denotes the j th entry of the design. Here, $\delta_1 = (110000\dots 0)'$ and $\delta_2 = (000101\dots 0)'$. We then form $\mathbf{w}_q = \delta_q \otimes [1, \mathbf{0}'_{m_{ISI}-1}]'$, where \otimes is the Kronecker product, $\mathbf{0}_a$ is an a -by-1 zero vector, and $m_{ISI} = (\tau_{ISI}/\Delta T) = 3$; $q = 1, 2$. The vector \mathbf{w}_q indicates the onset times of the q th-type stimuli in units of ΔT . The number of elements of \mathbf{w}_q is adjusted to $(m_{TR})T$ by leaving out the last few elements or adding zeros at the end (recall that T is the number of MR scans and $m_{TR} = (\tau_{TR}/\Delta T)$, which is 2 in this example). The design matrix for the q th-type stimulus is then $\mathbf{X}_q = [\mathbf{I}_T \otimes (1, \mathbf{0}'_{m_{TR}-1})][\mathbf{w}_q, \mathbf{B}\mathbf{w}_q, \dots, \mathbf{B}^{k-1}\mathbf{w}_q]$, where \mathbf{I}_T is the $T \times T$ identity matrix, $k = 1 + \lfloor 32/\Delta T \rfloor = 33$, and

$$\mathbf{B} = \left[\begin{array}{c|c} \mathbf{0}' & 0 \\ \hline \mathbf{I}_{(m_{TR})T-1} & \mathbf{0} \end{array} \right].$$

In particular, parts of \mathbf{X}_1 and \mathbf{X}_2 are

$$\mathbf{X}_1 = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & \dots \\ 0 & 0 & 1 & 0 & 0 & \dots \\ 0 & 1 & 0 & 0 & 1 & \dots \\ 0 & 0 & 0 & 1 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \dots \end{bmatrix}; \quad \mathbf{X}_2 = \begin{bmatrix} \mathbf{O}_{5,33} \\ 0 & 1 & 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & 1 & 0 & \dots \\ 0 & 0 & 0 & 0 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \dots \end{bmatrix},$$

where $\mathbf{O}_{m,n}$ is the $m \times n$ matrix of zeros. For example, the third row of \mathbf{X}_1 corresponds to the third scan, which occurs at 4 s. At that time, the HRF for the stimulus of type 1 with an onset at 0 s is at its height at 4 s. The HRF for the stimulus of type 1 with an onset of 3 s is at its height at 1 s. Therefore, in the third row of \mathbf{X}_1 , we see a 1 in positions 2 (HRF height at 1 s) and 5 (HRF height at 4 s), so that $\mathbf{X}_1\mathbf{h}^*$ is the sum of these two heights. The third row in \mathbf{X}_2 is zero since there has not yet been an onset of a type 2 stimulus at time 4 s. □

If $\tau_{ISI} = \tau_{TR}$, then the $(j + 1)$ st column of \mathbf{X}_q can be obtained by shifting the elements of the j th column one position down and adding a 0 at the top. For other cases, including Example 25.3, the design matrix is more complicated.

The nuisance term $\mathbf{S}\boldsymbol{\gamma}$ in Model (25.3) attempts to model systemic shifts in the signals that are of no scientific interest. It is not uncommon to use a second-order polynomial for this (Liu 2004; Maus et al. 2010b, 2011), although there are alternative models; see also Lund et al. (2006); Sarkka et al. (2012); and Churchill et al. (2012). In addition, the noise $\boldsymbol{\epsilon}$ is often assumed to follow an AR1 or AR2 process (Lindquist 2008). There will thus be

unknown autocorrelation coefficients; see Section 25.5.2 for a further discussion. While other correlation structures may be considered, the autoregressive processes tend to render satisfactory results (e.g., Worsley et al. 2002; Lenoski et al. 2008; Maus et al. 2010b).

With Model (25.3), the focus may be on $C_\theta \theta$, where $\theta = (\theta_1, \theta_2, \dots, \theta_Q)'$ and C_θ is a user-specified matrix of coefficients corresponding to linear combinations of interest. When $C_\theta = I_Q$, the objective is on studying the strength of brain activation due to each stimulus type. For the comparison of the activation strengths between stimulus types, C_θ consists of coefficients for pairwise comparisons; for example, $C_\theta = [1 \ -1]$ for $Q = 2$.

25.3.1.2 Linear Model for Estimation

For estimating the HRF $h_q(t)$, a second model is usually considered (e.g., Dale 1999; Liu and Frank 2004). In contrast to Model (25.3), the model for estimation assumes nearly no prior knowledge about the HRF. An unknown parameter vector h_q is used to represent all the HRF heights that contribute to MR signals. The model for estimation is then

$$y = \sum_{q=1}^Q X_q h_q + S\gamma + \epsilon, \tag{25.5}$$

where $h_q = (h_{q1}, \dots, h_{qk})'$ with h_{qj} representing the unknown HRF height $h_q((j - 1)\Delta T)$, $j = 1, \dots, k (= 1 + \lfloor 32/\Delta T \rfloor)$, and $q = 1, \dots, Q$; the remaining terms are as in Model (25.3). Interest would focus on $C_h h = C_h(h'_1, \dots, h'_Q)'$ for a coefficient matrix C_h . For example, when $C_h = I_{Qk}$, the main interest is on the estimation of the heights in h .

25.3.1.3 Design Selection Criteria

Optimal fMRI designs yield, in some sense, the best least squares estimates of the parametric functions of interest. To define what we mean by best, we focus on appropriate covariance matrices. For detection, the covariance matrix of $C_\theta \hat{\theta}$ is proportional to

$$\Sigma_\theta(d) \equiv \frac{\text{Cov}(C_\theta \hat{\theta})}{\sigma^2} = C_\theta [(I_Q \otimes h^*)' X' V' (I_T - P_{VS}) V X (I_Q \otimes h^*)]^{-1} C'_\theta, \tag{25.6}$$

where $X = (X_1, \dots, X_Q)$; $P_A = A(A'A)^{-1}A'$ is the orthogonal projection matrix on the space spanned by the columns of A with A^- being a generalized inverse of A . Thus, P_{VS} projects orthogonally onto the column space of VS , the product of V and S , where V is a whitening matrix so that $\text{Cov}(V\epsilon) = \sigma^2 I_T$ and S is as in (25.3) and (25.5). Similarly, the covariance matrix of $C_h \hat{h}$ is proportional to

$$\Sigma_h(d) \equiv \frac{\text{Cov}(C_h \hat{h})}{\sigma^2} = C_h [X' V' (I_T - P_{VS}) V X]^{-1} C'_h. \tag{25.7}$$

Note that the matrices $\Sigma_\theta(d)$ and $\Sigma_h(d)$ depend on design d through the choice of X .

Dale (1999) and Friston et al. (1999) worked on estimation and detection problems, respectively, and proposed to evaluate designs by considering the reciprocal of the trace of the appropriate covariance matrix. This corresponds to the A -optimality criterion and

was subsequently applied in many studies on fMRI designs (e.g., Liu et al. 2001; Wager and Nichols 2003; Liu 2004; Liu and Frank 2004; Kao et al. 2009a). For this criterion, we will write

$$\begin{aligned} \phi_{\theta}(d) &= \frac{r_{\theta}}{\text{trace}[\Sigma_{\theta}(d)]}, & \text{for detection;} \\ \phi_h(d) &= \frac{r_h}{\text{trace}[\Sigma_h(d)]}, & \text{for estimation.} \end{aligned}$$

Here, r_{θ} and r_h denote the number of rows of C_{θ} and C_h , respectively.

In addition to A -optimality, Wager and Nichols (2003) included the D -optimality criterion in their MATLAB program for selecting fMRI designs. This criterion can be written as

$$\begin{aligned} \phi_{\theta}(d) &= \det[\Sigma_{\theta}(d)]^{-1/r_{\theta}}, & \text{for detection;} \\ \phi_h(d) &= \det[\Sigma_h(d)]^{-1/r_h}, & \text{for estimation.} \end{aligned}$$

Maus et al. (2010b) applied the genetic algorithm of Kao et al. (2009a) to obtain both A - and D -optimal fMRI designs and evaluated the robustness of the two sets of designs against the choice of optimality criterion. In their empirical studies, A -optimal designs tend to be more robust than D -optimal designs. Specifically, the relative D -efficiencies (relative to the D -optimal designs) of the A -optimal designs were at least as good as the relative A -efficiencies (relative to the A -optimal designs) of the D -optimal designs. The computational tools described in Section 25.3.3 facilitate the use of criteria that include A - and D -optimality, as well as multiobjective criteria.

While statistical criteria are important, psychological effects and practical concerns may also need to be taken into account at the design stage. This includes avoiding designs with patterns that can be predicted by a subject based on the stimuli presented (e.g., Dale 1999; Liu et al. 2001; Wager and Nichols 2003). For example, with a design that repeats the same pattern, such as {111000}, a subject will quickly guess which stimulus comes next. Wager and Nichols (2003) proposed an R th-order counterbalancing criterion to avoid using designs with such patterns. The version of this criterion presented here is based on a slight improvement proposed by Kao et al. (2009a):

$$\phi_c(d) = \sum_{r=1}^R \sum_{i=1}^Q \sum_{j=1}^Q \left| \left| n_{ij}^{(r)} - (n-r)f_i f_j \right| \right|. \tag{25.8}$$

The value for this criterion is computed from a design $d_{(0)}$ obtained from d by removing all its zeros. The number of elements of $d_{(0)}$ is denoted by n ; $n_{ij}^{(r)}$ is the number of times that a type i stimulus is followed with a lag r by a type j stimulus in $d_{(0)}$, $r = 1, \dots, R$. The value of R controls the *degree* of counterbalancing and is typically set to 3; f_i is a user-specified desired proportion for type- i stimuli in $d_{(0)}$. With no preference, we may take $f_i = 1/Q$, $i = 1, \dots, Q$.

For the simplest case where $R = 1$, $Q = 2$, and $f_1 = f_2 = 1/2$, this criterion helps to achieve a design $d_{(0)}$ in which the four possible pairs of stimulus types, namely, (1, 1),

(1, 2), (2, 1), and (2, 2), appear nearly equally often. With a larger value for R , nearly counterbalanced designs become even harder to predict. However, even a small R can help to prevent subjects from detecting any patterns; see also Liu and Frank (2004).

Brendel et al. (2010) used a design with more stimuli of one type than of the other types; see also Kao et al. (2012). To aim for a design with user-specified proportions of occurrences for the various types of stimuli, the following criterion may be used:

$$\phi_f(d) = \sum_{i=1}^Q [|n_i - nf_i|], \tag{25.9}$$

where n_i is the number of occurrences of the type- i stimulus in $d_{(0)}$ and n and f_i are as in (25.8).

