

# ***Bayesian Approaches to Clinical Trials and Health-Care Evaluation***

**David J. Spiegelhalter**

MRC Biostatistics Unit, Cambridge, UK

**Keith R. Abrams**

University of Leicester, UK

**Jonathan P. Myles**

Cancer Research UK, London, UK



John Wiley & Sons, Ltd



***Bayesian Approaches to Clinical Trials  
and Health-Care Evaluation***

## STATISTICS IN PRACTICE

*Advisory Editor*

**Stephen Senn**

University of Glasgow, UK

*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.

# ***Bayesian Approaches to Clinical Trials and Health-Care Evaluation***

**David J. Spiegelhalter**

MRC Biostatistics Unit, Cambridge, UK

**Keith R. Abrams**

University of Leicester, UK

**Jonathan P. Myles**

Cancer Research UK, London, UK



John Wiley & Sons, Ltd

The work is based on an original NHS Health Technology Assessment funded project (93/50/05). Adapted with kind permission of the National Coordinating Centre for Health Technology Assessment.

Copyright © 2004 John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ, England

Telephone (+44) 1243 779777

Email (for orders and customer service enquiries): [cs-books@wiley.co.uk](mailto:cs-books@wiley.co.uk)  
Visit our Home Page on [www.wileyeurope.com](http://www.wileyeurope.com) or [www.wiley.com](http://www.wiley.com)

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, scanning or otherwise, except under the terms of the Copyright, Designs and Patents Act 1988 or under the terms of a licence issued by the Copyright Licensing Agency Ltd, 90 Tottenham Court Road, London W1T 4LP, UK, without the permission in writing of the Publisher. Requests to the Publisher should be addressed to the Permissions Department, John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ, England, or emailed to [permreq@wiley.co.uk](mailto:permreq@wiley.co.uk), or faxed to (44) 1243 770620.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the Publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

### ***Other Wiley Editorial Offices***

John Wiley & Sons Inc., 111 River Street, Hoboken, NJ 07030, USA

Jossey-Bass, 989 Market Street, San Francisco, CA 94103-1741, USA

Wiley-VCH Verlag GmbH, Boschstr. 12, D-69469 Weinheim, Germany

John Wiley & Sons Australia Ltd, 33 Park Road, Milton, Queensland 4064, Australia

John Wiley & Sons (Asia) Pte Ltd, 2 Clementi Loop #02-01, Jin Xing Distripark, Singapore 129809

John Wiley & Sons Canada Ltd, 22 Worcester Road, Etobicoke, Ontario, Canada M9W 1L1

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

### ***British Library Cataloguing in Publication Data***

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

ISBN 0-471-49975-7

Typeset in 10/12 pt Photina from L<sup>A</sup>T<sub>E</sub>X files supplied by the author, processed by Kolam Information Services Pvt. Ltd, Pondicherry, India.

Printed and bound in Great Britain by Antony Rowe Ltd, Chippenham, Wiltshire.

This book is printed on acid-free paper responsibly manufactured from sustainable forestry in which at least two trees are planted for each one used for paper production.

# Contents

Preface	xi
List of examples	xiii
<b>1 Introduction</b>	<b>1</b>
1.1 What are Bayesian methods?	1
1.2 What do we mean by ‘health-care evaluation’?	2
1.3 A Bayesian approach to evaluation	3
1.4 The aim of this book and the intended audience	3
1.5 Structure of the book	4
<b>2 Basic Concepts from Traditional Statistical Analysis</b>	<b>9</b>
2.1 Probability	10
2.1.1 What is probability?	10
2.1.2 Odds and log-odds	12
2.1.3 Bayes theorem for simple events	13
2.2 Random variables, parameters and likelihood	14
2.2.1 Random variables and their distributions	14
2.2.2 Expectation, variance, covariance and correlation	16
2.2.3 Parametric distributions and conditional independence	17
2.2.4 Likelihoods	18
2.3 The normal distribution	20
2.4 Normal likelihoods	22
2.4.1 Normal approximations for binary data	23
2.4.2 Normal likelihoods for survival data	27
2.4.3 Normal likelihoods for count responses	30
2.4.4 Normal likelihoods for continuous responses	31
2.5 Classical inference	31
2.6 A catalogue of useful distributions*	34
2.6.1 Binomial and Bernoulli	34
2.6.2 Poisson	35
2.6.3 Beta	36
2.6.4 Uniform	38
2.6.5 Gamma	39
2.6.6 Root-inverse-gamma	40
2.6.7 Half-normal	41

2.6.8	Log-normal	42
2.6.9	Student's $t$	43
2.6.10	Bivariate normal	44
2.7	Key points	46
	Exercises	46
<b>3</b>	<b>An Overview of the Bayesian Approach</b>	<b>49</b>
3.1	Subjectivity and context	49
3.2	Bayes theorem for two hypotheses	51
3.3	Comparing simple hypotheses: likelihood ratios and Bayes factors	54
3.4	Exchangeability and parametric modelling*	56
3.5	Bayes theorem for general quantities	57
3.6	Bayesian analysis with binary data	57
3.6.1	Binary data with a discrete prior distribution	58
3.6.2	Conjugate analysis for binary data	59
3.7	Bayesian analysis with normal distributions	62
3.8	Point estimation, interval estimation and interval hypotheses	64
3.9	The prior distribution	73
3.10	How to use Bayes theorem to interpret trial results	74
3.11	The 'credibility' of significant trial results*	75
3.12	Sequential use of Bayes theorem*	79
3.13	Predictions	80
3.13.1	Predictions in the Bayesian framework	80
3.13.2	Predictions for binary data*	81
3.13.3	Predictions for normal data	83
3.14	Decision-making	85
3.15	Design	90
3.16	Use of historical data	90
3.17	Multiplicity, exchangeability and hierarchical models	91
3.18	Dealing with nuisance parameters*	100
3.18.1	Alternative methods for eliminating nuisance parameters*	100
3.18.2	Profile likelihood in a hierarchical model*	102
3.19	Computational issues	102
3.19.1	Monte Carlo methods	103
3.19.2	Markov chain Monte Carlo methods	105
3.19.3	WinBUGS	107
3.20	Schools of Bayesians	112
3.21	A Bayesian checklist	113
3.22	Further reading	115
3.23	Key points	116
	Exercises	117
<b>4</b>	<b>Comparison of Alternative Approaches to Inference</b>	<b>121</b>
4.1	A structure for alternative approaches	121
4.2	Conventional statistical methods used in health-care evaluation	122
4.3	The likelihood principle, sequential analysis and types of error	124
4.3.1	The likelihood principle	124
4.3.2	Sequential analysis	126
4.3.3	Type I and Type II error	127
4.4	$P$ -values and Bayes factors*	127



