

Robust Bayesian reliability for complex systems under prior-data conflict

Citation for published version (APA):

Walter, G., & Coolen, F. P. A. (2018). Robust Bayesian reliability for complex systems under prior-data conflict. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, 4(3), Article 04018025. <https://doi.org/10.1061/AJRUA6.0000974>

DOI:

[10.1061/AJRUA6.0000974](https://doi.org/10.1061/AJRUA6.0000974)

Document status and date:

Published: 01/09/2018

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

ROBUST BAYESIAN RELIABILITY FOR COMPLEX SYSTEMS UNDER PRIOR-DATA CONFLICT

Gero Walter¹, Ph.D.
and Frank P.A. Coolen², Ph.D.

Abstract

In reliability engineering, data about failure events is often scarce. To arrive at meaningful estimates for the reliability of a system, it is therefore often necessary to also include expert information in the analysis, which is straightforward in the Bayesian approach by using an informative prior distribution. A problem called prior-data conflict then can arise: observed data seem very surprising from the viewpoint of the prior, i.e., information from data is in conflict with prior assumptions. Models based on conjugate priors can be insensitive to prior-data conflict, in the sense that the spread of the posterior distribution does not increase in case of such a conflict, thus conveying a false sense of certainty. An approach to mitigate this issue is presented, by considering sets of prior distributions to model limited knowledge on Weibull distributed component lifetimes, treating systems with arbitrary layout using the survival signature. This approach can be seen as a robust Bayesian procedure or imprecise probability method that reflects surprisingly early or late component failures by wider system reliability bounds.

Keywords: System Reliability, Imprecise Probability, Survival Signature, Robust Bayesian Methods, Remaining Useful Life

1 INTRODUCTION

In reliability engineering, a central task is to describe the reliability of a complex system. This is usually done by determining the *reliability function* $R(t)$, in other contexts also known as the *survival function*, giving the probability that the system has not failed by time t :

$$R_{\text{sys}}(t) = P(T_{\text{sys}} \geq t), \quad (1)$$

where T_{sys} is the random variable giving the failure time of the system. Based on the distribution of T_{sys} , which can also be expressed in terms of the cumulative distribution function $F_{\text{sys}}(t) = 1 - R_{\text{sys}}(t)$, decisions about, e.g., scheduling of maintenance work can be made.

Often, there is no failure data for the system itself (e.g., if the system is a prototype, or the system is used under unique circumstances), but there is some information about failure times for the components the system is made of. The proposed method allows to analyse systems of arbitrary system structures, i.e., any combination and nesting of series, parallel, k-out-of-n, or bridge-type arrangements, by use of the survival signature (Coolen and Coolen-Maturi 2012). In doing so, the present paper extends similar previous work which focused on

¹Postdoctoral Researcher, School of Industrial Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands. E-mail: g.m.walter@tue.nl. Corresponding author.

²Professor, Department of Mathematical Sciences, Durham University, Durham, United Kingdom. E-mail: frank.coolen@durham.ac.uk.

a simple parallel system with homogeneous components (Walter et al. 2015). In this paper, we assume that components can be divided into K different groups, and components within each group k ($k = 1, \dots, K$) can be assumed to be exchangeable, i.e., to follow the same failure time distribution. Components of group k are denoted as *type k components*, and are assumed to be independent from components of other types. Type k component lifetimes T_i^k ($i = 1, \dots, n_k$) are assumed as Weibull distributed with type-specific parameters, where the shape parameter β_k is known and the scale parameter λ_k is unknown.

Focusing on a running system, we assume that observations consist solely of the failure times of components in this system up to time t_{now} , such that the failure times of components that have not failed by t_{now} are right-censored, and calculate $R_{\text{sys}}(t \mid t > t_{\text{now}}) = P(T_{\text{sys}} \geq t \mid t > t_{\text{now}})$, which can be used to determine the remaining useful life of the system (in short RUL, see, e.g., Ahmadzadeh and Lundberg 2014).

The Bayesian approach allows to base estimation of the component failure distributions on both data and further knowledge not given by the data, the latter usually provided in the form of expert knowledge. This knowledge is encoded in form of a so-called prior distribution, which here, as the shape parameter β_k is assumed to be known, is a distribution on the scale parameter λ_k . Expert knowledge is especially important when there is very few data on the components, as only with its help meaningful estimates for the system reliability can be made. Due to the iterative nature of the Bayesian framework, it is also possible to use a set of posteriors based on component test data as the set of priors instead of a purely expert-based set of priors as discussed so far. The component test data may also contain right-censored observations, which can be treated in the same way as those from the running system (see Section 4.1 for details).