Note that both criteria (25.8) and (25.9) are of the smaller-the-better variety.

Designs that are very different in terms of the locations of zeros that they contain can lead to the same $d_{(0)}$. Such designs may perform very differently under other criteria, such as A - or D -optimality. Therefore, ϕ_c and ϕ_f are normally used in combination with other criteria. In that case a weighted-sum criterion may be considered for achieving a multiobjective design:

$$\phi(d) = w_1\phi_\theta^*(d) + w_2\phi_h^*(d) + w_3\phi_c^*(d) + w_4\phi_f^*(d), \tag{25.10}$$

where $w_i \in [0, 1]$ are user-specified weights with $\sum_{i=1}^4 w_i = 1$ and ϕ_j^* is a standardized form of ϕ_j , $j = \theta, h, c, f$.

Following Kao et al. (2009a), we use the standardization

$$\phi_j^* = \begin{cases} \frac{\phi_j - \min(\phi_j)}{\max(\phi_j) - \min(\phi_j)}, & j = \theta, h; \\ 1 - \frac{\phi_j - \min(\phi_j)}{\max(\phi_j) - \min(\phi_j)}, & j = c, f. \end{cases}$$

The maxima for ϕ_c and ϕ_f are obtained by a design that only contains the stimulus type with the smallest specified proportion f_j . The minima for these two criteria are zero, corresponding to (possibly hypothetical) optimal designs under these criteria. Similarly, the minimal values for ϕ_θ and ϕ_h are zero, corresponding to designs for which the parametric functions of interest are nonestimable. The maxima for ϕ_θ and ϕ_h are generally not available. An algorithm, such as that introduced in Section 25.3.3, may be used to approximate these values.

We now turn to a discussion of some traditional fMRI designs.

25.3.2 Efficient fMRI Designs

Different fMRI designs are recommended in the literature for different study objectives. Friston et al. (1999) studied detection problems with linear models and observed that so-called blocked designs can outperform other designs that they considered. Blocked designs are sometimes, perhaps confusingly, called block designs in the fMRI literature. They present stimuli of the same type in clusters or blocks. For example, repetitions of {111122220000} or {11112222} form blocked designs for $Q = 2$ with block size four. The

former includes blocks of the control and is usually recommended if individual stimulus effects ($C_\theta = I_Q$) are of interest. The blocked design without zeros tends to perform well for detecting pairwise differences ($\theta_i - \theta_j$) between stimulus types; see also Kao et al. (2008) and Maus et al. (2010a). Following Maus et al. (2010a), these two types of blocked designs will be said to have the pattern 'ABN' and 'AB', respectively. Block designs with the pattern 'ANBN', which would repeat a sequence like {1111000022220000} for $Q = 2$, can also perform well at times (Maus et al. 2010a). It has twice as many zeros as occurrences of any stimulus type. We will also consider blocked designs in which the blocks of zeros are about half as long as those for the other stimulus types and will denote these by 'ANBN/2'. For simplicity, we will also use this notation if $Q > 2$, even though it would then be more precise to make some adjustments, such as ABC, ABCN, ANBNCN, and ANBNCN/2 for $Q = 3$.

With the double-gamma function in (25.4) for h^* in Model (25.3), a design with a block duration of about 15 s can yield high ϕ_θ -values (Henson 2007). Consequently, with a τ_{ISI} of 4 s, say, a good blocked design will have blocks of, approximately, size four. Nevertheless, the optimal block duration can vary with the selection of the HRF shape $b(t)$ and the autocorrelation of the noise. Maus et al. (2010a) provide a computational method for approximating the optimal block duration. The genetic algorithm presented in Section 25.3.3 can also be used for obtaining an efficient design for detection, and it normally yields designs with slightly higher ϕ_θ -values than blocked designs. The designs obtained by the genetic algorithm tend to feature blocks of stimuli of the same type, and are thus similar to blocked designs.

For estimation of an HRF, Dale (1999) showed that blocked designs are inefficient. Buxton et al. (2000) demonstrated that random designs can yield high efficiencies in estimation, but do not perform well in detection. Their results also suggested that a design formed by a fraction of a random design and a fraction of a blocked design may be considered as a compromise when interested in both detection and estimation. Following these pioneering contributions, Liu et al. (2001) and Liu and Frank (2004) provided mathematical formulas to approximate the fundamental trade-off between detection and estimation. Liu (2004) studied not only blocked and random designs but also m -sequences, mixed designs, permuted blocked designs, and clustered m -sequences. An m -sequence is a linear recurrence sequence (LRS) over a finite field $GF(p)$, where p is a prime or a prime power. Specifically, an LRS is determined by a homogeneous linear recurrence relation, $d_{l+M} = \sum_{i=1}^M a_i d_{l+M-i}$, where M is a positive integer, $a_1, \dots, a_M \in GF(p)$ are coefficients, and d_l is the l th element of the sequence; $l = 1, \dots, L$. The recurrence relation for generating an m -sequence with Q types of stimuli can be obtained from a primitive polynomial over the Galois field $GF(Q+1)$ (MacWilliams and Sloane 1977; Godfrey 1993). For example, with $Q = 2$, an m -sequence {11012202} over $GF(3)$ can be obtained from the primitive polynomial $f(x) = x^2 - 2x - 1$. The selected $f(x)$ gives the recurrence relation $d_{l+2} = 2d_{l+1} + d_l$. Specifically, the coefficient of d_{l+j} on the right-hand side of the recurrence relation corresponds to the coefficient of x^j of $f(x)$; $j = 0, 1, \dots, \deg(f) - 1$, where $\deg(f)$ is the degree of $f(x)$. An m -sequence is then obtained with a specified initial status, say, $(d_1, d_2) = (1, 1)$. Buračas and Boynton (2002) proposed the use of m -sequences for estimating the HRF because of the low autocorrelation property of these designs. But, while efficient for estimation, m -sequences do not perform as well as blocked designs for detection.

A mixed design is a design that is obtained by concatenating part of a blocked design with part of an m -sequence or a random design. A permuted blocked design is constructed by permuting elements of a blocked design to make the design increasingly *random*. By contrast, a clustered m -sequence is generated by sequentially permuting elements of an

m -sequence to create blocks of stimuli of the same type. These designs can be generated by the MATLAB program of Liu (2004).

While the designs discussed so far offer sound choices for standard situations, fMRI experiments are complicated and expensive (\$200–\$1000 per hour scan), and experimenters tend to seek efficient designs best suited to their unique experiments. Designs for specific experimental settings are increasingly important for the ever more sophisticated fMRI experiments (Lindquist 2008). Algorithmic approaches, as discussed in Section 25.3.3, facilitate finding efficient fMRI designs for such situations and also tend to lead to more efficient designs than the traditional designs introduced in this section for standard situations.

25.3.3 Genetic Algorithms for Finding Efficient fMRI Designs

For a given L (number of elements) and Q (number of stimulus types), there are $(Q + 1)^L$ possible fMRI designs. Since L can easily be as large as several hundreds, finding an optimal design for a given criterion is a challenging optimization problem. An answer will depend on several factors, such as the interstimulus interval τ_{ISI} , time to repetition τ_{TR} , study objectives (detection, estimation, or both), model assumptions, parametric functions of interest, optimality criterion (A -optimality, D -optimality, or others), and possible psychological or practical constraints (see Section 25.3.1). Tackling this problem requires the use of a versatile algorithm.

To this end, Wager and Nichols (2003) proposed the use of a genetic algorithm (see Holland 1975, and Chapter 21). A genetic algorithm mimics Darwin's theory of evolution by moving through generations of chromosomes, which are representations of fMRI designs in our context. Each generation consists of a selected number of chromosomes, and the goodness of a chromosome is evaluated through the use of an appropriately chosen objective function. Good chromosomes of the current generation are used to reproduce chromosomes for the next generation. Following the survival-of-the-fittest principle, only chromosomes with better fit (higher objective function values) survive to pass on their *genes*. The process, when repeated, tends to preserve more and more good traits, resulting in better and better chromosomes (see also Ahn 2006).

Wager and Nichols (2003) treated fMRI designs as chromosomes, and segments of designs as genes. They used a weighted sum of normalized optimality criteria as their objective function and proposed a genetic algorithm to search for efficient fMRI designs based on user-specified experimental settings. Kao et al. (2009a) proposed an improved knowledge-based genetic algorithm, with a key innovation being the incorporation of traditional fMRI designs in the algorithm. This not only helped with a significant increase in speed but also resulted in designs with higher efficiencies. We briefly describe this improved algorithm in the succeeding text. A MATLAB program implementing this approach with a user's manual can be found in Kao (2009a).

With a user-specified constant g , the genetic algorithm of Kao et al. (2009a) starts with $2g$ initial designs, including blocked designs of various block sizes, m -sequences (provided they exist), random designs, and mixed designs. These $2g$ designs form the first generation, and their fit is evaluated through an objective function that is a single optimality criterion or, in the case of multiobjective studies, a weighted sum of standardized criteria. With probability proportional to the objective function value, g pairs of different designs are selected from the current generation with replacement. The selected pairs of designs are then used to generate g pairs of offspring designs via crossover and mutation. Specifically, the crossover operator exchanges corresponding subsequences of the paired designs

based on a randomly selected cut point. The mutation operator then randomly selects a portion of elements of the resulting designs and replaces these elements by integers randomly generated from the discrete uniform distribution over $0, 1, 2, \dots, Q$. At each generation, the algorithm also considers N immigrant designs. These are additionally generated designs consisting of random designs, blocked designs, and mixed designs. Immigrant designs help to assure that the algorithm does not get trapped at a local optimum. From the pool of the $2g$ designs in the current generation, the $2g$ offspring designs, and the N immigrants, we then select the $2g$ designs of the best fit; these selected designs form the next generation. The process is repeated until a stopping rule is met, for example, no significant improvement is made (for details, see Kao 2009a). The algorithm keeps track of the design with the best fit over all generations.

The objective function of the algorithm can be set to ϕ_θ , ϕ_h , or ϕ (see Section 25.3.3). With ϕ_θ (or ϕ_h), the result is an efficient design for detection (or estimation). The obtained design can be used to approximate the maximal value for ϕ_θ (or ϕ_h), which is required for calculating ϕ in (25.10). With ϕ as the objective function, and at least two user-specified positive weights, the algorithm searches for an efficient multiobjective design. The weights may be selected based on the importance of each objective.

As demonstrated in Kao et al. (2009a), this algorithm is fast and obtains designs that are more efficient than those obtained by the algorithm of Wager and Nichols (2003) and by the traditional designs introduced in Section 25.3.2. We now present examples for both a single- and multiobjective criterion.

