

Some useful features of the Bayesian setting while dealing with uncertainties in industrial practice

A. Pasanisi, E. de Rocquigny & N. Bousquet

EDF-R&D, Management of Industrial Risks, 6 Quai Watier, 78401 Chatou Cedex, France

E. Parent

ENGREF-AgroParisTech, 19 Avenue du Maine, 75732 Paris Cedex 15, France

ABSTRACT : In this paper, we put into evidence some particular features of the Bayesian paradigm that prove valuable to take into account uncertainty within industrial studies. First, Bayesian setting is in essence a level-2 probabilistic approach. The predictive distribution of variables of interest, integrating all possible uncertainty sources, can thus naturally be obtained as the marginal distribution of the joint probability law of parameters and observable variables. Secondly, in industrial practice sometimes few data or measurements are available for setting inputs' probability distributions but, on the other hand, engineers expertise is often available and in the Bayesian framework, it can be properly taken into account by a prior distribution. Moreover, the subjective interpretation of probabilities and credibility intervals underlying the Bayesian settings proves more flexible and possibly easier to communicate to non-specialists such as industrial engineers and decision-makers than the frequentist interpretation. These arguments will be illustrated by two exemplary case studies.

1 INTRODUCTION

Uncertainty treatment in physical, environmental or risk modeling is the subject of a long-standing theoretical literature rooted in fundamental statistical and economical thinking (Knight 1921, Savage 1954), and then developed in the risk and environmental assessment fields (Beck 1987, Granger Morgan & Henrion 1990, Helton 1993, Hamby 1994, Pathé-Cornell 1996). It recently gained large-scale industrial importance as a number of applications gradually included some uncertainty treatment in particular : (i) nuclear safety studies involving large scientific computing (thermohydraulics, mechanics, neutronics etc.) or Probabilistic Safety Analyses (PSA); (ii) advanced design in the aerospace, automotive or more generally mechanical industries ; (iii) oil exploration and underground waste disposal control ; (iv) environmental and natural risk studies. The ESReDA European network of industries and academics undertook a large review and methodological research resulting in a consistent generic framework applicable to many industrial studies (de Rocquigny et al. 2008).

For most of those applications, at least partially *probabilistic* modeling of the uncertainties is considered,

although deterministic uncertainty treatment or more elaborate non-probabilistic approaches keep their share. Then, apart from the important uncertainty propagation issues in the context of often complex and high-CPU demanding scientific computing, one of the key issues regards the quantification of the sources of uncertainties, i.e. choosing in an accountable manner statistical models for the input variables such as uncertain physical properties of the materials or industrial process or natural aleatory phenomena (wind, flood, temperature, etc.).

Two key issues are traditionally encountered at this stage : (a) how to deal with the highly-limited sampling information directly available on uncertain input variables in real-world industrial cases ; (b) how to account for quite different natures of uncertainties, such as the traditionally-distinguished intrinsic aleatory uncertainties and reducible epistemic uncertainties. Indeed, a long debate (Apostolakis 1990, Helton 1993, Pathé-Cornell 1996, Helton & Oberkampf 2004) originating from the first large risk assessments in the US nuclear domain has ended up with the importance of distinguishing two types of uncertainty : namely the epistemic (or reducible, lack of knowledge, by ignorance) type referring to

uncertainty that decreases with the injection of more data, modeling or number of runs and the aleatory (or irreducible, intrinsic, variability, by essence) type for which there is a variation of the true characteristics that may not be reduced by the increase of data or knowledge.

The classical backup to issue (a) is to involve expertise and choose directly the uncertainty distributions in such a more or less formalized way. A few historical examples of large-scale nuclear waste or environmental impact assessments did imply a structured elicitation of expertise (Granger Morgan & Henrion 1990, Cooke 2001) including prior training and calibration steps in order to retrieve the expertise in the form of probabilistic distributions, quantiles etc. and possibly organize posterior consensus.

The present paper will put forward the point that a Bayesian framework may be viewed as rather natural for tackling issues (a) and (b) altogether. Indeed, beyond the forceful epistemological and decision-theory features of a Bayesian approach, it includes by definition a double-level probabilistic model separating "epistemic" and "aleatory" components and offers a traceable process to mix the encoding of engineering expertise inside priors and the observations inside an updated epistemic layer that proves mathematically consistent even when dealing with very low-size samples.

Explicit Bayesian treatments in the industrial applications are already well-documented since the 1990s in reliability analysis of mostly lifetime distributions (Lannoy & Procaccia 1994) or more generally in the field of risk analysis, as reviewed for instance by Aven (2003). In other fields, such as physical or environmental modeling, practice remains scarce or more or less informal although a large potential may be anticipated. The remaining of the paper is structured as follows. Section 2 recalls the modeling structure of industrial uncertainty studies. Section 3 reminds some peculiarities of Bayesian statistical modeling. Sections 4 and 5 deals with two case studies : predicting the lifetime of an industrial component and designing a flood protection work. In Section 6 we sketch some conclusions and state again our point of view.