The choice of prior distribution to encode given expert knowledge is often debatable, and a specific choice of prior is difficult to justify. A way to deal with this is to employ sensitivity analysis, i.e., studying the effect of different choices of prior distribution on the quantities of interest (in our case, the system reliability function, which, in Bayesian terms, is a posterior predictive distribution). This idea has been explored in systematic sensitivity analysis, or robust Bayesian methods (for an overview on this approach, see, e.g., Berger et al. 1994 or Ríos Insua and Ruggeri 2000).

The work presented here can be seen as belonging to the robust Bayesian approach since it uses sets of priors. However, our focus and interpretation is slightly different, as we consider the result of our procedure, sets of reliability functions, as the proper result, while a robust Bayesian would base his analyses on a single reliability function from the set in case (s)he was able to conclude that quantities of interest are not ‘too sensitive’ to the choice of prior. In contrast, our viewpoint is rooted in the theory of imprecise or interval probability (Walley 1991; Augustin et al. 2014), where sets of distributions are used to express the precision of probability statements themselves: the smaller the set, the more precise the probability statement. Indeed, the system reliability function $R_{\text{sys}}(t)$ is a collection of probability statements, and a small set for $R_{\text{sys}}(t)$ will indicate that the reliability of the system can be quantified quite precisely, while a large set will indicate that available knowledge about T_{sys} is rather shaky. In line with imprecise or interval probability methods, the method provides, for each t , a lower reliability $\underline{R}_{\text{sys}}(t) = \underline{P}(T_{\text{sys}} \geq t)$, and an upper reliability $\overline{R}_{\text{sys}}(t) = \overline{P}(T_{\text{sys}} \geq t)$. Sections 3 and 4 will explain how these bounds

are obtained based on sets of prior distributions on the scale parameters of the component lifetime distributions.

The central merit of the proposed method is that it adequately reflects prior-data conflict (see, e.g., Evans and Moshonov 2006), i.e. the conflict that can arise between prior assumptions on component lifetimes and observed behaviour of components in the system under study. As will be shown in Section 2, when taking the standard choice of a conjugate prior, prior-data conflict is ignored, as the spread of the posterior distribution does not increase in case of such a conflict, ultimately conveying a false sense of certainty by communicating that the reliability of a system can be quantified quite precisely when in fact it can not. In contrast, the proposed method will indicate prior-data conflict by wider bounds for $R_{\text{sys}}(t)$. This behaviour is obtained by a specific choice for the set of priors (see Walter and Augustin 2009 and Walter 2013 §3.1.4) which leads to larger sets of posterior distributions when prior knowledge and data are in conflict (see Section 3 for more details). If a prior based on component test data is used, such prior-data conflict sensitivity furthermore allows to uncover a conflict between current observations in the running system and past observations from component tests.

The paper is organized as follows. In Section 2, a Bayesian analysis for Weibull component lifetimes with fixed shape parameter is described, illustrating the issue of prior-data conflict. Section 3 then details the use of sets of priors for the scale parameter λ_k of the Weibull distribution, showing how this mitigates the prior-data conflict issue for λ_k . Section 4 then develops a method to calculate system reliability bounds for a running system based on sets of priors for λ_k . Section 5 discusses elicitation, by giving guidelines on how to choose prior parameter bounds that reflect expert information on components. Section 6 contains examples illustrating the merits of our method, by studying the effect of surprisingly early or late component failures, showing that observations in conflict to prior assumptions indeed lead to more cautious system reliability predictions. Section 7 concludes the paper by summarizing results and indicating avenues for further research.

2 BAYESIAN ANALYSIS OF WEIBULL LIFETIMES

Consider a system with components of $k = 1, \dots, K$ different types; for each type k , there are n_k exchangeable components in the system. For each type k component, we assume for its lifetime T_i^k ($i = 1, \dots, n_k$, $k = 1, \dots, K$) a Weibull distribution with fixed shape parameter $\beta_k > 0$, in short $T_i^k | \lambda_k \sim \text{Wei}(\beta_k, \lambda_k)$, with density and cdf

$$f(t_i^k | \lambda_k) = \frac{\beta_k}{\lambda_k} (t_i^k)^{\beta_k - 1} e^{-\frac{(t_i^k)^{\beta_k}}{\lambda_k}}, \quad (2)$$

$$F(t_i^k | \lambda_k) = 1 - e^{-\frac{(t_i^k)^{\beta_k}}{\lambda_k}} = P(T_i^k \leq t_i^k | \lambda_k), \quad (3)$$

where $\lambda_k > 0$ and $t > 0$.