25.3.4 Illustrative Examples

We consider four choices for (Q, L) , namely, $(2, 242)$, $(3, 255)$, $(4, 624)$, and $(6, 342)$. These choices were also considered in Liu (2004) and Kao et al. (2009a). The design has $L = (Q + 1)^\ell - 1$ elements for an integer ℓ , which allows consideration of m -sequences. Both for Models (25.3) and (25.5), we assume a second-order polynomial drift in the response and AR1 noise with $\rho = 0.3$. More discussion about the value of ρ appears in Section 25.5.2. We consider two possible choices for the parametric functions of interest: (1) individual stimulus effects, with C_θ and C_h being identity matrices, and (2) pairwise comparisons, with C_θ and C_h containing coefficients for all possible pairwise comparisons between stimulus types. The combinations for (τ_{TR}, τ_{ISI}) that we consider are $(2, 2)$, $(2, 3)$, and $(2, 4)$, with $\Delta T = 2, 1$, and 2 s, respectively. Both A - and D -optimality are applied for each combination. We note that the number of elements, k , of the HRF parameter vector \mathbf{h}_q of (25.5) with $\Delta T = 1$ s is about twice as large as that with $\Delta T = 2$ s.

We use the genetic algorithm of Kao et al. (2009a) with the parameters of the algorithm set to their default values as presented in Kao (2009a). In particular, we use a population size $2g = 20$, a mutation rate of 1%, and the number of immigrants $N = 4$. The search is terminated when there is no significant improvement in the value of the objective function (for details, see Kao 2009a). The GA designs obtained by this algorithm will be compared to some of the traditional designs in Section 25.3.2 with respect to the criteria ϕ_θ , ϕ_h , and $\phi = 0.5\phi_\theta^* + 0.5\phi_h^*$. For the traditional designs, we employ the MATLAB program provided by Liu (2004) to obtain m -sequences, mixed designs, permuted blocked designs, and clustered m -sequences. In addition, blocked designs with blocks of (nearly) 15 s length are also included among the designs that we compare. In particular, when $Q = 2$, these blocked designs are formed by repetitions of $\{1111111222222200000000\}$ for $\tau_{ISI} = 2$ s, by repetitions of $\{11112222200000\}$ for $\tau_{ISI} = 3$ s, and by repetitions of $\{11122220000\}$ for $\tau_{ISI} = 4$ s.

The same block length is also considered for the construction of blocked designs for $Q = 3, 4,$ and 6 . In addition to designs with the pattern ABN, we also consider blocked designs with the patterns AB, ANBN, and ANBN/2 (Section 25.3.2). For ANBN/2, the block length of zeros is about 7.5 s.

Table 25.1 provides the relative efficiencies for the various scenarios of the best traditional design to the GA design obtained by the algorithm. With relative efficiencies of at least 95.7% for the best traditional designs under ϕ_θ , Table 25.1 shows that highly efficient traditional designs can be found for detection. As expected, these designs are blocked designs. For estimation of the HRF (the ϕ_h criteria), m -sequences can perform well for D -optimality. However, their efficiency relative to GA designs can be less impressive for A -optimality, especially when τ_{ISI} is not (a multiple of) τ_{TR} . As noted previously, when $(\tau_{TR}, \tau_{ISI}) = (2, 3)$, the number of HRF parameters is about twice that in the other cases. For comparing HRFs between stimulus types, designs obtained by the genetic algorithm clearly outperform traditional designs, for both A - and D -optimality. The same tends to be true for the multiobjective criteria that place equal weight on detection and estimation (the ϕ -criteria).

Table 25.2 provides the pattern of the efficient blocked designs whose relative ϕ_θ -values are reported in Table 25.1. When pairwise comparisons are of interest, the pattern that yields the highest ϕ_θ -values for both A - and D -optimality is AB. If interest is in estimating

TABLE 25.1

Relative Performance (%) of the Best Traditional Design versus the GA Design

Setting	Individual Stimulus Effect				Pairwise Comparison			
	Q = 2	Q = 3	Q = 4	Q = 6	Q = 2	Q = 3	Q = 4	Q = 6
<i>Detection</i> (ϕ_θ)								
A-opt., $\tau_{ISI} = 2$	97.6	95.7	98.2	98.3	99.6	98.7	99.0	98.9
A-opt., $\tau_{ISI} = 3$	98.4	98.5	99.9	99.8	99.6	99.0	99.0	98.7
A-opt., $\tau_{ISI} = 4$	96.2	97.7	98.5	99.5	99.8	97.5	100.0	97.6
D-opt., $\tau_{ISI} = 2$	99.8	99.2	99.2	99.5	99.7	98.7	99.1	99.1
D-opt., $\tau_{ISI} = 3$	100.0	99.1	99.4	100.0	99.6	99.0	99.0	98.8
D-opt., $\tau_{ISI} = 4$	97.6	97.8	98.0	98.3	99.8	97.5	100.0	97.7
<i>Estimation</i> (ϕ_h)								
A-opt., $\tau_{ISI} = 2$	93.9	91.9	89.4	88.4	65.6	73.9	79.3	86.5
A-opt., $\tau_{ISI} = 3$	89.3	82.3	87.1	75.0	62.8	68.9	77.6	73.4
A-opt., $\tau_{ISI} = 4$	97.6	94.4	91.5	87.8	67.1	75.4	80.3	87.0
D-opt., $\tau_{ISI} = 2$	98.7	99.3	99.9	100.0	68.4	76.1	81.0	89.3
D-opt., $\tau_{ISI} = 3$	94.1	93.1	96.2	92.4	63.1	70.6	79.6	80.2
D-opt., $\tau_{ISI} = 4$	98.6	98.9	99.5	99.8	67.0	75.6	80.5	87.1
<i>Biobjective</i> ($\phi = (\phi_\theta^* + \phi_h^*)/2$)								
A-opt., $\tau_{ISI} = 2$	90.5	86.9	84.3	82.9	65.2	71.2	76.7	80.7
A-opt., $\tau_{ISI} = 3$	89.4	85.1	86.8	78.3	62.7	69.7	76.3	77.3
A-opt., $\tau_{ISI} = 4$	95.1	93.0	90.2	86.9	66.0	74.3	79.4	84.1
D-opt., $\tau_{ISI} = 2$	93.2	93.8	94.5	93.9	66.8	71.7	76.6	81.7
D-opt., $\tau_{ISI} = 3$	92.0	92.2	94.4	92.4	63.6	70.3	76.9	79.6
D-opt., $\tau_{ISI} = 4$	96.1	96.9	97.6	97.3	65.9	73.9	79.0	84.2

TABLE 25.2

The Patterns of the Best Blocked Designs for Detection

Setting	Individual Stimulus Effect				Pairwise Comparison			
	Q = 2	Q = 3	Q = 4	Q = 6	Q = 2	Q = 3	Q = 4	Q = 6
A-opt., $\tau_{ISI} = 2$	ABN	ANBN/2	ANBN/2	ANBN/2	AB	AB	AB	AB
A-opt., $\tau_{ISI} = 3$	ANBN	ANBN/2	ANBN/2	ANBN/2	AB	AB	AB	AB
A-opt., $\tau_{ISI} = 4$	ANBN	ANBN/2	ANBN/2	ANBN/2	AB	AB	AB	AB
D-opt., $\tau_{ISI} = 2$	ABN	ABN	ABN	ABN	AB	AB	AB	AB
D-opt., $\tau_{ISI} = 3$	ABN	ABN	ABN	ABN	AB	AB	AB	AB
D-opt., $\tau_{ISI} = 4$	ABN	ABN	ABN	ABN	AB	AB	AB	AB

the vector θ , blocked designs with the pattern ABN are efficient for *D*-optimality, while the pattern ANBN/2 is best for *A*-optimality when $Q = 3, 4,$ or 6 . However, when $Q = 2$, designs with this pattern can be outperformed by blocked designs with other patterns. Specifically, for $Q = 2$ and *A*-optimality, the patterns ABN ($\tau_{ISI} = 2$ s) and ANBN ($\tau_{ISI} = 3$ and 4 s) are most efficient. However, the pattern ABN is also highly efficient in the latter cases and would thus be a good choice for all combinations that we studied for $Q = 2$ if the interest is in θ . This block pattern is also recommended by Maus et al. (2010a) for a maximin robustness criterion; they consider a linear drift and an AR1 error structure with a correlation coefficient $\rho \in [0, 0.5]$.

25.4 Designs for Nonlinear Models

The results in Section 25.3 are for the widely used general linear model approach. Not only does this approach use slightly different model assumptions for estimation and detection, but it also assumes a common HRF shape h^* for every brain voxel when considering detection. This assumption is sometimes seen as being unrealistic. For example, Handwerker et al. (2004) observed variable HRF shapes across voxels. Lindquist and Wager (2007) argue that a misspecified HRF shape may lead to incorrect conclusions about brain activation, so that an approach allowing for uncertainty in the HRF shape seems preferable; see also Lindquist et al. (2009).

Attempts to accommodate the possibility of uncertainty in the HRF shape have also entered the fMRI design literature. Focusing on design selection, Kao (2009b) considered the nonlinear model

$$y = \sum_{q=1}^Q X_q h(p) \theta_q + S\gamma + \epsilon. \tag{25.11}$$

Here, the vector $h(p)$, which is indexed by a vector p of unknown parameters, represents the HRF shape. As in Model (25.3), the same shape is assumed for every stimulus type. The parameter vector p is to be estimated for each voxel at the analysis stage, and different voxels may yield different estimates of p . However, the parameter values, and thus $h(p)$, are unknown at the design stage. Other terms in Model (25.11) are as in Model (25.3). A possible choice for $h(p)$ is by considering the double-gamma function of (25.4) with unknown parameters, p_1, p_2, \dots, p_6 to allow for various shapes:

$$g_0(t; p_1, p_2, \dots, p_6) = \frac{(t - p_6)^{p_1/p_3 - 1} e^{-(t - p_6)/p_3}}{\Gamma(p_1/p_3) p_3^{p_1/p_3}} - \frac{1}{p_5} \frac{(t - p_6)^{p_2/p_4 - 1} e^{-(t - p_6)/p_4}}{\Gamma(p_2/p_4) p_4^{p_2/p_4}}, \quad (25.12)$$

where $\Gamma(a) = \int_0^\infty t^{a-1} e^{-t} dt$ is the gamma function. Thus $g(t)$ of (25.4) is a special case of $g_0(t; p_1, p_2, \dots, p_6)$ with $(p_1, p_2, \dots, p_6) = (6, 16, 1, 1, 6, 0)$. We follow Wager et al. (2005) to keep the two most influential parameters, p_1 (time to peak) and p_6 (time to onset), as unknown parameters and fix the others at their default values of $(p_2, p_3, p_4, p_5) = (16, 1, 1, 6)$, although (some of) these latter p_i 's could also be regarded as unknowns. With a slight abuse of notation, for simplicity we will from hereon write $\mathbf{p} = (p_1, p_6)'$; the HRF shapes for various values of \mathbf{p} can be found in Figure 25.3. For a given \mathbf{p} , the j th element of $\mathbf{h}(\mathbf{p})$ is then equal to $g_0(t = (j - 1)\Delta T; \mathbf{p}) / \max_s g_0(s; \mathbf{p})$, and the number of elements of $\mathbf{h}(\mathbf{p})$ is $1 + \lceil 32/\Delta T \rceil$.