2 MODELLING FRAMEWORK

Quantitative uncertainty assessment in the industrial practice typically involves (de Rocquigny et al 2008) :

- a preexisting physical or industrial system or component lying at the heart of the study, that is represented by a preexisting model $G(\cdot)$,
- inputs, affected by various sources of uncertainty,

- an amount of data and expertise available to calibrate the model and/or assess its uncertainty,
- industrial stakes that motivate the uncertainty assessment more or less explicitly (safety and security, environmental control etc).

It is useful, in practice to split the inputs into two vectors \mathbf{x} and \mathbf{d} . Some input variables (noted \mathbf{x}) are uncertain, subject to randomness, variability, lack of knowledge, errors or any sources of uncertainty, while other inputs (noted \mathbf{d}) are fixed, either being well known or being considered to be of secondary importance with respect to the vector of output variables of interest \mathbf{z} :

$$\mathbf{z} = G(\mathbf{x}, \mathbf{d}) \quad (1)$$

Note that computation of the preexisting model for a point value (\mathbf{x}, \mathbf{d}) (i.e. not uncertain at all) may require a very variable CPU time : depending on the case, from 10^{-4} s to several days for a single run.

Within the potentially large number of raw outputs of the models, it is useful to subdistinguish the model output "variables of interest" that are eventually important for the decision criteria.

While literature has discussed to a large extent the pros and cons of various probabilistic or non-probabilistic paradigms (Helton & Oberkampf, 2004), it is assumed in this paper that probabilistic uncertainty treatment is acceptable for the cases considered. Decision criteria will generally involve the consideration of some peculiar "quantities of interest" on the distribution of \mathbf{Z} (e.g. variance, confidence intervals, quantiles).

Formally, in a probabilistic setting, the uncertain inputs are modeled as a random vector :

$$\mathbf{X} \sim f_X(\mathbf{x}|\boldsymbol{\theta}_X), \quad (2)$$

whereby f_X indicates a joint probability distribution function (pdf) associated to the vector of uncertain inputs \mathbf{X} (capital letters denoting random variables), defining the *aleatory* uncertainty model. Parameters of f_X are grouped within the vector $\boldsymbol{\theta}_X$.

An additional uncertainty model is built to represent the *epistemic* uncertainty (level-2 or lack of knowledge) in the precise values of parameters $\boldsymbol{\theta}_X$.

$$\boldsymbol{\theta}_X \sim \pi(\boldsymbol{\theta}_X|\zeta). \quad (3)$$

In other words, the random vector \mathbf{X} is defined to follow $f_X(\mathbf{x}|\boldsymbol{\theta}_X)$ conditionally to the vector of parameters $\boldsymbol{\theta}_X$. The vector of parameters itself is supposed random, following a known joint pdf $\pi(\boldsymbol{\theta}_X|\zeta)$.

Finally, in some cases, the output itself \mathbf{Z} may not be directly available without error. An additional observation layer is to be introduced within the previous setting.

$$\mathbf{Z}_{\text{obs}} \sim \pi(\mathbf{z}_{\text{obs}}|\boldsymbol{\eta}, \mathbf{z}). \quad (4)$$

For instance, a "true" avalanche volume Z , computed from the transfer model $G(\mathbf{x}, \mathbf{d})$, with \mathbf{x} denoting the avalanche uncertain inputs such as snow height (aleatory), the moisture degree of snow and the friction coefficient of the path (epistemic), can only be known through some observation Z_{obs} (estimation of the runoff volume by the rule of the thumb) with, say, an error variance η . As another example, the random error term between \mathbf{Z} and \mathbf{Z}_{obs} may be due to a simplified version of the function $G(\cdot)$ or to numerical approximations or round-off errors.

3 UNCERTAINTIES AND BAYESIAN SETTING

3.1 The Bayesian setting as a level-2 approach

Aven (2003), among the main authors in the area of risk analysis, gave a sharp review of classical and Bayesian approaches for dealing with uncertainties. The classical frequentist approach is shown to be unsatisfying. Indeed, probabilistic results based on this approach are considered similar to frequencies of events (for instance, the event $\mathbf{Z} > z_0$). To be valid, these frequencies should theoretically be defined with respect to a studied population of unlimited size. In industrial practice, this condition seems somewhat problematic.

