



HAL
open science

A Bayesian analysis of industrial lifetime data with Weibull distributions

Nicolas Bousquet

► **To cite this version:**

Nicolas Bousquet. A Bayesian analysis of industrial lifetime data with Weibull distributions. [Research Report] 2006, pp.24. inria-00115528v2

HAL Id: inria-00115528

<https://inria.hal.science/inria-00115528v2>

Submitted on 22 Nov 2006 (v2), last revised 23 Nov 2006 (v4)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

*A Bayesian analysis of industrial lifetime data with
Weibull distributions*

Nicolas Bousquet

N° ????

Novembre 2006

Thème COG



*R*apport
de recherche

A Bayesian analysis of industrial lifetime data with Weibull distributions

Nicolas Bousquet

Thème COG — Systèmes cognitifs
Projet SELECT

Rapport de recherche n° ???? — Novembre 2006 — 21 pages

Abstract: The context of our study is industrial reliability, where lifetime data are usually censored and in small number. Background information is available from experts. Our prior subjective knowledge is only about the lifetime of an industrial component and not about the parameters of a Weibull distribution which represents this lifetime. We propose to focus the discussion between the experts and the industrial analyst about the size of *virtual* data representing the variability of the expert opinion. Indeed, this size is one of the scarce indicators that both can understand. The prior calibration is made easy, and some methods and indicators including a default calibration method are proposed to help the Bayesian analyst (they can be extended to inferences on other distributions than Weibull). Besides, the posterior computation by importance sampling is simple and satisfying. Finally, through a real example, the flexibility of the elicitation is illustrated.

Key-words: reliability, durability, Bayesian analysis, Weibull, expert opinion, subjective prior, virtual data, censored data, Kullback-Leibler.

Une modélisation bayésienne pratique des paramètres de Weibull dans un cadre d'étude fiabiliste

Résumé : À faire

Mots-clés : fiabilité industrielle ; inférence bayésienne ; opinion d'expert ; modélisation *a priori* ; distance de Kullback-Leibler ; *a priori* non informatifs ; *a priori* objectifs ; échantillons d'apprentissage minimaux ; entropie ; calibration *a priori* ; modèle exponentiel ; modèle de Weibull

1 Introduction

In a reliability context, one of the most employed lifetime distributions is the Weibull distribution (Lawless 1982, chap. 4). This versatile $\mathcal{W}(\eta, \beta)$ distribution, with density function

$$f_W(t|\eta, \beta) = \frac{\beta}{\eta} \left(\frac{t}{\eta}\right)^{\beta-1} \exp\left\{-\left(\frac{t}{\eta}\right)^\beta\right\}$$

and hazard rate $h_W(t|\eta, \beta) = \frac{\beta}{\eta} \left(\frac{t}{\eta}\right)^{\beta-1}$ where $\eta, \beta > 0$, can be used for modelling infant mortality defects when the shape parameter $\beta < 1$, aging when $\beta > 1$ or accidental failure (with a constant rate) when $\beta = 1$. In this case, it reduces to an exponential distribution with scale parameter η .

In numerous industrial studies, reliability *feedback experience lifetime* (FEL) data $\mathbf{t}_n = t_1, \dots, t_n$ are often modelled with the Weibull distribution and the parameters are estimated by maximum likelihood inference using Newton-Raphson (NR) or EM algorithms. However, the good behavior of likelihood maximization is ensured only when the sample size n is large enough and when the lifetimes are little censored (Bacha 1996). Otherwise, using Bayesian inference techniques is relevant when prior knowledge is available (Robert 2001). This is the context of our study, where some experts provided informations about the lifetime T of an industrial component Σ . The main difficulty is to translate this expert opinion into information on (η, β) , through the choice of a prior distribution with density $\pi(\eta, \beta)$. Then the posterior density

$$\pi(\eta, \beta|\mathbf{t}_n) = \frac{\mathcal{L}(\mathbf{t}_n; \eta, \beta) \pi(\eta, \beta)}{\int_{\mathbb{R}} \int_{\mathbb{R}} \mathcal{L}(\mathbf{t}_n; \eta, \beta) \pi(\eta, \beta) d\eta d\beta}$$

allows to obtain estimates of reliability functions $h(\eta, \beta)$ which are of interest for the industrial analyst, as the survival function. Here $\mathcal{L}(\mathbf{t}_n; \eta, \beta)$ denotes the likelihood of the data.

This article is focused on the choice of $\pi(\eta, \beta)$. We consider that the expert(s) (considered as male(s) for simplicity) can express his (their) beliefs only about the marginal distribution \mathcal{M} of T with density

$$m(t) = \int_{\mathbb{R}} \int_{\mathbb{R}} f_W(t|\eta, \beta) \pi(\eta, \beta) d\eta d\beta$$

and not directly about (η, β) . We take profit of this feature to build a prior for the scale parametrization of the Weibull model (see Prop. 1). Typically, he (they) can give prior estimates of the mean, the median, the mode or percentiles of \mathcal{M} . However, he (they) can have difficulties to estimate his self uncertainty; indeed, an expert is often not a statistician and notions as standard deviation remain fuzzy. To obtain good estimations of this prior uncertainty, discussions between the expert and the Bayesian analyst (who gathers the prior information, possibly weight it and then infer both the data and the prior) are desirable. Hence the need of indicators which are understandable for both of them.

After some recalls and precisions about the meaning of the Weibull parameters in Section 2, we propose in Section 3 an improvement of the approach proposed by Berger and Sun (1993) and Bacha *et al.* (1998) to elicit prior distributions. The main point of this improvement is to consider the prior distribution as an approximation of the posterior distribution coming from a noninformative prior and *virtual data* whose size a is a relevant indicator of uncertainty.

According the specifications of the prior knowledge, we obtain a flexible hierarchical prior family for which we propose some strategies of calibration in Section 4. Especially, a default calibration method is suggested. Independently from the expert, the Bayesian analyst must often proceed to a final recalibration step. In this aim, an indicator measuring the effective size of the censored dataset \mathbf{t}_n is given to locate the strenght of the subjective information with respect to the strenght of the objective data information. Besides, bibliographical results are recalled and can be proceeded in a reliable and simple way. Then the consensus between several experts is briefly studied.

Finally, we show in Section 5 that the computation of the posterior distribution is simple using an importance sampling algorithm. In Section 6, a full Bayesian analysis is led for real data and two expert opinions to illustrate the main points of our elicitation method.

2 Data and parameters

2.1 The statistical context

For $n \in \mathbb{N}^*$, let $\mathbf{T}_n = T_1, \dots, T_n \sim \mathcal{W}(\eta, \beta)$ be independently and identically distributed real-or-vector-valued random variables in the sample space \mathbb{R}_+^n with probability density function (pdf) $f_W(t|\eta, \beta)$ and survival (reliability) function $S_W(t|\eta, \beta)$.

The available data are as follows: let $\mathbf{t}_n = (t_1, \dots, t_n)$ be an observed sample of n data. Usually \mathbf{t}_n contains r uncensored *i.i.d.* data $\mathbf{x}_r = (x_1, \dots, x_r)$ following $\mathcal{W}(\eta, \beta)$ and $n - r$ fixed (progressive type-I) right-censored values, denoted $\mathbf{c}_{n-r} = (c_1, \dots, c_{n-r})$. Thus, the observed likelihood can be written as

$$\mathcal{L}(\mathbf{t}_n; \eta, \beta) = \frac{\beta^r}{\eta^r} \prod_{i=1}^r \left(\frac{x_i}{\eta} \right)^{\beta-1} \exp \left\{ - \sum_{j=1}^n \left(\frac{t_j}{\eta} \right)^\beta \right\}.$$

2.2 Meaning of Weibull parameters

The usual Weibull parameters (η, β) have different senses. The scale parameter η is the 63rd percentile of the distribution and is homogeneous to T . But the shape parameter β has no dimension.

An expert viewpoint about β appears to be *qualitative*: the value of β reflects a knowledge about the *prospective* (or future) behavior of the component Σ . This knowledge is intrinsic to the component, independently of the *quantitative* knowledge of the failure time which is translated by a knowledge on functions of (η, β) as the mean, the median, etc. This qualitative meaning of β can be highlighted as follows: when the component Σ is submitted to aging, aging rate is given by

$$\frac{\partial h_W(t|\eta, \beta)}{\partial t} = \frac{\beta(\beta - 1)}{\eta^2} \left(\frac{t}{\eta} \right)^{\beta-2}$$

and aging acceleration can be measured by

$$\frac{\partial^2 h_W(t|\eta, \beta)}{\partial t^2} = \frac{\beta(\beta - 1)(\beta - 2)}{\eta^3} \left(\frac{t}{\eta} \right)^{\beta-3}.$$

Thus, as we said before, the exponential distribution ($\beta = 1$) can be used for modelling a component which is only submitted to accidental failure, without aging rate. Face to aging, engineers usually try

to slow down aging by preventive care (such that the estimated β remains between 1 and 2).

