

Probabilistic Sensitivity Analysis: Be a Bayesian

Hendriek C. Boshuizen, PhD, Pieter H. M. van Baal, PhD

National Institute of Public Health and the Environment (RIVM), Bilthoven, The Netherlands

ABSTRACT

Objective: To give guidance in defining probability distributions for model inputs in probabilistic sensitivity analysis (PSA) from a full Bayesian perspective.

Methods: A common approach to defining probability distributions for model inputs in PSA on the basis of input-related data is to use the likelihood of the data on an appropriate scale as the foundation for the distribution around the inputs. We will look at this approach from a Bayesian perspective, derive the implicit prior distributions in two examples (proportions and relative risks), and compare these to alternative prior distributions.

Results: In cases where data are sparse (in which case sensitivity analysis is crucial), commonly used approaches can lead to unexpected results. We

show that this is because of the prior distributions that are implicitly assumed, namely that these are not as “uninformative” or “vague” as believed. We propose priors that we believe are more sensible for two examples and which are just as easy to apply.

Conclusions: Input probability distributions should not be based on the likelihood of the data, but on the Bayesian posterior distribution calculated from this likelihood and an explicitly stated prior distribution.

Keywords: Bayesian methods, maximum likelihood estimation, prior probability distribution, probabilistic sensitivity analysis.

Background

In economic evaluation employing modeling techniques, the model typically contains several unknown parameters [1]. The outcome of a study will depend on the values that are postulated for these parameters. These parameters are seldom based on hard facts; in most cases, there is uncertainty about their magnitude.

Probabilistic sensitivity analysis (PSA) has become the state-of-the-art method for determining the uncertainty in the outcomes of cost-effectiveness calculations for health-care interventions because of the uncertainty in input parameters. For instance, in the UK, the National Institute for Clinical Excellence recommendations state that PSA should be employed in order to yield unbiased estimates of expected net monetary benefits, and more importantly, to characterize decision uncertainty [2–4].

In a PSA [5], the uncertainty in each parameter is quantified in terms of a probability distribution of this parameter. One then carries out a Monte Carlo simulation, in which one randomly draws one value for each parameter from its probability distribution and then calculates the outcome corresponding to the set of parameters drawn. This process is repeated M times, yielding M outcome values that represent the distribution of the outcome values (for a given choice of the distributions of the input parameters of the model). PSA is a conceptually simple and intuitive method, and as such has considerable appeal. It can be seen as an implementation of Bayesian statistics, as the view that parameters have a probability distribution is a hallmark of the Bayesian outlook. Moreover, the decision context in which economics evaluations are carried out is essentially Bayesian [6,7].

Parameter values usually come from data that are collected in a single study, studies that combine data from multiple studies (meta-analysis), expert opinion, or applying complex methods of Bayesian evidence synthesis [8]. An important step in performing

a PSA is defining the probability distribution to quantify the uncertainty in the input parameters. One guide in this field (Briggs et al. [9] hereafter called BCS) describes methods to fit distributions of parameter values “directly to the data.” In our experiences, this book, which grew out of a popular course on health economic modeling, is a popular guide for those carrying out cost-effectiveness analyses, and its methods are followed widely. Although it clearly states (and advocates) the Bayesian context of decision modeling, and describes the underlying theory, this guide’s final recommendations with respect to the choice of input probability distributions are not discussed from the viewpoint of the underlying Bayesian prior distributions. We are aware that a Bayesian perspective with respect to the choice of input probability distributions may scare some applied modelers.

However, as we will argue below, if seen from a Bayesian perspective, fitting parameter values “directly to the data” implies choices for prior distributions that need justification, as, in our opinion, more suitable alternative choices are possible. More importantly, we argue that in the case of some parameters, it is just as easy to estimate input probability distributions by assuming a more sensible alternative prior distribution. We will elaborate on two important types of parameters, namely a probability (or proportion), and a ratio (e.g., a relative risk [RR]). Although the prior distribution plays only a minor role whenever data are abundant, this is not always the case, especially given the current trend toward modeling of many specific subgroups [10]. Also, as our proposals are just as simple to use as those proposed in BCS, there are no practical reasons for not using them.

Our article is confined to the situation where the uncertainty of different parameters in the model is assumed to be mutually independent, for instance in cases where they are based on different sources. When multiple parameters are correlated, for instance because they are based on the same data source, the correlation between the uncertainty should also be taken into account. If not, the outcome of the PSA might be severely biased [11–13]. Also, in the case that input parameters are based on Bayesian evidence synthesis of trial data, there will be correlation between the estimates of individual parameters. In all these cases,

Address correspondence to: Hendriek C. Boshuizen, National Institute of Public Health and the Environment (RIVM), Department Statistics and Mathematical Modelling, PO Box 1, 3720 AA Bilthoven, The Netherlands. E-mail: Hendriek.Boshuizen@rivm.nl
 10.1111/j.1524-4733.2009.00590.x

the PSA should sample from the joint posterior distribution of these parameters. This topic, although important in many situations, is not dealt with here.