The shape parameter β_k determines whether the hazard rate is increasing ($\beta_k > 1$) or decreasing ($\beta_k < 1$) over time. For $\beta_k = 1$, one obtains the Exponential distribution with constant hazard rate as a special case. The value for β_k will often be clear from practical considerations.

The scale parameter λ_k can be interpreted through the relation

$$E[T_i^k | \lambda_k] = \lambda_k^{1/\beta_k} \Gamma(1 + 1/\beta_k). \quad (4)$$

For encoding expert knowledge about the reliability of the components, one needs to assign a prior distribution over the scale parameter λ_k . A convenient choice is to use the inverse Gamma distribution, commonly parametrized in terms of the hyperparameters $a_k > 0$ and $b_k > 0$:

$$f(\lambda_k | a_k, b_k) = \frac{(b_k)^{a_k}}{\Gamma(a_k)} \lambda_k^{-a_k-1} e^{-\frac{b_k}{\lambda_k}} \quad (5)$$

in short, $\lambda_k | a_k, b_k \sim \text{IG}(a_k, b_k)$. The inverse Gamma is convenient because it is a conjugate prior, i.e., the posterior obtained by Bayes' rule is again inverse Gamma and thus easily tractable; the prior parameters only need to be updated to obtain the posterior parameters.

In the standard Bayesian approach, one has to fix a prior by choosing values for a_k and b_k to encode specific prior information about component lifetimes. In our imprecise approach, we allow instead these parameters to vary in a set, this is advantageous also because expert knowledge is often vague, and it is difficult for the expert(s) to pin down precise hyperparameter values. For the definition of the hyperparameter set, we use however a parametrization in terms of $n^{(0)} > 1$ and $y^{(0)} > 0$ instead of a_k and b_k , where we drop the index k for the discussion about the prior model in the following, keeping in mind that each component type will have its own specific parameters. We use $n^{(0)} = a - 1$ and $y^{(0)} = b/n^{(0)}$, where $y^{(0)}$ can be interpreted as the prior guess for the scale parameter λ , as $\text{E}[\lambda | n^{(0)}, y^{(0)}] = y^{(0)}$. This parametrization also makes the nature of the combination of prior information and data through Bayes' rule more clear: After observing n component lifetimes $\mathbf{t} = (t_1, \dots, t_n)$, the updated parameters are

$$n^{(n)} = n^{(0)} + n, \quad y^{(n)} = \frac{n^{(0)}y^{(0)} + \tau(\mathbf{t})}{n^{(0)} + n}, \quad (6)$$

where $\tau(\mathbf{t}) = \sum_{i=1}^n (t_i)^\beta$. We thus have

$$\lambda | n^{(0)}, y^{(0)}, \mathbf{t} \sim \text{IG}(n^{(0)} + n + 1, n^{(0)}y^{(0)} + \tau(\mathbf{t})). \quad (7)$$

From the simple update rule (6), we see that $y^{(n)}$ is a weighted average of the prior parameter $y^{(0)}$ and the maximum likelihood (ML) estimator $\tau(\mathbf{t})/n$, with weights proportional to $n^{(0)}$ and n , respectively. $n^{(0)}$ can thus be interpreted as a prior strength or pseudocount, indicating how much our prior guess should weigh against the n observations. Furthermore, $\text{Var}[\lambda | n^{(0)}, y^{(0)}] = (y^{(0)})^2 / (1 - 1/n^{(0)})$, so for fixed $y^{(0)}$, the higher $n^{(0)}$, the more probability mass is concentrated around $y^{(0)}$.