In our study, the available expert opinions are mainly quantitative and give little information about β . This is a significant difference with the conditions required for instance by Singpurwalla and Song (1986) and Singpurwalla (1988), who proposed hierarchical prior buildings with numerous hyperparameters. Note that Berger and Sun (1993) assumed the same hypothesis of a prominent prior knowledge on β . Here, we shall simply consider that $\text{Supp}(\pi(\beta))$ is bounded for objective reasons.

In the following, we use for convenience the parametrization $\mu = \eta^{-\beta}$. More generally, we denote θ the parameters of the model, especially when we define or use techniques which can be applied in more general settings than Weibull inference.

3 The prior modelling

3.1 The prior distribution of the shape parameter

We assume that an objective knowledge of β is available such that the prior domain can be bounded on $[\beta_l, \beta_r] \subset (0, \infty)$. Following arguments from Erto (1982), Berger and Sun (1993) and especially Bacha *et al.* (1998), this is a reasonable assumption since β directs the kinetics of aging. On mechanical systems, for physical reasons, $\beta > 5$ is never estimated on real Weibull data (except when they come from a 3-parameter Weibull distribution and must be reduced from a burn-in time). Usually, β stays in $[1, 2.5]$. See Lannoy and Procaccia (2001) for an engineering viewpoint. Moreover, when aging is assumed, we have clearly $\beta_l = 1$.

Following Bacha (1996), we propose to use a Beta $\mathcal{B}_e(p, q)$ distribution for the prior on β ,

$$\pi(\beta) = \frac{\Gamma(p+q)}{\Gamma(p)\Gamma(q)} \frac{(\beta - \beta_l)^{p-1} (\beta_r - \beta)^{q-1}}{(\beta_r - \beta_l)^{p+q-1}} \mathbb{1}_{[\beta_l, \beta_r]}(\beta)$$

because of its flexibility. It can be calibrated in terms of variance and central value independently of (β_l, β_r) . Moreover, $E[\beta]$ depends only of p/q and can be chosen independently of $\text{Var}[\beta]$ (and similarly for the mode which depends only of $(p-1)/(q-1)$). Bacha *et al.* (1998) provided a methodology to calibrate $\pi(\beta)$ which lays on the specification of an observation β_m . Jenkinson (2005) gives a review of the calibration methods of a Beta prior distribution if some values of β can be built from past experiments or expert questioning. However, such methodologies remain *ad hoc*. Moreover, since experts are not statisticians, the questioning has to be indirect.

We propose the following method to obtain prior estimations of β . About the component \sum , ask to the experts the probabilities p_0 and p_1 to fall down before the times t_0 and $t_1 > t_0$, respectively. Then

$$\frac{1-p_1}{1-p_0} = \exp \left\{ -\mu x_0^\beta \left(\left(\frac{x_1}{x_0} \right)^\beta - 1 \right) \right\} = (1-p_0) \left(\frac{x_1}{x_0} \right)^{\beta-1}$$

and deduce the prior estimation

$$\beta_m = \log \left\{ \frac{\log(1-p_1)}{\log(1-p_0)} \right\} \log^{-1}(x_1/x_0). \quad (1)$$

In next subsection, we focus our study on the prior distribution of the scale parameter η . A hyperparameter a which can be interpreted as a virtual sample size is defined. When fixed, this virtual size will be used afterwards (in § 4.1.2) to propose a calibration method of $\pi(\beta)$ when no *credible* information about the uncertainty of β is available.

3.2 The prior distribution of the scale parameter

3.2.1 Extending the approach of Berger and Sun (1993)

We extend here the approach of Berger and Sun (1993). Conditionally to β , they choose for the scale parameter η the generalized inverse gamma distribution

$$\eta|\beta \sim \mathcal{GIG}(a, b, \beta)$$

with density

$$f(\eta|a, b, \beta) = \frac{b^a \beta}{\Gamma(a)} \frac{1}{\eta^{a\beta+1}} \exp\left(-\frac{b}{\eta^\beta}\right) 1_{[0;+\infty[}(\eta) \quad (2)$$

where $(a, b, \beta) > 0$, with moment

$$\mathbb{E}[\eta^k] = \frac{b^{k/\beta} \Gamma(a - k/\beta)}{\Gamma(a)} \quad \forall a\beta > k > 0,$$

and mode $M_d[\eta] = (b\beta/a\beta + 1)^{1/\beta}$. This family is closed by scale transformation (i.e. $X \sim \mathcal{GIG}(a, b, \beta) \Rightarrow \forall c > 0, cX \sim \mathcal{GIG}(a, bc^\beta, \beta)$) which makes it interesting to represent prior informations on a scale parameter. Moreover, if we consider the reparametrization $\mu = \eta^{-\beta}$, we obtain prior and posterior

$$\begin{aligned} \mu|\beta &\sim \mathcal{G}(a, b), \\ \mu|\beta, \mathbf{t}_n &\sim \mathcal{G}\left(a + r, b + \sum_{i=1}^n t_i\right). \end{aligned}$$

Thus the posterior computation (see § 5) is made easier. However, Berger and Sun (1993) do not provide a meaning to the hyperparameters a and b . This choice remains controversial since $\pi(\mu|\beta) = \pi(\mu)$. Then we propose an alternative choice.

Suppose that the quantitative expert opinion can be represented by a *virtual* Weibull sample $\tilde{\mathbf{x}}_m = (\tilde{x}_1, \dots, \tilde{x}_m)$. Let $\pi^J(\eta|\beta) \propto \eta^{-1}$ be the conditional Jeffreys prior. If our conditional prior on η can be represented by the posterior $\pi^J(\eta|\beta, \tilde{\mathbf{x}}_m)$, we obtain

$$\begin{cases} a &= m, \\ b &= b(a, \beta) = \sum_{i=1}^m \tilde{x}_i^\beta. \end{cases}$$

Thus the choice of b as a function of β appears to be natural. Besides, a takes the meaning of the size of a virtual sample yielding the same information as the expert opinion. Hence a is now a *calibration* hyperparameter since the Bayesian analyst can modulate the strength of the prior *quantitative* opinion through this simple hyperparameter. This sample size is an easy interpretable parameter and a good focus point for a discussion between the analyst and the expert. We propose some default calibration values for a and questioning suggestions in § 4.

Ideally, $\pi(\mu|\beta)$ has a natural hierarchical structure. Assuming $\tilde{\mathbf{x}}_m \stackrel{i.i.d.}{\sim} \mathcal{W}(\eta_e, \beta_e)$ where (η_e, β_e) are prior estimations, we get

$$\begin{aligned} b|\beta &\sim \mathcal{G}(a, \eta_e^{-\beta}), \\ \mu|\beta, b &\sim \mathcal{G}(a, b). \end{aligned}$$

However, for simplicity reasons we prefer to choose a determinist expression of b in function of the available prior specifications. This is done in the next paragraph.

3.2.2 Translating the quantitative expert opinion

The expert opinion is *quantitative* when he talks in terms of lifetime. Many authors in Bayesian litterature use expert opinions on the values of the parameter vector (see Bacha 1996, Lijoi 2003 or Wisse *et al.* 2005). It can be understandable when it concerns the exponential model, since the mean lifetime is the parameter. But usually the expert does not know the statistical model and the prior information on parameters corresponds more to a “transformation” of the expert knowledge by the Bayesian analyst.

