Volume 12 • Number 8 • 2009
VALUE IN HEALTH

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

  1098-3015/09/1210 1210–1214

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:

$$p(\theta|x) \propto p(x|\theta) \times p(\theta) \quad (1)$$

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 beta($a+1$, $b+1$) (assuming a uniform prior), or 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.

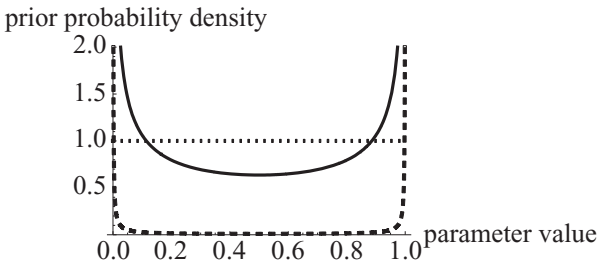

Figure 1 Possible prior probability for a binomial proportion: Haldane prior (dashed line), Jeffrey's prior (solid line), and uniform prior (dotted line). The Haldane prior has an infinitely large value at $P=0$ and $P=1$, and infinitely small values at P-values in between. Although all values in between 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: beta(0,b) or 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 a posteriori probability of dying in similar situations is taken to be zero with complete certainty. Common sense tells that this is not realistic.

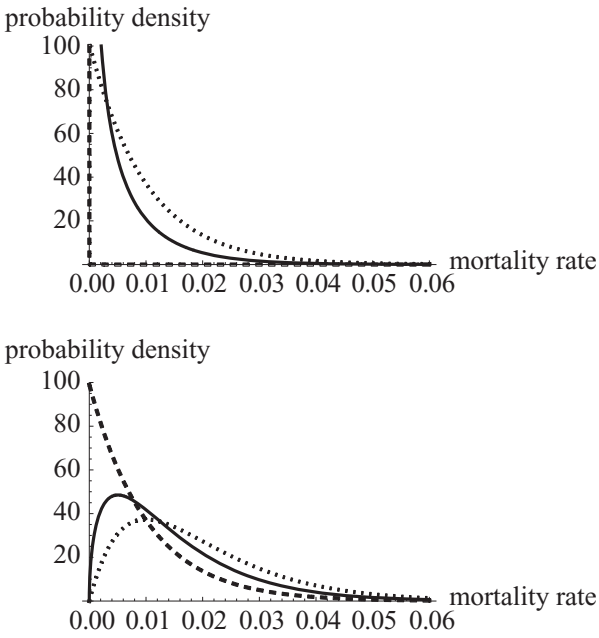

Figure 2 Resulting posterior probability density distributions for the data: 0 death in 100 participants (upper plot) and 1 death in 100 participants (lower plot) using a Haldane prior (dashed line), a Jeffrey's prior (solid line), or a uniform prior (dotted lines).

The advantage of using the Haldane prior is that the posterior mean is equal to the empirical rate. This resembles the advantage of the maximum likelihood estimator (MLE) for a proportion in statistics. An MLE estimate gives the parameter value at which the likelihood of the data is at its maximum and also yields a result equal to the empirical rate. The argument therefore might be made that using a Haldane prior conforms more to common methods of statistical analysis. We therefore will briefly review the arguments behind the MLE.

First, the MLE can be seen as the maximum of the posterior distribution given a uniform prior. This posterior is given by the dotted line in Figure 2. The likelihood (apart from a normalizing constant) also looks like the dotted line in Figure 2. The MLE is, in statistical terms, thus equal to the posterior mode of a posterior distribution based on a uniform prior distribution, rather than to the posterior mean of this posterior distribution.

The philosophy leading to using the mode rather than the mean of the likelihood in MLE estimation has (among others) to do with what is called asymptotic unbiasedness and efficiency: the first means that if one repeats the experiment an infinite number of times, the MLE will give the true answer on average. For instance, if the real mortality is 1 in 200, then repeating the experiment with N = 100 many times will give 60.6% trials with 0 cases, 30.4% trials with one case, 7.6% trials with two cases, and 1,4% trials with three or more cases. This implies that 60.6% of the MLEs are equal to 0, 30.4% are equal to 0.01, 7.6% are equal to 0.02, and 1.4% are equal to 0.03 or more. On average, the MLE than is 0.005.

