Implementing the Bayesian paradigm in risk analysis

T. Aven and J. T. Kvaløy

November 23, 2002

Abstract

The Bayesian paradigm comprises a unified and consistent framework for analysing and expressing risk. Yet, we see rather few examples of applications where the full Bayesian setting has been adopted with specifications of priors of unknown parameters. In this paper we discuss some of the practical challenges of implementing Bayesian thinking and methods in risk analysis, emphasising the introduction of probability models and parameters and associated uncertainty assessments. We conclude that there is a need for a pragmatic view in order to "successfully" apply the Bayesian approach, such that we can do the assignments of some of the probabilities without adopting the somewhat sophisticated procedure of specifying prior distributions of parameters. A simple risk analysis example is presented to illustrate ideas.

1 Introduction

Most engineers and risk analysts have been trained in the "classical" approach to risk analysis, where probability exists independent of the analyst - as a quantity characterizing the system being studied. This concept of probability is relative frequency based and the results of the risk analyses provide estimates of these "true" probabilities. This view was challenged about 25 years ago, when the need to quantify risk from large technological systems was recognized and resources were expended to produce numerical results (Apostolakis et al. [3], p. 247). The quantification of risk requires the quantification of the likelihood of rare accidental events, which normally cannot be done without employing engineering judgment. It has become evident that the problems of risk quantification cannot be handled with the methods of traditional statistics.

To us, the alternative is a Bayesian approach, where the concept of probability is used as the analyst's *measure of uncertainty* or *degree of belief*. This alternative approach has, however, not been commonly accepted; there is still a lot of scepticism among risk analysts when speaking about subjective probabilities. Most risk analysts do in fact use some subjective methods when carrying out risk analyses. For example, subjective probabilities are commonly developed for the branches of the event trees. But a total adoption of the Bayesian theory is surprisingly not very often seen among risk analysts. Perhaps one reason for this is lack of understanding of what this approach really means. When looking at different papers in the field, it is not clear how we should in fact implement this theory in a risk analysis. We find the Bayesian literature very technical and theoretical. The literature is to large extent concerned about mathematical and statistical aspects of the Bayesian paradigm. The more practical challenges of adopting the Bayesian approach is much more seldomly addressed.

In this paper we use a simple risk analysis example as a starting point for a discussion on how to think when using a Bayesian approach to risk analysis and we search for practical solutions. We primarily have in mind standard situations where the risk analysis is used in a decision making context, for example in the planning phase of a project, and where we have some relevant background information.

The understanding of what a Bayesian analysis means, varies a lot among risk analysts. The situation is somewhat confusing, and there is a need for clarifications. The purpose of this paper is to contribute to such clarifications.

2 An illustrative example.

We consider a process plant. As a part of a risk analysis of the plant, a separate study is to be carried out of the risk associated with the operation of the control room that is placed in a compressor module. Two persons operate the control room. The purpose of the study is to assess risk to the operators as a result of possible fires and explosions in the module and to evaluate the effect of implementing risk reducing measures. Based on the study a decision will be made on whether to move the control room out of the module or to implement some risk reducing measures. The risk is currently considered to be too high, but the management is not sure what is the overall best arrangement taking into account both safety and economy.

The management decides to conduct a risk analysis to support decision making. To simplify, suppose the analysis is based on one event tree as shown in Figure 1 below. The tree models the possible occurrence of gas leakages in a process plant during a period of time, say one year. The number of gas leakages, referred to as the initiating events, is denoted X. If an initiating event I occurs, it leads to N number of fatalities, where $N = 2$ if the events A and B occur, $N = 1$ if the events A and not B occur, and $N = 0$ if the event A does not occur. We may think of the event A as representing ignition of the gas and B as explosion. The total number of fatalities is denoted Y.