Thus, Sinpurwalla and Song (1986) then Sinpugpurwalla (1988) considered that an expert was able to speak about the median lifetime $m = \eta(\log 2)^{1/\beta}$: he is supposed to give an estimation of the prior mean $E[m]$. Fixing $\pi(\beta)$, they obtain a complete prior on both parameters. In our industrial context, to be more general, we consider that an expert can give an estimation of a quantity characterizing the *marginal distribution* \mathcal{M} of lifetime T with density

$$m_W(t) = \int_{\mathbb{R}} \int_{\mathbb{R}} f_W(t|\mu, \beta) \pi(\mu|\beta) \pi(\beta) d\mu d\beta. \quad (3)$$

An intensive litterature deals with the questioning of experts (see Daneskhah 2004 for a review). Some discussion techniques as the *bisection* method (Garthwaite *et al.* 2005) have been proposed with success to obtain from experts quantitative information about the behavior of a studied system. In our case, two questions are essential : “Can you give a representative value t_e of the lifetime T of component \sum ?” and “What is the probability α for the component \sum to fail before time t_e ?” (the context of both questions differ : the first one concerns *reliability* and the second one deals with *durability* ; see for instance Lawless 2000). Thus, we assume that the expert can be solicited to give a lifetime value t_e and the *specification* of t_e with respect to \mathcal{M} . This value can be perceived as the estimation

1. of the α order percentile, namely $P(T < t_e) = \alpha \in]0, 1[$;
2. of expectation $E[T]$;
3. of mode $M_d[T]$ (however, several private discussions with reliability specialists conclude to the very weak probability that an industrial expert can really specify a mode).

Through discussion, specifying precisely the nature of an expert opinion is essential. For instance, some authors have noticed that the experts tend to give a median value although they understand t_e as a mean (Schieren 1993, Lannoy and Procaccia 2001). The specification can be reinforced by some open questions as

1. When the age of a group of components reaches t_e , which proportion $1 - \alpha$ is still in use ?
2. Is the value t_e a mean of past events, taking into account extreme values or not ?
3. Is the value t_e the failure time which has been observed the most frequently in the past ?

Then, conditionally to the choice of a , the choice of b is given in the following proposition from the knowledge of the marginal distribution \mathcal{M} .

PROPOSITION 1. Let $t_e \in \mathbb{R}^{+*}$, $a > 0$ and $\alpha \in]0, 1]$. Denote

$$\begin{aligned} b_1(a, \beta) &= \left((1 - \alpha)^{-1/a} - 1 \right)^{-1} t_e^\beta, \\ b_2(a, \beta) &= \left(\frac{\Gamma(a)}{\Gamma(1 + 1/\beta)\Gamma(a - 1/\beta)} \right)^\beta t_e^\beta, \\ b_3(a, \beta) &= \frac{a\beta + 1}{\beta - 1} t_e^\beta. \end{aligned}$$

Denote B_i , $i = 1, 2, 3$ the induced prior modellings. Then, for all choice of $\pi(\beta)$ such that $\text{Supp}(\pi(\beta)) \subset [\beta_l, \beta_r]$,

$$\begin{aligned} (i) \quad P(T \leq t_e | B_1) &= \alpha, \\ (ii) \quad E[T | B_2] &= t_e \text{ if } a > \beta_l^{-1}, \\ (iii) \quad M_d[T | B_3] &= t_e \text{ if } \beta_l > 1. \end{aligned}$$

where the expectation is taken with respect to the marginal distribution \mathcal{M} with density (3).

Proof. See Appendix A.

Notice that an expert can give more than one specification. For instance, he can give a *credibility interval* $[t_{e1}, t_{e2}]$ and probabilities (α_1, α_2) such that $P(T \leq t_{ei}) = \alpha_i$ for $i = 1, 2$. We propose to reduce any expert opinion to “discrete” specifications $i = 1, \dots, P$ then to select couples (a_i, b_i) . Finally, we sum all specifications through the prior

$$\mu | \beta \sim \mathcal{G} \left(\sum_{i=1}^P P \gamma_i a_i, \sum_{i=1}^P b_i (P \gamma_i a_i, \beta) \right) \quad (4)$$

which is the posterior coming from successive Bayesian inferences on virtual samples. Weights $\gamma_i > 0$ (such that $\sum_{i=1}^P \gamma_i = 1$) are fixed in function of the relative trust in the specifications. Later in the article, an application will exemplify (and clarify) this choice. It is reinforced by the convexity properties of applications $\beta \mapsto b(a, \beta)$. Especially, it appears obviously that $b_1(a, \beta)$ (thus any convex sum of several $b_1(a, \beta)$) “mimics” the geometrical behavior of $\beta \mapsto \sum_{i=1}^a \tilde{x}_i^\beta$. It needs some weighty calculations (not given here) but can easily be numerically checked.

In next figures we display some densities $m(t)$ of the prior marginal distribution \mathcal{M} , where $t_e = 100$, $[\beta_l, \beta_r] = [1.1, 5]$, $(p, q) = (1.5, 1.5)$ and $a = 2$. Figure 1 shows $m(t)$ when t_e is successively considered as the prior median/mean/mode. Figure 2 shows the evolution of $m(t)$ when t_e is the α order percentile, for several values of α . The correlation between parameters induced by the prior modelling is revealed by the *convex hull* of the prior sampling ; in Figure 3, 95% of prior simulated (η, β) values are used for delimiting the convex hull. The form of the correlation remains near the form of the theoretical joint confidence area showed by Wu (2002) ; the inverse distribution of parameter values which is usually perceived is coherent with the classical behavior of frequentist estimations of the parameters (using unknown data): if β is overestimated then η is underestimated, and conversely.

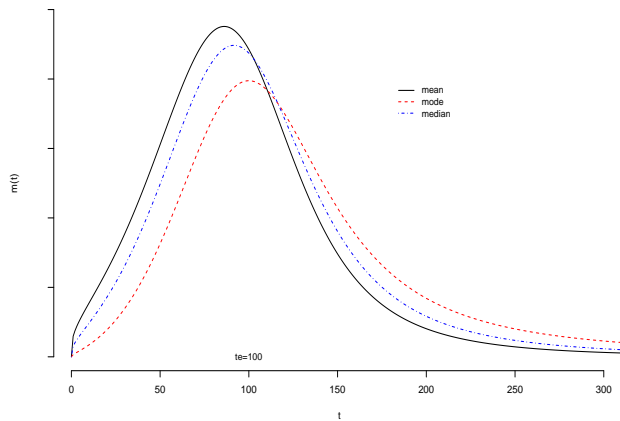


Figure 1: Densities $m(t)$ of the prior marginal distribution \mathcal{M} (indexed by the specification of $t_e = 100$).

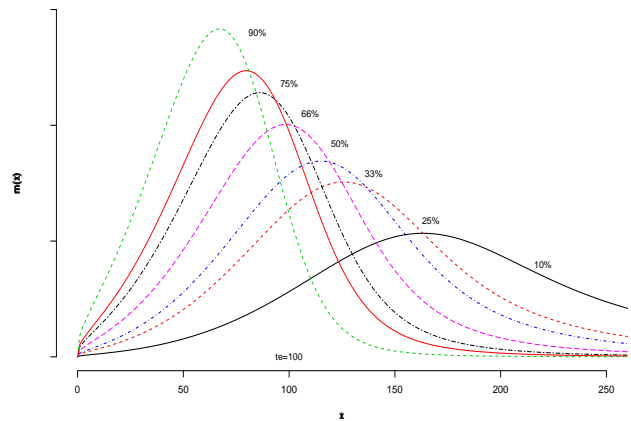


Figure 2: Densities $m(t)$ of the prior marginal distribution \mathcal{M} (indexed by the specification of α).

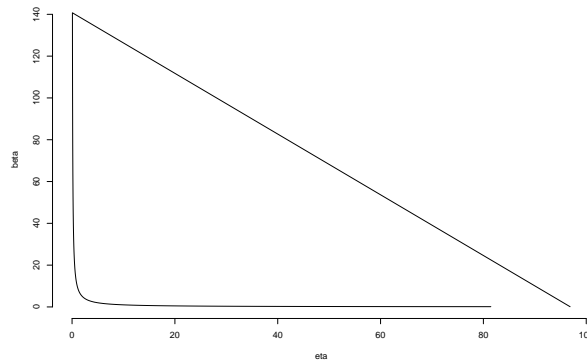


Figure 3: Convex hull of a prior sampling of parameters (η, β) .

4 Prior calibration

Some authors as Singpurwalla and Song (1986) build prior modellings with numerous hyperparameters, which allow to modify the expert information in location and uncertainty. As Lindley and Singpurwalla (1986) noticed, it has been observed that experts tend to produce “location and scale bias”. In this work, our opinion is that the full expert knowledge is given through t_e (and α if t_e is a percentile value). No objective criterion allows the Bayesian analyst to modify this expert knowledge. The only possibility is to modulate the expert uncertainty through the choice of a . In our sense, three steps of calibration should be considered.

1. A *default calibration* step; according to the number and the nature of specifications, we propose some default values for a . This is done in § 4.1. Besides, when a has been chosen, we propose

in § 4.1.2 a default method for calibrating $\pi(\beta)$ when no information is available about the uncertainty of β .