Efficiency means that the average amount that the estimate is "off target" (defined as the root mean square error) is as small as possible. This is related to the fact that if the size of the experiment increases (say you have 1000 patients with 0 or 10 deaths), the MLE stays in place (while the posterior mean will not). These characteristics, however, are not of much use in PSA: one does not have data from repeating the experiment many times: the knowledge that a particular estimate is "on average" unbiased and precise does not tell you whether it is right or precise in this particular case. Common sense will tell that an observation of 0 or 1 deaths in a group of 100 persons in most cases will not mean that the true mortality rate is exactly 0 or 0.01, but rather that it is compatible with many other "true" values beside 0.0 or 0.01. Also, although an observation of 0 deaths in 100 participants gives the same empirical rate as 0 in 10,000 participants, the true mortality rates that are compatible with these data will be lower in the last case.

In many cases, one is pretty sure that a particular event can occur in some patients and not in others, although its probability might be very low or very high. If you are sure that (based on evidence external to the collected data) one single case of the event has ever occurred, a possible event rate of zero is no longer possible. In this case, we would advocate a uniform prior (that assigns only a very small prior probability to the chance that the rate is zero or 1), and in this case the distribution to be used in PSA is beta($a + 1$, $b + 1$). Briggs et al. [14] also noted the problem of using beta(a,b) for data with zero rates, and proposed using distributions based on such priors here.

However, there might be situations where one seriously questions whether the event could occur at all. For example, one might model the probability of hair color changing to pink as a result of taking a particular red-colored drug, which it is rumored to do on the Internet. The outcome is included in the model, but we do not think this effect possible from a biological point of view. Still, the Haldane prior does not seem prudent in this case either as this would make a zero rate already a certainty based on a study with only a few subjects. A better choice in this case is the Jeffrey's prior (beta(0.5, 0.5)), an intermediate prior between the Haldane prior and uniform prior.

Ratio Measures

Ratio measures, like odds ratios (ORs) and (hazard) rate ratios (HRs), are generally modeled on the logarithmic (log) scale, using mostly logistic regression (modeling the log[OR]) for count data or proportional hazards regression (modeling the log HR) for survival data. These models yield parameter estimates and parameter (co)variances on the log scale. When such ratio measures are used as input in a health economic model, BCS advise that their uncertainty is modeled by a normal distribution on the log scale, using the parameter estimate (log OR or log HR) as mean, and the standard error of the log(OR) or log(HR) as standard deviation. From a Bayesian point of view, this normal distribution is the posterior distribution of the log(OR) or log(HR) given the data used in the logistic or proportional hazards model and assuming a uniform prior distribution on log(OR) or log(HR), that is, assuming that all values of log(OR) or log(HR) have equal probability.

In modeling effects of interventions, however, one uses the ratio itself, not the log of the ratio. For instance, event rates are modeled by multiplying baseline event rates with the RR itself, not with its logarithm. So, it is the ratio itself, and not the log(ratio) that is usually proportional to the outcome in a health economic setting.

The implicit uniform prior on log(OR) or log(HR) (uniform prior on the log scale) at first sight seems reasonable: all values of log(OR) or log (HR) between minus infinity and plus infinity are equally likely and the average of this prior on the log scale is 0 (as each positive value of log(OR) will cancel out a corresponding negative value), corresponding to an OR or HR of 1. However, an uninformative prior one scale can be informative on another scale. On the linear scale, this prior implies a prior that is proportional to 1/OR or 1/HR. Despite the fact that this prior assigns probabilities that decline with increasing OR or HR values, this prior has an average of infinity. Thus, with a limited amount of empirical information, the average of the posterior will partly reflect the infinitely high average of the implicit prior.

To illustrate this, we will use the example of using an RR on larynx cancer in current smokers compared to never smokers, as taken from a meta-analysis [17] (RR 6.76 with a 95% confidence interval of [2.86, 16.0]). The statistical methods used in meta-analysis yield an estimate of log(RR) and a standard error of the estimated log(RR). Figure 3 shows three different prior distributions that could be considered in this case and which are discussed below. Figure 4 displays the corresponding posterior distributions (the potential input distributions for PSA) for this RR on larynx cancer that are based on these three different prior distributions. Two of the plotted prior distributions in Figure 3 are improper priors, that is, their integral is infinite. This means that they cannot be rescaled to a probability density, which must integrate to 1. In other terms, their scaling constant is infinitely small and we can therefore not plot these prior distributions. Here, we therefore plotted a function that is proportional to these priors, choosing an arbitrary scaling constant. Note, however, that an improper prior still can have a finite mean.