Such a fictional notion of probability should be, for Aven (2003), replaced by a simple subjective characterization : following Lindley (1985), "probability is a subjective measure of uncertainty". Assuming this point, Aven considered the Bayesian alternative to the classical approach, where $(\mathbf{X}, \theta_{\mathbf{X}})$ are both perceived as random variables, as a valuable way of propagating *aleatory* (through the choice of a model in \mathbf{X} knowing $\theta_{\mathbf{X}}$) and *epistemic* uncertainties (through the choice of a model in $\theta_{\mathbf{X}}$). These last *epistemic* uncertainties can be reduced (by definition) through the Bayes theorem. However, Aven (2003) pointed :

- the need for the practitioner to keep in mind the importance of focusing on observable (or at least understandable) quantities : \mathbf{X} and not $\theta_{\mathbf{X}}$. Thus, by the expert point of view, the observable uncertainty is often yielded by the prior *predictive* cumulative distribution function (cdf) of \mathbf{X} (which is generally a unidimensional variable X when eliciting expertise) :

$$P(X < x) = \int P(X < x | \theta_{\mathbf{X}}) \pi_0(\theta_{\mathbf{X}}) d\theta_{\mathbf{X}}, \quad (5)$$

where $\pi_0(\theta_{\mathbf{X}})$ is the prior distribution of $\theta_{\mathbf{X}}$, which formally encodes the expert knowledge.

- the need for reducing subjectivity in $\pi_0(\theta_{\mathbf{X}})$, for instance through formal and systematic elicitation methods.

In our opinion, the most interesting feature of the Bayesian setting, while dealing with uncertainties, is that it is in essence a level-2 probabilistic approach.

Actually, Bayesian data analysis is, above all, based on setting up a full probability model, i.e. a joint probability distribution for all observable and unobservable variables of the problem (Gelman et al. 2004). Given that we have observed a sample \mathbf{x}_S , more refined bets than $\pi_0(\theta_{\mathbf{X}})$ can be taken about $\theta_{\mathbf{X}}$. Namely, $\pi(\theta_{\mathbf{X}} | \mathbf{x}_S)$ is the posterior distribution of $\theta_{\mathbf{X}}$ given by Bayes formula :

$$\pi(\theta_{\mathbf{X}} | \mathbf{x}_S) \propto \pi_0(\theta_{\mathbf{X}}) \cdot f_X(\mathbf{x}_S | \theta_{\mathbf{X}}). \quad (6)$$

With \mathbf{x}_S at hand, the uncertainty about \mathbf{X} will be subsequently quantified by the posterior predictive distribution. For instance, in case of unidimensional variable, Equation 5 is updated as :

$$P(X < x | x_S) = \int P(X < x | \theta_{\mathbf{X}}) \pi(\theta_{\mathbf{X}} | x_S) d\theta_{\mathbf{X}}. \quad (7)$$

The change of knowledge about future \mathbf{x} brought by the observation of the past x_S will, in turn, impact the distribution of the output variable \mathbf{Z} via the deterministic transfer function given by Equation 1. Girard & Parent (2004) particularly insist on the idea that the Bayesian analyst should focus on inferring on the posterior predictive distribution of observable variables rather than on model's parameters θ . As already pointed by Box (1980), parameters estimation is just the first step of the statistician's work (inductive phase) which must be followed by the deductive phase of statistical analysis, i.e. coming back from the conceptual world of models to the real problems.

3.2 Bayesian inference as sequential learning

A very powerful idea behind Bayesian inference is that statistical inference is simply updating a previous knowledge, assessed by a prior distribution. The obtained posterior distribution, which encodes the current state of knowledge, can be sequentially updated by adding more and more data. To exemplify this idea, let us consider a very simple problem : the Bayesian assessment of an exponential lifetime model, of failure rate λ :

$$f(t | \lambda) = \lambda \exp(-\lambda t). \quad (8)$$

We suppose that available data are in the form $S_1 = \{t_1, \dots, t_{n_1}, \delta_1, \dots, \delta_{n_1}\}$, where δ_i is one if t_i is a failure time or zero if t_i is right-censored. In this case, a very well known result states that sufficient statistics for the data set S_1 are :

$$a_1 = \sum_{i=1}^{n_1} \delta_i, \quad b_1 = \sum_{i=1}^{n_1} t_i, \quad (9)$$

i.e. the number a_1 of observed failures and the sum b_1 of observed operating times prior to failure. A conjun-

gate prior for the exponential model is a gamma distribution $\text{Ga}(a, b)$ for λ such that :

$$\pi_0(\lambda|a, b) = \frac{b^a}{\Gamma(a)} \lambda^{a-1} \exp(-b\lambda). \quad (10)$$

By applying Bayes rule, the posterior distribution of λ is still a gamma distribution, the parameters of which are $a + a_1$ and $b + b_1$ respectively. It has to be noticed that the prior parameters a and b , which play the same role of a_1 and b_1 respectively, can be interpreted as sufficient statistics of a fictitious initial data set. If a new data set S_2 is available, the new posterior of λ is a gamma distribution of parameters $a + a_1 + a_2$ and $b + b_1 + b_2$ and so on. The previous posterior $\text{Ga}(a + a_1, b + b_1)$ can be interpreted as the prior distribution of λ , before the observation of data set S_2 . The sequence of the successive posteriors shows the enhancement of the knowledge, by adding more and more data, as shown in the Figure 1. In this example, the prior knowledge about the failure rate is modeled by a gamma distribution with $a = 14$ and $b = 2.81E + 5$ and the two available data sets have the following sufficient statistics :