4.4.1	Criticism of $P$ -values	127
4.4.2	Bayes factors as an alternative to $P$ -values: simple hypotheses	128
4.4.3	Bayes factors as an alternative to $P$ -values: composite hypotheses	130
4.4.4	Bayes factors in preference studies	133
4.4.5	Lindley's paradox	135
4.5	Key points	136
	Exercises	136

## **5 Prior Distributions** **139**

5.1	Introduction	139
5.2	Elicitation of opinion: a brief review	140
5.2.1	Background to elicitation	140
5.2.2	Elicitation techniques	141
5.2.3	Elicitation from multiple experts	142
5.3	Critique of prior elicitation	147
5.4	Summary of external evidence*	148
5.5	Default priors	157
5.5.1	'Non-informative' or 'reference' priors:	157
5.5.2	'Sceptical' priors	158
5.5.3	'Enthusiastic' priors	160
5.5.4	Priors with a point mass at the null hypothesis ( <i>'lump-and-smear' priors</i> )*	161
5.6	Sensitivity analysis and 'robust' priors	165
5.7	Hierarchical priors	167
5.7.1	The judgement of exchangeability	167
5.7.2	The form for the random-effects distribution	168
5.7.3	The prior for the standard deviation of the random effects*	168
5.8	Empirical criticism of priors	174
5.9	Key points	176
	Exercises	177

## **6 Randomised Controlled Trials** **181**

6.1	Introduction	181
6.2	Use of a loss function: is a clinical trial for inference or decision?	182
6.3	Specification of null hypotheses	184
6.4	Ethics and randomisation: a brief review	187
6.4.1	Is randomisation necessary?	187
6.4.2	When is it ethical to randomise?	187
6.5	Sample size of non-sequential trials	189
6.5.1	Alternative approaches to sample-size assessment	189
6.5.2	'Classical power': hybrid classical-Bayesian methods assuming normality	193
6.5.3	'Bayesian power'	194
6.5.4	Adjusting formulae for different hypotheses	196
6.5.5	Predictive distribution of power and necessary sample size	201
6.6	Monitoring of sequential trials	202
6.6.1	Introduction	202
6.6.2	Monitoring using the posterior distribution	204
6.6.3	Monitoring using predictions: 'interim power'	211
6.6.4	Monitoring using a formal loss function	220

6.6.5	Frequentist properties of sequential Bayesian methods	221
6.6.6	Bayesian methods and data monitoring committees	222
6.7	The role of 'scepticism' in confirmatory studies	224
6.8	Multiplicity in randomised trials	227
6.8.1	Subset analysis	227
6.8.2	Multi-centre analysis	227
6.8.3	Cluster randomisation	227
6.8.4	Multiple endpoints and treatments	228
6.9	Using historical controls*	228
6.10	Data-dependent allocation	235
6.11	Trial designs other than two parallel groups	237
6.12	Other aspects of drug development	242
6.13	Further reading	244
6.14	Key points	245
	Exercises	247
<b>7</b>	<b>Observational Studies</b>	<b>251</b>
7.1	Introduction	251
7.2	Alternative study designs	252
7.3	Explicit modelling of biases	253
7.4	Institutional comparisons	258
7.5	Key points	262
	Exercises	263
<b>8</b>	<b>Evidence Synthesis</b>	<b>267</b>
8.1	Introduction	267
8.2	'Standard' meta-analysis	268
8.2.1	A Bayesian perspective	268
8.2.2	Some delicate issues in Bayesian meta-analysis	274
8.2.3	The relationship between treatment effect and underlying risk	278
8.3	Indirect comparison studies	282
8.4	Generalised evidence synthesis	285
8.5	Further reading	298
8.6	Key points	299
	Exercises	299
<b>9</b>	<b>Cost-effectiveness, Policy-Making and Regulation</b>	<b>305</b>
9.1	Introduction	305
9.2	Contexts	306
9.3	'Standard' cost-effectiveness analysis without uncertainty	308
9.4	'Two-stage' and integrated approaches to uncertainty in cost-effectiveness modelling	310
9.5	Probabilistic analysis of sensitivity to uncertainty about parameters: two-stage approach	312
9.6	Cost-effectiveness analyses of a single study: integrated approach	315
9.7	Levels of uncertainty in cost-effectiveness models	320
9.8	Complex cost-effectiveness models	322
9.8.1	Discrete-time, discrete-state Markov models	322
9.8.2	Micro-simulation in cost-effectiveness models	323

9.8.3	Micro-simulation and probabilistic sensitivity analysis	324
9.8.4	Comprehensive decision modelling	328
9.9	Simultaneous evidence synthesis and complex cost-effectiveness modelling	329
9.9.1	Generalised meta-analysis of evidence	329
9.9.2	Comparison of integrated Bayesian and two-stage approach	335
9.10	Cost-effectiveness of carrying out research: payback models	335
9.10.1	Research planning in the public sector	335
9.10.2	Research planning in the pharmaceutical industry	336
9.10.3	Value of information	337
9.11	Decision theory in cost-effectiveness analysis, regulation and policy	341
9.12	Regulation and health policy	343
9.12.1	The regulatory context	343
9.12.2	Regulation of pharmaceuticals	343
9.12.3	Regulation of medical devices	344
9.13	Conclusions	344
9.14	Key points	345
	Exercises	345