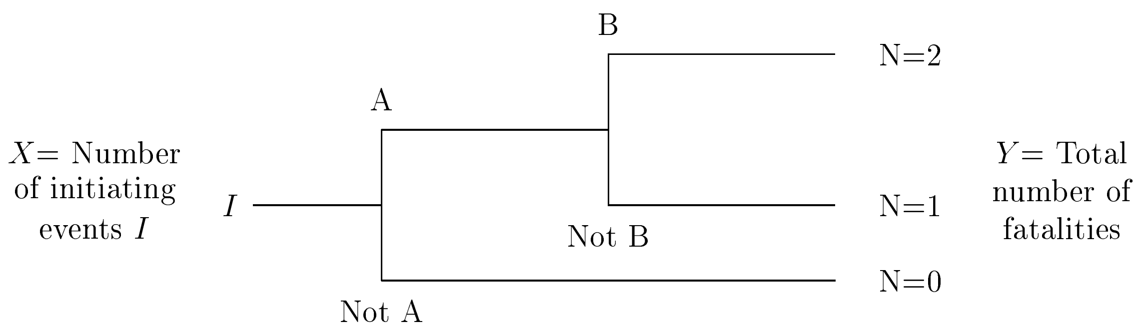

Figure 1: Event tree example

3 The prevailing Bayesian thinking

In this section we approach the problem presented in Section 2 using the prevailing "textbook" Bayesian thinking. We first present the calculations that will typically be performed, and then we discuss various interpretations of what the calculations really mean.

Let $\mathcal{B}$ denote the background information we have about the problem, i.e. all sorts of information which is relevant for the problem. This information might come from various sources, e.g. general information from similar situations, more or less relevant historical data from similar situations, expert judgments etc. The whole analysis is conditional on this background information.

We begin by specifying a probability model for the problem, consisting of probability distributions for random variables (quantities) and parameters. First consider the number of initiating events X. A typical choice is to use a Poisson distribution $p(x|\lambda)$, with parameter λ, where

$$p(x|\lambda) = \frac{\lambda^x}{x!}e^{-\lambda}.$$

Further a so called prior distribution $f(\lambda|\mathcal{B})$ for λ is specified, using for instance a gamma distribution. The gamma distribution is a mathematical convenient choice for prior distribution in this case by being a so called conjugate distribution. A prior distribution is called conjugate if it leads to a posterior distribution (posterior distributions are discussed later) in the same distribution class (in this case gamma), see for instance Bernardo & Smith [8] for more details.

Next we specify $\theta_1 = P(A|I)$ and $\theta_2 = P(B|A)$. Then

$$\begin{aligned}
P(N=2|I,\theta_1,\theta_2,\mathcal{B}) &= P(N=2|I,\theta_1,\theta_2) = \theta_1\theta_2 \\
P(N=1|I,\theta_1,\theta_2,\mathcal{B}) &= P(N=1|I,\theta_1,\theta_2) = \theta_1(1-\theta_2) \\
P(N=0|I,\theta_1,\theta_2,\mathcal{B}) &= P(N=0|I,\theta_1,\theta_2) = 1-\theta_1,
\end{aligned}$$

and prior distributions $f(\theta_1|\mathcal{B})$ and $f(\theta_2|\mathcal{B})$ are specified, using for instance beta distributions. Note that more generally a joint prior distribution $f(\lambda,\theta_1,\theta_2|\mathcal{B})$ of all parameters could be specified, but it is most common to use independent prior distributions, i.e. $f(\lambda,\theta_1,\theta_2|\mathcal{B}) = f(\lambda|\mathcal{B})f(\theta_1|\mathcal{B})f(\theta_2|\mathcal{B})$.

Further a conditional distribution of the total number of fatalities can be specified as

$$\begin{aligned}
P(Y=y|X=x,\theta_1,\theta_2,\lambda,\mathcal{B}) &= P(Y=y|X=x,\theta_1,\theta_2) \\
&= P(\sum_{i=1}^{x} N_i = y|X=x,\theta_1,\theta_2).
\end{aligned}$$