Data set	Number of failure	Sum of obs.times
S_1	12	344,000 hours
S_2	21	400,000 hours

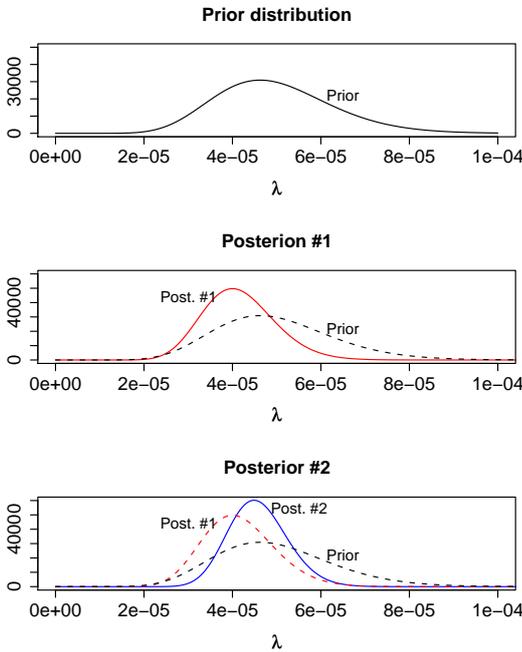


FIGURE 1: Sequential Inference on a life-time model.

More generally, one can quantify the enhancement in terms of posterior variance of $\theta_{\mathbf{X}}$ when updating the prior knowledge. Indeed, Gelman et al. (2004) pointed the fact that the posterior variance $V[\theta_{\mathbf{X}}|\mathbf{X}]$ is such that :

$$E[V[\theta_{\mathbf{X}}|\mathbf{X}]] = V[\theta_{\mathbf{X}}] - V[E[\theta_{\mathbf{X}}|\mathbf{X}]]. \quad (11)$$

Namely, on average on \mathbf{X} , the posterior variance is smaller than the prior variance. In other terms, the formula above states that posterior epistemic uncertainty decreases when adding data to prior knowledge.

3.3 Easily understandable results

Credibility intervals help to sum up posterior distribution of unknowns and posterior predictive distribution of observables. For instance, the 90% credibility interval for the failure rate in the previous example is the shortest interval containing 90% of the gamma distribution $\text{Ga}(a + a_i, b + b_i)$. Relying on Equation (7), a credibility interval for the future failure time t , or any function of it, can also be derived. Credibility intervals give probabilistic judgments about the unknowns. Their frequentist counterparts, i.e. confidence intervals, do not, although most people unaware of statistical jargon fallaciously believe they do (Leconte 1997)! Confidence intervals state reliability properties of the statistical procedure used to recover the unknown, not the unknown itself : for instance, changing the max-likelihood estimation technique for the method of moments will change the bounds of the frequentist interval !

An other difference between Bayesians and frequentists lies in the theory of tests : in the Bayesian setting, testing hypotheses via Bayes factor (Raftery 1995) is naturally encompassed in the general theory of statistical decision (Berger 1985), which seems very natural for engineering oriented audiences. And finally, the result of the Bayesian test can be simply expressed in terms of probability that a given hypothesis, between many others, is the best one, instead of refusing or not a null hypothesis. The bridge between the two approaches (and the point where they break apart) can be found in (Wald 1950).

3.4 Taking into account expert advice

Incorporating expert opinion through formal prior representation is a difficult task. General questioning methods have been reviewed by O'Hagan et al. (2006) and lay as well on psychological aspects as on past quantifiable experience from experts. In general, prior modeling, as in Garthwaite & Dickey (1996), often remains highly depending on the context and the nature of expert information, which can rather express his or her beliefs about a function of parameters, or unconditionally to the parametric model that he or she can know or not (usually not).

The main issue is taking into account the irreducible uncertainty emanating from subjective opinions, in terms clearer than "pure" statistical indicators like standard deviations (thus, easier to calibrate). The exponential example in section 3.2 shows that one could simply consider the expert as a furnisher of "old" data, which are only known through exhaustive statistics.

In this case, calibrating the prior would be similar to select the size of these data. Such an approach, focalized on "understandable" indicators of uncertainty, has increased in applied works in the recent years (Morita et al. 2007). Another issue is certainly the elicitation of a unique prior parametric structure itself among the variety of priors which can fit the experts' requirements. Conjugate models, like exponential-gamma in Section 3.2, appear quite convenient. However it also appears too restrictive since the class of models for which conjugate structures can be found for any sample size merely coincides with the traditional family of exponential distributions. In Kass and Wasserman (1996), various formal ways are presented, including famous maximal entropy methods, which however can lead to priors with defects (like strong dissymmetry).

According to the Bayesian version of the central limit theorem (see Berger 1985, Chapter 4, p. 224), the posterior distribution turns to a Normal pdf when the sample size gets large, providing that the prior distribution does not exclude any value of the parameter. The mean of that asymptotic approximation is the maximum likelihood estimate and its variance is the opposite of the inverse of the Fisher information matrix. In practice, this is to be interpreted that the information conveyed by the data overweights prior beliefs for large sample sizes : experts with different priors would then no longer disagree and frequentist maximum likelihood and Bayesian analyses would come to the same end.