## **10 Conclusions and Implications for Future Research 349**

10.1	Introduction	349
10.2	General advantages and problems of a Bayesian approach	349
10.3	Future research and development	350

## **A Websites and Software 353**

A.1	The site for this book	353
A.2	Bayesian methods in health-care evaluation	353
A.3	Bayesian software	354
A.4	General Bayesian sites	355

References	357
------------	-----

Index	381
-------	-----



# *Preface*

This book began life as a review of Bayesian methods in health technology assessment commissioned by the UK National Health Service Research and Development Programme, which appeared as Spiegelhalter *et al.* (2000). It was then thought to be a good idea to change the review into a basic introduction to Bayesian methods which also tried to cover the field of clinical trials and health-care evaluation. We did not realise the amount of work this would involve.

We are very grateful to all those who have read all or part of the manuscript and given such generous comments, particularly David Jones, Laurence Freedman, Mahesh Parmar, Tony Ades, Julian Higgins, Nicola Cooper, Cosetta Minelli, Alex Sutton and Denise Kendrick. Unfortunately, by tradition, we must take full responsibility for all errors and idiosyncrasies. Our particular thanks go to Daniel Farewell for writing the BANDY program, and Nick Freemantle for providing data. The University of Leicester provided the second author with study leave, during which part of this work was carried out. Finally, we must thank Rob Calver and Siân Jones at Wiley for being so patient with the repeated excuses for delay: in the words of Douglas Adams (1952–2001), “I love deadlines. I especially like the whooshing sound they make as they go flying by”. We hope it has been worth the wait.



# List of Examples

<b>Example 2.1</b>	<i>Dice: Illustration of rules of probability</i>	11
<b>Example 2.2</b>	<i>Prognosis: Marginalisation and extending the conversation</i>	12
<b>Example 2.3</b>	<i>Prognosis (continued): Bayes theorem for single events</i>	13
<b>Example 2.4</b>	<i>Response: Combining Bernoulli likelihoods</i>	19
<b>Example 2.5</b>	<i>GREAT: Normal likelihood from a <math>2 \times 2</math> table</i>	26
<b>Example 2.6</b>	<i>Power: Choosing the sample size for a trial</i>	32
<b>Example 3.1</b>	<i>Diagnosis: Bayes theorem in diagnostic testing</i>	52
<b>Example 3.2</b>	<i>Drug: Binary data and a discrete prior</i>	58
<b>Example 3.3</b>	<i>Drug (continued): Binary data and a continuous prior</i>	60
<b>Example 3.4</b>	<i>SBP: Bayesian analysis for normal data</i>	63
<b>Example 3.5</b>	<i>SBP (continued): Interval estimation</i>	67
<b>Example 3.6</b>	<i>GREAT (continued): Bayesian analysis of a trial of early thrombolytic therapy</i>	69
<b>Example 3.7</b>	<i>False positives: ‘The epidemiology of clinical trials’</i>	74
<b>Example 3.8</b>	<i>Credibility: Sumatriptan trial results</i>	77
<b>Example 3.9</b>	<i>GREAT (continued): Sequential use of Bayes theorem</i>	79
<b>Example 3.10</b>	<i>Drug (continued): Making predictions for binary data</i>	82
<b>Example 3.11</b>	<i>GREAT (continued): Predictions of continuing the trial</i>	84
<b>Example 3.12</b>	<i>Neural tube defects: Making personal decisions about preventative treatment</i>	87
<b>Example 3.13</b>	<i>Magnesium: Meta-analysis using a sceptical prior</i>	95
<b>Example 3.14</b>	<i>Coins: A Monte Carlo approach to estimating tail areas of distributions</i>	103
<b>Example 3.15</b>	<i>Drug (continued): Using WinBUGS to implement Markov chain Monte Carlo methods</i>	108
<b>Example 4.1</b>	<i>Stopping: The likelihood principle in action</i>	124
<b>Example 4.2</b>	<i>Preference: P-values as measures of evidence</i>	128
<b>Example 4.3</b>	<i>Preference (continued): Bayes factors in preference studies</i>	134
<b>Example 4.4</b>	<i>GREAT (continued): A Bayes factor approach</i>	136
<b>Example 5.1</b>	<i>CHART: Eliciting subjective judgements before a trial</i>	143
<b>Example 5.2</b>	<i>GUSTO: Using previous results as a basis for prior opinion</i>	153
<b>Example 5.3</b>	<i>CHART (continued): Sceptical priors</i>	160
<b>Example 5.4</b>	<i>Urokinase: ‘lump and smear’ prior distributions</i>	163