Calculating this probability is in principle straightforward, but for large x and y, tedious. The simplest cases are

$$P(Y=0|X=x,\theta_1,\theta_2) = (1-\theta_1)^x$$

$$\begin{aligned} P(Y=1|X=x,\theta_1,\theta_2) &= x(1-\theta_1)^{x-1}\theta_1(1-\theta_2) \\ P(Y=2|X=x,\theta_1,\theta_2) &= x(1-\theta_1)^{x-1}\theta_1\theta_2 + \frac{x(x-1)}{2}(1-\theta_1)^{x-2}\theta_1^2(1-\theta_2)^2 \end{aligned}$$

A conditional joint distribution of X and Y can now be specified as

$$p(x,y|\theta_1,\theta_2,\lambda,\mathcal{B}) = p(x,y|\theta_1,\theta_2,\lambda) = p(y|x,\theta_1,\theta_2)p(x|\lambda)$$

and unconditionally

$$p(x,y|\mathcal{B}) = \int_{\theta_1}\int_{\theta_2}\int_{\lambda} p(y|x,\theta_1,\theta_2)p(x|\lambda)f(\lambda,\theta_1,\theta_2|\mathcal{B})d\lambda d\theta_1 d\theta_2.$$

Finally the unconditional distribution of the total number of fatalities becomes

$$p(y|\mathcal{B}) = \sum_x \int_{\lambda}\int_{\theta_1}\int_{\theta_2} p(y|x,\theta_1,\theta_2)p(x|\lambda)f(\lambda,\theta_1,\theta_2|\mathcal{B})d\lambda d\theta_1 d\theta_2. \qquad (1)$$

If additional information in the form of observations $\mathbf{w} = (x_1, y_1, \ldots, x_n, y_n)$ of X and Y (not previously included in $\mathcal{B}$) is or becomes available, the likelihood function is given as

$$L(\theta_1,\theta_2,\lambda|\mathbf{w},\mathcal{B}) = \prod_{i=1}^{n} p(x_i,y_i|\theta_1,\theta_2,\lambda,\mathcal{B}),$$

or if data on number of fatalities for each initiating event is available, i.e. $\mathbf{w} = (n_1, \ldots, n_x)$, then the likelihood function becomes

$$L(\theta_1,\theta_2,\lambda|\mathbf{w},\mathcal{B}) = \prod_{i=1}^{x} P(N_i = n_i|\theta_1,\theta_2,I)P(X=x|\lambda).$$

The prior distributions can now be updated to posterior distributions by using Bayes theorem. For instance

$$f(\lambda|\mathbf{w},\mathcal{B}) \propto L(\theta_1,\theta_2,\lambda|\mathbf{w},\mathcal{B})f(\lambda|\mathcal{B})$$

where the constant of proportionality ensures that the posterior distribution is a proper distribution, i.e. that it integrates to one integrated over λ. The posterior distributions $f(\theta_1|\mathbf{w},\mathcal{B})$ and $f(\theta_2|\mathbf{w},\mathcal{B})$ are calculated similarly. θ_2 and their posterior distributions.

So far nothing has been said about interpretations. The calculations presented above are basically the calculations which have to be done according to most texts on Bayesian statistics, but the interpretation of parameters, distributions and what the calculations really means differs in the literature.

One interpretation which is often seen is what we will call the combined classical Bayesian approach, discussed for instance in Aven [5] and Aven & Rettedal [6], and used for instance by Martz & Waller [11], Sander & Badoux [13] and many others. In this approach there are assumed to exist true underlying probability distributions and true unknown parameters thought to represent some true underlying states of nature. In other words, a true underlying risk is assumed to exist. This is the classical part, which is in line with the classical relative frequency approach to risk analysis. Then subjective probabilities are used to express the analysts uncertainties regarding what the true value of the underlying probabilities and parameters are.

The combined classical Bayesian approach is also referred to as the "probability of frequency" framework or approach, see Apostolakis and Wu[2] and Kaplan[9]. In this framework the concept probability is used for the subjective probability and frequency for the "objective", relative frequency based probability.