However, prior-data conflicts must be avoided or, at least, detected and explained before the inference. Such conflicts can appear when prior and data favour areas of the parameters space far from each other, or when the subjective source (expert) becomes largely predominant. Dedicated criteria have been recently proposed by Evans and Moshonov (2006) then Bousquet (2008). Of course, many scientist do not feel at ease with the concept of prior distribution. Although its mathematical existence is grounded of five axioms of rational behavior (Savage, 1954), one can object that the man in the street does not obey this strict axiomatic corpus (Munier & Parent, 1995) or that encoding expert knowledge into probability distribution looks like an impossible mission.

However, not taking into account expert advice is against common good sense and modeling prior knowledge is part of the very general tricky task of modeling, i.e a never ending compromise between simplicity and realism.

4 CASE STUDY 1 : PREDICTING THE LIFETIME OF AN INDUSTRIAL COMPONENT

The first case study concerns the prediction of the lifetime T of an industrial component. The uncer-

tainty over T is modeled, as usual in industrial reliability problems, by a Weibull model $W(\eta, \beta)$, where η and β are the scale and the shape parameters respectively :

$$f(t|\eta, \beta) = \frac{\beta}{\eta} \left(\frac{t}{\eta}\right)^{\beta-1} \exp\left(-\left(\frac{t}{\eta}\right)^\beta\right). \quad (12)$$

Note that this model turns to be the exponential one (see Section 3.2) when $\beta = 1$. Many variables of interests can be imagined, depending on the underlying industrial problem. Sometimes, the variable of interest is the lifetime itself. For instance, one can be interested in evaluating the Mean Time To Failure (MTTF) or the probability that lifetime T is lower than a given value t^* . Other typical variables of interest can be the cost of a given age-based preventive maintenance strategy or the cost of a warranty post-sale policy. An interesting review of decision models involving Weibull models is given by Murthy & al. (2004).

The example shown hereby concerns the Bayesian analysis of the lifetime of a specific component Ω of the secondary water circuit of French nuclear power plants (Bousquet & Celeux 2006) :

Real failure times (months) :	134.9, 152.1, 133.7, 114.8, 110.0 129.0, 78.7, 72.8, 132.2, 91.8
-------------------------------	---

In addition to these data, expert knowledge is available over the future behavior of Ω . An expert provided an estimate of a reference lifetime value $t_e = 250$ months (interpreted as a median value) and ensures Ω is submitted to decelerated ageing. Doing this, the expert gives indirect and direct information on the prior $\pi_0(\eta, \beta)$ through a feature of the predictive prior distribution with cdf given by Equation (5), and a propriety of the marginal prior $\pi_0(\beta)$ since it is well known β quantifies the intrinsic dynamics of ageing of Ω (Lannoy & Procaccia 1994). Decelerated ageing means that β evolves between 1 and 2.

For setting a proper prior distribution $\pi(\eta, \beta)$ we use the hierarchical elicitation method proposed by Bousquet (2006). Consider reparametrization $\mu = \eta^{-\beta}$, then set a priori :

$$\beta \sim U(1, 2), \quad \mu|\beta \sim \text{Ga}(a, b(a, \beta)). \quad (13)$$

where $\text{Ga}(\cdot, \cdot)$ and $U(\cdot, \cdot)$ denote gamma and uniform pdf's respectively. In this prior setting, the (unconditional) prior distribution for β is chosen uniform to maximize the entropy over the bounded domain $[1, 2]$. Concerning μ , its prior, conditional to β , is specified by the function $b(a, \beta)$.

In this case, we will put : $b(a, \beta) = (2^{1/a} - 1)^{-1} t_e^\beta$. With this choice, we could show that the median of the predictive prior distribution of T is t_e , i.e. the expert value.

The parameter a is interpreted as the size of a virtual sample yielding a similar information to that given by the experts, and should not be too much greater than real sample size (in this case 10 observed failure times) to ensure that any posterior decision will be guided by the objective data information. Actually, a plays the role of a weight of the available expertise with respect to data.

Results

We are interested in computing the predictive MTTF and the failures probabilities before $t = (100, 200, 400)$ months, defined with respect to the posterior predictive distribution. We will also show how these predictions depend on the ratio $\rho = a/(a + 10)$, which is an easily understandable measure of the weight of prior information.

Bayesian inference has been performed with Markov Chain Monte Carlo (MCMC) techniques, using the OpenBUGS software via the BRugs package of R (Spiegelhalter et al. 2009). For $\rho = 25\%$ and $\rho = 50\%$, corresponding prior and posterior predictive densities are displayed in Figure 2. In this case, expert opinion is clearly optimistic with respect to data. Thus increasing ρ has the effect of predicting a posterior lifetime more and more longer than which could have been estimated using the data only, with a decreasing uncertainty.

