

THIRD EDITION

# Applied Mixed Models in Medicine



HELEN BROWN  
ROBIN PRESCOTT

STATISTICS IN PRACTICE

WILEY



# Applied Mixed Models in Medicine

# STATISTICS IN PRACTICE

*Series Advisors*

## **Human and Biological Sciences**

Stephen Senn

*CRP-Santé, Luxembourg*

## **Earth and Environmental Sciences**

Marian Scott

*University of Glasgow, UK*

## **Industry, Commerce and Finance**

Wolfgang Jank

*University of Maryland, USA*

## **Founding Editor**

Vic Barnett

*Nottingham Trent University, UK*

---

Statistics in Practice is an important international series of texts which provide detailed coverage of statistical concepts, methods and worked case studies in specific fields of investigation and study.

With sound motivation and many worked practical examples, the books show in down-to-earth terms how to select and use an appropriate range of statistical techniques in a particular practical field within each title's special topic area.

The books provide statistical support for professionals and research workers across a range of employment fields and research environments. Subject areas covered include medicine and pharmaceuticals; industry, finance and commerce; public services; the earth and environmental sciences, and so on.

The books also provide support to students studying statistical courses applied to the above areas. The demand for graduates to be equipped for the work environment has led to such courses becoming increasingly prevalent at universities and colleges.

It is our aim to present judiciously chosen and well-written workbooks to meet everyday practical needs. Feedback of views from readers will be most valuable to monitor the success of this aim.

A complete list of titles in this series appears at the end of the volume.

# Applied Mixed Models in Medicine

Third Edition

**Helen Brown**

*The Roslin Institute*

*University of Edinburgh, UK*

**Robin Prescott**

*Centre for Population Health Sciences*

*University of Edinburgh, UK*

**WILEY**

This edition first published 2015  
© 2015 John Wiley & Sons, Ltd

*Registered office*

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ,  
United Kingdom

For details of our global editorial offices, for customer services and for information about how to apply for permission to reuse the copyright material in this book please see our website at [www.wiley.com](http://www.wiley.com).

The right of the author to be identified as the author of this work has been asserted in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by the UK Copyright, Designs and Patents Act 1988, without the prior permission of the publisher.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The publisher is not associated with any product or vendor mentioned in this book.

**Limit of Liability/Disclaimer of Warranty:** While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. It is sold on the understanding that the publisher is not engaged in rendering professional services and neither the publisher nor the author shall be liable for damages arising herefrom. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

*Library of Congress Cataloging-in-Publication Data*

Brown, Helen, 1962- , author.

Applied mixed models in medicine / Helen Brown, Robin Prescott. – Third edition.  
p. : cm.

Includes bibliographical references and index.

ISBN 978-1-118-77825-8 (cloth)

I. Prescott, Robin, author. II. Title.

[DNLN: 1. Medicine. 2. Statistics as Topic—methods. 3. Models, Statistical. WA 950]

R853.S7

610.72'7—dc23

2014026811

A catalogue record for this book is available from the British Library.

ISBN: 9781118778258

Set in 10/12pt Photina by Laserwords Private Limited, Chennai, India

1 2015

# Contents

<b>Preface to third edition</b>	<b>xiii</b>
<b>Mixed models notation</b>	<b>xvii</b>
<b>About the Companion Website</b>	<b>xix</b>
1 Introduction	1
1.1 The use of mixed models	1
1.2 Introductory example	3
1.2.1 Simple model to assess the effects of treatment (Model A)	4
1.2.2 A model taking patient effects into account (Model B)	6
1.2.3 Random effects model (Model C)	7
1.2.4 Estimation (or prediction) of random effects	12
1.3 A multi-centre hypertension trial	12
1.3.1 Modelling the data	13
1.3.2 Including a baseline covariate (Model B)	15
1.3.3 Modelling centre effects (Model C)	16
1.3.4 Including centre-by-treatment interaction effects (Model D)	17
1.3.5 Modelling centre and centre-treatment effects as random (Model E)	18
1.4 Repeated measures data	19
1.4.1 Covariance pattern models	19
1.4.2 Random coefficients models	21
1.5 More about mixed models	22
1.5.1 What is a mixed model?	22
1.5.2 Why use mixed models?	23
1.5.3 Communicating results	24
1.5.4 Mixed models in medicine	25
1.5.5 Mixed models in perspective	25

1.6	Some useful definitions	27
1.6.1	Containment	28
1.6.2	Balance	29
1.6.3	Error strata	31
2	Normal mixed models	34
2.1	Model definition	34
2.1.1	The fixed effects model	35
2.1.2	The mixed model	37
2.1.3	The random effects model covariance structure	39
2.1.4	The random coefficients model covariance structure	42
2.1.5	The covariance pattern model covariance structure	44
2.2	Model fitting methods	46
2.2.1	The likelihood function and approaches to its maximisation	47
2.2.2	Estimation of fixed effects	50
2.2.3	Estimation (or prediction) of random effects and coefficients	51
2.2.4	Estimation of variance parameters	53
2.3	The Bayesian approach	57
2.3.1	Introduction	58
2.3.2	Determining the posterior density	59
2.3.3	Parameter estimation, probability intervals and $p$ -values	60
2.3.4	Specifying non-informative prior distributions	62
2.3.5	Evaluating the posterior distribution	67
2.4	Practical application and interpretation	69
2.4.1	Negative variance components	69
2.4.2	Accuracy of variance parameters	73
2.4.3	Bias in fixed and random effects standard errors	74
2.4.4	Significance testing	75
2.4.5	Confidence intervals	78
2.4.6	Checking model assumptions	79
2.4.7	Missing data	81
2.4.8	Determining whether the simulated posterior distribution has converged	83
2.5	Example	84
2.5.1	Analysis models	84
2.5.2	Results	86
2.5.3	Discussion of points from Section 2.4	88
3	Generalised linear mixed models	113
3.1	Generalised linear models	114
3.1.1	Introduction	114