While Model (25.11) allows for different HRF shapes for different voxels, the nonlinearity of the model (through $\mathbf{h}(\mathbf{p})\theta_q$) also introduces some new challenges. The covariance matrix of the least squares estimate $\hat{\theta}$ of θ can be approximated by

$$\begin{aligned} \sigma^2 \Sigma_\theta(d; \theta, \mathbf{p}) &= \sigma^2 [E(\mathbf{p})' (\mathbf{I}_T - \mathbf{P}_{L(\theta, \mathbf{p})}) E(\mathbf{p})]^{-1}; \text{ where} \\ E(\mathbf{p}) &= (\mathbf{I}_T - \mathbf{P}_{V_S}) \mathbf{VX} (\mathbf{I}_Q \otimes \mathbf{h}(\mathbf{p})); \\ L(\theta, \mathbf{p}) &= [L_1, L_6]; L_i = (\mathbf{I}_T - \mathbf{P}_{V_S}) \mathbf{VX} (\mathbf{I}_Q \otimes \partial \mathbf{h}(\mathbf{p}) / \partial p_i) \theta; i = 1, 6. \end{aligned}$$

Here, \mathbf{X} is as in (25.6) and $\partial \mathbf{h}(\mathbf{p}) / \partial p_i$ is the partial derivative of $\mathbf{h}(\mathbf{p})$ with respect to p_i .

This approximation is obtained by a linearization of Model (25.11) via a Taylor expansion (Fedorov and Hackl 1997). A difficulty in finding an fMRI design optimizing some function




FIGURE 25.3
The HRF shapes for various values of $\mathbf{p} = (p_1, p_6)$.

of $\Sigma_\theta(d; \theta, \mathbf{p})$ is that the answer will depend on unknown model parameters. One way to tackle this issue is by using locally optimal designs (Chernoff 1953), which are optimal for a best guess for θ and \mathbf{p} . However, locally optimal designs are not satisfactory since fMRI experiments involve multiple brain voxels with, possibly, different values for these parameters. This calls for designs that are efficient for multiple values of θ and \mathbf{p} .

With a prior distribution for the model parameters, Kao (2009b) adapted the genetic algorithm of Section 25.3.3 to obtain Bayesian optimal designs. In cases where only a range of the parameter values is available, Maus et al. (2012) searched for maximin robust designs under the same model as in Kao (2009b). Their primary focus was on cases with one stimulus type ($Q = 1$) under D -optimality. Kao et al. (2013) also studied maximin designs for Model (25.11) and proposed an approach that is also applicable for larger Q and for optimality criteria that are invariant under simultaneous permutations of rows and columns of the information matrix $E(\mathbf{p})'(I_T - P_{L(\theta, \mathbf{p})})E(\mathbf{p})$. Both the A - and D -optimality criteria satisfy this invariance property.

The main focus of Kao et al. (2013) is on finding designs d that maximize

$$\min_{(\theta, \mathbf{p}) \in \Theta \times \mathcal{P}} \phi_\theta(d; \theta, \mathbf{p}), \tag{25.13}$$

where Θ and \mathcal{P} are parameter spaces for θ and \mathbf{p} , respectively. The function ϕ_θ is chosen as in Section 25.3.1 with $\Sigma_\theta(d)$ replaced by $\Sigma_\theta(d; \theta, \mathbf{p})$. Based on empirical results (see Rosen et al. 1998; Lindquist 2008), the parameter space \mathcal{P} for \mathbf{p} is set to $\{(p_1, p_6) \mid p_1 \in [6, 9], p_6 \in [0, 2]\}$. Without precise information about the possible value of θ , Θ is temporarily set to the entire Q -dimensional space \mathbb{R}^Q . We note that a negative θ_q corresponds to *deactivation* described in Friston et al. (1998). Kao et al. (2013) made use of the following observations to reduce the parameter space and computational burden.

Lemma 25.1 $\Sigma_\theta(d; \mathbf{0}, \mathbf{p}) \leq \Sigma_\theta(d; \theta, \mathbf{p})$ in Löwner ordering for any θ, \mathbf{p} , and design d that makes $\Sigma_\theta(d; \theta, \mathbf{p})$ nonsingular.

Lemma 25.2 $\Sigma_\theta(d; c\theta, \mathbf{p}) = \Sigma_\theta(d; \theta, \mathbf{p})$ for any scalar $c \neq 0$.

Lemma 25.3 Let $\mathcal{G} = \{G_1, \dots, G_M\}$ be a set of $Q \times Q$ permutation matrices. Suppose $\Theta_0 \subset \Theta$ is such that $\Theta = \bigcup_{m=1}^M \Theta_m$, where $\Theta_m = G_m \Theta_0 = \{\theta_m \mid \theta_m = G_m \theta_0, \theta_0 \in \Theta_0\}$ for $m = 1, 2, \dots, M$. If d_0^* is a maximin design for $\Theta_0 \times \mathcal{P}$ and $\min_{\Theta_0 \times \mathcal{P}} \phi_\theta(d_0^*; \theta_0, \mathbf{p}) = \min_{\Theta_m \times \mathcal{P}} \phi_\theta(d_0^*; \theta_m, \mathbf{p})$ for any $m = 1, 2, \dots, M$, then d_0^* is also a maximin design for $\Theta \times \mathcal{P}$.

Proofs of these properties can be found in Kao et al. (2013). We note that, with Lemma 25.2, we can slightly extend Lemma 25.3.

Corollary 25.1 The result in Lemma 25.3 still holds after replacing Θ_m by $\Theta_m^* = G_m^c \Theta_0 = \{c(\theta_m)\theta_m \mid \theta_m = G_m \theta_0, \theta_0 \in \Theta_0, c(\theta_m) \text{ is a nonzero scalar that may vary with } \theta_m\}$; $m = 1, 2, \dots, M$.

From Lemma 25.1, $\phi_\theta(d; \theta, \mathbf{p}) \leq \phi_\theta(d; \mathbf{0}, \mathbf{p})$ for any $\theta \neq \mathbf{0}$. Consequently, the minimal ϕ_θ -value can be achieved within $(\mathbb{R}^Q - \{\mathbf{0}\}) \times \mathcal{P}$. In addition, Lemma 25.2 helps to further reduce the parameter space for θ . Specifically, when $Q = 1$, we have $\phi_\theta(d; \theta_1, \mathbf{p}) = \phi_\theta(d; 1, \mathbf{p})$ for $\theta_1 \neq 0$. For this case, the search algorithm introduced in Section 25.3.3 can be adapted to search for a design maximizing $\min_{\mathcal{P}} \phi_\theta(d; 1, \mathbf{p})$, where the minimum is evaluated by using a fine grid on \mathcal{P} .

When $Q > 1$, Lemmas 25.1 and 25.2 allow us to focus on the surface of a Q -dimensional unit hemisphere centered at the origin. For example, when $Q = 2$, the parameter space of θ can be reduced to $\Theta = \{(\cos \varphi_1, \sin \varphi_1) \mid \varphi_1 \in (-\pi/2, \pi/2]\}$. Similarly, $\Theta = \{(\cos \varphi_1, \sin \varphi_1 \cos \varphi_2, \sin \varphi_1 \sin \varphi_2) \mid \varphi_i \in (-\pi/2, \pi/2]\}$ for $Q = 3$. In general,

$$\Theta = \left\{ (\theta_1, \dots, \theta_Q) \mid \theta_1 = \cos \varphi_1; \theta_q = \cos \varphi_q \prod_{i=1}^{q-1} \sin \varphi_i, q = 2, 3, \dots, Q - 1; \right. \\ \left. \theta_Q = \prod_{i=1}^{Q-1} \sin \varphi_i; \varphi_1, \dots, \varphi_{Q-1} \in (-\pi/2, \pi/2] \right\}.$$

We may now evaluate the minimum of (25.13) over a grid on the reduced parameter space.

Based on Lemma 25.3, Kao et al. (2013) proposed a strategy for finding a maximin design for $Q \geq 2$. The key idea is to focus on an even smaller judiciously selected subset Θ_0 of Θ when comparing the minimal ϕ_θ -values of candidate designs. If $\Theta = \bigcup_{m=0}^M \Theta_m$, where $\Theta_m = \mathbf{G}_m \Theta_0$ for permutation matrices $\mathbf{G}_1, \dots, \mathbf{G}_M$, then it follows from Lemma 25.3 that a design d_0^* that is maximin over $\Theta_0 \times \mathcal{P}$ is also a maximin design for $\Theta \times \mathcal{P}$ if $\mathcal{R}_m = 1$ for all m , where

$$\mathcal{R}_m = \frac{\min_{\Theta_m \times \mathcal{P}} \phi_\theta(d_0^*; \theta_m, \mathbf{p})}{\min_{\Theta_0 \times \mathcal{P}} \phi_\theta(d_0^*; \theta_0, \mathbf{p})}. \tag{25.14}$$

In addition, if d_0^* and d^* are maximin designs for $\Theta_0 \times \mathcal{P}$ and $\Theta \times \mathcal{P}$, respectively, then it can easily be seen that (Kao et al. 2013)

$$\min_m \mathcal{R}_m \leq \frac{\min_{\Theta \times \mathcal{P}} \phi_\theta(d_0^*; \theta, \mathbf{p})}{\min_{\Theta \times \mathcal{P}} \phi_\theta(d^*; \theta, \mathbf{p})}.$$

Consequently, when $\min_m \mathcal{R}_m$ is not 1 but very close to 1, design d_0^* , while not maximin over $\Theta \times \mathcal{P}$, is quite efficient compared to the maximin design d^* . We note that, with Corollary 25.1, the same is true by replacing Θ_m with $\Theta_m^c = \mathbf{G}_m^c \Theta_0, m = 1, 2, \dots, M$.

Based on our experience, the grid method is convenient and efficient for evaluating the minimum ϕ_θ -value for each candidate design, especially when $\Theta_0 \times \mathcal{P}$ is relatively small. The genetic algorithm of Kao et al. (2009a) can then be adapted to search for a maximin design. In the next section, we provide an illustrative example for implementing the proposed strategy.