	ρ		
	25%	50%	75%
MTTF	160	210	252
$P(T < 100)$	0.33	0.24	0.19
$P(T < 200)$	0.71	0.55	0.46
$P(T < 400)$	0.97	0.90	0.82

TABLE 1: Case-study 1 results.

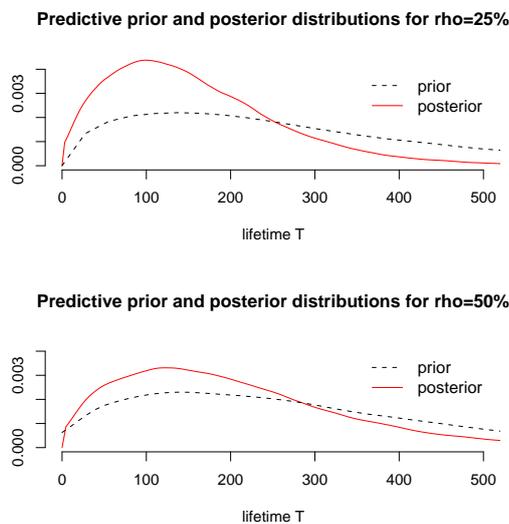


FIGURE 2: Predictive Weibull densities.

5 CASE STUDY 2 : FLOOD PROTECTION DESIGN

The second case-study deals with the design of a flood protection dike. The variable of interest is here the maximal water level Z_c of the river in a given wection. For the exemplary purposes of this study, we suppose that Z_c is related by a simple analytical formula to several input variables :

$$Z_c = Z_v + \left(\frac{Q}{K_s \cdot B \cdot \sqrt{(Z_m - Z_v)/L}} \right)^{3/5} \quad (14)$$

where :

- Q is the yearly maximal water discharge (m³/s),
- K_s is the Strickler friction coefficient,
- Z_m and Z_v are the riverbed levels (m asl) at the upstream and downstream respectively part of the river under investigation,
- B and L are the width and the length respectively of the river part (m).

In the formula above, the term in brackets is the water level H , measured in m above the riverbed level. Possible quantities of interest for Z_c are extreme quantiles, corresponding to very high probabilities (e.g. 0.99 or 0.999). These values allows to properly design a protection dike to cover flood risk. The case shown hereby is to be interpreted a *toy example* and the simplified analytical model given by Equation 14 is not representative of the models used by EDF to assess hydrological risk.

Among the input variables of the model, the width B and the length L of the part of river under investigation are considered as deterministic and their values are set to 300 and 5000 m respectively. The Strickler friction coefficient K_s is affected by epistemic uncertainty. Assessing the uncertainty of K_s is difficult as, in practice, even if this coefficient is strongly related to the morphology of the river, it cannot be measured. When a sample of couples (water level, discharge) is available, a way for properly assessing a pdf for K_s could be using inverse probabilistic methods, as in Celeux et al. (2007), for instance. We consider here that the probability distribution of K_s is already known and that it is normal with mean and standard deviation equal to 30 and 7.5 respectively. The riverbed levels upstream and downstream, Z_m and Z_v respectively, are uncertain and their uncertainty will be quantified by a bivariate normal distribution $N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Indeed, as upstream and downstream section are quite close it seems more reasonable to model them as possibly dependent variables. 29 couples of data $(z_m^{(i)}, z_v^{(i)})$ are available to perform Bayesian inference and setting the posterior distribution of $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$. The prior distributions for both components of vector $\boldsymbol{\mu}$ are normal with means equal to 56 and 50 m for μ_1 and μ_2 respectively and

standard deviations equal to 1 m. This prior translates a vague knowledge around two reference values. Concerning the prior of Σ , a classical choice is the inverse-Wishart distribution :

$$\pi_0(\Sigma|\Lambda, \nu) = \frac{|\Lambda|^{\nu/2} |\Sigma|^{-(\nu+p+1)/2} \cdot e^{(-\text{tr}(\Lambda \cdot \Sigma^{-1})/2)}}{2^{\nu p/2} \Gamma_p(\nu/2)}$$

where Σ is a $p \times p$ symmetric positive definite matrix (here $p = 2$), Λ is a $p \times p$ positive definite matrix, $\nu \geq p$ are the degrees of freedom and $\Gamma_p(\cdot)$ is the p-variate Gamma function.

A quite common choice for the parameter of this prior distribution is to set $\nu = p$. Concerning Λ , we put $\Lambda(1, 1) = \Lambda(2, 2) = 2$ and $\Lambda(1, 2) = \Lambda(2, 1) = -1$. According to this (slightly informative) choice the prior mean of Σ has a positive covariance term, which physically makes sense.