<b>Example 5.5</b>	<i>GREAT (continued): Criticism of the prior</i>	175
<b>Example 6.1</b>	<i>CHART (continued): Clinical demands for new therapies</i>	185
<b>Example 6.2</b>	<i>Bayesian power: Choosing the sample size for a trial</i>	194
<b>Example 6.3</b>	<i>Bayesian power (continued): Choosing the sample size for a trial</i>	196
<b>Example 6.4</b>	<i>Gastric: Sample size for a trial of surgery for gastric cancer</i>	197
<b>Example 6.5</b>	<i>Uncertainty: Predictive distribution of power</i>	201
<b>Example 6.6</b>	<i>CHART (continued): Monitoring trials using sceptical and enthusiastic priors</i>	207
<b>Example 6.7</b>	<i>B-14: Using predictions to monitor a trial</i>	214
<b>Example 6.8</b>	<i>CALGB: Assessing whether to perform a confirmatory randomised clinical trial</i>	224
<b>Example 6.9</b>	<i>ECMO: incorporating historical controls</i>	231
<b>Example 6.10</b>	<i>N of 1: pooling individual response studies</i>	237
<b>Example 6.11</b>	<i>CRM: An application of the continual reassessment method</i>	242
<b>Example 7.1</b>	<i>OC: interpreting case-control studies in pharmacoepidemiology</i>	255
<b>Example 7.2</b>	<i>IVF: estimation and ranking of institutional performance</i>	259
<b>Example 8.1</b>	<i>ISIS: Prediction after meta-analyses</i>	271
<b>Example 8.2</b>	<i>EFM: meta-analyses of trials with rare events</i>	275
<b>Example 8.3</b>	<i>Hyper: Meta-analyses of trials adjusting for baseline rates</i>	279
<b>Example 8.4</b>	<i>Blood pressure: Estimating effects that have never been directly measured</i>	283
<b>Example 8.5</b>	<i>Screen: generalised evidence synthesis</i>	288
<b>Example 8.6</b>	<i>Maple: estimating complex functions of parameters</i>	292
<b>Example 8.7</b>	<i>HIV: synthesising evidence from multiple sources and identifying discordant information</i>	295
<b>Example 9.1</b>	<i>Anakinra: Two-stage approach to cost-effectiveness analysis</i>	313
<b>Example 9.2</b>	<i>TACTIC: integrated cost-effectiveness analysis</i>	316
<b>Example 9.3</b>	<i>HIPS: Cost-effectiveness analysis using discrete-time Markov models</i>	325
<b>Example 9.4</b>	<i>HIPS (continued): Integrated generalised evidence synthesis and cost-effectiveness analysis</i>	331
<b>Example 9.5</b>	<i>HIV (continued): Calculating the expected value of perfect information</i>	339



# Introduction

## 1.1 WHAT ARE BAYESIAN METHODS?

Bayesian statistics began with a posthumous publication in 1763 by Thomas Bayes, a Nonconformist minister from the small English town of Tunbridge Wells. His work was formalised as *Bayes theorem* which, when expressed mathematically, is a simple and uncontroversial result in probability theory. However, specific uses of the theorem have been the subject of continued controversy for over a century, giving rise to a steady stream of polemical arguments in a number of disciplines. In recent years a more balanced and pragmatic perspective has developed and this more ecumenical attitude is reflected in the approach taken in this book: we emphasise the benefits of Bayesian analysis and spend little time criticising more traditional statistical methods.

The basic idea of Bayesian analysis is reasonably straightforward. Suppose an unknown quantity of interest is the median years of survival gained by using an innovative rather than a standard therapy on a defined group of patients: we shall call this the ‘treatment effect’. A clinical trial is carried out, following which conventional statistical analysis of the results would typically produce a *P*-value for the null hypothesis that the treatment effect is zero, as well as a point estimate and a confidence interval as summaries of what this particular trial tells us about the treatment effect. A Bayesian analysis supplements this by focusing on how the trial should change our opinion about the treatment effect. This perspective forces the analyst to explicitly state

- a reasonable opinion concerning the plausibility of different values of the treatment effect *excluding* the evidence from the trial (known as the prior distribution),
  - the support for different values of the treatment effect based *solely* on data from the trial (known as the likelihood),
- and to combine these two sources to produce
- a final opinion about the treatment effect (known as the posterior distribution).

The final combination is done using Bayes theorem, which essentially weights the likelihood from the trial with the relative plausibilities defined by the prior distribution. This basic idea forms the entire foundation of Bayesian analysis, and will be developed in stages throughout the book.

One can view the Bayesian approach as a formalisation of the process of learning from experience, which is a fundamental characteristic of all scientific investigation. Advances in health-care typically happen through incremental gains in knowledge rather than paradigm-shifting breakthroughs, and so this domain appears particularly amenable to a Bayesian perspective.

## **1.2 WHAT DO WE MEAN BY ‘HEALTH-CARE EVALUATION’?**

Our concern is with the evaluation of ‘health-care interventions’, which is a deliberately generic term chosen to encompass all methods used to improve health, whether drugs, medical devices, health education programmes, alternative systems for delivering care, and so on. The appropriate evaluation of such interventions is clearly of deep concern to individual consumers, health-care professionals, organisations delivering care, policy-makers and regulators: such evaluations are commonly called ‘health-technology assessments’, but we feel this term carries connotations of ‘high’ technology that we wish to avoid.

A wide variety of research designs have been used in evaluation, and it is not the purpose of this book to argue the benefits of one design over another. Rather, we are concerned with appropriate methods for analysing and interpreting evidence from one or multiple studies of possibly varying designs. Many of the standard methods of analysis revolve around the classical randomised controlled trial (RCT): these include power calculations at the design stage, methods for controlling Type I error within sequential monitoring, calculation of *P*-values and confidence intervals at the final analysis, and meta-analytic techniques for pooling the results of multiple studies. Such methods have served the medical research community well.

The increasing sophistication of evaluations is, however, highlighting the limitations of these traditional methods. For example, when carrying out a clinical trial, the many sources of evidence and judgement available beforehand may be inadequately summarised by a single ‘alternative hypothesis’, monitoring may be complicated by simultaneous publication of related studies, and multiple subgroups may need to be analysed and reported. Randomised trials may not be feasible or may take a long time to reach conclusions. A single clinical trial will also rarely be sufficient to inform a policy decision, such as embarking or continuing on a research programme, regulatory approval of a drug or device, or recommendation of a treatment at an individual or population level. Standard statistical methods are designed for summarising the evidence from single studies or pooling evidence from similar studies, and have difficulties dealing with the pervading complexity of multiple sources of evidence. Many have argued that a fresh, Bayesian, approach is worth investigating.