3.1.2	Distributions	115
3.1.3	The general form for exponential distributions	117
3.1.4	The GLM definition	118
3.1.5	Fitting the GLM	121
3.1.6	Expressing individual distributions in the general exponential form	123
3.1.7	Conditional logistic regression	125
3.2	Generalised linear mixed models	125
3.2.1	The GLMM definition	126
3.2.2	The likelihood and quasi-likelihood functions	127
3.2.3	Fitting the GLMM	129
3.3	Practical application and interpretation	133
3.3.1	Specifying binary data	133
3.3.2	Uniform effects categories	133
3.3.3	Negative variance components	134
3.3.4	Presentation of fixed and random effects estimates	135
3.3.5	Accuracy of variance parameters and potential bias	136
3.3.6	Bias in fixed and random effects standard errors	137
3.3.7	The dispersion parameter	139
3.3.8	Significance testing	140
3.3.9	Confidence intervals	141
3.3.10	Model checking	142
3.3.11	Determining whether the simulated posterior distribution has converged	142
3.4	Example	143
3.4.1	Introduction and models fitted	143
3.4.2	Results	144
3.4.3	Discussion of points from Section 3.3	148
4	Mixed models for categorical data	168
4.1	Ordinal logistic regression (fixed effects model)	168
4.2	Mixed ordinal logistic regression	173
4.2.1	Definition of the mixed ordinal logistic regression model	173
4.2.2	Residual variance matrix	174
4.2.3	Likelihood and quasi-likelihood functions	176
4.2.4	Model fitting methods	177
4.3	Mixed models for unordered categorical data	177
4.3.1	The G matrix	180
4.3.2	The R matrix	180
4.3.3	Fitting the model	180
4.4	Practical application and interpretation	180
4.4.1	Expressing fixed and random effects results	181
4.4.2	The proportional odds assumption	181

4.4.3	Number of covariance parameters	181
4.4.4	Choosing a covariance pattern	182
4.4.5	Interpreting covariance parameters	182
4.4.6	Checking model assumptions	182
4.4.7	The dispersion parameter	182
4.4.8	Other points	183
4.5	Example	183
5	Multi-centre trials and meta-analyses	197
5.1	Introduction to multi-centre trials	197
5.1.1	What is a multi-centre trial?	197
5.1.2	Why use mixed models to analyse multi-centre data?	198
5.2	The implications of using different analysis models	198
5.2.1	Centre and centre-treatment effects fixed	198
5.2.2	Centre effects fixed, centre-treatment effects omitted	199
5.2.3	Centre and centre-treatment effects random	200
5.2.4	Centre effects random, centre-treatment effects omitted	201
5.3	Example: a multi-centre trial	202
5.4	Practical application and interpretation	209
5.4.1	Plausibility of a centre-treatment interaction	209
5.4.2	Generalisation	210
5.4.3	Number of centres	210
5.4.4	Centre size	210
5.4.5	Negative variance components	211
5.4.6	Balance	211
5.5	Sample size estimation	211
5.5.1	Normal data	212
5.5.2	Binary data	215
5.5.3	Categorical data	217
5.6	Meta-analysis	217
5.7	Example: meta-analysis	218
5.7.1	Analyses	219
5.7.2	Results	220
5.7.3	Treatment estimates in individual trials	221
6	Repeated measures data	231
6.1	Introduction	231
6.1.1	Reasons for repeated measurements	231
6.1.2	Analysis objectives	232
6.1.3	Fixed effects approaches	232
6.1.4	Mixed models approaches	234
6.2	Covariance pattern models	234
6.2.1	Covariance patterns	235

6.2.2	Choice of covariance pattern	239
6.2.3	Choice of fixed effects	241
6.2.4	General points	242
6.3	Example: covariance pattern models for normal data	244
6.3.1	Analysis models	245
6.3.2	Selection of covariance pattern	245
6.3.3	Assessing fixed effects	247
6.3.4	Model checking	248
6.4	Example: covariance pattern models for count data	254
6.4.1	Analysis models	255
6.4.2	Analysis using a categorical mixed model	258
6.5	Random coefficients models	262
6.5.1	Introduction	262
6.5.2	General points	264
6.5.3	Comparisons with fixed effects approaches	266
6.6	Examples of random coefficients models	267
6.6.1	A linear random coefficients model	267
6.6.2	A polynomial random coefficients model	270
6.7	Sample size estimation	286
6.7.1	Normal data	286
6.7.2	Binary data	288
6.7.3	Categorical data	288
7	Cross-over trials	289
7.1	Introduction	289
7.2	Advantages of mixed models in cross-over trials	290
7.3	The AB/BA cross-over trial	290
7.3.1	Example: AB/BA cross-over design	292
7.4	Higher order complete block designs	297
7.4.1	Inclusion of carry-over effects	297
7.4.2	Example: four-period, four-treatment cross-over trial	298
7.5	Incomplete block designs	302
7.5.1	Example: Three treatment two-period cross-over trial	303
7.6	Optimal designs	305
7.6.1	Example: Balaam's design	305
7.7	Covariance pattern models	307
7.7.1	Structured by period	308
7.7.2	Structured by treatment	308
7.7.3	Example: four-way cross-over trial	308
7.8	Analysis of binary data	317
7.9	Analysis of categorical data	321
7.10	Use of results from random effects models in trial design	325
7.10.1	Example	325
7.11	General points	326