25.4.1 Illustrative Example for the Nonlinear Model

Suppose $(Q, L) = (2, 242)$, $\tau_{ISI} = 4$ s, and $\tau_{TR} = 2$ s. We consider model (25.11) with a second-order polynomial drift and AR1 noise with $\rho = 0.3$. The parameter space for the model parameters θ and $\mathbf{p} = (p_1, p_6)'$ is $\mathbb{R}^2 \times \mathcal{P}$, where $\mathcal{P} = [6, 9] \times [0, 2]$. The focus is on $\theta = (\theta_1, \theta_2)'$, that is, on detection, and designs are compared via a maximin criterion as in (25.13) with ϕ_θ chosen as the function corresponding to A -optimality. The genetic algorithm in Section 25.3.3 is adapted to search for a design maximizing this criterion.

We first consider an off-the-shelf optimization tool, namely, the MATLAB function *fmincon*, for approximating the minimum of (25.13) for each candidate design during the genetic algorithm search. This naive approach can take a large amount of CPU time and yields unsatisfactory results. We thus apply Lemmas 25.1 and 25.2 to reduce the parameter space to $\Theta \times \mathcal{P}$ with $\Theta = \{(\cos \varphi_1, \sin \varphi_1) \mid \varphi_1 \in (-\pi/2, \pi/2)\}$ and obtain $\min_{\Theta \times \mathcal{P}} \phi_\theta(d; \theta, \mathbf{p})$ for each design over a grid on $\Theta \times \mathcal{P}$; we use grid intervals of 0.2 and 0.1π for \mathbf{p} and φ_1 , respectively. By evaluating $\phi_\theta(d; \theta, \mathbf{p})$ at every (θ, \mathbf{p}) -value on the grid, this grid search tends to provide a better approximation of the minimum than the previously mentioned off-the-shelf optimization tool. We also note that, while we work on a reduced parameter space, the maximin design d^* is still optimal for $\mathbb{R}^2 \times \mathcal{P}$.

To further reduce the computational burden, the strategy outlined in this section is applied and the parameter space is further reduced by taking

$$G_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

and $\Theta_0 = \{(\cos \varphi_1, \sin \varphi_1) \mid \varphi_1 \in [-\pi/4, \pi/4]\}$. The genetic algorithm is adapted to search for a design d_0^* maximizing $\min_{\Theta_0 \times \mathcal{P}} \phi_\theta(d; \theta, \mathbf{p})$, where the minimal value is each time obtained over a grid on $\Theta_0 \times \mathcal{P}$ with grid intervals of 0.2 and 0.1π for \mathbf{p} and φ_1 , respectively. For this example, $\Theta = \Theta_0 \cup \Theta_1^*$ with $\Theta_1^* = \{\text{sign}(\theta_{11})\theta_1 \mid \theta_1 = (\theta_{11}, \theta_{21})' = G_1\theta_0, \theta_0 \in \Theta_0\}$; $\text{sign}(\theta_{11})$ is the sign of the first element of $\theta_1 = G_1\theta_0$. Following the results in the previous section, d_0^* is efficient if $\mathcal{R}_1^* = \min_{\Theta_1^* \times \mathcal{P}} \phi_\theta(d_0^*; \theta, \mathbf{p}) / \min_{\Theta_0 \times \mathcal{P}} \phi_\theta(d_0^*; \theta, \mathbf{p})$ is close to 1. We demonstrate that this strategy can help to reduce CPU time and achieve efficient maximin designs.

On a desktop computer with a 3.4 GHz Core i7-2600 processor, the genetic algorithm with *fmincon* took about 2.56 h to search for a maximin design. With the grid method on the reduced parameter space $\Theta \times \mathcal{P}$, the genetic algorithm used about 15.36 min to obtain d^* . While Lemmas 25.1 and 25.2 help to reduce computing time significantly, the strategy based on Lemma 25.3 and Corollary 25.1 provides an even more efficient way to obtain a design that yields performance similar to d^* . For this example, the genetic algorithm with this strategy took about 6.35 min to obtain a maximin design d_0^* for $\Theta_0 \times \mathcal{P}$. The time spent for d_0^* is slightly less than half of that for d^* ; this is mainly because the size of Θ_0 is half of that of Θ . This time reduction is greater for a larger Q since the area of Θ_0 is in general $1/Q!$ of that of Θ . While the computational burden can increase with Q , the strategy helps to reduce the CPU time by a factor of $1/Q!$; see Kao et al. (2013) for other examples and further discussions.

The performance of various designs is presented in Figure 25.4, where $\min(\phi_\theta) = \min_{\Theta_0 \times \mathcal{P}} \phi_\theta(d; \theta, \mathbf{p})$ is evaluated over $\Theta \times \mathcal{P}$ with grid intervals of 0.1 and 0.05π for \mathbf{p} and φ_1 , respectively; the same grid intervals will be used throughout this example for comparing and evaluating designs. The design obtained with *fmincon* does not perform well; it yields a maximin efficiency of 86.83% relative to d^* . On the other hand, the maximin design d_0^* for




FIGURE 25.4

The $\min(\phi_\theta)$ -values over $\Theta \times \mathcal{P}$ of some traditional designs and the maximin designs. For random designs, the mean and the standard deviation of $\min(\phi_\theta)$ of 100 randomly generated designs are presented.

$\Theta_0 \times \mathcal{P}$ is very efficient and has a similar $\min(\phi_\theta)$ -value as d^* . Although we obtain d^* in this example to help evaluate the performance of d_0^* , this is unnecessary in practice, especially when obtaining d^* is time consuming (e.g., with a large Q). To evaluate the performance of d_0^* , we calculate the previously mentioned \mathcal{R}_1 , which does not rely on d^* . The \mathcal{R}_1 -value for d_0^* is 99.32%, indicating that d_0^* is very efficient in terms of the maximin criterion (25.13). A direct comparison with d^* yields a maximin efficiency of 99.55% for d_0^* relative to d^* .

Figure 25.4 also presents the $\min(\phi_\theta)$ -values of a blocked design and an m -sequence, and the average $\min(\phi_\theta)$ -value of 100 random designs. From the figure, while blocked designs are highly recommended for detection with linear models (Section 25.3.2), they do not perform well in terms of the maximin criterion for this nonlinear model. In particular, the blocked design has a maximin efficiency of 44.88% relative to d^* . As argued in Kao et al. (2013), blocked designs are a poor choice for estimation of the HRF and are therefore limited in their performance with the nonlinear model, which treats detection and estimation in a single model. Stated in a different way, with Model (25.11), a reasonably precise estimate of p is needed in order to obtain a precise estimate of θ .

The blocked design and maximin designs are presented in Figure 25.5; different shades of gray represent different stimulus types and a white bar corresponds to a zero. In contrast to the blocked design, maximin designs d^* and d_0^* are rather random in appearance and do not seem to have perceivable patterns. An m -sequence also looks random (not shown), it achieved 84.05% of the maximin efficiency of d^* . Designs that are randomly generated attain similar maximin efficiencies as the m -sequence (Figure 25.4). However, none of these traditional designs perform as well as d_0^* or d^* .

25.5 Additional Considerations

In this section, we provide a brief overview of some other issues related to fMRI experimental designs, including violation of the assumption of additivity of overlapping HRFs,

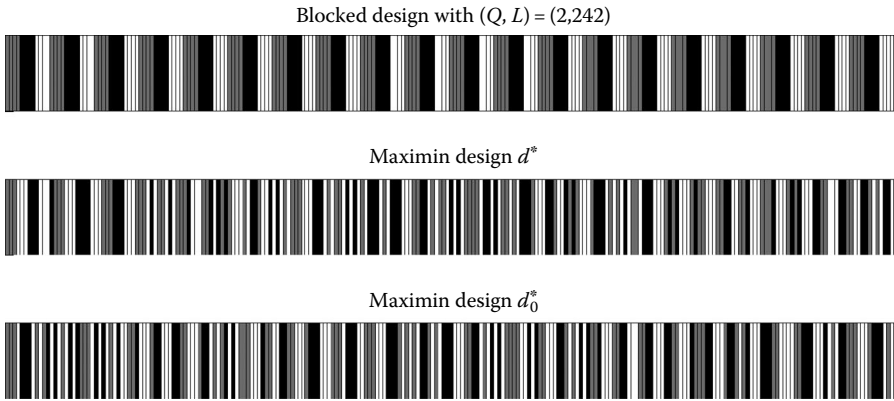


FIGURE 25.5

A blocked design for $Q = 2$ formed by repetitions of $\{000011112222\}$, a maximin design d^* for $\Theta \times \mathcal{P}$, and a maximin design d_0^* for $\Theta_0 \times \mathcal{P}$. Different shades of gray indicate different types of stimuli and white means 0.

misspecification of the correlation parameter for the noise, and use of multiple scanning sessions. Our main focus is on discussing results for the linear models introduced in Section 25.3. For a nonlinear model, the first issue, nonadditive HRFs, has also been discussed in Maus et al. (2012); this will be briefly described in the next section. While the other design issues may also be relevant for the nonlinear model, they have not been systematically studied under that model.

25.5.1 Nonadditive HRFs

The statistical models in Section 25.3 assume that the heights of overlapping HRFs accumulate additively to form a component of MR measurements. This additivity assumption is (roughly) valid in some cases (Boynton et al. 1996; Dale and Buckner 1997), but violations have been observed in other cases. For example, Wager et al. (2005) studied the effect of one single brief (250 ms) stimulus and that of multiple (2, 5, 6, 10, or 11) brief stimuli of the same type that are separated by 1 s. By comparing the MR signals evoked by a single stimulus and those by multiple stimuli, they found that the accumulated HRFs under the assumed system tend to overestimate the observed ones. This phenomenon had already been reported in Wager and Nichols (2003), and the accumulated heights of the HRFs of rapidly presented stimuli are said to *saturate* at a certain level of intensity. Such nonadditivity of the HRFs is still under active investigation; see Huettel (2012) and references therein. While not fully understood, methods for addressing this issue at the design stage have already been proposed. Although the nonadditivity issue may also exist in Model (25.5) for estimation, the proposed methods are primarily for Model (25.3) for detection.

Based on their empirical results, Wager and Nichols (2003) proposed to impose a ceiling value on the accumulated HRF heights to account for nonadditivity. With our notation in (25.3), Wager and Nichols set an upper bound of 2 for each of the elements of the vector Xh^* . Thus, any element of Xh^* that exceeds 2 is replaced by 2. This replacement is more likely to occur when times between stimulus onsets are short or the HRF shape h^* has an extended width (e.g., due to an extended neuronal activity). The method assumes a constant HRF amplitude θ_q (strength of activation) even with *saturated* accumulations of HRFs. It provides a simple, first-order adjustment for nonadditive HRFs. In the following example, we present and discuss the effect of this adjustment on design performance.