### 1.3 A BAYESIAN APPROACH TO EVALUATION

We may define a Bayesian approach as ‘the explicit quantitative use of external evidence in the design, monitoring, analysis, interpretation and reporting of a health-care evaluation’. The argument of this book is that such a perspective can be more *flexible* than traditional methods in that it can adapt to each unique situation, more *efficient* in using all available evidence, more *useful* in providing predictions and inputs for making decisions for specific patients, for planning research or for public policy, and more *ethical* in both clarifying the basis for randomisation and fully exploiting the experience provided by past patients.

For example, a Bayesian approach allows evidence from diverse sources to be pooled through assuming that their underlying probability models (their likelihoods) share parameters of interest: thus the ‘true’ underlying effect of an intervention may feature in models for both randomised trials and observational data, even though there may be additional adjustments for potential biases, different populations, crossovers between treatments, and so on.

Attitudes have changed since Feinstein (1977) claimed that ‘a statistical consultant who proposes a Bayesian analysis should therefore be expected to obtain a suitably informed consent from the clinical client whose data are to be subjected to the experiment’. Increasing attention to the Bayesian approach is shown by the medical and statistical literature, the popular scientific press, pharmaceutical companies and regulatory agencies. However, many important outstanding questions remain: in particular, to what extent will the scientific community, or the regulatory authorities, allow the explicit introduction of evidence that is not totally derived from observed data, or the formal pooling of data from studies of differing designs? Indeed, Berry (2001) warns that ‘There is as much Bayesian junk as there is frequentist junk. Actually, there’s probably more of the former because, to the uninitiated, the Bayesian approach seems like it provides a free lunch’. External evidence must therefore be introduced with caution, and used in a clear, explicit and transparent manner that can be challenged by those who need to critique any analysis: this balanced approach should help resolve these complex questions.

### 1.4 THE AIM OF THIS BOOK AND THE INTENDED AUDIENCE

This book is intended to provide:

- a review of the essential ideas of Bayesian analysis as applied to the evaluation of health-care interventions, without obscuring the essential message with undue technicalities;
- a suggested ‘template’ for reporting a Bayesian analysis;

## 4 Introduction

- a critical commentary on similarities and differences between Bayesian and conventional approaches;
- a structured review of published work in the areas covered;
- a wide range of stand-alone examples of Bayesian methods applied to real data, mainly in a common format, with accompanying software which will allow the reader to reproduce all analyses;
- a guide to potential areas where Bayesian methods might be particularly valuable, and where further research may be necessary;
- an indication of appropriate methods that may be applied in different contexts (although this is not intended as a ‘cookbook’);
- a range of exercises suitable for use in a course based on the material in this book.

Our intended audience comprises anyone with a good grasp of quantitative methods in health-care evaluation, and whose mathematical and statistical training includes basic calculus and probability theory, use of normal tables, clinical trial design, and familiarity with hypothesis testing, estimation, confidence intervals, and interpretation of odds and hazard ratios, up to the level necessary to use standard statistical packages. Bayesian statistics has a (largely deserved) reputation for being mathematically challenging and difficult to put into practice, although we recommend O’Hagan and Luce (2003) as a good non-technical preliminary introduction to the basic ideas. In this book we deliberately try to use the simplest possible analytic methods, largely based on normal distributions, without distorting the conclusions: more technical aspects are placed in starred sections that can be omitted without loss of continuity. There is a steady progression throughout the book in terms of analytic complexity, so that by the final chapters we are dealing with methods that are at the research frontier. We hope that readers will find their own level of comfort and make some effort to transcend it.

## 1.5 STRUCTURE OF THE BOOK

We have struggled to decide on an appropriate structure for the material in this book. It could be ordered by *stage of evaluation* and so separate, for example, initial observational studies, RCTs possibly for licensing purposes, cost-effectiveness analysis and monitoring interventions in routine use. Alternatively, we might structure by *study design*, with discussion of randomised trials, databases, case-control studies, and so on. Finally, we could identify the *modelling issue*, for example prior distributions, alternative forms for likelihoods, and loss functions. We have, after much deliberation, made a compromise and used aspects of all three proposals, using extensive examples to weave together analytic techniques with evaluation problems.

Chapter 2 is a brief *revision* of important aspects of traditional statistical analysis, covering issues such as probability distributions, normal tables, parameterisation of outcomes, summarising results by estimates and confidence intervals, hypothesis testing and sample-size assessment. There is a particular emphasis on normal likelihoods, since they are an important prerequisite for much of the subsequent Bayesian analysis, but we also provide a fairly detailed catalogue of other distributions and their use.

Chapter 3 forms the core of the book, being an *overview* of the main features of the Bayesian approach. Topics include the subjective interpretation of probability, use of prior to posterior analysis in a clinical trial, assessing the evidence in reported clinical trial results, comparing hypotheses, predictions, decision-making, exchangeability and hierarchical models, and computation: these topics are then applied to substantive problems in later chapters. Differing perspectives on prior distributions and loss functions are shown to lead to different schools of Bayesianism. A proposed checklist for reporting Bayesian health-care evaluations forms the basis for all further examples in the book.

Chapter 4 briefly critiques the ‘classical’ statistical approach to health-care evaluation and makes a *comparison* with the Bayesian approach. Hypothesis tests, *P*-values, Bayes factors, stopping rules and the ‘likelihood principle’ are discussed with examples. This chapter can be skipped without loss of continuity.

Chapter 5 deals in detail with sources of *prior distributions*, such as expert opinion, summaries of evidence, ‘off-the-shelf’ default priors and hierarchical priors based on exchangeability assumptions. The criticism of prior opinions in the light of data is featured, and a detailed taxonomy provided of ways of using historical data as a basis for prior opinion.