We start with discussing the prior that is implicit in the meta-analytic model, as recommended by BCS (the solid lines in Figs. 3 and 4). The implicit prior in this case is a uniform prior on the scale on which the analysis was carried out, in this case on the log(RR). The posterior probability of the RR in this case follows the lognormal distribution given by the solid curve in Figure 4.

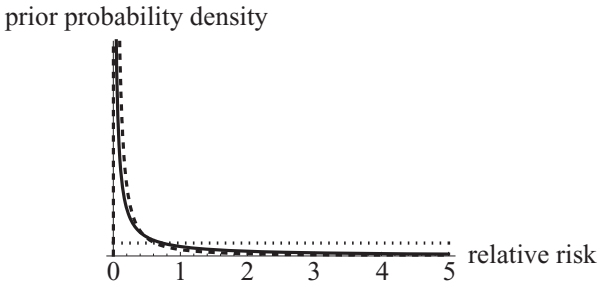

Figure 3 Possible prior probabilities for a ratio: uniform prior on the log-ratio scale (solid line), uniform prior on the ratio scale (dotted line), or lognormal prior with $\mu = -0.5$ σ^2 (here, we plotted $\sigma^2 = 100$; dashed line). The second prior has an infinitely large value at ratio = 0, while the lognormal prior is equal to 0 at ratio = 0.

Here, the median of this posterior distribution is equal to 6.8, the original point estimate. If this posterior distribution is used in PSA, in a model where the outcome value would just be a copy of the input, the mean outcome will be approximately 10% higher (7.4) than 6.8, reflecting the fact that the mean of a lognormal distribution is higher than its median. The less precise the estimate of the RR, the higher the mean of the lognormal distribution of the RR relative to its median. Generally, however, RRs used in modeling are fairly precise, and the difference between the median and the mean is smaller than in our example. However, in models where several RRs are multiplied with each other (adding also nonlinearity to the model), the overall effect could become large enough to be of practical interest.

An alternative to the BCS recommendations is to assume a uniform prior (on the interval zero to infinity) on the RR itself instead of on log(RR) (dotted line in Fig. 3). This leads to the posterior distribution given by the dotted line in Figure 3. As can be seen, the mode of this function (the MLE) is 6.8 (the figure given by the meta-analysis), but its mean is even higher than when using a uniform prior on log(RR) (the solid line). This is because assuming a uniform prior (between zero and infinity) on the RR itself has an average that is even higher than that of a uniform prior on log(RR). This approach therefore does not solve the problem of having a prior distribution with an average of infinity.

To us, assuming a prior distribution for the RR that has an average of infinity does not seem realistic, and we would prefer a prior that converges to a sensible finite value. A prior that has this characteristic and that is also mathematically convenient is a lognormal distribution, characterized by its parameters μ_{prior} and $sigma_{prior}$.

Such a lognormal prior on RR yields a posterior for RR (dashed line in Fig. 4) that is a lognormal distribution with

$$\mu_{post} = \mu_{lik} - (\mu_{lik} - \mu_{prior}) \frac{\sigma^2_{lik}}{\sigma^2_{lik} + \sigma^2_{prior}}$$

and

$$\sigma^2_{post} = \frac{\sigma^2_{prior}\sigma^2_{lik}}{\sigma^2_{prior} + \sigma^2_{lik}}$$

where μ_{post} is the posterior mean of log(RR), μ_{lik} the point estimate for log(RR) from the meta-analysis, and $sigma_{lik}$ the standard error of log(RR) from the meta-analysis. By choosing a sufficiently large value for $sigma_{prior}$, one could make sure that the prior is sufficiently vague, so it would not play an important role. A conservative choice for μ_{prior} would set the average prior RR to 1, implying

$$\mu_{prior} = -\frac{1}{2}\sigma^2_{prior}$$

Thus,

$$\mu_{post} = \mu_{lik} - \left(\mu_{lik} + \frac{\sigma^2_{prior}}{2}\right)\frac{\sigma^2_{lik}}{\sigma^2_{lik} + \sigma^2_{prior}} = \mu_{lik}\frac{\sigma^2_{prior}}{\sigma^2_{lik} + \sigma^2_{prior}} - \frac{1}{2}\sigma^2_{post}$$

$$\text{if } \sigma^2_{prior} \gg \sigma^2_{lik}$$

$$\sigma^2_{post} \approx \sigma^2_{lik}$$

$$\mu_{post} \approx \mu_{lik} - \frac{1}{2}\sigma^2_{lik}$$

This prior is given by the dashed line in Figure 3. These are exactly the values for μ_{post} and σ_{post} that BCS advise to use if one wants the expectation of the probabilistic distribution to correspond to the point estimate from a log link generalized linear model. It can be seen that this advice can be justified from using the vague prior described earlier.