8	Other applications of mixed models	329
8.1	Trials with repeated measurements within visits	329
8.1.1	Covariance pattern models	330
8.1.2	Example	335
8.1.3	Random coefficients models	341
8.1.4	Example: random coefficients models	343
8.2	Multi-centre trials with repeated measurements	349
8.2.1	Example: multi-centre hypertension trial	349
8.2.2	Covariance pattern models	350
8.3	Multi-centre cross-over trials	355
8.4	Hierarchical multi-centre trials and meta-analysis	356
8.5	Matched case-control studies	357
8.5.1	Example	358
8.5.2	Analysis of a quantitative variable	359
8.5.3	Check of model assumptions	360
8.5.4	Analysis of binary Variables	362
8.6	Different variances for treatment groups in a simple between-patient trial	370
8.6.1	Example	371
8.7	Estimating variance components in an animal physiology trial	374
8.7.1	Sample size estimation for a future experiment	375
8.8	Inter- and intra-observer variation in foetal scan measurements	380
8.9	Components of variation and mean estimates in a cardiology experiment	381
8.10	Cluster sample surveys	384
8.10.1	Example: cluster sample survey	384
8.11	Small area mortality estimates	386
8.12	Estimating surgeon performance	389
8.13	Event history analysis	391
8.13.1	Example	392
8.14	A laboratory study using a within-subject $4 \times 4$ factorial design	394
8.15	Bioequivalence studies with replicate cross-over designs	397
8.15.1	Example	402
8.16	Cluster randomised trials	411
8.16.1	Example: A trial to evaluate integrated care pathways for treatment of children with asthma in hospital	412
8.16.2	Example: Edinburgh randomised trial of breast screening	414
8.17	Analysis of bilateral data	418

8.18	Incomplete block designs	425
8.18.1	Introduction	425
8.18.2	Balanced incomplete block (BIB) designs	426
8.18.3	Studies where only comparisons to a particular treatment are of interest	429
8.18.4	Designs to produce lower variances for a specific treatment pair	433
8.18.5	Design to produce lower variances for more than one treatment pair	438
9	Software for fitting mixed models	452
9.1	Packages for fitting mixed models	452
9.2	PROC MIXED	454
9.3	Using SAS to fit mixed models to non-normal data	476
9.3.1	PROC GLIMMIX	477
9.3.2	PROC GENMOD	480
9.4	PROC MCMC	483
	<b>Glossary</b>	<b>489</b>
	<b>References</b>	<b>492</b>
	<b>Index</b>	<b>497</b>
	<b>SIP Series Titles</b>	<b>513</b>



# Preface to third edition

Analysis of variance and regression has for many years been the mainstay of statistical modelling. These techniques usually have as a basic assumption that the residual or error terms are independently and identically distributed. Mixed models are an important approach to modelling, which allows us to relax the independence assumption and take into account more complicated data structures in a flexible way. Sometimes, this interdependence of observations is modelled directly in a mixed model. For example, if a number of repeated measurements are made on a patient, then mixed models allow us to specify a pattern for the correlation between these measurements. In other contexts, such as the cross-over clinical trial, specifying that patient effects are normally distributed, rather than fixed as in the classical approach, induces observations on the same patient to be correlated.

There are many benefits to be gained from using mixed models. In some situations, the benefit will be an increase in the precision of our estimates. In others, we will be able to make wider inferences. We will sometimes be able to use a more appropriate model that will give us greater insight into what underpins the structure of the data. However, it is only the availability of software in versatile packages such as SAS<sup>®</sup> that has made these techniques widely accessible. It is now important that suitable information on their use becomes available so that they may be applied confidently on a routine basis.

Our intention in this book is to put all types of mixed models into a general framework and to consider the practical implications of their use. We aim to do this at a level that can be understood by applied statisticians and numerate scientists. Greatest emphasis is placed on skills required for the application of mixed models and interpretation of the results. An in-depth understanding of the mathematical theory underlying mixed models is not essential to gain these, but an awareness of the practical consequences of fitting different types of mixed models is necessary. While many publications are available on various aspects of mixed models, these generally relate to specific types of model and often differ in their use of terminology. Such publications are not always readily comprehensible

to the applied statisticians who will be the most frequent users of the methods. An objective of this book is to help overcome this deficit.

Examples given will primarily relate to the medical field. However, the general concepts of mixed models apply equally to many other areas of application, for example, social sciences, agriculture, veterinary science and official statistics. (In the social sciences, mixed models are often referred to as ‘multi-level’ models.) Data are becoming easier to collect, with the consequence that datasets are now often large and complex. We believe that mixed models provide useful tools for modelling the complex structures that occur in such data.

The third edition of this book retains the structure of the first two, but there are further changes to reflect the continued evolution of SAS. This edition fully incorporates features of SAS up to version 9.3. Compared to what was available at the time of the previous edition, enhancements to SAS include improved graphical facilities. Importantly, there is also a new procedure, PROC MCMC, which facilitates Bayesian analysis. This has led to extensive changes in our coverage of Bayesian methods. SAS 9.3 and later versions now provide output both in text format from the output window and, additionally, as an HTML file in the results viewer. There have been accompanying minor changes in the details of outputs and graphs, such as labelling. Our approach to reporting SAS outputs in this edition has been to change our presentation from earlier editions only when we wish to highlight features that have changed substantially and, importantly, to facilitate the reader’s use of mixed models, whatever their version of SAS.