Chapter 6 attempts to structure the substantial work on Bayesian approaches to all aspects of RCTs, including design, monitoring, reporting, and interpretation. The many worked examples emphasise the need for analysis of sensitivity to alternative prior assumptions.

Chapter 7 covers *observational studies*, such as case-control and other non-randomised designs. Particular aspects emphasised include the explicit modelling of potential biases with such designs, and non-randomised comparisons of institutions including ranking into ‘league tables’.

Chapter 8 considers the *synthesis of evidence* from multiple studies, starting from ‘standard’ meta-analysis and then considering various extensions such as potential dependence of treatment effects on baseline risk. We particularly focus on examples of ‘generalised evidence synthesis’, which might feature studies of different designs, or ‘indirect’ comparison of treatments that have never been directly compared in a trial.

Chapter 9 examines how Bayesian analyses may be used to inform *policy*, including cost-effectiveness analysis, research planning and regulatory affairs. The view of alternative stakeholders is emphasised, as is the integration of evidence synthesis and cost-effectiveness in a single unified analytic model.

Chapter 10 includes a final summary, general discussion and some suggestions for future research. Appendix A briefly describes available software and Internet sites of interest.

Most of the chapters finish with a list of key points and questions/exercises, and some have a further guide to the literature.

This structure will inevitably mean some overlap in methodological questions, such as the appropriate form of the prior distribution, and whether it is reasonable to adopt an explicit loss function. For example, a particular issue that arises in many contexts is the appropriate means of including historical data. This will be introduced as a general issue and a list of different approaches provided (Section 3.16), and then these approaches will be illustrated in four different contexts in which one might wish to use historical data: first, obtaining a prior distribution from historical studies (Section 5.4); second, historical controls in randomised trials (Section 6.9); third, modelling the potential biases in observational studies (Section 7.3), and fourth, pooling data from many sources in an evidence synthesis (Section 8.2). This overlap means that a considerable amount of cross-referencing is inevitable and ideally there would be hypertext links, but a traditional book format forces us into a linear structure.

Different audiences may want to focus on different parts of the book. The material up to Chapter 5 comprises a basic short course in Bayesian analysis, suitable for both students and researchers. After that, Chapter 6 may be of more interest to statisticians working with clinical trials in the pharmaceutical industry or the public sector, while Chapters 7–9 may be more appropriate for those exploring policy decisions. However, there are no clear boundaries and we hope that most of the material is relevant for much of the potential readership.

In order to avoid disappointment, we should make clear what this book does *not* contain:

- There is almost no guidance on data analysis, model checking and many other essential ingredients of professional statistical practice. Our discussion of study design is limited to sample-size calculations, and there is little contribution to the debate concerning the relative importance of observational and randomised studies.
- There is no rigorous mathematical or philosophical development of the Bayesian approach, and the technical development is limited entirely to the level required for the examples.
- The examples are almost all taken from published work by ourselves and others, and although they deal with real problems and use real data, there is necessarily a degree of simplification in the presentation. In addition, while the Bayesian approach emphasises the formal use of substantive knowledge and subjective opinion, it is inevitable that judgements are introduced in a somewhat stylised manner into such ‘second-hand’ examples. We should also point out that numbers given in the text have been rounded, and the accompanying programs should be used for a more accurate analysis.

- There is limited development of the decision-theoretic approach to evaluation, and many will feel this is a serious omission. This bias arises from two related reasons. First, our personal experience has been almost entirely concerned with problems of inference, and so that is what we feel qualified to write about. Second, it will become clear that we have some misgivings concerning the application of decision theory in this context, and so prefer to emphasise the more immediately relevant material.
- There is very limited exploration of more general Bayesian approaches to modelling data that arise in health-care evaluations, such as applications to survival analysis, longitudinal models, non-compliance in trials, drop-outs and other missing data, and so on.

The accompanying website will be found at <http://www.mrc-bsu.cam.ac.uk/bayeseval/>, which provides code for most of the examples in the book, either using the BANDY spreadsheet program for simple analysis of odds and hazard ratios, or WinBUGS code for more complex examples. The website will also contain a list of any errors detected.

Finally, we should emphasise that this book is not intended as a polemic in favour of Bayesianism – there have been enough of those – and we shall try to avoid making exaggerated claims as to the benefits of this new ‘treatment’ for statistical problems. Our hope is that we can contribute to the responsible use of Bayesian methods and hence help in a small way towards the development of cost-effective health-care.





# ***Basic Concepts from Traditional Statistical Analysis***

The Bayesian approach, to a considerable extent, supplements rather than replaces the kind of analyses traditionally carried out in assessing health-care interventions, and in this chapter we shall briefly review some of the basic ideas that will subsequently be found useful. In particular, probability theory is fundamental to Bayesian analysis, and we therefore revise the basic concepts with a natural emphasis on Bayes theorem. We also consider random variables and probability distributions with particular emphasis on the normal distribution, which plays a vital role in summarising what the observed data can tell us about unknown quantities of interest. A particularly important practical aspect is the transformation of output from standard statistical packages into a form amenable to Bayesian interpretation.

Bayesian analysis makes a much wider use of probability distributions than traditional statistical methods, in that not only are sampling distributions required for summaries of data, but also a wide range of distributions are used to represent prior opinion about proportions, event rates, and other unknown quantities. The *shapes* of distributions therefore become particularly important, as they are intended to represent the plausibility of different values, and so we shall provide (in starred sections) extensive graphical displays as well the usual formulae.

Most of the issues addressed in this chapter are covered in a concise and readable manner in standard textbooks such as Altman (2001) and Berry *et al.* (2001b). In addition, Clayton and Hills (1993) consider a likelihood-based approach to many of the models that are frequently encountered in epidemiology and health-care evaluation.

