

Preface

This book originates from the work that we have done, at different times and in different capacities, in the area of statistical modelling for health economic evaluation. In our view, this is a very interesting and exciting area for statisticians: despite the strong connotation derived by its name, health economic evaluation is just as much (if not more!) about statistics than it is about healthcare or economics. Statistical modelling is a fundamental part of any such evaluation and as models and the data that are used to populate them become bigger, more complex and representative of a complicated underlying reality, so do the skills required by a modeller.

Broadly speaking, the objective of publicly funded health care systems (such as the UK's) is to maximise health gains across the general population, given finite monetary resources and a limited budget. Bodies such as the National Institute for Health and Care Excellence (NICE) provide guidance on decision-making on the basis of health economic evaluation. This covers a suite of analytical approaches (usually termed “cost-effectiveness analysis”) for combining costs and consequences of intervention(s) compared to a control, the purpose of which is to aid decision making associated with resource allocation. To this aim, much of the recent research has been oriented towards building the health economic evaluation on sound and advanced statistical decision-theoretic foundations.

Historically, cost-effectiveness analysis has been based on modelling often performed in specialised commercial packages (such as *TreeAge*) or even more frequently spreadsheet calculators (almost invariably *Microsoft Excel*). The “party-line” for why this is the case is that these are “easy to use, familiar, readily available and easy to share with stakeholders and clients”. Possibly, in addition to these, another crucial factor for the wide popularity of these tools is the fact that often modellers are not statisticians by training (and thus less familiar with general-purpose statistical packages such as *SAS*, *Stata* or *R*). Even more interestingly, it is often the case that cost-effectiveness models are based on existing templates (usually developed as *Excel* spreadsheets, for example for a specific country or drug) and then “adapted” to the situation at hand.

Luckily, we are not alone (although perhaps not in the majority) in arguing that many of these perceived advantages require a serious rethink. In our view, there are several limitations to the current state of modelling in health economics: firstly, the process often implies a separation of the different steps required for the evaluation. This potentially increases the risk of human errors and confusion, because the results of the intermediate steps (*e.g.* the statistical analysis of data collected in a randomised trial) are usually copied and pasted in Excel to populate cells and formulæ (see for instance our discussion in §1.4 and §4.2). Secondly, in an Excel file calculations are usually spread over several sheets that are linked by formulæ or cross references. While in the case of simple models this is actually a neat way of structuring the work, it can become unwieldy and difficult to track modifications for more complex models, based on a combination of different datasets and thus analyses (which of course is increasingly the norm!).

The idea of the R package BCEA has evolved naturally by the need of replicating some types of analyses when post-processing the output of the models we were developing in our applied work, while overcoming the limitations of the “standard” work flow based on spreadsheets. It felt natural to make the effort of systematising the functions we were using to do standard analyses and as we started doing so, we realised that there was much potential and interesting work to be done. The main objective of this book is to aid statisticians and modellers in health economics with the “easier” part of the process — making sense of their model results and help them reproduce the analysis that is, more or less, ubiquitous in the relevant output (be it a research paper, or a dossier to be submitted to a regulatory agency such as NICE).

To this aim, the book is structured as follows. First, in chapter 1 we introduce the main concepts underlying the Bayesian approach and the basics of health economic evaluation, with particular reference to the relevant statistical modelling. Again, linking the two is natural to us as we are of a very strong Bayesian persuasion. In addition to this, however, it is interesting to note that Bayesian methods are extremely popular in this area of research, since they are particularly useful in modelling composite sources of information (often termed “evidence synthesis”) and effectively underlie the important concept of Probabilistic Sensitivity Analysis (PSA, see for instance chapter 4).

Chapter 2 presents the two case studies we use throughout the book. In particular, we introduce the statistical modelling and notation, describe the whole process of running the analysis and obtaining the relevant output (in the form of posterior distributions) and then the extra modelling required to compute the quantities of interest for the economic analysis. This process is performed under a fully Bayesian approach and is based on a combination of R and BUGS/JAGS, the *de facto* standard software to perform Markov Chain Monte Carlo analysis.

Chapter 3 introduces the R package BCEA and its basic functionalities by means of the two running examples. The very nature of BCEA is to follow a full Bayesian analysis of the statistical model used to estimate the economic

quantities required for the cost-effectiveness analysis, but we make here (and later in the book) the point that it can also be used in the case where the modelling is done using frequentist arguments — provided suitable inputs (*e.g.* in the forms of simulations for the effects and costs of the interventions under investigation).

Chapter 4 is perhaps the most technical and, as mentioned above, it introduces the concept of PSA and its application. It is also there that we make the point that increasingly popular and important analyses (*e.g.* based on the evaluation of the value of information) simply *cannot* be performed in spreadsheet or other sub-optimal software.

Finally, Chapter 5 presents an extension of BCEA, which can be turned into a web-based application using R Shiny. To a R user, this is perhaps not very useful — all that BCEAweb can do, BCEA in R can do even better (while the reverse is not totally true). However, we do feel that using web-interfaces is indeed very important to disseminate the message and convince practitioners of the supremacy of R over Excel or other specialised software. The main argument is that BCEAweb allows the user an intermediate step between the “standard” Excel based modelling and the “ideal” (at least to our mind) situation in which all the analysis is performed in R. BCEAweb can also be used to produce a graphical interface to help the “translation” of the model in simpler, may be graphical terms. This will probably overcome the complaints that clients (*e.g.* pharmaceutical companies commissioning cost-effectiveness analysis for their products) or stakeholders (*e.g.* reviewers and committee members in regulatory agencies) have: they want to be able to use menu-bars and sliders to modify the models in an easy and intuitive way. Tools such as BCEAweb will allow this.

The final version of the book has benefitted from comments, suggestions and discussion with many friends and colleagues. Among them, we would like to mention Mark Strong, Nicky Welton, Chris Jackson, Matthew Bending, James Jarrett, Andreas Karabis, Petros Pechlivanoglou and Katrin Haeussler. Polina Hadjipanayiotou has worked extensively on BCEAweb as part of her MSc dissertation at University College London and has contributed to the write up of Chapter 5. One particular thought goes to Richard Nixon, a brilliant researcher and a lovely person — it has been a pleasure to have met him and picked his brains on how to make BCEA better. Eva Hiripi at Springer has been extremely patient and helpful in her support, while we were preparing the book (we were often joking that we would publish it in 2036, so we are pretty pleased that in fact we are 20 years earlier than predicted!). Similarly, the comments of two anonymous reviewers have been very useful and have contributed to making the final version a better and, hopefully, clearer one.

London,
December 2016

Gianluca Baio
Andrea Berardi
Anna Heath