However, the weighted average structure for $y^{(n)}$ is behind the problematic behaviour in case of prior-data conflict. Assume that from expert knowledge we expect to have a mean component lifetime of 9 weeks. Using (4), with $\beta = 2$ we obtain $y^{(0)} = 103.13$. We choose $n^{(0)} = 2$, so our prior guess for the mean component lifetime counts like having two observations with this mean. If we now have a sample of two observations with surprisingly early failure times $t_1 = 1$ and $t_2 = 2$, using (6) we get $n^{(2)} = 4$ and $y^{(2)} = \frac{1}{4}(2 \cdot 103.13 + 1^2 + 2^2) = 52.82$, so our posterior expectation for the scale parameter λ is 52.82, equivalent to a mean component lifetime of 6.44 weeks. The posterior standard deviation (sd) for λ is 60.99. Compared to the prior standard deviation of 145.85, the posterior expresses




Figure 1. Prior and posterior cdf for λ given surprising observations; the conflict between prior assumptions and data is averaged out, with a more pointed posterior giving a false sense of certainty.

now more confidence that mean lifetimes are around $y^{(2)} = 52.82$ than the prior had about $y^{(0)} = 103.13$. This irritating conclusion is illustrated in Figure 1; the posterior cdf is shifted halfway towards the values for λ that the two observations suggest (the ML estimator for λ would be 2.5), and is steeper than the prior (so the pdf is more pointed), thus conveying a false sense of certainty about λ . We would obtain almost the same posterior distribution if we had assumed the mean component lifetime to be 7 weeks (so $y^{(0)} = 62.39$), and observed lifetimes $t_1 = 6$, $t_2 = 7$ in line with our expectations. It seems unreasonable to make the same probability statements on component lifetimes in these two fundamentally different scenarios.

Note that this is a general problem in Bayesian analysis with canonical conjugate priors. For such priors, the same update formula (6) applies, and so conflict is averaged out, for details see Walter and Augustin (2009) and Walter (2013, §3.1.4 and §A.1.2).

3 MODELS REFLECTING SURPRISING DATA

Despite this issue of ignoring prior-data conflict, the tractability of the update step is a very attractive feature of the conjugate setting. As was shown by Walter and Augustin (2009) (see also Walter 2013, §3.1), it is possible to retain tractability and to have a meaningful reaction to prior-data conflict when using sets of priors generated by varying both $n^{(0)}$ and $y^{(0)}$. Then, the magnitude of the set of posteriors, and with it the precision of posterior probability statements, will be sensitive to the degree of prior-data conflict, i.e. leading to more cautious probability statements when prior-data conflict occurs.

Instead of a single prior guess $y^{(0)}$ for the mean component lifetimes, we will now assume a range of prior guesses $[\underline{y}^{(0)}, \bar{y}^{(0)}]$, and also a range $[\underline{n}^{(0)}, \bar{n}^{(0)}]$ of pseudocounts, i.e., we now consider the set of priors

$$\mathcal{M}^{(0)} := \{f(\lambda \mid n^{(0)}, y^{(0)}) \mid n^{(0)} \in [\underline{n}^{(0)}, \bar{n}^{(0)}], y^{(0)} \in [\underline{y}^{(0)}, \bar{y}^{(0)}]\} \quad (8)$$

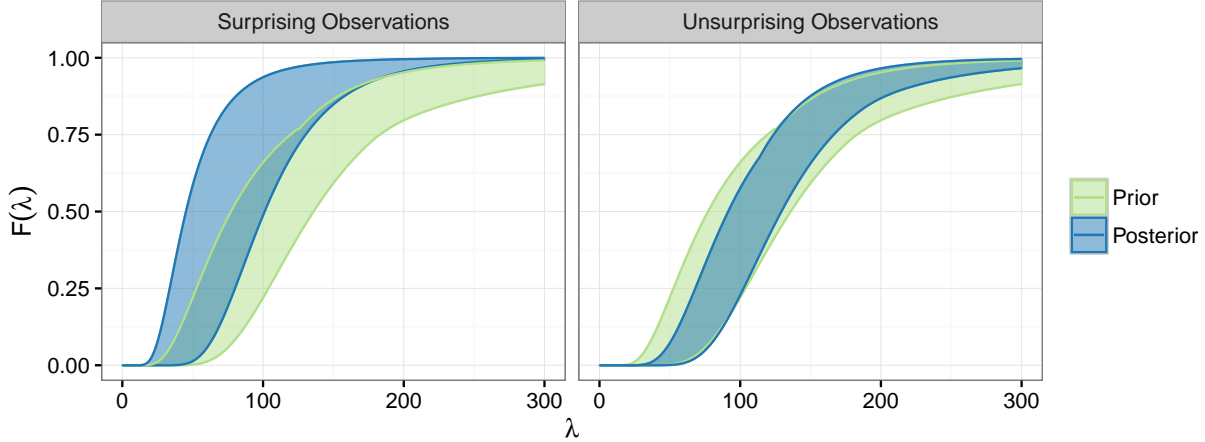


Figure 2. Set of prior and posterior cdfs for λ for two surprising observations $t_1 = 1$, $t_2 = 2$ (left) and two unsurprising observations $t_1 = 10$, $t_2 = 11$ (right).