## 2.1 PROBABILITY

### 2.1.1 What is probability?

Suppose  $a$  is some event which may or may not take place, such as the next toss of a coin coming up heads. Although we may casually speak of the ‘probability’ of  $a$  occurring, and give it a mathematical notation  $p(a)$ , it is perhaps remarkable that there is no universally agreed definition of what this term means. Perhaps the currently most accepted interpretation is the following:  $p(a)$  is the proportion of times  $a$  will occur in an infinitely long series of repeated identical situations. This is known as the ‘frequentist’ perspective, as it rests on the frequency with which specific events occur. However, a number of other interpretations of probability have been made throughout history, and we shall consider a different, ‘subjective’, definition in Section 3.1.

There is little dispute, however, about the mathematical properties of probability. Let  $a$  and  $b$  be events, and  $H$  represent the context in which  $a$  and  $b$  might arise, and let  $p(a|H)$  denote the probability of  $a$  given the context  $H$ : the vertical line represents ‘conditioning’. Then  $p(a|H)$  is a number that satisfies the following three basic rules:

1. *Bounds.*

$$0 \leq p(a|H) \leq 1,$$

where  $p(a|H) = 0$  if  $a$  is impossible and  $p(a|H) = 1$  if  $a$  is certain in the context  $H$ .

2. *Addition rule.* If  $a$  and  $b$  are mutually exclusive (i.e. one at most can occur),

$$p(a \text{ or } b|H) = p(a|H) + p(b|H).$$

(We note that, for technical reasons, it is helpful if Rule 2 is taken as holding for an infinite set of mutually exclusive events.)

3. *Multiplication rule.* For any events  $a$  and  $b$ ,

$$p(a \text{ and } b|H) = p(a|b,H)p(b|H).$$

We say that  $a$  and  $b$  are independent if  $p(a \text{ and } b|H) = p(a|H)p(b|H)$  or equivalently  $p(a|b,H) = p(a|H)$ ; thus the fact that  $b$  has occurred does not alter the probability of  $a$ . The multiplication rule can equivalently be expressed as the definition of conditional probability,

$$p(a|b, H) = \frac{p(a \text{ and } b|H)}{p(b|H)},$$

provided  $p(b|H) \neq 0$ .

The explicit introduction of the context  $H$  is unusual in standard texts and we shall subsequently drop it to avoid accusations of pedantry: however, it is always useful to keep in mind that *all probabilities are conditional* and so, if the situation changes, then probabilities may change. We shall see in Section 3.1 that this notion forms the basis of *subjective probability*, in which  $H$ , the context, represents the information on which an individual bases their *own* subjective assessment of the *degree of belief*, i.e. probability, of an event occurring.

Example 2.1 illustrates that these rules can be given an immediate intuitive justification by comparison with a standard experiment.

**Example 2.1** *Dice: Illustration of rules of probability*

Suppose  $H$  denotes the roll of two perfectly balanced six-sided dice, and let ' $\equiv$ ' denote 'is equivalent to'.

*Rule 1.* For a single die: if  $a \equiv$  'throw 7', then  $p(a) = 0$ ; if  $a \equiv$  'throw  $\leq 6$ ', then  $p(a) = 1$ . If  $c$  is the sum of the two dice: then if  $c \equiv$  '13', then  $p(c) = 0$ ; if  $c \equiv$  ' $\leq 12$ ', then  $p(c) = 1$ .

*Rule 2.* For a single die: if  $a \equiv$  'throw 3',  $b \equiv$  'throw 4', then

$$\begin{aligned} p(a \text{ or } b) &= p(a) + p(b) \text{ since } a \text{ and } b \text{ are mutually exclusive} \\ &= 1/6 + 1/6 = 1/3. \end{aligned}$$

*Rule 3.* If we throw two dice: if  $a \equiv$  'first die throw 2',  $b \equiv$  'second die throw 5', then

$$\begin{aligned} p(a \text{ and } b) &= p(a)p(b) \text{ since } a \text{ and } b \text{ are independent} \\ &= 1/6 \times 1/6 = 1/36. \end{aligned}$$

If  $a \equiv$  'total score of the two throws is greater than or equal to 6',  $b \equiv$  'first die throw 1', then

$$\begin{aligned} p(a \text{ and } b) &= p(a|b)p(b) \\ &= 1/3 \times 1/6 = 1/18. \end{aligned}$$

Suppose we also consider the events ' $a$  and  $b$ ' and ' $a$  and  $\bar{b}$ ', where  $\bar{b}$  represents the event 'not  $b$ '. Then ' $a$  and  $b$ ' and ' $a$  and  $\bar{b}$ ' are mutually exclusive and together form the event  $a$ , and hence, using Rule 2, we have the identity

$$p(a) = p(a \text{ and } b) + p(a \text{ and } \bar{b}) \quad (2.1)$$

which is known as 'marginalisation'. Further, by using Rule 3, we obtain

$$p(a) = p(a|b)p(b) + p(a|\bar{b})p(\bar{b}), \quad (2.2)$$

which is known by the curious title of ‘extending the conversation’ (or ‘extending the argument’). Example 2.2 shows these expressions follow naturally from considering the full ‘joint’ distribution over all possible combinations of events.

**Example 2.2** *Prognosis: Marginalisation and extending the conversation*

Suppose we wish to determine the probability of survival (up to a specified point in time) following a particular cancer diagnosis, given that it depends on the stage of disease at diagnosis amongst other factors. Whilst directly specifying the probability of surviving, denoted  $b$ , may be difficult, by extending the conversation to include whether the cancer was at an early stage, denoted  $a$ , or not, denoted  $\bar{a}$ , we obtain from (2.1),

$$p(b) = p(b|a)p(a) + p(b|\bar{a})p(\bar{a}).$$

For example, suppose patients with early stage disease have a good prognosis, say  $p(b|a) = 0.80$ , but for late stage it is poor, say  $p(b|\bar{a}) = 0.20$ , and that of new diagnoses the majority, 90%, are early stage, *i.e.*  $p(a) = 0.90$  and  $p(\bar{a}) = 0.10$ . Then the marginal probability of surviving is  $p(b) = 0.80 \times 0.90 + 0.20 \times 0.10 = 0.74$ .