In our example, the RR was larger than 1. If it would have been smaller than 1 (for instance, in case of a treatment effect), then using a uniform prior on the log scale would yield a posterior with a mean closer to 1 than the original point estimate (representing a smaller treatment effect). The alternative prior again would yield a mean equal to the original point estimate.

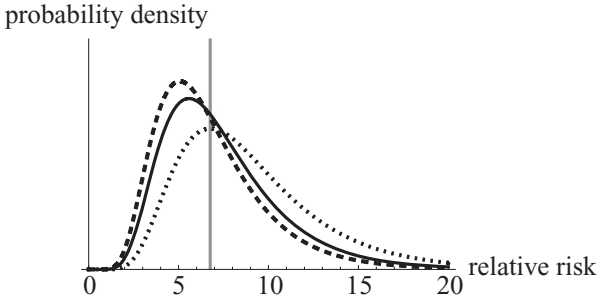

Figure 4 Posterior distribution for a ratio using a uniform prior on the ratio itself (dotted line), using a uniform prior on the log ratio (solid line), and using the alternative prior described in the text (dashed line). The reference line gives the value of the point estimate of the ratio as calculated in the meta-analysis.

Discussion

The importance of a Bayesian outlook in health economic modeling has been stated before [6,18]. We feel that this should be extended to the topic of defining input distributions as used in sensitivity analysis. The reasoning that such input distributions should be founded on data only is tempting. Such reasoning has brought us likelihood-based statistics, but, unfortunately, for purposes other than summarizing evidence from data, the reasoning is deceptive. When working with the uncertainty of an estimated parameter, one uses the Bayesian concept of posterior probability, which is always based on an assumed prior. The choice of such a prior therefore should be discussed.

One important option of the Bayesian outlook is that one can specify "informative" prior distributions for the parameters of interest, using any relevant information available and without being constrained to a prespecified form for the posterior distributions. The resulting posterior distributions can then be used in

the economic analysis, which then is based on all existing relevant information. This is surely the optimal approach from a scientific point of view, but in economic evaluations, where financial stakes are often high, one prefers to use information based on hard facts as much as possible, in order to avoid the semblance of having manipulated the data.

In health economical modeling expectancies rather than modes or medians are the summarizing entities of interest. If input parameters have skewed distributions, it generally is the mean of this distribution rather than its mode or median that is related to the expected value of the outcome. It is therefore intuitive to "doctor" the distribution of an input parameter in a way that its mean is equal to our "best estimate" for that parameter (mostly an MLE estimator). We show that in case of a ratio measure, the justification of this method is that one implicitly chooses a prior distribution on the RR that has an average of 1 (which is conservative, but sensible), instead of the standard prior which has an average of zero on the log ratio scale, but an average of infinity on the ratio scale. The latter, in our opinion, does not make a sensible prior, as infinity is an unrealistic value.

For proportions, however, we do not advocate the approach of using a beta distribution with parameters equal to the number of positives (a) and negatives (b) in the data, despite its advantage of similarly forcing the distribution into having an expected value equal to the MLE. The reason is that it prescribes complete certainty in cases where zero positives or negatives are observed. Especially when events are rare and/or data sets are split up in many subpopulations, such situations are not uncommon.

Instead, we propose using a uniform prior distribution in those cases where one is sure that values of 0% and 100% are extremely unlikely, implying the use of a beta distribution with parameters $a + 1$ and $b + 1$. If values of 0% or 100% are likely, Jeffrey's prior is a good choice, implying a beta-distribution with parameters $a + 0.5$, $b + 0.5$.

Although these recommendations are in the spirit of BCS, they nevertheless differ subtly from the final recommendations made in BCS. Its relevance to the outcome, however, in many cases will be minimal. First, in all cases where data are ample, the influence of the type of prior that has been chosen will vanish. Second, the uncertainty of modeling is only partly due to parameter uncertainty [20], and parameter uncertainty in turn is only partly due to the type of uncertainty on which we have focused here. The estimates of uncertainty given by a statistical model only reflect the uncertainty for populations that are completely similar to the study population in which the data were observed. In reality, there often is considerable uncertainty on whether study results can be generalized to the population of interest in the health economic problem. Ignoring this uncertainty influences results of PSA more than the subtle changes because of following the recommendations given here. Nevertheless, the recommendations given here are just as simple and just as easy to implement, and so there is no real reason not to use them.