During the drafting of this edition, SAS 9.4 became available. It is not fully incorporated into this book because its new features are focused more on the SAS high performance procedures than on improvements to the SAS/STAT procedures. These high performance procedures ‘provide predictive modelling tools that have been specially developed to take advantage of parallel processing in both multithread single-machine mode and distributed multi-machine mode’. Typically, the high performance procedures such as PROC HPLMIXED have a greatly reduced range of options compared to PROC MIXED and, consequently, are peripheral to the aims of this book. We do, however, consider some of the small modifications to improve procedures such as GLIMMIX and MCMC that are available in SAS/STAT® 12.1 and later versions.

Chapter 1 provides an introduction to the capabilities of mixed models, defines general concepts and gives their basic statistical properties. Chapter 2 defines models and fitting methods for normally distributed data. Chapter 3 first introduces generalised linear models that can be used for the analysis of data that are binomial or Poisson or from any other member of the exponential family of distributions. These methods are then extended to incorporate mixed models concepts under the heading of generalised linear mixed models. The fourth chapter examines how mixed models can be applied when the variable to be analysed is categorical. The main emphasis in these chapters, and indeed in the whole book, is on classical statistical approaches to inference, based on significance tests and confidence intervals. However, the Bayesian approach is



also introduced in Chapter 2, since it has several potential advantages and its use is becoming more widespread. Although the overall emphasis of the book is on the application of mixed models techniques, these chapters can also be used as a reference guide to the underlying theory of mixed models.

Chapters 5–7 consider the practical implications of using mixed models for particular designs. Each design illustrates a different feature of mixed models.

Multi-centre trials and meta-analyses are considered in Chapter 5. These are examples of hierarchical data structures, and the use of a mixed model allows for any additional variation in treatment effects occurring between centres (or trials) and hence makes results more generalisable. The methods shown can be applied equally to any type of hierarchical data.

In Chapter 6, the uses of covariance pattern models and random coefficients models are described using the repeated measures design. These approaches take into account the correlated nature of the repeated observations and give more appropriate treatment effect estimates and standard errors. The material in this chapter will apply equally to any situation where repeated observations are made on the same units.

Chapter 7 considers cross-over designs where each patient may receive several treatments. In this design, more accurate treatment estimates are often achieved by fitting patient effects as random. This improvement in efficiency can occur for any dataset where a fixed effect is ‘crossed’ with a random effect.

In Chapter 8, a variety of other designs and data structures is considered. These either incorporate several of the design aspects covered in Chapters 5–7 or have structures that have arisen in a more unplanned manner. They help to illustrate the broad scope of application of mixed models. This chapter includes two new sections. We have added a section on the analysis of bilateral data, a common structure in some areas of medical research, but one that we had not previously addressed. There is also a substantial new section on incomplete block designs.

Chapter 9 gives information on software available for fitting mixed models. Most of the analyses in the book are carried out using PROC MIXED in SAS, supplemented by PROC GENMOD, PROC GLIMMIX, and PROC MCMC. This chapter introduces the basic syntax for these procedures. This information should be sufficient for fitting most of the analyses described, but the full SAS documentation should be referenced for those who wish to use more complex features. The SAS code used for most of the examples is supplied within the text. In addition, the example datasets and SAS code may be obtained electronically from [www.wiley.com/go/brown/applied\\_mixed](http://www.wiley.com/go/brown/applied_mixed).

This book has been written to provide the reader with a thorough understanding of the concepts of mixed models, and we trust it will serve well for this purpose. However, readers wishing to take a shortcut to the fitting of normal mixed models should read Chapter 1 for an introduction, Section 2.4 for practical details, and the chapter relevant to their design. To fit non-normal or categorical mixed models, Section 3.3 or Section 4.4 should be read in addition to Section 2.4. In an attempt to make this book easier to use, we have presented at the beginning of

the text a summary of the notation we have used, while at the end, we list some key definitions in a glossary.

Our writing of this book has been aided in many ways. The first edition evolved from a constantly changing set of course notes that accompanied a 3-day course on the subject, run regularly over the previous 6 years. The second edition was helped by many individuals who were kind enough to comment on the first edition, including the identification of some errors that had slipped in, and by further participants at our courses who have contributed to discussions and have thereby helped to shape our views. This process has continued with the third edition. We are also grateful to many other colleagues who have read and commented on various sections of the manuscript and especially to our colleagues who have allowed us to use their data. We hope that readers will find the resulting book a useful reference in an interesting and expanding area of statistics.

***Helen Brown***  
***Robin Prescott***  
*Edinburgh*

# Mixed models notation

The notation below is provided for quick reference. Models are defined more fully in Sections 2.1, 3.1 and 4.1.

## Normal mixed model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\alpha} + \mathbf{Z}\boldsymbol{\beta} + \mathbf{e},$$

$$\boldsymbol{\beta} \sim N(\mathbf{0}, \mathbf{G}),$$

$$\text{var}(\mathbf{e}) = \mathbf{R},$$

$$\text{var}(\mathbf{y}) = \mathbf{V} = \mathbf{Z}\mathbf{G}\mathbf{Z}' + \mathbf{R}.$$

## Generalised linear mixed model

$$\mathbf{y} = \boldsymbol{\mu} + \mathbf{e},$$

$$g(\boldsymbol{\mu}) = \mathbf{X}\boldsymbol{\alpha} + \mathbf{Z}\boldsymbol{\beta},$$

$$\boldsymbol{\beta} \sim N(\mathbf{0}, \mathbf{G}),$$

$$\text{var}(\mathbf{e}) = \mathbf{R},$$

$$\text{var}(\mathbf{y}) = \mathbf{V} = \text{var}(\boldsymbol{\mu}) + \mathbf{R},$$

$$\approx \mathbf{B}\mathbf{Z}\mathbf{G}\mathbf{Z}'\mathbf{B} + \mathbf{R} \quad (\text{a first-order approximation}),$$

where

$\mathbf{y}$  = dependent variable,

$\mathbf{e}$  = residual error,

$\mathbf{X}$  = design matrix for fixed effects,

$\mathbf{Z}$  = design matrix for random effects,

$\boldsymbol{\alpha}$  = fixed effects parameters,

$\boldsymbol{\beta}$  = random effects parameters,

$\mathbf{R}$  = residual variance matrix,

$\mathbf{G}$  = matrix of covariance parameters,

$\mathbf{V} = \text{var}(\mathbf{y})$  variance matrix,

$\boldsymbol{\mu}$  = expected values,

$g$  = link function,

$\mathbf{B}$  = diagonal matrix of variance terms (e.g.  $\mathbf{B} = \text{diag}\{\mu_i(1 - \mu_i)\}$  for binary data).

### Ordered categorical mixed model

$$\mathbf{y} = \boldsymbol{\mu} + \mathbf{e},$$

$$\text{logit}(\boldsymbol{\mu}^{[c]}) = \mathbf{X}\boldsymbol{\alpha} + \mathbf{Z}\boldsymbol{\beta},$$

$$\boldsymbol{\beta} \sim N(\mathbf{0}, \mathbf{G}),$$

$\text{var}(\mathbf{y})$  is defined as in the GLMM,

where

$$\boldsymbol{\mu} = (\mu_{11}, \mu_{12}, \mu_{13}, \mu_{21}, \mu_{22}, \mu_{23}, \dots, \mu_{n1}, \mu_{n2}, \mu_{n3})',$$

$\mu_{ij}$  = probability observation  $i$  is in category  $j$ ,

$$\boldsymbol{\mu}^{[c]} = (\mu_{11}^{[c]}, \mu_{12}^{[c]}, \mu_{13}^{[c]}, \mu_{21}^{[c]}, \mu_{22}^{[c]}, \mu_{23}^{[c]}, \dots, \mu_{n1}^{[c]}, \mu_{n2}^{[c]}, \mu_{n3}^{[c]})',$$

$$\mu_{ij}^{[c]} = \text{probability } (y_i \leq j) = \sum_{k=1}^j \mu_{ik}.$$

# About the Companion Website

This book is accompanied by a companion website:

**[www.wiley.com/go/brown/applied\\_mixed](http://www.wiley.com/go/brown/applied_mixed)**

This website includes SAS codes and datasets for most of the examples. In the future, updates and further materials may be added.



# 1

## Introduction

At the start of each chapter, we will ‘set the scene’ by outlining its content. In this introductory chapter, we start Section 1.1 by describing some situations where a mixed models analysis will be particularly helpful. In Section 1.2, we describe a simplified example and use it to illustrate the idea of a statistical model. We then introduce and compare fixed effects and random effects models. In the next section, we consider a more complex ‘real-life’ multi-centre trial and look at some of the variety of models that could be fitted (Section 1.3). This example will be used for several illustrative examples throughout the book. In Section 1.4, the use of mixed models to analyse a series of observations (repeated measures) is considered. Section 1.5 broadens the discussion on mixed models and looks at mixed models with a historical perspective of their use. In Section 1.6, we introduce some technical concepts: containment, balance and error strata.

We will assume in our presentation that the reader is already familiar with some of the basic statistical concepts as found in elementary statistical textbooks.

### 1.1 The use of mixed models

In the course of this book, we will encounter many situations in which a mixed models approach has advantages over the conventional type of analysis, which would be accessible via introductory texts on statistical analysis. Some of them are introduced in outline in this chapter and will be dealt in detail later on.

*Example 1: Utilisation of incomplete information in a cross-over trial* Cross-over trials are often utilised to assess treatment efficacy in chronic conditions, such as asthma. In such conditions, an individual patient can be tested for response to a succession of two or more treatments, giving the benefit of a ‘within-patient’ comparison. In the most commonly used cross-over design, just two treatments

are compared. If, for generality, we call these treatments A and B, then patients will be assessed either on their response to treatment A, followed by their response to treatment B, or vice versa. If all patients complete the trial, and both treatments are assessed, then the analysis is fairly straightforward. However, commonly, patients drop out during the trial and may have a valid observation from only the first treatment period. These incomplete observations cannot be utilised in a conventional analysis. In contrast, the use of a mixed model will allow all of the observations to be analysed, resulting in more accurate comparisons of the efficacy of treatment. This benefit, of more efficient use of the data, applies to all types of cross-over trial where there are missing data.

*Example 2: Cross-over trials with fewer treatment periods than treatments* In cross-over trials, for logistical reasons, it may be impractical to ask a patient to evaluate more than two treatments (e.g. if the treatment has to be given for several weeks). Nevertheless, there may be the need to evaluate three or more treatments. Special types of cross-over design can be used in this situation, but a simple analysis will be very inefficient. Mixed models provide a straightforward method of analysis, which fully uses the data, resulting again in more precise estimates of the effect of the treatments.

*Example 3: A surgical audit* A surgical audit is to be carried out to investigate how different hospitals compare in their rates of postoperative complications following a particular operation. As some hospitals carry out the operation commonly, while other hospitals perform the operation rarely, the accuracy with which the complication rates are estimated will vary considerably from hospital to hospital. Consequently, if the hospitals are ordered according to their complication rates, some may appear to be outliers compared with other hospitals, purely due to chance variation. When mixed models are used to analyse data of this type, the estimates of the complication rates are adjusted to allow for the number of operations, and rates based on small numbers become less extreme.

*Example 4: Analysis of a multi-centre trial* Many clinical trials are organised on a multi-centre basis, usually because there is an inadequate number of suitable patients in any single centre. The analysis of multi-centre trials often ignores the centres from which the data were obtained, making the implicit assumption that all centres are identical to one another. This assumption may sometimes be dangerously misleading. For example, a multi-centre trial comparing two surgical treatments for a condition could be expected to show major differences between centres. There could be two types of differences. First, the centres may differ in the overall success, averaged over the two surgical treatments. More importantly, there may be substantial differences in the relative benefit of the two treatments across different centres. Surgeons who have had more experience with one operation (A) may produce better outcomes with A, while surgeons with more experience with the alternative operation (B) may obtain better results



with B. Mixed models can provide an insightful analysis of such a trial by allowing for the extent to which treatment effects differ from centre to centre. Even when the difference between treatments can be assumed to be identical in all centres, a mixed model can improve the precision of the treatment estimates by taking appropriate account of the centres in the analysis.

*Example 5: Repeated measurements over time* In a clinical trial, the response to treatment is often assessed as a series of observations over time. For example, in a trial to assess the effect of a drug in reducing blood pressure, measurements might be taken at two, four, six and eight weeks after starting treatment. The analysis will usually be complicated by a number of patients failing to appear for some assessments or withdrawing from the study before it is complete. This complication can cause considerable difficulty in a conventional analysis. A mixed models analysis of such a study does not require complete data from all subjects. This results in more appropriate estimates of the effect of treatment and their standard errors (SEs). The mixed model also gives great flexibility in analysis, in that it can allow for a wide variety of ways in which the successive observations are correlated with one another.

## 1.2 Introductory example

We consider a very simple cross-over trial using artificial data. In this trial, each patient receives each of treatments A and B for a fixed period. At the end of each treatment period, a measurement is taken to assess the response to that treatment. In the analysis of such a trial, we commonly refer to treatments being *crossed* with patients, meaning that the categories of ‘treatments’ occur in combination with those of ‘patients’. For the purpose of this illustration, we will suppose that the response to each treatment is unaffected by whether it is received in the first or second period. The table shows the results from the six patients in this trial.

Patient	Treatment		Difference A – B	Patient mean
	A	B		
1	20	12	8	16.0
2	26	24	2	25.0
3	16	17	–1	16.5
4	29	21	8	25.0
5	22	21	1	21.5
6	24	17	7	20.5
Mean	22.83	18.67	4.17	20.75

### 1.2.1 Simple model to assess the effects of treatment (Model A)

We introduce in this section a very simple example of a statistical model using this data. A model can be thought of as an attempt to describe quantitatively the effect of a number of factors on each observation. Any model we describe is likely to be a gross oversimplification of reality. In developing models, we are seeking ones which are as simple as possible but which contain enough truth to ask questions of interest. In this first simple model, we will deliberately be oversimplistic in order to introduce our notation. We just describe the effect of the two treatments. The model may be expressed as

$$y_{ij} = \mu + t_j + e_{ij},$$

where

- $j = A$  or  $B$ ,
- $y_{ij}$  = observation for treatment  $j$  on the  $i$ th patient,
- $\mu$  = overall mean,
- $t_j$  = effect of treatment  $j$ ,
- $e_{ij}$  = error for treatment  $j$  on the  $i$ th patient.

The constant  $\mu$  represents the overall mean of the observations.  $\mu + t_A$  corresponds to the mean in the treatment group A, while  $\mu + t_B$  corresponds to the mean in the treatment group B. The constants  $\mu$ ,  $t_A$  and  $t_B$  can thus be estimated from the data. In our example, we can estimate the value of  $\mu$  to be 20.75, the overall mean value. From the mean value in the first treatment group, we can estimate  $\mu + t_A$  as 22.83, and hence our estimate of  $t_A$  is  $22.83 - 20.75 = 2.08$ . Similarly, from the mean of the second treatment group, we estimate  $t_B$  as  $-2.08$ . The term  $t_j$  can therefore be thought of as a measure of the relative effect that treatment  $j$  has had on our outcome variable.