Thus in the present example according to the probability of frequency approach, λ represents the true expected number of leakages while θ_1 and θ_2 represent true probabilities. Further, probabilities like $p(x|\lambda)$ and $p(x, y|\theta_1, \theta_2, \lambda)$ are interpreted as a probability distribution describing the true inherent, or aleatory, variation of random variables. The prior distributions $f(\lambda|\mathcal{B})$, $f(\theta_1|\mathcal{B})$ and $f(\theta_2|\mathcal{B})$ are subjective epistemic uncertainty distributions expressing the analysts uncertainty, or lack of knowledge, regarding the true value of the parameters λ, θ_1 and θ_2. Then finally $p(y|\mathcal{B})$ is a mix of aleatory and subjective epistemic uncertainties, and must thus basically be interpreted as a subjective distribution function. The importance of the aleatory versus subjective contributions will vary according to amount of "objective" data available, strength of background information etc. Similarly the posterior distribution will also be a mix of aleatory and subjective epistemic uncertainties, and much focus and interest will typically be on these posterior distributions.

The above interpretation can hardly be interpreted as proper Bayesian. In a proper Bayesian setting the terms "true probabilities" and "true risks" are meaningless. Let us see how the modern Bayesian theory, as it is presented in the literature, see e.g. Barlow[7], Bernardo and Smith[8], Lindley[10], Singpurwalla[14] and Singpurwalla and Wilson[15], would interpret the various elements of the set-up. As we will see from the brief summary below, the Bayesian thinking is in fact not that different from the probability of frequency analysis described above.

Probabilities are always conditioned on the background information, $\mathcal{B}$. To specify the probabilities related to a random quantity like X, a direct assignment, $P(X \leq x|\mathcal{B})$, could be used, based on everything we know. Since this knowledge is often complex, of high dimension, and much in $\mathcal{B}$ may be irrelevant to X, this approach is often replaced by the use of probability models, which is a way of abridging $\mathcal{B}$ so that it is manageable. A probability model, $p(x|\lambda)$, expresses the probability distribution of the unknown quantity X, given a parameter λ. This parameter λ is unknown, and our uncertainty related to its value is specified through a prior distribution $f(\lambda|\mathcal{B})$. According to the law of total probability

$$P(X \leq x|\mathcal{B}) = \int p(x|\lambda, \mathcal{B}) df(\lambda|\mathcal{B}) = \int p(x|\lambda) df(\lambda|\mathcal{B}),$$

where $p(x|\lambda, \mathcal{B}) = p(x|\lambda)$ implies that λ summarizes all the information we have about X in $\mathcal{B}$, or if we knew λ we would judge X to be independent of $\mathcal{B}$.

Thus to summarize, $p(x|\lambda)$, also called a probability model, is interpreted as a distribution reflecting our uncertainty regarding X if we knew λ. The interpretation is similar for distributions like $p(x, y|\theta_1, \theta_2, \lambda)$. Prior distributions like $f(\lambda|\mathcal{B})$, $f(\theta_1|\mathcal{B})$ and $f(\theta_2|\mathcal{B})$ are subjective epistemic uncertainty distributions expressing the analysts uncertainty, or lack of knowledge, regarding the value of the parameters λ, θ_1 and θ_2. Finally $p(y|\mathcal{B})$ is a subjective distribution expressed via probability models and prior distributions. If more data become available, the prior distribution is updated to the posterior distribution using Bayes theorem.

What about the parameters? The interpretation is that a parameter is the limit of a function of exchangeable observations. This interpretation follows from the so called representation theorem, see e.g. Bernardo & Smith [8], chap. 4. A sequence of observations is exchangeable if the joint distribution of the observations is invariant under permutations of the order of the observations. Thus in our example λ is interpreted as the limit of the average of observations of the number of leakages, while θ_1 and θ_2 are interpreted as limiting frequencies of 0-1 events.

Note that the Bayesian approach, as presented here, allows for fictional parameters, based on thought experiments. Such parameters are introduced and uncertainty of these assessed. Thus, from a practical point of view, an analyst would probably not see much difference between this framework and the probability of frequency approach described above. Of course, Bayesians would not speak about true, objective risk and probabilities, and the predictive form is seen as the most important one. However, in practice, Bayesian parametric analysis is often seen as an end-product of a statistical analysis. The use of and understanding of probability models gives focus on limiting values of quantities constructed through a thought experiment, which are very close to the mental constructions of probability and risk used in the classical relative frequency approach.