The variable Q is intrinsically random. For quantifying the uncertainty of Q , a sample of 149 annual maximal value is available. The extreme value theory gives the general theoretical framework for properly fitting a probability distribution to the maxima of iid samples. We choose to model the annual maxima of the river discharges with a Gumbel distribution :

$$f(q|\alpha, \beta) = \frac{1}{\beta} \exp \left[-\exp \left(\frac{q - \alpha}{\beta} \right) \right] \exp \left(\frac{\alpha - q}{\beta} \right)$$

with location parameter α and scale parameter β . Prior parameters for α and β are chosen uniform over a quite large support : $[500, 2000]$ for α and $[100, 1000]$ for β . That give rise to the predictive prior distribution for Q , shown in Figure 3.

variable	mean	st.dev.	2.5%	97.5%
α	1014.0	48.7	919.2	1110.0
β	565.4	36.7	498.7	642.7
μ_1	55.0	0.086	54.9	55.2
μ_2	50.2	0.073	50.0	50.3
σ_1	0.46	0.063	0.36	0.60
σ_2	0.39	0.053	0.30	0.51
$r_{1,2}$	0.64	0.109	0.39	0.82
Q	1340	719	286	3098
Z_c	52.8	1.07	51.0	55.1

TABLE 2: Case-study 2 results.

Results

As in the first case-study, Bayesian inference has been performed using Open BUGS software. Table 2 shows mean, standard deviation and 95% credibility interval bounds of marginal posteriors of model's parameters (α , β , μ , Σ) and of predictive distributions of Q and Z_c . The results concerning Σ are given in terms of standard deviations $\sigma_i = \sqrt{\Sigma(i, i)}$ and Pearson's correlation coefficient

$r_{i,j} = \Sigma(i, j)/(\sigma_i \sigma_j)$, more easily understandable in practice than variances and covariances. The values of the predictive distribution of Z_c corresponding to the probabilities 0.99 and 0.999, i.e. to returns periods of 100 and 1000 years respectively, are 55.8 and 57.9 m.

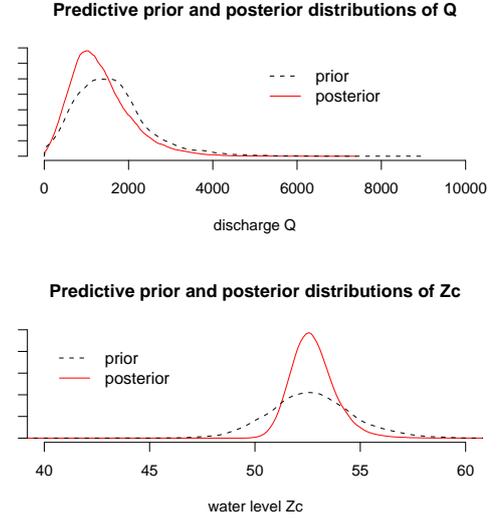


FIGURE 3: Case study 2 predictive distributions.

6 CONCLUSION

We will conclude this paper by stating again that Bayesian framework is very appropriate for dealing with uncertainty in industrial practice. The more powerful feature of Bayesian uncertainty setting is that a unique joint distribution exists for observable and unobservable variables. Working with conditional or marginal distributions allows to isolate or aggregate respectively uncertainty sources.

Moreover, as a consequence of working with conditional probabilities that force them to change their point of view back and forth from data to model, Bayesians are less prone to fall in love with their model, which helps to step back, discuss hypotheses and entertain the cycle of statistical analysis (Box, 1980). The possibility to take into account expert knowledge, the more natural way of interpreting probabilities, credibility intervals, statistical tests are other interesting features in practical applications. In spite of that, it cannot be denied that Bayesian setting is not very common in the industrial practice, at least less common than the frequentist approach.

This can be explained on the one hand by the fact that practitioners are more familiar with frequentist statistics, more largely taught in engineering courses, on the other hand by a persisting old-fashioned and rather negative vision of Bayesian statistics.

Classical reasons often advanced to explain the relative mistrust in Bayesian methods are the presumed

difficulty in putting inference into practice and the "arbitrary side" of prior information. If the first argument is clearly out-of-date, as nowadays fast calculations can be easily made by using Monte Carlo methods, deployed in friendly softwares as R or BUGS, it is true that the proper choice of prior distribution is not trivial and must be regarded as a real modeling problem. Indeed, properly methods exist to elicit expertise, when available, and measure possible conflicts between data and this prior knowledge. If no information is available, using non-informative or low-informative priors guarantees that the impact of the prior on the posterior will be very weak.