2. A *combined calibration* step; because a can be understood as a virtual sample size, it is a cgood parameter for a discussion between the expert and the Bayesian analyst. First values of a proposed by the default calibration step can be high when the expert is very precise ; the intent is to make the expert react to his self recklessness and correct a .
3. A *final recalibration* step; the Bayesian analyst can have some prior knowledge about the reliability of the expert opinion. A huge litterature is dedicated to this analysis (see Cooke and Goossens 2001 for a review). This analysis can be based on other subjective beliefs (which are not explored here) and the comparison with objective indicators. Here, if $a < n$ and if the data \mathbf{t}_n are uncensored, the posterior distribution is ensured to be more dependent from the objective data knowledge. In § 4.2, a similar upper bound for a is defined when the data \mathbf{t}_n are censored. Besides, some bibliographical results are given to moderate the expert opinions and the consensus of several expert opinions is studied.

4.1 Default prior calibration

4.1.1 Default calibration of the scale parameter

This calibration can be separated according to the kind of specification. The value of a that we proposed is the minimal number a_{\min} of past virtual data that seem necessary to assess this specification.

Mean, median or mode ? A confusion between those three specification is easy (especially the two first ones). Therefore a default choice must be done. We suggest to choose $a_{\min} = 1$ for the mean and $a_{\min} = 2$ for the median. Specifying the mode requires intuitively more past virtual data (and the certain knowledge of aging). We suggest $a_{\min} = 3$ since it is the minimal number of data to define the modal class of a continuous distribution. In Table 1, we give estimations of the standard deviation σ and the *skewness* γ of \mathcal{M} , using the prior choices done for displaying Figures 1 and 2, using several values of a . The mode specification appears as giving heavier tails and favorizing higher values of X . We recommend to use the median specification by default since it ensures a larger uncertainty than the mean specification.

a	mean		mode		median	
	σ	γ	σ	γ	σ	γ
2	67.4	4.5	360.7	31.3	94.2	5.5
3	53.8	2.4	185.0	12.5	66.0	2.7
4	50.0	1.91	151.6	9.8	58.2	1.94
5	48.1	1.45	136.7	7.5	55.8	1.46
10	44.0	1.26	110.7	6.3	50.3	1.13

Table 1: Estimations of standard deviation σ and skewness γ of distribution \mathcal{M} .

Quantile specification. Our working hypothesis is that the expert has perceived as much past virtual data as it is necessary to obtain the precision α when he specifies t_e as a percentile. For instance, the Bayesian analyst can propose $a = 10$ when $\alpha = 10\%$. This is similar to the numerous *histogram* methods that have been developed for the asking of experts. Thus, Van Noortwijk *et al.*

(1992) propose to segment \mathbb{R}_+ in separate intervals, choosing boundaries near to real observed data (or censored data). Since those lifetimes can be representative for the expert (even he is supposed not to know the data), he can be asked about his probability of failure. Next example illustrates our way of thinking.

EXAMPLE 1. Let $(50, 80, 90, 170)$ be some (possibly censored) observed failure times. Divide the lifetime scale in $D_1 = [0, 90)$ and $D_2 = [90, \infty)$. Let (n_1, n_2) be the numbers of virtual data in D_1 and D_2 , respectively. Suppose that the expert gives 66% chance of breaking down in D_1 . Then $n_1 = 2n_2$. Divide now D_1 in $D_{1,a} = [0, 50)$ and $D_{1,b} = [50, 90)$. Indicate to the expert that the effective lifetime is contained in D_1 . Suppose that he indicates 25% chance of breaking down in $D_{1,b}$; to a similar question on a partition of D_2 , $D_{1,a}$ or $D_{1,b}$, suppose that he is unable to answer. He is supposed to have “perceived” at least $n_1 = 4$ virtual data. Finally, he can be given $n_1 + n_2 = 6$ virtual data. ■

Credibility interval specification. Suppose to obtain from an expert the interval $[t_{e1}, t_{e2}]$ and probabilities (α_1, α_2) such that $P(T \leq t_{ei}) = \alpha_i$ for $i = 1, 2$. Adding two percentile specifications needs to normalize a_1, a_2 such that $a = \gamma_1 a_1 + \gamma_2 a_2$ is the maximal size attainable by the separate specifications. Thus, specifications (50%, 90%) will give the same default size $a = 10$ than specifications (10%, 90%). Moreover, if the prior domain for β has been elicited, the coherence of the credibility interval with the Weibull model must be checked using formula (1). If

$$\beta_e = \log \left\{ \frac{\log(1 - \alpha_2)}{\log(1 - \alpha_1)} \right\} \log^{-1}(t_{e2}/t_{e1}) \tag{5}$$

is not in $[\beta_l, \beta_r]$ the orders should be modified such that $\beta_e = \beta_0$ (for instance the middle or a bound of $[\beta_l, \beta_r]$). Solving (5) = β_0 with a Newton-Raphson algorithm is simple (see Bousquet 2006 for details). Orders are weighted with the constant convergence rate $\gamma_1/(1-\gamma_1)$. Thus, if $\gamma_1 \rightarrow 1$, α_1 remains stable.

Algorithm 1. Weighting the credibility α .

1. Let $0 < \alpha_1^0 < \alpha_2^0 < 1$ and $\alpha_0 = (\alpha_1^0, \alpha_2^0)$. Denote $\ell_0 = (x_{e,2}/x_{e,1})^{\beta_0}$, $\ell(\alpha) = \frac{\log(1 - \alpha_2)}{\log(1 - \alpha_1)}$ and choose a precision ε . Fix $0 < \rho \ll 1$.

2. Step $k = 0, \dots, K$:

- compute the vector $\delta_k = \begin{pmatrix} (1 - \alpha_1^k) \log(1 - \alpha_1^k) (\ell(\alpha_k) - \ell_0) / \ell(\alpha_k) \\ -(1 - \alpha_2^k) \log(1 - \alpha_1^k) (\ell(\alpha_k) - \ell_0) \end{pmatrix}$;
- compute $h_k = -\frac{(1 - \alpha_1^k)}{(1 - \alpha_2^k) \ell(\alpha_k)}$.
- compute $\alpha^{k+1} = \begin{cases} \alpha_1^k - \rho h_k (\alpha_2^k - \alpha_1^k), \\ \alpha_2^k - \rho h_k (\alpha_2^k - \alpha_1^k) \frac{\gamma_1}{1 - \gamma_1} \end{cases}$
- stop when $\|\delta_k\| \leq \varepsilon$

EXAMPLE 2. Choose $\beta_0 = 3$ and $(t_{e,1}, t_{e,2}) = (200, 300)$. Fixing $(\alpha_1^0, \alpha_2^0) = (0.05, 0.95)$, we obtain $\beta_e \simeq 10.03$. The expert opinion induces an unrealistic parameter shape. By default, fix $\gamma = (1/2, 1/2)$. Then we obtain $\alpha = (0.3, 0.7)$. If we choose now $\alpha_1^0 = 0.25$, we obtain $\beta_e \simeq 5.78$. The induced aging remains unreasonable. With equal weights, we find $\alpha = (0.4, 0.8)$. Assuming much more credibility in the lower bound, we fix $\gamma = (0.95, 0.05)$. Then $\alpha = (0.26, 0.64)$. Conversely, with $\gamma = (0.05, 0.95)$, we obtain $\alpha = (0.55, 0.93)$. ■

4.1.2 Default calibration of the shape parameter

Once a has been chosen, it can be used to calibrate $\pi(\beta)$ if no real information about the uncertainty of β is available. Ideally, since we see $\pi(\eta|\beta)$ as a reference posterior density with respect to virtual data $\tilde{\mathbf{x}}_{\mathbf{a}}$, $\pi(\beta)$ should be chosen as

$$\pi^J(\beta|\tilde{\mathbf{x}}_{\mathbf{a}}) \propto \beta^{a-1} \frac{\left(\prod_{i=1}^a \tilde{x}_i\right)^\beta}{\left(\sum_{j=1}^a \tilde{x}_j^\beta\right)^a}, \quad (6)$$

with $\pi^J(\beta)$ the noninformative reference prior. Because this prior distribution is not tractable (it is well known that the Weibull distribution does not admit any continuous conjugate, cf. Soland 1969), we made an arbitrary choice of $\pi(\beta)$. But the Beta prior $\pi(\beta|p, q)$ can be elicited as the minimizer of the relative entropy

$$\int_{\beta_l}^{\beta_r} \pi^J(\beta|\tilde{\mathbf{x}}_{\mathbf{a}}) \log \frac{\pi^J(\beta|\tilde{\mathbf{x}}_{\mathbf{a}})}{\pi(\beta|p, q)} d\beta \quad (7)$$

under a constraint on the mean or the mode, which is tractable using (1). However, the values of the virtual sample $\tilde{\mathbf{x}}_{\mathbf{a}}$ remain unknown. Therefore we propose to sample virtual samples from the $\mathcal{W}(\eta_e, \beta_e)$ distribution, where η_e and β_e are prior estimations. Thus we replace $\pi^J(\beta|\tilde{\mathbf{x}}_{\mathbf{a}})$ in (7) by the *expected posterior prior*

$$\pi_a^J(\beta) = \int_{\mathbb{R}_+^a} \pi^J(\beta|\tilde{\mathbf{x}}_{\mathbf{a}}) \prod_{i=1}^a f_W(\tilde{x}_i|\eta_e, \beta_e) d\tilde{x}_1 \dots d\tilde{x}_a.$$

Then minimizing (7) in (p, q) is similar to minimize

$$\log \frac{\Gamma(p)\Gamma(q)}{\Gamma(p+q)} + (p+q) \log(\beta_r - \beta_l) - (p+q) \mathbb{E} \left[\log \frac{\beta_r - \beta}{\beta - \beta_l} \right] \quad (8)$$

where the mean $\mathbb{E}[\cdot]$ is with respect to $\pi_a^J(\beta)$. When the prior mean (or mode) of $\pi(\beta)$ is chosen, the solution is unique. This minimization can be done using Monte Carlo estimations of the extreme right term of (8). In our applications, this minimization gave promising results. But it needs more work to ensure that the Kullback-Leibler projection does not lead to an overestimation of the prior information.

4.2 Recalibration and consensus

4.2.1 Comparing subjective and objective knowledge

In this subsection we give the definition of an indicator \tilde{n} which measures the effective size of the available censored data $\mathbf{t}_{\mathbf{n}}$. The virtual size a should be compared to \tilde{n} to locate the subjective information with respect to the objective data information. A similar indicator can be easily defined for other models than the Weibull distribution.

Denote $\mathbf{x}_{\mathbf{m}}$ an *i.i.d.* Weibull sample of size m yielding the same data information of $\mathbf{t}_{\mathbf{n}}$. Formally, denoting

$$\mathcal{D}^{J_c, J}(\mathbf{t}_{\mathbf{n}}, \mathbf{x}_{\mathbf{m}}) = \int_{\Theta} \pi_c^J(\theta|\mathbf{t}_{\mathbf{n}}) \log \frac{\pi_c^J(\theta|\mathbf{t}_{\mathbf{n}})}{\pi^J(\theta|\mathbf{x}_{\mathbf{m}})} d\theta \quad (9)$$

the Kullback-Leibler divergence between two reference posterior distributions, we define our indicator of effective size by

$$\tilde{n} = \arg \min_m \mathbb{E}_{X_m} \mathcal{D}^{J_c, J}(\mathbf{t}_n, \mathbf{x}_m). \tag{10}$$

From Lin *et al.* (2006, Prop. 1), the existence and unicity of \tilde{n} is ensured. In (9), π^J and π_c^J are noninformative priors with good frequentist coverage of the posterior Bayes intervals for \mathbf{x}_m and \mathbf{t}_n , respectively. Indicator \tilde{n} is of interest since it improves the conservative choice r (the number of uncensored data) which underestimates the effective data information. Especially, when all available data are censored ($r = 0$) we have $\tilde{n} > 0$. Of course, when \mathbf{t}_n is uncensored, we have $\tilde{n} = n$.

Usually π^J and π_c^J are said *coverage matching priors* ; see Datta (1996) and Ghosal (1999) for reviews and precisions. The definition follows: denoting $\theta_n(\alpha)$ the posterior α -quantile of θ based on observations \mathbf{t}_n (i.e. $P_c^J(\theta \leq \theta_n(\alpha) | \mathbf{t}_n) = \alpha$), it means that

$$P_\theta(\theta \leq \theta_n(\alpha)) = P_c^J(\theta \leq \theta_n(\alpha) | \mathbf{t}_n) + \mathcal{O}(n^{-i/2})$$

where the left-hand side is the frequentist probability, where θ is fixed in Θ and $\theta_n(\alpha)$ is random, and i is (ideally) the highest attainable value (the coverage order). In our applications, π^J was chosen as the *reference prior* (Berger and Bernardo 1992) which is at least of order two according to Sun (1993). Since \mathbf{t}_n is censored, we chose π_c^J as the special Jeffreys prior (including censored times) which can be derived from De Santis *et al.* (2001). The authors show it has better posterior coverage than the standard Jeffreys prior. From Bousquet (2006, Chapter 3), it is defined as follows, when \mathbf{t}_n contains $n - r$ right-censored fixed values c_1, \dots, c_{n-r} .

Let γ be the Euler constant ($\gamma \simeq 0.57722$). Let $\gamma_1 = \pi^2/6 + \gamma^2 - 2\gamma > 0$ and $\gamma_2 = -2(1 - \gamma)$. Denote

$$\begin{aligned} \tilde{\delta}(\mu, \beta) &= \delta^2(\mu, \beta) + [\delta(\mu, \beta) - 1] (\gamma_1 + \gamma_2 \log \mu + \log^2 \mu) + \pi^2/6 - 1, \\ \delta(\mu, \beta) &= n - \sum_{i=1}^{n-r} \exp(-\mu c_i^\beta). \end{aligned}$$

Then the special Jeffreys prior for the Weibull parametrization (μ, β) is

$$\pi_c^J(\mu, \beta) \propto (\mu\beta)^{-1} \sqrt{\tilde{\delta}(\mu, \beta)}.$$

4.2.2 Percentile orders correction

Suppose that the discussion between the expert and the analyst came to an agreement about the order α of a percentile value t_e . Then, independently from the expert, the analyst can correct this order using the results averaged by numerous authors about the real sense of the percentile prior estimations. Face to the results of various experiments in reliability, those estimations are given credibility orders that usually are very overestimated. Illustrated by Meyer and Booker (1987) and Lannoy and Procaccia (2001), a tacit rule is summarized in next table, modifying the order of percentile then giving us upper bounds a^* for a .

expert opinion	reality	a^*
5%	25%	4
20%	33%	3
25%	40%	2
75%	60%	2
80%	66%	3
95%	75%	4

Table 2: Tacit reduction of percentile order between expert opinion and reality.

4.2.3 Consensus between experts

When priors $\pi_1(\theta), \dots, \pi_M(\theta)$ are modelling several available (independent) expert opinions, the *convex weighted combination*

$$\pi(\theta) = \frac{\prod_{i=1}^M \pi_i^{\varpi_i}(\theta)}{\int_{\Theta} \prod_{i=1}^M \pi_i^{\varphi_i}(\theta) d\theta},$$

where $\sum_{i=1}^M \varphi_i = 1$, minimizes the Kullback-Leibler information loss

$$KL(\pi; \pi_1, \dots, \pi_M | \beta) = \sum_{i=1}^M \varpi_i \int_{\mathbb{R}} \pi(\theta) \log \frac{\pi(\theta)}{\pi_i(\theta)} d\theta$$

and carries out an optimal *consensus* of the opinions. See Liisberg (1991) ou Alturazza *et al.* (2004) for more precisions about this elicitation. The relative importance of the experts is judged through the choices of weights ϖ_i . Cooke *et al.* (1988) or Budescu and Rantilla (2000) among others propose several criteria, like the past error rate, to fix the weights through methodologies whose most famous is probably the Delphi method (see Linstone and Turoff 2002 for a review). Fortunately we obtain

$$\begin{aligned} \mu | \beta &\sim \mathcal{G} \left(\sum_{i=1}^M \varpi_i a_i, \sum_{i=1}^M \varpi_i b_i(a_i, \beta) \right), \\ \beta &\sim \mathcal{B}_e \left(\sum_{i=1}^M \varpi_i p_i, \sum_{i=1}^M \varpi_i q_i \right) \end{aligned}$$

when $\pi_i(\beta)$ is a Beta density $\mathcal{B}_e(p_i, q_i)$ defined on a common domain $[\beta_l, \beta_r]$. The consensus distribution on μ appears as the posterior distribution coming from a consensus virtual sample whose size is the weighted sum of all sizes. If experts can not be considered as independent (usually when $M > 2$), correlations must be add to the modelling. Face to this issue, O'Hagan (2003, 2005) gives numerous arguments to define a consensus (by discussion means) before the modelling. But we obtain the same kind of prior modelling.

5 Posterior computation

Berger and Sun (1993) provided steps of Gibbs sampling when $\pi(\beta)$ was log-concave. Our choice could be similar. Alternatively, since $\text{Supp}(\pi(\beta))$ is bounded, it is easy to use *importance sampling* to

estimate the posterior mean of a function of interest $h(\theta)$

$$\begin{aligned} I &= \int_{\Theta} h(\theta)\pi(\theta|\mathbf{t}_n) d\theta, \\ &= \int_{\Theta} h(\theta)\frac{\pi(\theta|\mathbf{t}_n)}{\rho(\theta)} \rho(\theta) d\theta \end{aligned}$$

where $\rho(\theta)$ is any function such that $\text{Supp}(\pi(\theta|\mathbf{t}_n)) \subset \text{Supp}(\rho(\theta))$ (Robert and Casella 2004). Choosing ρ as a density, it is easy to estimate I by

$$\hat{I}_M = \sum_{i=1}^M \omega_i h(\theta_i)$$

where the θ_i are simulated from $\rho(\theta)$ and the weights $\omega_i = \omega'_i / \sum_{j=1}^M \omega'_j$ with

$$\omega'_i = \frac{\pi(\theta_i)\mathcal{L}(\mathbf{t}_n; \theta_i)}{\rho(\theta_i)}.$$

Under mild conditions, a limit central theorem ensures the convergence of \hat{I}_M to I when $M \rightarrow \infty$. The main difficulty is that the tails of π must be heavier than the tails of $\pi(\theta|\mathbf{t}_n)$ (and ideally, $\rho(\theta)$ should be close to $\pi(\theta|\mathbf{t}_n)$). In our case, with $\theta = (\mu, \beta)$, we obtain

$$\begin{aligned} \mu|\beta, B, \mathbf{t}_n &\sim \mathcal{G}\left(a+r, b(a, \beta) + \sum_{i=1}^n t_i^\beta\right), \\ \pi(\beta|B, \mathbf{t}_n) &\propto \beta^r \left(\prod_{j=1}^r x_j\right)^\beta (\beta - \beta_l)^{p-1} (\beta_r - \beta)^{q-1} b^a(a, \beta) \left\{b(a, \beta) + \sum_{i=1}^n t_i^\beta\right\}^{-(a+r)} \mathbb{1}_{[\beta_l, \beta_r]}(\beta). \end{aligned}$$

Thus, choosing $\rho(\mu, \beta) = \pi(\mu|B, \beta, \mathbf{t}_n) \mathbb{1}_{\{\beta_l \leq \beta \leq \beta_r\}} / (\beta_r - \beta_l)$, we respect the conditions of a satisfying importance sampling and obtain the theoretical (unnormalized) weight

$$\omega'(\beta|B, \mathbf{t}_n) = \beta^r \left(\prod_{i=1}^r x_i\right)^\beta (\beta - \beta_l)^{p-1} (\beta_r - \beta)^{q-1} \frac{b^a(a, \beta)}{\left(b(a, \beta) + \sum_{i=1}^n t_i^\beta\right)^{a+r}} \mathbb{1}_{\{\beta_l \leq \beta \leq \beta_r\}}.$$

Especially, an industrial reliabilist is interested in the computation of the following *predictive* quantities:

1. the *mean lifetime*

$$\mathbb{E}[T|B, \mathbf{t}_n] = \int_{\mathbb{R}} \Gamma(1 + 1/\beta) \left\{b(a, \beta) + \sum_{i=1}^n t_i^\beta\right\}^{1/\beta} \frac{\Gamma(a+r-1/\beta)}{\Gamma(a+r)} \pi(\beta|B, \mathbf{t}_n) d\beta;$$

2. the *survival at time t_0*

$$\begin{aligned} S(t_0|B, \mathbf{t}_n) &= P(T > t_0|B, \mathbf{t}_n), \\ &= \int_{\mathbb{R}} \left(1 + \frac{t_0^\beta}{\left\{b(a, \beta) + \sum_{i=1}^n t_i^\beta\right\}}\right)^{-(a+r)} \pi(\beta|B, \mathbf{t}_n) d\beta; \end{aligned}$$

3. the *residual lifetime after time* t_0 (see for instance Finkelstein 2006)

$$\begin{aligned} \text{MRTF}(t_0|B, \mathbf{t}_n) &= \text{E}[T - t_0 | T > t_0, B, \mathbf{t}_n], \\ &= \frac{1}{S(t_0|B, \mathbf{t}_n)} \int_{t_0}^{\infty} S(x|B, \mathbf{t}_n) dx. \end{aligned}$$

Thus Monte Carlo estimations of these quantities need simply posterior sampling of β , or uniform importance sampling on (β_l, β_r) .

6 A numerical example

We consider the right-censored real lifetime data \mathbf{t}_n ($n = 18$) from Table 3. They correspond to failure times or stopping times collected on some similar devices \sum belonging to the secondary water circuit of nuclear plants. Lifetimes are given in months. For physical reasons and according to a large consensus, those data are assumed to arise from a Weibull distribution $\mathcal{W}(\eta, \beta)$. The maximum likelihood estimator (MLE) is $(\hat{\eta}_n, \hat{\beta}_n) = (140.8, 4.51)$ with estimated standard deviations $\hat{\sigma}_n = (7.3, 1.8)$. The high value of $\hat{\beta}_n$ is unexpected and suggests a Bayesian estimation.

real failure times:	134.9, 152.1, 133.7, 114.8, 110.0, 129.0, 78.7, 72.8, 132.2, 91.8
right-censored times :	70.0, 159.5, 98.5, 167.2, 66.8, 95.3, 80.9, 83.2

Table 3: Lifetimes (months) of nuclear components (from secondary water circuits).

Two prior opinions on the lifetime are available on device \sum , given by independent experts \mathcal{E}_1 and \mathcal{E}_2 . They are summarized in Table 4. The \mathcal{E}_1 opinion is much more informative than \mathcal{E}_2 and both are right-shifted with respect to the data. Moreover the experts are not asked at the same level of precision. \mathcal{E}_1 is a nuclear operator and speaks for a particular component while \mathcal{E}_2 can be seen as a component producer whose opinion takes into account a variety of running conditions. Thus the expert opinions can be considered as independent.

	credibility intervals (5%,95%)	median value
expert \mathcal{E}_1	[200,300]	250
expert \mathcal{E}_2	[100,500]	250

Table 4: Expert opinions about the lifetime T .

Aging is assumed: we choose $\beta_l = 1$. For technical reasons we choose $\beta_r = 5$. Using (1), the underlying prior estimates of β proposed by the experts take values in $\{8.02, 10.0, 11.7\}$ (expert \mathcal{E}_1) and $\{2.11, 2.53, 2.84\}$ (expert \mathcal{E}_2). Thus, the first expert opinion seems dubious since it induces an unreasonable aging. Note that when we replace orders (5%,95%) by (33%,66%) (found using Algorithm 1) we obtain prior values of β near to 2.5. For this reason, we prefer these corrected orders. Then, for each prior density $\pi(\beta|(p, q)_{\mathcal{E}_i})$ ($i = 1, 2$), we specify for the *mode* the same value $\beta^* = 2.5$. Indeed, we necessarily have $(p, q)_{\mathcal{E}_i} > 1$ which allows not to obtain flat priors. This is the starting

point of the elicitation method in Bacha *et al.* (1998).

Now let (a_1, a_2) be the virtual size for each expert. A default prior for the combination of both expert opinions is

$$\begin{aligned}\mu|\beta &\sim \mathcal{G}(a, b(\beta)), \\ \beta &\sim \mathcal{B}_e(p, q)\end{aligned}$$

with

$$\begin{cases} a &= \varpi_1 a_{\mathcal{E}_1} + (1 - \varpi_1) a_{\mathcal{E}_2}, \\ b(\beta) &= \varpi_1 b_{\mathcal{E}_1}(\beta) + (1 - \varpi_1) b_{\mathcal{E}_2}(\beta), \\ p &= \varpi_1 p_{\mathcal{E}_1} + (1 - \varpi_1) p_{\mathcal{E}_2}, \\ q &= \varpi_1 q_{\mathcal{E}_1} + (1 - \varpi_1) q_{\mathcal{E}_2}, \\ \\ a_{\mathcal{E}_1} &= 3(\gamma_1 a_1 + \gamma_2 a_2 + \gamma_3 a_3), \\ a_{\mathcal{E}_2} &= 3(\gamma'_1 a'_1 + \gamma'_2 a'_2 + \gamma'_3 a'_3), \\ q_{\mathcal{E}_i} &= 1 + (\delta - 1)(p_{\mathcal{E}_i} - 1) \text{ for } i = 1, 2, \\ \delta &= (\beta_r - \beta_l)/(\beta^* - \beta_l), \end{cases}$$

and

$$\begin{aligned}b_{\mathcal{E}_1}(\beta) &= \left(\frac{1}{0.66^{1/3\gamma_1 a_1}} - 1\right)^{-1} 200^\beta + \left(\frac{1}{0.5^{1/3\gamma_2 a_2}} - 1\right)^{-1} 250^\beta + \left(\frac{1}{0.33^{1/3\gamma_3 a_3}} - 1\right)^{-1} 300^\beta, \\ b_{\mathcal{E}_2}(\beta) &= \left(\frac{1}{0.95^{1/3\gamma'_1 a'_1}} - 1\right)^{-1} 100^\beta + \left(\frac{1}{0.5^{1/3\gamma'_2 a'_2}} - 1\right)^{-1} 250^\beta + \left(\frac{1}{0.05^{1/3\gamma'_3 a'_3}} - 1\right)^{-1} 500^\beta.\end{aligned}$$

We have no precision about the relative legitimacy of the percentiles. Thus we fix $\gamma_i = \gamma'_i = 1/3$. Similarly, we have no objective criterion allowing us to favor any expert. Thus we fix $\varpi_1 = 1/2$.

The minimal number of virtual data to obtain the percentiles (25%, 50%, 75%) is 4. To specify percentiles (25%, 75%), we need $a_1 = a_3 = 2a_2$. By normalization, we obtain the default values $a_1 = a_3 = 1.6$ and $a_2 = 0.8$. For the second expert opinion, by the same method, we obtain $a'_1 = a'_3 = 10a'_2$ and $a'_1 + a'_2 + a'_3 = 20$. Then $a'_2 = 20/21 \simeq 0.95$ and $a'_1 = a'_3 \simeq 9.5$. Note that the censored dataset \mathbf{t}_n yields approximatively as much information as $\tilde{n} \simeq 11$ *i.i.d.* data (see § 4.2.1). Thus, the second posterior distribution will be more dependent from the prior than the data (in an approximate ratio of 66%) if a_2 is not modified (for instance using Table 2).

Then, since a_1 and a_2 are fixed, the prior uncertainty on β is elicited as explained in § 4.1.2. We use $(\eta_e, \beta_e) = (290, 2.5)$ to sample virtual data, since the median of $\mathcal{W}(\eta_e, \beta_e)$ is 250. We obtained prior variances $\sigma_1^2 = 1.07$ and $\sigma_2^2 = 0.67$, respectively. We obtained finally $(p, q)_{\mathcal{E}_1} = (1.19, 1.31)$ for the first expert, and $(p, q)_{\mathcal{E}_2} = (3.13, 4.56)$ for the second expert.

Finally, the complete prior corresponds to a virtual sample of size $a = 12$, which implies that the posterior distribution is approximatively as much dependent from subjective knowledge than frequentist knowledge. Separate prior densities on parameters (η, β) and the marginal density $m(t)$ are displayed on Figures 4, 5 and 6. Marginal densities $m(t)$ get good compromises between the specifications. For the expert \mathcal{E}_1 , empirical percentiles of order (33%, 50%, 66%) are (197, 250, 298). For the expert \mathcal{E}_2 , empirical percentiles of order (5%, 50%, 95%) are (98, 255, 494) (these results are computed on 10^5 sampled particles). Posterior survival functions $S(t)$ are displayed in Figure 7.

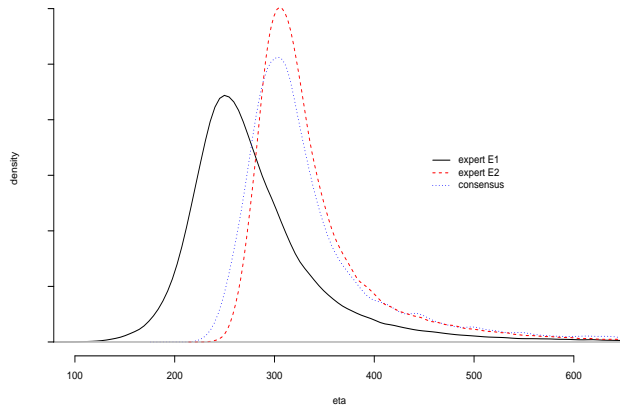


Figure 4: Default prior densities $\pi(\eta)$ for separate and consensus expert opinions.

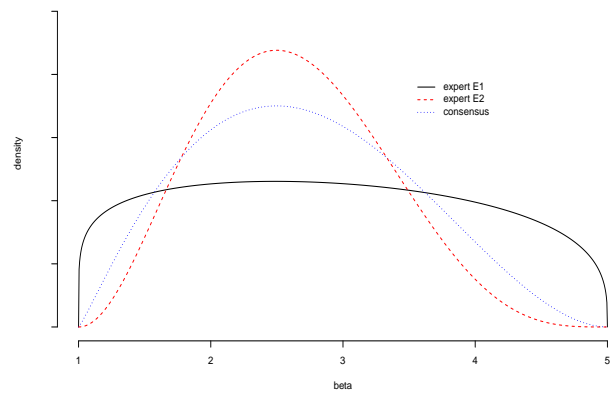


Figure 5: Default prior densities $\pi(\beta)$ for separate and consensus expert opinions.

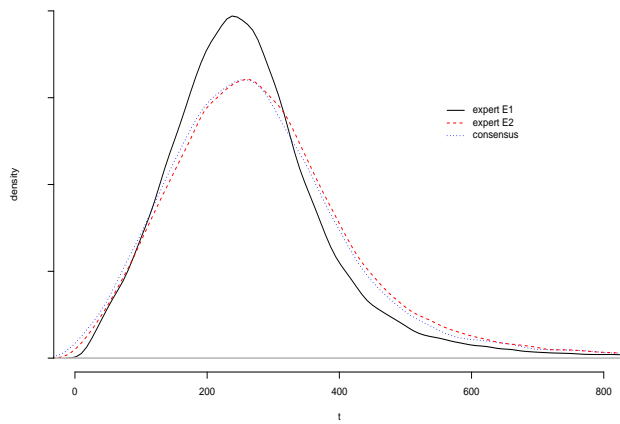


Figure 6: Default marginal prior densities $m(t)$ for separate and consensus expert opinions.

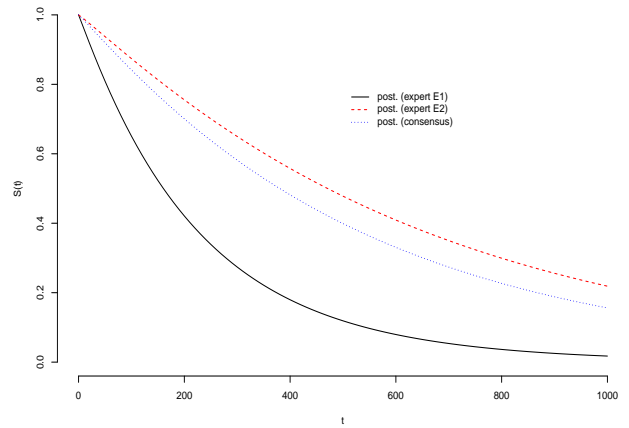


Figure 7: Posterior survival functions $S(t)$ for separate and consensus expert opinions.

7 Conclusion

In this article, we have provided a prior modelling of the Weibull parameters which is practical to use. Some strategies of calibration have been proposed when the prior uncertainty is difficult to assess. Considering the size of a virtual sample, yielding an information which is approximately the same than the expert information, is practical and subjective and objective data informations can be easily compared. An indicator \tilde{n} is defined to replace the size n when data \mathbf{t}_n are censored, in a more general setting than the Weibull analysis. Finally, the posterior computation remains simple.