To us, the hypothetical population thinking in the above Bayesian interpretations is a blind alley. As in the present example, risk analyses of real life complex systems are typically conducted in the planning phase, and the situations considered are not repeated a number of times under similar conditions - the situation is rather unique. In other words, no actual population exist. Consider for instance the parameter λ in the example. This parameter is interpreted as the true expected number of leakages or as the limit of the average of the number of leakages, i.e., λ is interpreted as a parameter of a thought hypothetical population. The same is the case for θ_1 and θ_2 which are interpreted as true probabilities or limiting frequencies of 0-1 events. There is no way that we can accurately measure these quantities, and introducing them produce uncertainties. And what does it then mean to have uncertainty distributions on fictitious parameters like λ, θ_1 and θ_2?

The interpretation problems related to the hypothetical population thinking are common to both approaches discussed above. For the probability of frequency approach, problems related to model uncertainty also arise. The true risk numbers generated by the event tree with associated probabilities represents a model in this approach, and consequently its uncertainty should be addressed. But that is not obvious how to do, and in practice it is ignored. See Aven [4] and Nilsen & Aven [12] for a further discussion of this issue.

In addition to the problem related to interpretation, a serious problem with the probability of frequency approach and the full Bayesian approach is the complexity of the calculations. The presented example is an extremely simple case compared to most real life problems, still some of the calculations are complicated. Depending on the choice of prior distributions, calculation of for instance the mean of the posterior distributions could require Markov chain Monte Carlo (McMC) simulations, and calculation of (1) would certainly require McMC simulations and/or numerical methods. In more difficult problems the calculations involved becomes very complicated, McMC simulations are required and the calculations becomes very time consuming. In fact, a complete application of the standard Bayesian approach, without doing a number of simplifications, is seldom done in practice.

4 How to make the Bayesian thinking work

As we have discussed in Section 3, there are several problems with applying the Bayesian approach in practice, it is too complex and there are interpretational problems. A complete adoption of textbook Bayesian thinking is seldom done in practice. A rethinking is needed for "successfully" adopting the Bayesian approach. Our starting point is that we want to establish a framework based on subjective probabilities that works in practice and has a clear interpretation and focus, i.e. an approach which is easy to use in communication and decision-making.

The first problem pointed out in the discussion of the Bayesian approach in Section 3 is that focus tends to be on fictitious parameters of the problem like λ, θ_1 and θ_2, and inference on these. Focus should primary be on the actual outcomes, i.e. the observable quantities, such as X and Y in the example, representing the number of leakages and the number of fatalities. These quantities should be predicted and associated uncertainties assessed.

The previous section identified an interpretational problem when adopting the prevailing Bayesian thinking, related to the use of parametric probability models, parameters and the use of hypothetical population thinking. We believe that there is a need for a new look on how to approach the analysis on these points. Let us start from scratch. Probabilities are subjective. They are measures of our personal uncertainties or beliefs. There is only one type of uncertainty, stemming from lack of knowledge about the future value of observable quantities, i.e. the uncertainty is epistemic. Further, hypothetical populations and related parameters should not be introduced.

To be more specific, let us again consider the event tree example. The focus is on Y, the number of fatalities. To predict this number and to assess uncertainties, we develop a deterministic model, which is the event tree model presented in Figure 1. The aim of the modeling is to get insight into the uncertainties and reduce uncertainties. Given the model, the remaining uncertainties are related to the observable quantities X, A and B. The next step is then to assess uncertainties of these quantities. Let us look at X first.

We would like to predict X and assess uncertainties. How should we do this? We distinguish between two cases:

1. Data from situations "similar" to the one analysed, are available, and let us assume for the sake of simplicity that these are of the form $x_1, x_2, ..., x_n$, where x_i is the number of initiating events during one year. These data are considered relevant for the situation being studied.