Example 25.4

We consider the same settings as in the illustrative example of Section 25.3.4, including a time to repetition τ_{TR} of 2 s. However, we take the interstimulus interval τ_{ISI} to be 1 s, corresponding to a rapid presentation of the stimuli. Consequently, ΔT is 1 s. The focus is on detection with Model (25.3), but the nonadditivity of the HRF is accounted for using the method of Wager and Nichols (2003). The HRF shape h^* is obtained from the double-gamma function of (25.4), normalized to have a maximum of 1 (Section 25.3.1). Individual stimulus effects and pairwise comparisons between stimulus types are separately investigated, and both *A*- and *D*-optimality criteria are applied. We use the genetic algorithm of Kao et al. (2009a) to search for optimal designs for detection. For comparison purposes, we also use the algorithm to find optimal designs when possible nonadditivity is ignored. The obtained designs are compared via $\phi_\theta(d)$ computed by using the model with the nonadditivity adjustment.

Table 25.3 presents relative efficiencies based on ϕ_θ -values under the model with the nonadditivity assumption. The relative efficiencies are for designs obtained when ignoring the nonadditivity relative to designs obtained when taking the nonadditivity into account. The results suggest that the former designs can suffer significant efficiency loss when the HRFs are nonadditive. We also note that while the obtained designs with additive HRFs are close to blocked designs, the designs with nonadditive HRFs contain random components. Intuitively, with additive HRFs, the MR signals (or accumulated HRFs) evoked by a block of multiple stimuli can be strong, making blocked designs efficient for detection. However, such an advantage of blocked designs can be suppressed when the MR signals are bounded above; adding stimuli to a block does not help to increase the signal intensity if the upper bound has already been attained.

We also present in Table 25.3 the performance of designs obtained with the nonadditivity assumption when the HRFs are additive. As seen from the table, using designs obtained under the assumption of nonadditive HRFs can also lead to considerable loss of efficiency when the HRFs are actually additive. It would thus be good to know at the design stage which of these assumptions applies. \square

Maus et al. (2010a) also considered the nonadditivity issue when studying the performance of blocked designs; each block is formed by brief stimuli of the same type separated by a specific time interval ($\tau_{ISI} = 1, 2$ or 3 s). With our notation, they reduced the heights of the HRF shape h^* by 19% for blocks with $\tau_{ISI} = 1$ s; no adjustments were made for blocks with $\tau_{ISI} > 1$ s. In Maus et al. (2012), where a nonlinear model is considered, this method is adapted again. Specifically, they reduced the maximal height of the HRF shape by 30%

TABLE 25.3

Relative Efficiencies (%) for Detection Under (I) Nonadditive HRFs and (II) Additive HRF of Designs Obtained with the Wrong Assumption to Designs Obtained with the Correct Assumption

Parametric Function	A-Optimality				D-Optimality			
	Q = 2	Q = 3	Q = 4	Q = 6	Q = 2	Q = 3	Q = 4	Q = 6
<i>(I) Correct assumption: Nonadditive HRFs</i>								
Individual effects	80.1	74.2	73.2	67.4	82.4	79.1	74.6	69.2
Pairwise comparisons	78.7	85.1	83.6	73.3	84.2	84.2	82.6	74.3
<i>(II) Correct assumption: Additive HRFs</i>								
Individual effects	85.2	71.2	60.5	54.4	79.6	74.6	65.3	60.5
Pairwise comparisons	83.7	88.6	68.2	48.7	72.5	84.7	65.0	49.8

for $\tau_{ISI} = 2.5$ s and by 21% for $\tau_{ISI} = 5$ s. The reductions on the HRF height are based on empirical results from previous studies. This adjustment provides an alternative way for accommodating nonadditive HRFs.

25.5.2 Uncertain Correlation Coefficient

Observations obtained from a brain voxel are often temporally correlated, and thus, correlated noise is assumed when analyzing fMRI data. The assumption of an AR1 error process seems to provide satisfactory results; see also Zhang and Yu (2008) and Worsley et al. (2002). While the autocorrelation coefficient ρ of an AR1 process can be estimated at the analysis stage, it is often unknown at the design stage. In addition, the value of ρ may vary across brain voxels (Maus et al. 2010b).

To take the uncertainty about ρ into account, Maus et al. (2010b) considered a maximin procedure for finding designs that protect against the worst case over a specified range for ρ . To emphasize the dependence of the optimality criterion in (25.10) on ρ , we write

$$\phi(d; \rho) = w_1 \phi_0^*(d; \rho) + w_2 \phi_i^*(d; \rho) + w_3 \phi_c^*(d) + w_4 \phi_f^*(d).$$

Maus et al. (2010b) studied cases with $(Q, L) = (3, 255)$ and $\tau_{ISI} = \tau_{TR} = 2$ s. They studied Models (25.3) and (25.5) with a second-order polynomial drift and AR1 noise for detection and estimation, respectively. For the optimality criterion they used third-order counterbalancing, taking $R = 3$ for ϕ_c^* in (25.8), and set $f_1 = f_2 = f_3 = 1/3$ for ϕ_f^* in (25.9). For selected weights w_i , they considered the maximin criterion

$$\min_{\rho \in [0, 0.5]} \frac{\phi(d; \rho)}{\phi(d^*_\rho; \rho)} \tag{25.15}$$

for evaluating designs. Here, d^*_ρ is a design maximizing $\phi(d; \rho)$ for a given ρ . Maus et al. (2010b) used the genetic algorithm of Kao et al. (2009a) (see Section 25.3.3) to generate 51 d^*_ρ designs for ρ ranging from 0 to 0.5 in steps of size 0.01 to help calculate (25.15). Among these 51 designs, they then selected the design that maximizes the criterion in (25.15).

We summarize their findings in Table 25.4 in which the last two columns provide the ρ -value that yielded the maximin design. We also study and report the results for the cases whose best ρ -value is not (precisely) reported by Maus et al. (2010b). The reported ρ -values yield designs with very high maximin efficiencies (>97%). Nevertheless, one could adapt the genetic algorithm again to search for (possibly) better designs by using (25.15) as the fitness function, except that the minimization would only be over the 51 ρ -values for which d^*_ρ has been obtained.

From Table 25.4, the assumption of $\rho = 0.3$ used in the previous sections does not seem unreasonable, especially for the A -optimality criterion. Additional results based on Maus et al. (2010b) and our own investigations show that the d^*_ρ design with $\rho = 0.3$ for D -optimality is also highly efficient. However, the results available are only for some particular experimental settings with $Q = 3$. Additional investigations could also consider an AR2 correlation structure for the noise.

TABLE 25.4

Maximin Designs for Uncertainty in the Correlation Coefficient ρ

Criterion	Parametric Function	ρ -Value of Maximin Design d_ρ^*	
		A-Optimality	D-Optimality
ϕ_θ (Detection)	Individual effects	$\rho = 0.30$	$\rho = 0.45$
	Pairwise comparisons	$\rho = 0.30$	$\rho = 0.32$
ϕ_h (Estimation)	Individual effects	$\rho = 0.28$	$\rho = 0.21$
	Pairwise comparisons ^a	$\rho = 0.24$	$\rho = 0.12$
$(\phi_\theta^* + \phi_h^*)/2$	Individual effects ^a	$\rho = 0.30$	$\rho = 0.36$
	Pairwise comparisons ^a	$\rho = 0.27$	$\rho = 0.24$
$(\phi_\theta^* + \phi_h^* + \phi_c^* + \phi_f^*)/4$	Individual effects ^a	$\rho = 0.31$	$\rho = 0.31$
	Pairwise comparisons ^a	$\rho = 0.29$	$\rho = 0.33$

^a Cases whose results are not (precisely) reported by Maus et al. (2010b).

25.5.3 Multiple Scanning Sessions

For experiments with a long duration, experimenters may use multiple short scanning sessions rather than a single long one (e.g., Brown et al. 2008). Each short session may last 5 min, say. The experimental subject can rest between sessions, which should help to maintain comparable performance throughout the experiment.

For design selection with multiple scanning sessions, Kao et al. (2009b) generalized Models (25.3) and (25.5) by allowing different drift or trend effects for different sessions. For detection with B scanning sessions, the model becomes

$$y = \begin{bmatrix} X^{(1)} \\ X^{(2)} \\ \vdots \\ X^{(B)} \end{bmatrix} h^* \theta + [I_B \otimes S] \begin{bmatrix} \gamma^{(1)} \\ \gamma^{(2)} \\ \vdots \\ \gamma^{(B)} \end{bmatrix} + \begin{bmatrix} \epsilon^{(1)} \\ \epsilon^{(2)} \\ \vdots \\ \epsilon^{(B)} \end{bmatrix}, \tag{25.16}$$

where, for the b th scanning session, $b = 1, \dots, B$, $X^{(b)} = [X_1^{(b)} \dots X_Q^{(b)}]$ is the 0–1 design matrix obtained from the design; $\gamma^{(b)}$ is the parameter vector for the drift/trend; $\epsilon^{(b)}$ represents the noise. The remaining terms are as in Model (25.3).

We again follow a common practice to assume the same HRF shape h^* for every stimulus type and focus only on (some parametric functions of) the possibly different unknown HRF amplitudes θ_q , $q = 1, \dots, Q$. The construction of $X_q^{(b)}$ is the same as that of X_q in Model (25.3); see Section 25.3.1. The *rest periods* between sessions are assumed to be sufficiently long (e.g., >30 s) so that the stimulus effects from a previous session do not carry over to the next session. In addition, all sessions are assumed to be of the same length, and $\epsilon^{(1)}, \dots, \epsilon^{(B)}$ are independent and have the same covariance matrix, although these assumptions are not essential at all.