Table 2.1 shows all possible combinations of events and their probabilities, as well as the marginal probabilities that, appropriately, appear in the margin of the table. The joint probabilities of events have been obtained by Rule 2 so that, for example,  $p(b \text{ and } a) = p(b|a)p(a) = 0.80 \times 0.90 = 0.72$ .

**Table 2.1** Probabilities of all combinations of survival and stage, including marginal probabilities.

	Early stage $a$	Late stage $\bar{a}$	
Survive $b$	0.72	0.02	0.74
Not survive $\bar{b}$	0.18	0.08	0.26
	0.90	0.10	1.00

### 2.1.2 Odds and log-odds

Any probability  $p$  can also be expressed in terms of ‘odds’  $O$ , where

$$O = \frac{p}{1-p}$$

and

$$p = \frac{O}{1+O},$$

so that, for example, a probability of 0.20 (20% chance) corresponds to odds of  $O = 0.20/0.80 = 0.25$  or, in betting parlance, '4 to 1 against'. Conversely, betting odds of '7 to 4 against' correspond to  $O = 4/7$ , or a probability of  $p = 4/11 = 0.36$ .

The natural logarithm (denoted  $\log$ ) of the odds is termed the 'logit', so that

$$\text{logit}(p) = \log \left[ \frac{p}{1-p} \right].$$

### 2.1.3 Bayes theorem for simple events

A number of properties can immediately be derived from Rules 1 to 3 of Section 2.1.1. Since  $p(b \text{ and } a) = p(a \text{ and } b)$ , Rule 3 implies that  $p(b|a)p(a) = p(a|b)p(b)$ , or equivalently

$$p(b|a) = \frac{p(a|b)}{p(a)} \times p(b). \quad (2.3)$$

We have proved Bayes theorem! In words, this vital result tells us how an initial probability  $p(b)$  is changed into a conditional probability  $p(b|a)$  when taking into account the event  $a$  occurring: it should be clear by this description that we are interpreting Bayes theorem as providing a formal mechanism for learning from experience.

Equation (2.3) also holds for  $\bar{b}$ , so that

$$p(\bar{b}|a) = \frac{p(a|\bar{b})}{p(a)} \times p(\bar{b}), \quad (2.4)$$

and dividing (2.3) by (2.4) we obtain the *odds form* for Bayes theorem:

$$\frac{p(b|a)}{p(\bar{b}|a)} = \frac{p(a|b)}{p(a|\bar{b})} \times \frac{p(b)}{p(\bar{b})}. \quad (2.5)$$

Thus  $p(b)/p(\bar{b}) = p(b)/(1-p(b))$ , the odds on  $b$  before taking into account the event  $a$ , which is changed into the new odds  $p(b|a)/p(\bar{b}|a)$  after conditioning on  $a$ . Equation (2.5) shows how Bayes theorem accomplishes this transformation without even explicitly calculating  $p(a)$ , and this insight is exploited in Section 3.2.

---

#### Example 2.3 Prognosis (continued): Bayes theorem for single events

Suppose we were given Table 2.1, and wanted to use Bayes theorem to tell us how knowing the stage of the disease at diagnosis revises our probability for survival  $a$ . Initially, before we know the stage,  $p(b) = 0.74$  from the

marginal probability in Table 2.1. Suppose we find out that the disease is at an early stage, *i.e.*  $a$ , where we know from Table 2.1 that  $p(a|b) = 0.72/0.74 = 0.97$  and  $p(a) = 0.9$ . Hence from (2.3) we obtain a revised probability of survival

$$p(b|a) = \frac{0.97}{0.9} \times 0.74 = 0.80,$$

matching what, in fact, we knew already.

To use the odds form of Bayes theorem (2.5) we first require the initial odds for survival, *i.e.*  $p(b)/p(\bar{b}) = 0.74/0.26 = 2.85$ , and the ratio  $p(a|b)/p(a|\bar{b}) = 0.97/0.69 = 1.405$ . Then from (2.5) we obtain the final odds on survival as  $2.85 \times 1.41 = 4.01$ , corresponding to a probability  $p(b|a) = 0.80$  (up to rounding error).

The two forms of Bayes theorem both give the required results and can be thought of as a means of moving from a marginal probability in a table to a conditional probability having taken into account some evidence. As we shall see in Section 3.2, it is this use of Bayes theorem that is used in many diagnostic testing situations without any controversy.

---

## 2.2 RANDOM VARIABLES, PARAMETERS AND LIKELIHOOD

### 2.2.1 Random variables and their distributions

Random variables have a somewhat complex formal definition, but it is sufficient to think of them as unknown quantities that may take on one of a set of values: traditionally a random variable is denoted by a capital Latin letter, say  $Y$ , before being observed and by a lower-case letter  $y$  as a specific observed value. This convention tends to be broken in Bayesian analysis, in which all unknown quantities are considered as random variables, but we shall try to keep to it where it clarifies the exposition.

Loosely speaking,  $p(y)$  denotes the probability of a random variable  $Y$  taking on each of its possible values  $y$ .  $p(y)$  is formally known as the *probability density function*, and the probability that  $Y$  does not exceed  $y$ ,  $P(Y \leq y)$ , is termed the *probability distribution function*. We shall tend to use ‘probability distribution’ as a generic term, hopefully without causing confusion.

Probability distributions may be:

**Binary.** When  $Y$  can take on one of two values, we shall generally use the notation  $Y = 1$  for when an event of interest occurs, and  $Y = 0$  when it does not: this is