2. We have no relevant data available for this situation.

First, let us consider case 1. The data allow a prediction simply by using the mean $\overline{x}$ of the observations $x_1, x_2, ..., x_n$. But what about uncertainty in this prediction? How should we express uncertainty related to X and the prediction of X? Suppose the observations $x_1, x_2, ..., x_n$ are 1,1,2,0,1, so that $n = 5$ and the observed mean is equal to 1. In this case we have rather strong background information, and we suggest to use the Poisson distribution with mean 1 as our uncertainty distribution of X. How can this uncertainty distribution be "justified"? Well, if this distribution reflects our uncertainty about X, it *is* justified, and there is nothing more to say. This is a subjective probability distribution and there is no need for further justification. But is a Poisson distribution with mean 1 "reasonable", given the background information? We note that this distribution has a variance not larger than 1. By using this distribution, 99% of the mass is on values less than 4.

Adopting the prevailing Bayesian thinking, as outlined above, using the Poisson distribution with mean 1, means that we have no uncertainty about the parameter λ, which is interpreted as the long run average number of failures when considering an infinite number of exchangeable random quantities, representing similar systems as the one being analyzed. According to the Bayesian theory, ignoring the uncertainty about λ, gives misleading over-precise inference statements about X, *cf. e.g.* Bernardo and Smith[8], p. 483. This reasoning is of course valid if we work within the standard Bayesian setting, considering an infinite number of exchangeable random quantities. In our case, however, we just have one X, so what do we gain by making a reference to limiting quantities of a sequence of similar hypothetical Xs? The point is that given the observations $x_1, x_2, ..., x_5$, the choice of the Poisson distribution with mean 1, is in fact reasonable. Consider the following argumentation.

Suppose that we divide the year $[0, T]$ into time periods of length T/k, where k is for example 1000. Then we may ignore the possibility of having two events occurring in one time period, and we assign an event probability of $1/k$ for the first time period, as we predict one event in the whole interval $[0, T]$. Suppose that we have observations related to $i-1$ time periods. Then for the next time period we should take these observations into account—using independence means ignoring available information. A natural way of balancing the prior information and the observations is to assign an event probability of $(d_i + 1 \cdot n)/((i-1) + nk)$, where d_i is equal to the total number of events that occurred in $[0, T(i-1)/k]$; i.e., we assign a probability which is equal to the total number of events occurred per unit of time. It turns out that this assignment process gives an approximate Poisson distribution for X. This can be shown for example by using Monte Carlo simulation. The Poisson distribution is justified as long as the background information dominates the uncertainty assessment of the number of events occurring in a time period. Thus from a practical point of view, there is no problem in using the Poisson distribution with mean 1. The above reasoning provides a "justification" of the Poisson distribution, even with not more than one or two years of observations.

Alternatively, the Poisson approximation follows by studying the predictive distribution of X in a full Bayesian analysis, assuming that $x_1, x_2, ..., x_5$ are observations coming from a Poisson distribution, given the mean λ and using a suitable (for example a non-informative) prior distribution on λ. Restricting attention to observable quantities only, a procedure as specified in Barlow[7], Chapter 3, can be used. This procedure, in which the multinomial distribution is used to establish the Poisson distribution, is based on exact calculation of the conditional probability distribution of the number of events in subintervals, given the observed number of events for the whole interval.

Note that for the direct assignment procedure using the k time periods, the observations $x_1, x_2, ..., x_5$ are considered a part of the background information, meaning that this procedure does not involve any modeling of these data. In contrast, the more standard Bayesian approach requires that we model $x_1, x_2, ..., x_5$ as observations coming from a Poisson distribution, given the mean λ.

We conclude that a Poisson distribution with mean 1 can be used to describe the analyst's uncertainty with respect to X in this case. The background information is sufficiently strong.

Now consider the case 2 with no historical data. Then we will probably find the direct use of the Poisson distribution as described above to have too small variance. The natural approach is then to implement a full parametric Bayesian procedure with specification of prior distributions for unknown parameters. But how should we interpret the various elements of the set-up? Should we speak about λ having a true value (the limit of an infinite sequence of exchangeable random quantities), introduce an uncertainty distribution over λ, and refer to the Poisson distribution as a model?

No, as long as λ is fictional, not a state of the world (the nature) and thus not observable, a true value of λ does not exist and the Poisson distribution is not a representation of the world. Instead we suggest the following interpretation.

The Poisson probability distribution $p(x|\lambda)$ is a candidate for our subjective probability for the event $X = x$, and $f(\lambda|\mathcal{B})$ is a confidence measure, reflecting for a given value of λ, the confidence we have in $p(x|\lambda)$ for being able to predict X. If we have several X_is, similar to X, and λ is our choice, we believe that about $p(x|\lambda) \cdot 100\%$ of the X_is will take a value equal to x, and $f(\lambda|\mathcal{B})$ reflects for a given value of λ, the confidence we have in $p(x|\lambda)$ for being able to predict the number of X_is taking the value x. Following this interpretation, note that $p(x|\lambda)$ is not a model, and $f(\lambda|\mathcal{B})$ is not an uncertainty measure. We refer to this as the confidence interpretation.

The above analysis is a tool for predicting X and assessing associated uncertainties. When we have little data available, modeling is required to get insight into the uncertainty related to X and hopefully reduce the uncertainty. The modeling also makes it possible to see the effects of changes in the system and to identify risk contributors.

We now turn to how to assess uncertainties for A and B. For these events we just need to assign two probabilities, $\theta_1 = P(A|I, \mathcal{B})$, expressing our uncertainty related to occurrence of ignition, and $\theta_2 = P(B|A, \mathcal{B})$, expressing our uncertainty regarding an explosion given an ignition. The basis for the probability assignments would be "hard" data and expert opinions. These probabilities are not "true underlying probabilities" or limiting frequencies of 0-1 events, they just represent our subjective uncertainties regarding the observable events A and B, expressed as probabilities. This is a difference from what is common in Bayesian analyses where prior distributions expressing the uncertainties regarding the "true values" of θ_1 and θ_2 are usually specified. Why introduce such hypothetical limiting quantities and associated prior distributions, when we can easily assess our uncertainties regarding what would happen by the single numbers θ_1 and θ_2?

What now remains is using probability calculus for calculating the predictive uncertainty distribution for the total number of fatalities Y. This distribution is now straightforwardly calculated as

$$p(y|\mathcal{B}) = \sum_x p(y|x, \theta_1, \theta_2)p(x|\mathcal{B}), \qquad (2)$$

where $p(x|\mathcal{B})$ is the uncertainty distribution of X. Comparing (2) with (1) we see that the computational simplifications by not having prior distributions on θ_1 and θ_2 are considerable.

So the end product of the analysis is simply the predictive uncertainty distribution $p(y|\mathcal{B})$ in (2), expressing our uncertainty regarding the future value of Y. There are no further "uncertainties of uncertainties" or similar confusing notions. The uncertainty distribution regarding the "top level" quantity, here Y, is calculated by first focusing on the observable quantities on a more "detailed level", in this case X, and A and B, establishing uncertainty distributions for these, and then using probability calculus to propagate this into an uncertainty distribution for the "top level" quantity Y.

5 Summary and final remarks

To summarize, the key principles we recommend should be followed to make the Bayesian thinking work in practice in risk analyses are the following: Focus is on so-called observable quantities, that is, quantities expressing states of the "world" or the nature, that are unknown at the time of the analysis but will (or could) become known in the future; these quantities are predicted in the risk analysis and probability is used as a measure of uncertainty related to the true values of these quantities. The emphasis of these principles gives a framework which is easy to understand and use in a decision-making context. As for the practical calculations, compared to the common approaches we see no need for "prior distributions" on quantities like θ_1 and θ_2 representing uncertainties. If we have strong background information we also suggest that "priors" could be skipped for parameters like λ, otherwise a distribution $f(\lambda|\mathcal{B})$ expressing the confidence we have in distribution functions with different values of λ for expressing our uncertainty could be introduced.