Similarly, for estimation, the model accommodating multiple scanning sessions can be written as

$$y = \begin{bmatrix} X^{(1)} \\ X^{(2)} \\ \vdots \\ X^{(B)} \end{bmatrix} h + [I_B \otimes S] \begin{bmatrix} \gamma^{(1)} \\ \gamma^{(2)} \\ \vdots \\ \gamma^{(B)} \end{bmatrix} + \begin{bmatrix} \epsilon^{(1)} \\ \epsilon^{(2)} \\ \vdots \\ \epsilon^{(B)} \end{bmatrix}. \tag{25.17}$$

As in Model (25.5), the unknown HRF parameter vector $\mathbf{h} = (h'_1, \dots, h'_Q)'$ for the Q types of stimuli or some linear combinations of \mathbf{h}_q , $q = 1, \dots, Q$, is the main interest. With Models (25.16) and (25.17), the covariance matrices of the estimators $C_\theta \hat{\theta}$ and $C_h \hat{\mathbf{h}}$ are proportional to

$$\Sigma_\theta(d) = C_\theta \left[\sum_{b=1}^B (\mathbf{I}_Q \otimes \mathbf{h}^*) \mathbf{X}^{(b)'} \mathbf{V}' (\mathbf{I}_{T_B} - \mathbf{P}_{V_S}) \mathbf{V} \mathbf{X}^{(b)} (\mathbf{I}_Q \otimes \mathbf{h}^*) \right]^{-1} C'_\theta; \text{ and}$$

$$\Sigma_h(d) = C_h \left[\sum_{b=1}^B \mathbf{X}^{(b)'} \mathbf{V}' (\mathbf{I}_{T_B} - \mathbf{P}_{V_S}) \mathbf{V} \mathbf{X}^{(b)} \right]^{-1} C'_h,$$

respectively. Here, $T_B = T/B$ is the number of observations collected in each scanning session, and \mathbf{V} is the assumed whitening matrix making $\mathbf{V}\epsilon^{(b)}$ white noise for $b = 1, \dots, B$. The optimality criteria ϕ_θ and ϕ_h can then be adapted for evaluating the quality of designs in detection and estimation for experiments with multiple scanning sessions.

A design for multiple scanning sessions can be written as $d = \{d^{(1)}d^{(2)} \dots d^{(B)}\}$, where $d^{(b)}$ is a design sequence with $L_B = L/B$ elements for the b th scanning session. To achieve high-quality designs, Kao et al. (2009b) proposed to adapt the genetic algorithm of Kao et al. (2009a). In particular, one may treat the optimality criterion defined for multiple scanning sessions as the objective function and then use the algorithm to search for a design d maximizing the criterion. When applying this approach, the genetic operators (crossover and mutation) work on the entire design d by ignoring multiple sessions; the session effect is then taken into account when evaluating the quality of designs. Another approach proposed by Kao et al. (2009b) is to utilize the search algorithm to search over a restricted design class, in which $d^{(b+1)}$ is obtained by replacing the symbol q of $d^{(b)}$ with $q + 1$ and setting $Q + 1$ to 1; $q = 1, \dots, Q$. For example, if $d^{(1)} = \{01210132\}$ for $Q = 3$, then $d^{(2)} = \{02320213\}$ and $d^{(3)} = \{03130321\}$. That is, $d^{(i+1)}$ is obtained by cyclically permuting the labels of stimulus types of $d^{(i)}$, and $d^{(i)}$ is used again in the $(Q + i)$ th scanning session. As demonstrated in Kao et al. (2009b), this latter approach can reduce the computing time in obtaining efficient designs for multiple scanning sessions.

25.6 Summary and Discussion

Selecting a design for an fMRI study requires consideration of statistical efficiency and of psychological and practical constraints. Finding an optimal or near-optimal design often involves a multiobjective optimization and calls for the use of a computationally efficient, versatile algorithmic approach that can accommodate a wide spectrum of experimental settings. In this chapter, we have presented some existing results and approaches for selecting efficient fMRI designs. For detecting brain activation with a linear model, blocked designs are often recommended. For the estimation problem, m -sequences can achieve high efficiencies and are advocated by Buračas and Boynton (2002). When both detection and estimation are of interest, designs generated from blocked designs and

m -sequences may be considered; these designs include mixed designs, permuted blocked designs, or clustered m -sequences. While these traditional designs are useful in some cases, designs obtained from an efficient optimization algorithm can significantly outperform these designs in many cases. Such algorithms include the genetic algorithm of Kao et al. (2009a), which is useful for taking both statistical efficiencies and practical constraints into account. Such an algorithm can also be modified to handle extensions or variations, such as nonadditivity or multiple scanning sessions as in Section 25.5 or a nonlinear model as in Section 25.4.

Although previous research on fMRI designs provides useful tools and guidance, there are additional topics that deserve closer investigation. One such topic concerns the assumed correlation structure of the noise. Based on empirical results provided in Lenoski et al. (2008), assuming AR2 noise rather than AR1 noise tends to improve the analysis results. Approaches that allow experimenters to obtain efficient designs for such a more complicated situation will be of a great practical value. One possibility is to extend the genetic algorithm of Kao et al. (2009a) to accommodate AR2 noise. However, the computation is expected to be very expensive and can easily become infeasible with an increased size and complexity of the problem such as the design problem discussed in Section 25.4. Developing an efficient algorithm for this problem is an interesting research topic. Another topic concerns the nature of the stimuli. The previous studies mainly focus on *simple stimuli*, each consisting of only one single component. In practice, the use of compound stimuli involving two or more components (e.g., a cue followed by a mental task) is not uncommon (e.g., Pochon et al. 2001; Serences 2004; Silver et al. 2007). In such situations, each compound stimulus evokes multiple HRFs rather than a single one. When the focus is on one of the components of each stimulus, empirical results suggest that blocked designs, where compound stimuli of the same type are clustered, are not as efficient as some random designs. This is observed even when detecting brain activation is of primary interest, which is contrary to what we saw for simple stimuli. A possible reason is that, with a blocked design, the columns for the different components in the design matrix tend to have a large inner product. A novel, sophisticated approach for dealing with this situation is needed. Moreover, both for simple and compound stimuli, the stimulus (or one of the components of a compound stimulus) could be a level combination of several cognitive factors whose main effects and interactions are of interest (e.g., Friston et al. 1996). Finding efficient designs for such a factorial fMRI experiment is also an interesting research topic.

In addition, current studies about selecting fMRI designs only consider statistical models for single-voxel analysis, thereby ignoring spatial correlation across voxels. How will design selection be affected if such spatial correlation is to be taken into account? Developing tools to search for efficient designs that allow such spatial correlation is an interesting research direction. Moreover, while detecting activation and estimating the HRF are common objectives, other study objectives, such as the assessment of connectivity or interaction between brain regions, are becoming increasingly popular (Lindquist 2008). Obtaining efficient designs in the context of such newly emerging objectives is an open research problem.

So far, the research on fMRI designs mainly focuses on the development of efficient computational methods. While computer-generated designs are crucially important, theoretical and analytical results that provide insights and guidance for the fMRI design selection problem can also be of great value. Not much work has been done in this direction.

Due to the complexity and usefulness of fMRI studies, research on fMRI designs is arguably an important and useful area. It will continue to offer challenging design

questions, as also indicated by the following quote from Lindquist (2008): “As research hypotheses ultimately become more complicated, the need for more advanced experimental designs will only increase further and this is clearly an area where statisticians can play an important role.”

Acknowledgments

The research for this chapter of Ming-Hung Kao was in part supported by the National Science Foundation grant DMS-13-52213 and that of John Stufken by the National Science Foundation grant DMS-10-07507 and DMS-14-06760.

References

- Ahn, C. W. (2006), *Advances in Evolutionary Algorithms: Theory, Design and Practice*, Studies in computational intelligence, Berlin, Germany: Springer.
- Blamire, A. M. (2012), The Yale experience in first advancing fMRI, *NeuroImage*, 62, 637–640.
- Boynton, G. M., Engel, S. A., Glover, G. H., and Heeger, D. J. (1996), Linear systems analysis of functional magnetic resonance imaging in human V1, *Journal of Neuroscience*, 16, 4207–4221.
- Brendel, B., Hertrich, I., Erb, M., Lindner, A., Riecker, A., Grodd, W., and Ackermann, H. (2010), The contribution of mesiofrontal cortex to the preparation and execution of repetitive syllable productions: An fMRI study, *Neuroimage*, 50, 1219–1230.
- Brown, M. R. G., Vilis, T., and Everling, S. (2008), Isolation of saccade inhibition processes: Rapid event-related fMRI of saccades and nogo trials, *NeuroImage*, 39, 793–804.
- Buračas, G. T. and Boynton, G. M. (2002), Efficient design of event-related fMRI experiments using m-sequences, *NeuroImage*, 16, 801–813.
- Buxton, R. B., Liu, T. T., Martinez, A., Frank, L. R., Luh, W. M., and Wong, E. C. (2000), Sorting out event-related paradigms in fMRI: The distinction between detecting an activation and estimating the hemodynamic response, *NeuroImage*, 11, S457.
- Chernoff, H. (1953), Locally optimal designs for estimating parameters, *Annals of Mathematical Statistics*, 24, 586–602.
- Churchill, N. W., Yourganov, G., Spring, R., Rasmussen, P. M., Lee, W., Ween, J. E., and Strother, S. C. (2012), PHYCAA: Data-driven measurement and removal of physiological noise in BOLD fMRI, *Neuroimage*, 59, 1299–1314.
- Culham, J. C. (2006), Functional neuroimaging: Experimental design and analysis, in *Handbook of Functional Neuroimaging of Cognition*, eds. Cabeza, R. and Kingstone, A., Cambridge, MA: MIT Press, 2nd ed., pp. 53–82.
- Dale, A. M. (1999), Optimal experimental design for event-related fMRI, *Human Brain Mapping*, 8, 109–114.
- Dale, A. M. and Buckner, R. L. (1997), Selective averaging of rapidly presented individual trials using fMRI, *Human Brain Mapping*, 5, 329–340.
- Fedorov, V. V. and Hackl, P. (1997), *Model-Oriented Design of Experiments*, New York: Springer.
- Friston, K. J., Fletcher, P., Josephs, O., Holmes, A., Rugg, M. D., and Turner, R. (1998), Event-related fMRI: Characterizing differential responses, *NeuroImage*, 7, 30–40.
- Friston, K. J., Price, C. J., Fletcher, P., Moore, C., Frackowiak, R. S. J., and Dolan, R. J. (1996), The trouble with cognitive subtraction, *NeuroImage*, 4, 97–104.
- Friston, K. J., Zarahn, E., Josephs, O., Henson, R. N. A., and Dale, A. M. (1999), Stochastic designs in event-related fMRI, *NeuroImage*, 10, 607–619.