The structure of the article is the following. First, we will discuss the implicit prior distributions for ratio and proportion parameters obtained by fitting distributions “directly to the data” (as described by BCS) from a Bayesian perspective. Next, we derive alternative prior distributions and show that their use leads to more intuitive input distributions for PSA.

Being a Bayesian

From a Bayesian point of view, the distributions that enter as input into the PSA are themselves “posterior distributions” based both on a “prior distribution” and on the data according to the following central formula in Bayesian statistics:

$$p(\theta|x) \propto p(x|\theta) \times p(\theta) \tag{1}$$

This formula states that the posterior distribution of the input parameter ($p(\theta|x)$) is proportional to the product of the likelihood of the data ($p(x|\theta)$) and the prior distribution of the parameter θ . The posterior distribution ($p(\theta|x)$) is the distribution we want to use in our PSA, as this gives the probability distribution of the parameter after we take the data into account.

Fitting input distributions directly to the data has considerable appeal as it seems to avoid the potentially messy business of having to choose a prior. However, this is deceiving: the methods as described in BCS implicitly assume a particular prior distribution. In their chapter on choosing distributions for input parameters, BCS give guidelines for the choice of distributions for (among others) proportions and ratios. In the case of a proportion, other possible priors are discussed by BCS in their technical appendix to that chapter, and distributions based on other priors are also applied in another article from Briggs et al. [14]. However, if we consider the underlying prior distributions, their final recommendations would not be our preference. In the case of a ratio, we propose an alternative prior that, as far as we know, has not yet been discussed in the health economic literature. Although we did not come up with this prior for that particular purpose, this alternative prior also obviates the shortcoming of the expected value of the input distribution not being equal to the point estimate as computed from the data. As the binomial proportion lends itself well to explaining the Bayesian method, we will discuss this first.

Binomial Proportion

For the binomial proportion, BCS advise using a beta distribution characterized by two parameters a and b , and propose to use the number of positive outcomes observed in the data for a , and the number of negative outcomes for b . This approach implies a so-called Haldane prior distribution, which is proportional to $p^{-1}(1-p)^{-1}$ (where p is the proportion). Alternative distributions to use in PSA are $\text{beta}(a+1, b+1)$ (assuming a uniform prior), or $\text{beta}(a+0.5, b+0.5)$ (assuming a Jeffrey’s prior, proportional to $p^{-1/2}(1-p)^{-1/2}$) [15,16]. Figure 1 displays the probability density functions for these different prior distributions. To illustrate the differences between using alternative priors, we take the following simple example: In a trial, there are two arms, each with 100 patients. The object of study is (among other outcomes) the overall mortality, which in our example is rare: there is only one death in arm A, while there are 0 deaths in arm B. Figure 2 displays the posterior probabilities based on these data using the three different priors.




Figure 1 Possible prior probability for a binomial proportion: Haldane prior (dashed line), Jeffrey’s prior (solid line), and uniform prior (dotted line). The Haldane prior has an infinitely large value at $P=0$ and $P=1$, and infinitely small values at P -values in between. Although all values in between 0 and 1 are infinitely small, they are not all equal, as shown by our plotted function.

If we use these posterior distributions in a PSA, using (for illustrational purposes) a model that just copies the input to the output, we get a distribution of output values which would look exactly like Figure 2. Using the distributions based on the Haldane prior, the means of the output values from PSA are equal to the empirical rates, (0 and 0.01, respectively), while this is not the case when using the posteriors based on the other priors. This can be seen as an advantage of using the Haldane prior. The use of the Haldane prior, however, has also a serious drawback: in the case that either a or b is zero (i.e., the empirical proportion is 0 or 100%), the resulting distribution is no longer a proper probability distribution: $\text{beta}(0,b)$ or $\text{beta}(a,0)$ represents complete certainty that the value really is 0% or 100%, also in the case where data are scarce. So, if none of two patients die in a particular situation, it means that the posterior probability of dying in similar situations is taken to be zero with complete certainty. Common sense tells that this is not realistic.





Figure 2 Resulting posterior probability density distributions for the data: 0 death in 100 participants (upper plot) and 1 death in 100 participants (lower plot) using a Haldane prior (dashed line), a Jeffrey’s prior (solid line), or a uniform prior (dotted lines).

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات