

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.