- Godfrey, K. (1993), *Perturbation Signals for System Identification*, New York: Prentice Hall.
- Grinband, J., Wager, T. D., Lindquist, M., Ferrera, V. P., and Hirsch, J. (2008), Detection of time-varying signals in event-related fMRI designs, *Neuroimage*, 43, 509–520.
- Handwerker, D. A., Ollinger, J. M., and D'Esposito, M. (2004), Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses, *Neuroimage*, 21, 1639–1651.
- Henson, R. N. A. (2007), Efficient experimental design for fMRI, in *Statistical Parametric Mapping: The Analysis of Functional Brain Images*, eds. Friston, K. J., Ashburner, J. T., Kiebel, S. J., Nichols, T. E., and D., P. W., London, U.K.: Academic, 1st ed., pp. 193–210.
- Holland, J. H. (1975), *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*, Ann Arbor, MI: University of Michigan Press.
- Huettel, S. A. (2012), Event-related fMRI in cognition, *Neuroimage*, 62, 1152–1156.
- Josephs, O. and Henson, R. N. A. (1999), Event-related functional magnetic resonance imaging: Modelling, inference and optimization, *Philosophical Transactions of the Royal Society of London Series B-Biological Sciences*, 354, 1215–1228.
- Josephs, O., Turner, R., and Friston, K. (1997), Event-related fMRI, *Human Brain Mapping*, 5, 243–248.
- Kao, M.-H. (2009a), Multi-objective optimal experimental Designs for ER-fMRI using Matlab, *Journal of Statistical Software*, 30, 1–13.
- Kao, M.-H. (2009b), Optimal experimental designs for event-related functional magnetic resonance imaging, PhD thesis, University of Georgia, Athens, GA.
- Kao, M.-H., Majumdar, D., Mandal, A., and Stufken, J. (2013), Maximin and maximin-efficient event-related fMRI designs under a nonlinear model, *Annals of Applied Statistics*, 7, 1940–1959.
- Kao, M.-H., Mandal, A., and Stufken, J. (2008), Optimal design for event-related functional magnetic resonance imaging considering both individual stimulus effects and pairwise contrasts. *Statistics and Applications*, 6, 235–256.
- Kao, M.-H., Mandal, A., Lazar, N., and Stufken, J. (2009a), Multi-objective optimal experimental designs for event-related fMRI studies, *NeuroImage*, 44, 849–856.
- Kao, M.-H., Mandal, A., and Stufken, J. (2009b), Efficient designs for event-related functional magnetic resonance imaging with multiple scanning sessions, *Communications in Statistics-Theory and Methods*, 38, 3170–3182.
- Kao, M.-H., Mandal, A., and Stufken, J. (2012), Constrained multi-objective designs for functional MRI experiments via a modified nondominated sorting genetic algorithm, *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 61, 515–534.
- Lazar, N. A. (2008), *The Statistical Analysis of Functional MRI Data*, Statistics for Biology and Health, New York: Springer.
- Lenoski, B., Baxter, L. C., Karam, L. J., Maisog, J., and Debbins, J. (2008), On the performance of autocorrelation estimation algorithms for fMRI analysis, *IEEE Journal of Selected Topics in Signal Processing*, 2, 828–838.
- Lindquist, M. A. (2008), The statistical analysis of fMRI data, *Statistical Science*, 23, 439–464.
- Lindquist, M. A., Loh, J. M., Atlas, L. Y., and Wager, T. D. (2009), Modeling the hemodynamic response function in fMRI: Efficiency, bias and mis-modeling, *NeuroImage*, 45, S187–S198.
- Lindquist, M. A. and Wager, T. D. (2007), Validity and power in hemodynamic response modeling: A comparison study and a new approach, *Human Brain Mapping*, 28, 764–784.
- Liu, T. T. (2004), Efficiency, power, and entropy in event-related fMRI with multiple trial types: Part II: Design of experiments, *NeuroImage*, 21, 401–413.
- Liu, T. T. and Frank, L. R. (2004), Efficiency, power, and entropy in event-related fMRI with multiple trial types: Part I: Theory, *NeuroImage*, 21, 387–400.
- Liu, T. T., Frank, L. R., Wong, E. C., and Buxton, R. B. (2001), Detection power, estimation efficiency, and predictability in event-related fMRI, *NeuroImage*, 13, 759–773.
- Loh, J. M., Lindquist, M. A., and Wager, T. D. (2008), Residual analysis for detecting mis-modeling in fMRI, *Statistica Sinica*, 18, 1421–1448.

- Lu, Y., Bagshaw, A. P., Grova, C., Kobayashi, E., Dubeau, F., and Gotman, J. (2006), Using voxel-specific hemodynamic response function in Eeg-fMRI data analysis, *NeuroImage*, 32, 238–247.
- Lund, T. E., Madsen, K. H., Sidaros, K., Luo, W. L., and Nichols, T. E. (2006), Non-white noise in fMRI: Does modelling have an impact? *Neuroimage*, 29, 54–66.
- MacWilliams, F. J. and Sloane, N. J. A. (1977), *The Theory of Error Correcting Codes*, Amsterdam, the Netherlands: Elsevier/North-Holland.
- Maus, B., van Breukelen, G. J. P., Goebel, R., and Berger, M. P. F. (2010a), Optimization of blocked designs in fMRI studies, *Psychometrika*, 75, 373–390.
- Maus, B., van Breukelen, G. J. P., Goebel, R., and Berger, M. P. F. (2010b), Robustness of optimal design of fMRI experiments with application of a genetic algorithm, *Neuroimage*, 49, 2433–2443.
- Maus, B., van Breukelen, G. J. P., Goebel, R., and Berger, M. P. F. (2011), Optimal design of multi-subject blocked fMRI experiments, *NeuroImage*, 56, 1338–1352.
- Maus, B., van Breukelen, G. J. P., Goebel, R., and Berger, M. P. F. (2012), Optimal design for nonlinear estimation of the hemodynamic response function, *Human Brain Mapping*, 33, 1253–1267.
- Miezin, F. M., Maccotta, L., Ollinger, J. M., Petersen, S. E., and Buckner, R. L. (2000), Characterizing the hemodynamic response: Effects of presentation rate, sampling procedure, and the possibility of ordering brain activity based on relative timing, *NeuroImage*, 11, 735–759.
- Pochon, J. B., Levy, R., Poline, J. B., Crozier, S., Lehericy, S., Pillon, B., Deweer, B., Le Bihan, D., and Dubois, B. (2001), The role of dorsolateral prefrontal cortex in the preparation of forthcoming actions: An fMRI study, *Cerebral Cortex*, 11, 260–266.
- Poldrack, R. A., Mumford, J. A., and Nichols, T. E. (2011), *Handbook of functional MRI Data Analysis*, New York: Cambridge University Press.
- Raichle, M. E. (2006), Functional neuroimaging: A historical and physiological perspective, in *Handbook of Functional Neuroimaging of Cognition*, eds. Cabeza, R. and Kingstone, A., Cambridge, MA: MIT Press, 2nd ed., pp. 1–20.
- Rosen, B. R., Buckner, R. L., and Dale, A. M. (1998), Event-related functional MRI: Past, present, and future, *PNAS*, 95, 773–780.
- Sarkka, S., Solin, A., Nummenmaa, A., Vehtari, A., Auranen, T., Vanni, S., and Lin, F. H. (2012), Dynamic retrospective filtering of physiological noise in BOLD fMRI: Drifter, *Neuroimage*, 60, 1517–1527.
- Serences, J. T. (2004), A comparison of methods for characterizing the event-related BOLD time series in rapid fMRI, *NeuroImage*, 21, 1690–1700.
- Silver, M. A., Ress, D., and Heeger, D. J. (2007), Neural correlates of sustained spatial attention in human early visual cortex, *Journal of Neurophysiology*, 97, 229–237.
- Smith, A. M., Lewis, B. K., Ruttimann, U. E., Ye, F. Q., Sinnwell, T. M., Yang, Y., Duyn, J. H., and Frank, J. A. (1999), Investigation of low frequency drift in fMRI signal, *NeuroImage*, 9, 526–533.
- Song, A. W., Huettel, S. A., and McCarthy, G. (2006), Functional neuroimaging: Basic principles of functional MRI, in *Handbook of Functional Neuroimaging of Cognition*, eds. Cabeza, R. and Kingstone, A., Cambridge, MA: MIT Press, 2nd ed., pp. 21–52.
- The Wellcome Trust Centre for Neuroimaging (2003), SPM—Statistical Parametric Mapping. The Wellcome Trust Centre for Neuroimaging at the University College London, London, U.K. <http://www.fil.ion.ucl.ac.uk/spm/>.
- Wager, T. D. and Nichols, T. E. (2003), Optimization of experimental design in fMRI: A general framework using a genetic algorithm, *NeuroImage*, 18, 293–309.
- Wager, T. D., Vazquez, A., Hernandez, L., and Noll, D. C. (2005), Accounting for nonlinear BOLD effects in fMRI: Parameter estimates and a model for prediction in rapid event-related studies, *NeuroImage*, 25, 206–218.
- Worsley, K. J. and Friston, K. J. (1995), Analysis of fMRI time-series revisited—Again, *NeuroImage*, 2, 173–181.
- Worsley, K. J., Liao, C. H., Aston, J., Petre, V., Duncan, G. H., Morales, F., and Evans, A. C. (2002), A general statistical analysis for fMRI data, *NeuroImage*, 15, 1–15.
- Zhang, C. M. and Yu, T. (2008), Semiparametric detection of significant activation for brain fMRI, *Annals of Statistics*, 36, 1693–1725.

Chapman & Hall/CRC

Handbooks of Modern Statistical Methods

Handbook of Design and Analysis of Experiments provides a detailed overview of the tools required for the optimal design of experiments and their analyses. The handbook gives a unified treatment of a wide range of topics, covering the latest developments.

This carefully edited collection of 25 chapters in seven parts synthesizes the state of the art in the theory and applications of designed experiments and their analyses. Written by leading researchers in the field, the chapters offer a balanced blend of methodology and applications.

The first part presents a historical look at experimental design and the fundamental theory of parameter estimation in linear models. The second part deals with settings such as response surfaces and block designs in which the response is modeled by a linear model, the third part covers designs with multiple factors (both treatment and blocking factors), and the fourth part presents optimal designs for generalized linear models, other nonlinear models, and spatial models. The fifth part addresses issues involved in designing various computer experiments. The sixth part explores “cross-cutting” issues relevant to all experimental designs, including robustness and algorithms. The final part illustrates the application of experimental design in recently developed areas.

Features

- Covers many recent advances in the field, including designs for nonlinear models, the design and analysis of computer experiments, and algorithms applicable to a wide variety of design problems
- Describes block designs, crossover and repeated measurement designs, designs for estimating response surfaces, optimal designs for linear models, fractional factorial designs, multifactor designs, spatial models, and more
- Explores the extensive use of experimental designs in marketing, the pharmaceutical industry, engineering, medicine, and other areas
- Provides further reading suggestions on related topics



CRC Press

Taylor & Francis Group
an informa business

www.crcpress.com

6000 Broken Sound Parkway, NW
Suite 300, Boca Raton, FL 33487

711 Third Avenue
New York, NY 10017

2 Park Square, Milton Park
Abingdon, Oxon OX14 4RN, UK