We restricted the topic of this article to two simple, but frequently occurring types of input parameters. We did not discuss topics as including uncertainty on the distributional form of the input parameters [18,19] or using more complex methods of Bayesian evidence synthesis [7], mostly using MCMC. In the latter case, the posterior joint distribution from these methods can directly be used as input for PSA.

Summarizing, considering the (implicit) priors used in constructing input distributions for PSA, we recommend using a lognormal distribution for ratios, with median $\exp(\log(RR) - 0.5*se^2)$ and standard error equal to se, and a beta distribution with parameters $a + 1$ and $b + 1$ for proportions. Only in cases where a real proportion of 0 or 1 is anticipated, a beta distribution with parameters $a + 0.5$ and $b + 0.5$ might be preferable.

Source of financial support: The work was carried out as part of the regular work of the institute (no external funding).

References

1 Buxton MJ, Drummond MJ, Van Hout BA, et al. Modelling in economic evaluation: an unavoidable fact of life. Health Econ 1997;6:217–27.
2 Claxton K. Exploring uncertainty in cost-effectiveness analysis. Pharmacoeconomics 2008;26:781–98.
3 Claxton K, Sculpher M, McCabe C, et al. Probabilistic sensitivity analysis for NICE technology assessment: not an optional extra. Health Econ 2005;14:339–47.
4 Claxton K, Sculpher M, Drummond M. A rational framework for decision making by the National Institute for Clinical Excellence (NICE). Lancet 2002;360:711–15.
5 Oakley JE., O'Hagan A. Probabilistic sensitivity analysis of complex models: a Bayesian approach. J R Stat Soc Series B 2004;66:751–69.
6 Claxton K. The irrelevance of inference: a decision-making approach to the stochastic evaluation of health care technologies. J Health Econ 1999;18:341–64.
7 O'Hagan A, McCabe C, Akehurst R, et al. Incorporation of uncertainty in health economic modelling studies. Pharmacoeconomics 2005;23:529–36.
8 Sutton AJ, Abrams KR. Bayesian methods in meta-analysis and evidence synthesis. Stat Methods Med Res 2001;10:277–303.
9 Briggs A, Claxton K, Sculpher M. Decision Modelling for Health Economic Evaluation. Oxford: Oxford University Press, 2006.
10 Sculpher M. Subgroups and heterogeneity in cost-effectiveness analysis. Pharmacoeconomics 2008;26:799–806.
11 Greenland S. Sensitivity analysis, Monte Carlo risk analysis, and Bayesian uncertainty assessment. Risk Anal 2001;21:579–83.
12 Ades AE, Claxton K, Sculpher M. Evidence synthesis, parameter correlation and probabilistic sensitivity analysis. Health Econ 2006;15:373–81.
13 Ades AE, Lu G. Correlations between parameters in risk models: estimation and propagation of uncertainty by Markov Chain Monte Carlo. Risk Anal 2003;23:1165–72.
14 Briggs AH, Ades AE, Price MJ. Probabilistic sensitivity analysis for decision trees with multiple branches: use of the Dirichlet distribution in a Bayesian framework. Med Decis Making 2003;23:341–50.
15 Spiegelhalter DJ, Abrams KR, Myles JP. Bayesian Approaches to Clinical Trials and Health-Carre Evaluation. Chichester: John Wiley & Sons, 2004.
16 Gelman A, Carlin JB, Stern HS, Rubin DB. Bayesian Data Analysis. London: Chapman & Hall/CRC Press, 2004.
17 Gandini S, Botteri E, Iodice S, et al. Tobacco smoking and cancer: a meta-analysis. Int J Cancer 2008;122:155–64.
18 Claxton K. Exploring uncertainty in cost-effectiveness analysis. Pharmacoeconomics 2008;26:781–98.
19 Chick SE. Input distribution selection for simulation experiments: accounting for input uncertainty. Operations Res 2001;49:744–58.
20 Walker WE, Harremoes P, Rotmans J, et al. Defining uncertainty: a conceptual basis for uncertainty management in model-based decision support. Integr Assess 2003;4:5–17.