The error term,  $e_{ij}$ , or *residual* is what remains for each patient in each period when  $\mu + t_j$  is deducted from their observed measurement. This represents random variation about the mean value for each treatment. As such, the residuals can be regarded as the result of drawing random samples from a distribution. We will assume that the distribution is Gaussian or normal, with standard deviation  $\sigma$ , and that the samples drawn from the distribution are independent of each other. The mean of the distribution can be taken as zero, since any other value would simply cause a corresponding change in the value of  $\mu$ . Thus, we will write this as

$$e_{ij} \sim N(0, \sigma^2),$$

where  $\sigma^2$  is the variance of the residuals. In practice, checks should be made to determine whether this assumption of normally distributed residuals is reasonable. Suitable checking methods will be considered in Section 2.4.6. As individual

observations are modelled as the sum of  $\mu + t_j$ , which are both constants, and the residual term, it follows that the variance of individual observations equals the residual variance:

$$\text{var}(y_{ij}) = \sigma^2.$$

The covariance of any two separate observations  $y_{ij}$  and  $y_{i'j'}$  can be written as

$$\begin{aligned} \text{cov}(y_{ij}, y_{i'j'}) &= \text{cov}(\mu + t_i + e_{ij}, \mu + t_{i'} + e_{i'j'}) \\ &= \text{cov}(e_{ij}, e_{i'j'}) \text{ (since other terms are constants).} \end{aligned}$$

Since all the residuals are assumed independent (i.e. uncorrelated), it follows that

$$\text{cov}(y_{ij}, y_{i'j'}) = 0.$$

The residual variance,  $\sigma^2$ , can be estimated using a standard technique known as analysis of variance (ANOVA). The essence of the method is that the total variation in the data is decomposed into components that are associated with possible causes of this variation, for example, that one treatment may be associated with higher observations, with the other being associated with lower observations. For this first model, using this technique, we obtain the following ANOVA table:

<b>Source of variation</b>	<b>Degrees of freedom</b>	<b>Sums of squares</b>	<b>Mean square</b>	<b>F</b>	<b>p</b>
Treatments	1	52.08	52.08	2.68	0.13
Residual	10	194.17	19.42		

*Note:*  $F$  = value for the  $F$  test (ratio of mean square for treatments to mean square for residual).

$p$  = significance level corresponding to the  $F$  test.

The residual mean square of 19.42 is our estimate of the residual variance,  $\sigma^2$ , for this model. The key question often arising from this type of study is as follows: ‘do the treatment effects differ significantly from each other?’ This can be assessed by the  $F$  test, which assesses the null hypothesis of no mean difference between the treatments (the larger the treatment difference, the larger the treatment mean square and the higher the value of  $F$ ). The  $p$  value of 0.13 is greater than the conventionally used cutoff point for statistical significance of 0.05. Therefore, we cannot conclude that the treatment effects are significantly different. The difference between the treatment effects and the SE of this difference provides a measure of the size of the treatment difference and the accuracy with which it is estimated:

$$\text{difference} = t_A - t_B = 2.08 + 2.08 = 4.16.$$

The SE of the difference is given by the formula

$$\begin{aligned} \text{SE}(t_A - t_B) &= \sqrt{\sigma^2(1/n_A + 1/n_B)} \\ &= \sqrt{(2 \times \sigma^2/6)} = \sqrt{6.47} = 2.54. \end{aligned}$$

Note that a  $t$  test can also be constructed from this difference and SE, giving  $t = 4.16/2.54 = 1.63$ . This is the square root of our  $F$  statistic of 2.68 and gives an identical  $t$  test  $p$  value of 0.13.

**1.2.2 A model taking patient effects into account (Model B)**

Model A as discussed previously did not utilise the fact that pairs of observations were taken on the same patients. It is possible, and indeed likely, that some patients will tend to have systematically higher measurements than others, and we may be able to improve the model by making allowance for this. This can be done by additionally including patient effects into the model:

$$y_{ij} = \mu + p_i + t_j + e_{ij},$$

where  $p_i$  are constants representing the  $i$ th patient effect.

The ANOVA table arising from this model is as follows:

Source of variation	Degrees of freedom	Sums of squares	Mean square	$F$	$p$
Patients	5	154.75	30.95	3.93	0.08
Treatments	1	52.08	52.08	6.61	0.05
Residual	5	39.42	7.88		

The estimate of the residual variance,  $\sigma^2$ , is now 7.88. It is lower than in Model A because it represents the ‘within-patient’ variation, as we have taken account of patient effects. The  $F$  test  $p$  value of 0.05 indicates that the treatment effects are now significantly different. The difference between the treatment effects is the same as in Model A, 4.16, but its SE is now as follows:

$$\text{SE}(t_A - t_B) = \sqrt{(2 \times \sigma^2/6)} = \sqrt{2.63} = 1.62.$$

(Note that the SE of the treatment difference could alternatively have been obtained directly from the differences in patient observations.)

Model B is perhaps the ‘obvious’ one to think of for this dataset. However, even in this simple case, by comparison with Model A we can see that the statistical

modeller has some flexibility in his/her choice of model. In most situations, there is no single 'correct' model, and, in fact, models are rarely completely adequate. The job of the statistical modeller is to choose that model which most closely achieves the objectives of the study.

### 1.2.3 Random effects model (Model C)

In the Models A and B, the only assumption we made about variation was that the residuals were normally distributed. We did not assume that patient or treatment effects arose from a distribution. They were assumed to take constant values. These models can be described as *fixed effects models*, and all effects fitted within them are *fixed effects*.

An alternative approach available to us is to assume that some of the terms in the model, instead of taking constant values, are realisations of values from a probability distribution. If we assumed that patient effects also arose from independent samples from a normal distribution, then the model could be expressed as

$$\begin{aligned}y_{ij} &= \mu + p_i + t_j + e_{ij}, \\e_{ij} &\sim N(0, \sigma^2) \\p_i &\sim N(0, \sigma_p^2).\end{aligned}$$

The  $p_i$  are now referred to as *random effects*. Such models, which contain a mixture of fixed and random effects, provide an example of a *mixed model*. In this book, we will meet several different types of mixed model, and we describe in Section 1.5 the common feature that distinguishes them from fixed effects models. To distinguish the class of models we have just met from those we will meet later, we will refer to this type of model as a *random effects model*.