Regardless of the approach taken, sensitivity analyses are useful, showing how sensitive the results are to variations in input quantities, for example as a result of changes in assumptions and suppositions. Such analyses are also useful for identifying potential areas for improvement.

Of course, if an infinite (or large) population of similar situations can be defined, then the parameters represent a state of nature they are observable quantities and we can speak about our uncertainty of these quantities. The question is whether such a population of similar situations should be introduced. In our view, as a general rule, the analyst should avoid introducing fictitious populations. In the example considered in this paper, it is not obvious how to define an infinite or large population of similar situations, and without a precise understanding of the population, the uncertainty assessments become difficult to perform and it introduces an element of arbitrariness. What is a fictitious population and what is a real population is a matter for the analyst to decide, but the essential point we are making here is that the analyst should think first before he/she introduces such a population. The full Bayesian set-up, covering the introduction of a parametric distribution class, specification of a prior distribution for the parameters, Bayesian updating to establish the posterior distribution, and the calculation of the predictive distributions, is a useful tool for coherent analysis of statistical data, but should not be used when not needed.

The presented predictive Bayesian approach to risk analysis has been discussed in other contexts in for instance Aven [5], Aven & Rettedal [6], Apeland et al. [1].

We have used a standard process risk analysis example to illustrate our points. Other examples, including examples from Structural Reliability Analyses, are presented in Aven [4], see also Aven and Rettedal [6].

Acknowledgements

The authors are grateful to Bjørn Auestad and a referee for valuable comments and suggestions.

References

[1] Apeland, S., Nilsen, T., Aven, T. 2001. Quantifying uncertainty under a predictive, fully subjective approach to risk analysis. *Reliability Engineering and System Safety*, **75**, 93-102.

[2] Apostolakis, G., Wu, J.S. 1993. In *Reliability and Decision Making*, Eds. Barlow, R.E, Clarotti, C.A. 311-322. London: Chapman & Hall.

[3] Apostolakis, G.E. et al. (eds.). 1988. *Reliability Engineering and System Safety, The Interpretation of Probability in Probabilistic Safety Assessments* **23**.

[4] Aven, T. 2002 *How to Approch Risk and Uncertainty to Support Decision-Making*, New York: Wiley. To appear.

[5] Aven, T. 2000. Risk analysis - A tool for expressing and communicating uncertainty. In Proceedings of *European Safety and Reliability Conference (ESREL)* 21-28. Edinburgh.

[6] Aven, T., Rettedal, W. 1998. Bayesian frameworks for integrating QRA and SRA methods. *Structural Safety* **20** 155-165.

[7] Barlow, R.E. 1998. *Engineering Reliability.* Philadelphia: SIAM.

[8] Bernardo, J.M., Smith, A. 1994. *Bayesian Theory.* Chichester: Wiley & Sons.

[9] Kaplan, S. 1992. Formalisms for handling phenomenological uncertainties: the concepts of probability, frequency, variability, and probability of frequency. *Nuclear Technology*, **102**, 137-142.

[10] Lindley, D.V. 2000. The philosophy of statistics. *The Statistician*, **49**, 293-337.

[11] Martz, H.F., Waller, R.A. 1982. *Bayesian Reliability Analysis.* New York: Wiley & Sons.

[12] Nilsen, T., Aven, T. 2001. Models and model uncertainty in the context of risk analysis. Sumitted for publication in *Reliability Engineering and System Safety.*

[13] Sander, P., Badoux, R. (eds.). 1991. *Bayesian Methods in Reliability.* Dordrecht: Kluwer Academic Publishers.

[14] Singpurwalla, N.D. 1988. Foundational issues in reliability and risk analysis. *SIAM Review*, **30**, 264-282.

[15] Singpurwalla, N.D., Wilson, S.P. 1999. *Statistical Methods in Software Engineering*. New York: Springer Verlag.

[16] Winkler, R.L. 1996. Uncertainty in probabilistic risk assessment. *Reliability Engineering and System Safety* **54**, 127-132.