7 REFERENCES

- Apostolakis G. 1990. The Concept of Probability in Safety Assessments of Technological Systems. *Science* 250 : 1359-1364.
- Aven, T. 2003. *Foundations of Risk Analysis*, Chichester : J. Wiley & Sons.
- Beck, M.B. 1987. Water Quality Modelling : A Review of the Analysis of Uncertainty, *Wat. Res. Research* 23(8) : 1393-1442.
- Berger, J.O. 1985. *Statistical Decision Theory and Bayesian Analysis*. New York : Springer-Verlag.
- Bousquet, N. 2006. *Analyse bayésienne de la dureté de vie de composants industriels*. Ph.D. Thesis, Univ. Paris-Sud.
- Bousquet, N. 2008. Diagnostics of prior-data agreement in applied Bayesian analysis. *J. Appl. Statist.* 35 : 1011-1029.
- Bousquet, N. & Celeux, G. 2006. Measures of Bayesian discrepancy between prior beliefs and data knowledge. In Guedes Soares, C. & Zio E. (eds.) *Proceedings of the European safety and reliability conference ESREL 2006*. London : Taylor & Francis.
- Box, G.E.P. 1980. Sampling and Bayes' inference in scientific modelling and robustness (with discussion). *J. Roy. Stat. Society A* 143 : 383-430.
- Celeux, G., Grimaud, A., Lefebvre Y. & de Rocquigny, E. 2007. *Identifying intrinsic variability in multivariate systems through linearised inverse methods*. INRIA Research Report 6400.
- Evans, M. & Moshonov, H. 2006. Checking for prior-data conflict. *Bayesian Analysis* 1 : 893-914.
- Garthwaite, P. & Dickey, J. M. 1996. Quantifying and using expert opinion for variable-selection problems in regression. *Chemometrics and Intelligent Laboratory Systems* 35 : 1-26.
- Gelman, A., Carlin, J.B., Stern, S. & Rubin, D.B. (2004) *Bayesian Data Analysis, 2nd edition*. New York : Chapman & Hall.
- Girard, P. & Parent, E. 2004. The deductive phase of statistical analysis via predictive simulations : test, validation and control of a linear model with autocorrelated errors representing a food process. *J. Statist. Plann. Infer.* 124 : 99-120.
- Granger Morgan, M. & Henrion M., 1990, *Uncertainty : A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis* Cambridge : Cambridge University Press.
- Hamby, D.M. 1994. A Review of Techniques for Parameter Sensitivity Analysis of Environmental Models. *Environmental Monitoring and Assessment* 32(2) : 135-154.
- Helton, J., 1993, Uncertainty and sensitivity analysis techniques for use in performance assessment for radioactive waste disposal. *Rel. Eng. & Syst. Saf.* 42 : 327-367.
- Helton, J.C. & Oberkampf, W.L. (eds.) 2004, Alternative Representations of Epistemic Uncertainty. *Rel. Eng. & Syst. Saf.* (special issue) 85 (1-3).
- Kass, R. & Raftery, A. 1995. Bayes factors. *J. Amer. Stat. Assoc.*, 90 : 773-795.
- Kass, R.E. & Wasserman, L. 1996. The Selection of Prior Distributions by Formal Rules. *it J. Amer. Stat. Assoc.* 91 : 1343-1370.
- Knight, F.H. 1921. *Risk, Uncertainty and Profit*, Chicago : University of Chicago Press.
- Lannoy, A. & Procaccia, H. 1994. *Méthodes avancées d'analyse des bases de données du retour d'expérience industriel*. Paris : Eyrolles.
- Lecoutre, B. 1997. C'est bon à savoir ! et si vous étiez un bayésien qui s'ignore. *Revue du Modulad* 18 : 81-87.
- Lindley, D.V. (1985). *Making Decisions*. New York : J. Wiley & Sons.
- Morita, S., Thall, P.F. & Mueller, P. 2007. *Determining the Effective Sample Size of a Parametric Prior*. Univ. Texas Working Paper Series 36.
- Munier, B. & Parent, E. 1998. Le développement récent des sciences de la décision : un regard critique sur la statistique décisionnelle bayésienne, In Parent, E. & al. (eds.), *Bayesian Methods in Hydrology Sciences*. Paris : UNESCO Publishing.
- Murthy, D.N., Xie, M. & Jiang, R. 2004. *Weibull Models*. Hoboken (NJ) : J. Wiley & Sons.
- O'Hagan, A., Buck, C.E., Daneshkhan, A., Eiser, J.R., Garthwaite, P.H., Jenkinson, D.J., Oakley, J.E. & Rakow, T. 2006. *Uncertain Judgements : Eliciting Experts' Probabilities*. New York : J. Wiley & Sons.
- Pathé-Cornell, ME. 1996. Uncertainties in Risk Analysis : Six Levels of Treatment. *Rel. Eng. & Syst. Saf.* 54(2-3) : 95-111.
- de Rocquigny, E. Devictor, N. & Tarantola S. (eds.) 2008. *Uncertainty in industrial practice. A guide to quantitative uncertainty management*. Chichester : J. Wiley & Sons.
- Savage, L.H. 1954. *The Foundations of Statistics*. New York : Dover Publications.
- Spiegelhalter, D., Best, N., Lunn, D. & Thomas, A. 2009 *Bayesian Analysis using the BUGS language : A Practical Introduction*, Chapman & Hall (in press).
- Sundberg, R. 2001. Comparison of confidence procedures for type I censored exponential lifetimes. *Lifetime Data Analysis* 7 : 393-413.
- Wald, A. 1950. *Statistical Decision Functions*. New York : J. Wiley & Sons..