Each random effect in the model gives rise to a *variance component*. This is a model parameter that quantifies random variation due to that effect only. In this model, the patient variance component is  $\sigma_p^2$ . We can describe variation at this level (between patients) as occurring within the patient *error stratum* (see Section 1.6 for a full description of the error stratum). This random variation occurs in addition to the residual variation (the residual variance can also be defined as a variance component.)

Defining the model in this way causes some differences in its statistical properties compared with the fixed effects model met earlier.

The variance of individual observations in a random effects model is the sum of all the variance components. Thus,

$$\text{var}(y_{ij}) = \sigma_p^2 + \sigma^2.$$

This contrasts with the fixed effects models where we had

$$\text{var}(y_{ij}) = \sigma^2.$$

The effect on the covariance of pairs of observations in the random effects model is interesting and perhaps surprising. Since  $y_{ij} = \mu + p_i + t_j + e_{ij}$ , we can write

$$\begin{aligned} \text{cov}(y_{ij}, y_{i'j'}) &= \text{cov}(\mu + p_i + t_j + e_{ij}, \mu + p_{i'} + t_{j'} + e_{i'j'}) \\ &= \text{cov}(p_i + e_{ij}, p_{i'} + e_{i'j'}). \end{aligned}$$

When observations from different patients are being considered (i.e.  $i \neq i'$ ), because of the independence of the observations,  $\text{cov}(y_{ij}, y_{i'j'}) = 0$ . However, when two samples from the same patient are considered (i.e.  $i = i'$ ), then

$$\begin{aligned} \text{cov}(y_{ij}, y_{i'j'}) &= \text{cov}(p_i + e_{ij}, p_i + e_{ij'}) \\ &= \text{cov}(p_i, p_i) = \sigma_p^2. \end{aligned}$$

Thus, observations on the same patient are correlated and have covariance equal to the patient variance component, while observations on different patients are uncorrelated. This contrasts with the fixed effects models where the covariance of any pair of observations is zero.

The ANOVA table for the random effects model is identical to that for the fixed effects model. However, we can now use it to calculate the patient variance component using results from the statistical theory that underpins the ANOVA method. The theory shows the expected values for each of the mean square terms in the ANOVA table, in terms of  $\sigma^2$ ,  $\sigma_p^2$  and the treatment effects. These are tabulated in the following table. We can now equate the expected value for the mean squares expressed in terms of the variance components to the observed values of the mean squares to obtain estimates of  $\sigma^2$  and  $\sigma_p^2$ .

Source of variation	Degrees of freedom	Sums of squares	Mean square	E(MS)
Patients	5	154.75	30.95	$2\sigma_p^2 + \sigma^2$
Treatments	1	52.08	52.08	$\sigma^2 + 6\Sigma t_i^2$
Residual	5	39.42	7.88	$\sigma^2$

Note: E(MS) = expected mean square.

Thus, from the residual line in the ANOVA table,  $\hat{\sigma}^2 = 7.88$ . In addition, by subtracting the third line of the table from the first we have:

$$2\hat{\sigma}_p^2 = (30.95 - 7.88), \quad \text{and} \quad \hat{\sigma}_p^2 = 11.54.$$

(We are introducing the notation  $\hat{\sigma}_p^2$  to denote that this is an estimate of the unknown  $\sigma_p^2$ , and  $\hat{\sigma}^2$  is an estimate of  $\sigma^2$ .)

In this example, we obtain identical treatment effect results to those for the fixed effects model (Model B). This occurs because we are, in effect, only using within-patient information to estimate the treatment effect (since all information on treatment occurs in the within-patient residual error stratum). Again, we obtain the treatment difference as  $-4.16$  with a SE of  $1.62$ . Thus, in this case, it makes no difference at all to our conclusions about treatments whether we fit patient effects as fixed or random. However, had any of the values in the dataset been missing, this would not have been the case. We now consider this situation.

**Dataset with missing values**

We will now consider analysing the dataset with two of the observations set to missing.

<b>Patient</b>	<b>Treatment</b>		<b>Difference A – B</b>	<b>Patient mean</b>
	<b>A</b>	<b>B</b>		
1	20	12	8	16.0
2	26	24	2	25.0
3	16	17	-1	16.5
4	29	21	8	25.0
5	22	–	–	22.0
6	–	17	–	17.0
Mean			4.25	

As shown previously, there are two ways we can analyse the data. We can base our analysis on a model where the patient effects are regarded as fixed (Model B) or can regard patient effects as random (Model C).

**The fixed effects model** For this analysis, we apply ANOVA in the standard way, and the result of that analysis is summarised as follows:

<b>Source of variation</b>	<b>Degrees of freedom</b>	<b>Sums of squares</b>	<b>Mean square</b>	<b>F</b>	<b>p</b>
Patients	5	167.90	33.58	3.32	0.18
Treatments	1	36.13	36.13	3.57	0.16
Residual	3	30.38	10.12		

In the fitting of Model B, it is interesting to look at the contribution that the data from patient 5 are making to the analysis. The value of 22 gives us information that will allow us to estimate the level in that patient, but it tells us nothing at all