A Appendix : proof of proposition 1

Denote $F_W(t)$ the Weibull distribution function. We have

$$\begin{aligned} P(T \leq t_e | B_1) &= \int_{\mathbb{R}} \int_{\mathbb{R}} F_W(t_e | \mu, \beta) \pi(\mu | B_1, \beta) \pi(\beta) \, d\mu d\beta, \\ &= 1 - \int_{\mathbb{R}} \left(\frac{1}{1 + \frac{x_e^\beta}{b_1(\beta)}} \right)^a \pi(\beta) \, d\beta, \\ &= 1 - \int_{\mathbb{R}} (1 - \alpha) \pi(\beta) \, d\beta = \alpha. \end{aligned}$$

This results holds even if $\text{Supp}(\pi(\beta))$ is not compact. Assuming $a > \beta_l^{-1}$,

$$\begin{aligned} E[T | B_2] &= \int_{\mathbb{R}} \int_{\mathbb{R}} \mu^{-1/\beta} \Gamma(1 + 1/\beta) \pi(\mu | B_2, \beta) \pi(\beta) \, d\mu d\beta, \\ &= \int_{\mathbb{R}} \Gamma(1 + 1/\beta) b_2^{1/\beta} (\beta) \frac{\Gamma(a - 1/\beta)}{\Gamma(a)} \pi(\beta) \, d\beta = t_e. \end{aligned}$$

This result holds if $\text{Supp}(\pi(\beta)) \subset [\beta_l, \infty[$ for any $\beta_l > 0$. Finally, a mode of distribution \mathcal{M} is located in $t_e \neq 0$ such that the derivative $m'_W(t)$ be zero. We have

$$m_W(t) = \int_{\mathbb{R}} \frac{ab^a}{(b + t^\beta)^{a+1}} \beta t^{\beta-1} \pi(\beta) \, d\beta.$$

With $\text{Supp}(\pi(\beta))$ compact, the derivative is defined and

$$m'_W(t) = - \int_{\mathbb{R}} \frac{ab^a t^{\beta-2}}{(b + t^\beta)^{a+2}} [t^\beta (a\beta + 1) - b(\beta - 1)] \pi(\beta) \, d\beta. \quad (11)$$

When $a > \beta_l^{-1}$, choosing $b = b_3$ visibly allows to obtain $m'(t_e) = 0$. Besides, $\forall x > 0$, we have

$$m'_W(t | B_3) = \int_{\mathbb{R}} \frac{ab_3^a t^{\beta-2}}{(b_3 + t^\beta)^{a+2}} (a\beta + 1) \pi(\beta) \{t_e^\beta - t^\beta\} \, d\beta$$

whose sign is the same than $t_e - t$. Then the unicity of the mode is ensured.

References

- Bacha, M. (1996). "Inférence statistique pour des modèles de durées de vie et applications", *Thèse de doctorat*, Université de Rouen.
- Bacha, M., Celeux, G., Idée, E., Lannoy, A. & Vasseur, D. (1998). *Estimation de modèles de durées de vie fortement censurées*, Eyrolles.
- Berger, J.O & Bernardo, J.M. (1992). On the development of reference priors (with discussion). In: J.M. Bernardo, J.O. Berger, A.P. Dawid and A.F.M. Smith, Eds., *Bayesian Statistics 4*, Oxford University Press, pp. 35-60.
- Bertholon, H., Bousquet, N. & Celeux, G. (2006). An alternative competing risk model to the Weibull distribution for modelling aging in lifetime data analysis, *Lifetime Data Analysis*, accepted.
- Bousquet, N. (2006). Une méthodologie d'analyse bayésienne pour la prévision de la durée de vie de composants industriels, *Ph.D. thesis manuscript*, Université Paris-XI.
- Budescu, D.V. & Rantilla, A. K. (2000). Confidence in aggregation of expert opinions. *Acta Psychologica*, **104**, pp. 371-398.
- Consonni, G., Veronese, P. & Gutierrez-Pena, E. (2004) Reference priors for natural exponential families having a simple quadratic variance function, *J. Multivariate Analysis*, **88**, pp. 335-364.
- Cooke, R.M., Mendel, M. & Thijs, W. (1988). Calibration and Information in Expert Resolution, *Automatica*, **24**, pp. 87-94.
- Cooke, R.M. & Goossens, L.H.J. (2001). *Expert judgement elicitation in risk assessment*. Nederland: Kluwer Academic Publishers.
- Daneshkhah, A.R. (2004). Psychological Aspects Influencing Elicitation of Subjective Probability, *research report*, University of Sheffield.
- Datta, G.S. (1996). On priors providing frequentist validity for Bayesian inference for multiple parametric functions, *Biometrika*, **83**, pp. 287-298.
- Dawid, A.P. & Lauritzen, S. (2000). Compatible prior distribution, in *Bayesian Methods with Application to Science Policy and Official Statistics, ISBA proceedings*, pp. 109-118.
- De Santis, F., Mortera, J. & Nardi, A. (2001). Jeffreys priors for survival models with censored data, *J. Statis. Planning and Inference*, **99**, pp. 193-209.
- Erto, P. (1982). New practical Bayes estimators for the 2-parameter Weibull distribution, *IEEE Transactions on Reliability*, **31**, pp. 194-197.
- Finkelstein, M. (2006). Aging: damage accumulation versus increasing mortality rate, *ALT 2006*, Angers.
- Garthwaite, P.H., Kadane, J.B. & O'Hagan, A. (2005). Statistical methods for eliciting probability distributions, *JASA*, **100**, pp. 680-701.
- Ghoshal, S. (1999). Probability matching priors for non-regular cases, *Biometrika*, **86**, pp. 956-964.
- Jenkinson, D. (2005). The elicitation of probabilities - A review of the statistical literature, *research report*, Open University and University of Sheffield.

- Lannoy, A. & Procaccia, H. (2001). *L'utilisation du jugement d'expert en sûreté de fonctionnement*, Tec & Doc.
- Lawless, J.F. (1982). *Statistical Models and Methods for Lifetime Data*, Wiley.
- Lawless, J.F. (2000). Statistics in Reliability, *Journal of the American Statistical Association*, **95**, pp. 989-992.
- Liisberg, C. (1991). Possible Low-Priced, robust expert systems using neural networks and minimal entropy coding, *Expert Systems with Applications*, **3**, pp. 249-257.
- Lijoi, A. (2003). Approximating priors by finite mixtures of conjugate distributions for an exponential family, *J. Statist. Planning and Inference*, **113**, pp. 419-435.
- Lin, X., Pittman, J. & Clarke, B. (2006). Bayesian Effective Samples and Parameter Size, *submitted*.
- Lindley, D., and Singpurwalla, N.D. (1986). Reliability (and Fault Tree) Analysis using expert opinions, *JASA*, **393**, pp. 87-91.
- Linstone, H.A. and Turoff, M. (eds) (2002). *The Delphi method: techniques and applications (electronic edition)*, New Jersey Institute of Technology.
- Marin, J.-M. (2006). Conjugate compatible prior distributions, *submitted*.
- Meyer, M. & Booker, J.M. (1987). *Source of correlation between experts: empirical results from two extremes*, NRC report, NUREG/CR-4814.
- O'Hagan, A. (2003). HSSS model criticism (with discussion). In: *Highly Structured Stochastic Systems*, P. J. Green, N. L. Hjort and S. T. Richardson (eds), Oxford University Press, pp. 423-453.
- O'Hagan, A. (2005). Elicitation, *Significance*, June, pp. 84-86.
- Robert, C.P. (2001). *The Bayesian Choice. A Decision-Theoretic Motivation* (second edition), Springer.
- Schieren, G.A. (1993). Median Worklife, Mean Age at Final Separation, or Transition Probabilities to calculate Expected Lost Earnings?, *Journal of Forensic Economics*, **1**, pp. 103-109.
- Singpurwalla, N.D. & Song, M.S. (1986). An analysis of Weibull lifetime data incorporating expert opinion, in *Probability and Bayesian Statistics* (R.Viertl ed.), Plenum Pub.Corp., pp. 431-442.
- Singpurwalla, N.D. (1988). An interactive PC-Based procedure for reliability assessment incorporating expert opinion and survival data, *JASA*, **83**, pp. 43-51.
- Soland, R. (1969). Bayesian analysis of the Weibull process with unknown scale and shape parameters, *IEEE Transactions on Reliability*, **18**, pp. 181-184.
- Sun, D. (1997). A note on noninformative priors for Weibull distributions, *J. Statist. Planning and Inference*, **61**, pp. 319-338.
- van Noortwijk, J.M., Dekker, R., Cooke, R.M. & Mazzuchi, T.A. (1992). Expert judgment in Maintenance Optimization, *IEEE Transactions on Reliability*, **41**, pp. 427-431.
- Wisse, B., Bedford, T. & Quigley, J. (2005). Combining Expert Judgements in the Bayes Linear Methodology, *Proceedings of the Workshop on the Use of Expert Judgement in Decision-Making*.
- Wu, S.J. (2002). Estimations of the the parameters of the Weibull distribution with progressively censored data, *J. Japan Statis. Soc.*, **32**, pp. 155-163.



Unité de recherche INRIA Futurs
Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399