to express our prior knowledge about component lifetimes. Each of the priors $f(\lambda | n^{(0)}, y^{(0)})$ is then updated to the posterior $f(\lambda | n^{(0)}, y^{(0)}, \mathbf{t}) = f(\lambda | n^{(n)}, y^{(n)})$ by using (6), such that the set of posteriors $\mathcal{M}^{(n)}$ can be written as $\mathcal{M}^{(n)} = \{f(\lambda | n^{(n)}, y^{(n)}) | n^{(0)} \in [\underline{n}^{(0)}, \bar{n}^{(0)}], y^{(0)} \in [\underline{y}^{(0)}, \bar{y}^{(0)}]\}$. This procedure of using Bayes' Rule element by element is seen as self-evident in the robust Bayesian literature, but can be formally justified as being *coherent* (a self-consistency property) in the framework of imprecise probability, where it is known as *Generalized Bayes' Rule* (Walley 1991, §6.4).

Technically, it is crucial to consider a range of pseudocounts $[\underline{n}^{(0)}, \bar{n}^{(0)}]$ along with the range of prior guesses $[\underline{y}^{(0)}, \bar{y}^{(0)}]$, as only then $\tau(\mathbf{t})/n \notin [\underline{y}^{(0)}, \bar{y}^{(0)}]$ leads to the set of posteriors being larger and hence reflecting prior-data conflict.

Continuing the example from Section 2 and Figure 1, assume now for the mean component lifetimes the range 9 to 11 weeks, this corresponds to $[\underline{y}^{(0)}, \bar{y}^{(0)}] = [103.13, 154.06]$. Choosing $[\underline{n}^{(0)}, \bar{n}^{(0)}] = [2, 5]$ means to value this information on mean component lifetimes as equivalent to having seen two to five observations. Compare now the set of posteriors obtained from observing $t_1 = 1$, $t_2 = 2$ (as before), see Figure 2 (left), and $t_1 = 10$, $t_2 = 11$, see Figure 2 (right). There is now a clear difference between the two scenarios of observations in line with expectations and observations in conflict. In the prior-data conflict case, the set of posteriors (blue) is shifted towards the left, but has about the same size as the set of priors (yellow), and so posterior quantification of reliability has the same precision, despite having seen two observations. Instead, in the no conflict case, the set of posteriors is smaller than the set of priors, such that the two observations have increased the precision of reliability statements.

As each posterior in $\mathcal{M}^{(n)}$ corresponds to a predictive distribution for T_{sys} , we will have a set of reliability functions $R_{\text{sys}}(t)$. The derivation of $R_{\text{sys}}(t)$ for systems with K component types and arbitrary layout will be given in Section 4 below. This will include, in contrast to previous studies using sets of priors of type (8), the treatment of censored observations. More specifically, we consider the case of *non-informative right censoring*, where the censoring process is independent of the failure process.

4 ROBUST RELIABILITY FOR COMPLEX SYSTEMS VIA THE SURVIVAL SIGNATURE

Consider now a system of arbitrary layout, consisting of components of K types. This system is observed until time t_{now} , leading to censored observation of lifetimes of components within the system, and we will explain in Section 4.1 how the scale parameter λ_k for component type k can be estimated in this situation. Section 4.2 describes then how the system reliability function can be efficiently obtained using the survival signature. For this, we need the posterior predictive distribution of the number of components that function at times $t > t_{\text{now}}$, which we will derive in Section 4.3. Finally, in Section 4.4 we describe how the lower and upper bound for the system reliability function are obtained when prior parameters $(n_k^{(0)}, y_k^{(0)})$ vary in sets $\Pi_k^{(0)} = [\underline{n}_k^{(0)}, \bar{n}_k^{(0)}] \times [\underline{y}_k^{(0)}, \bar{y}_k^{(0)}]$, defining sets $\mathcal{M}_k^{(0)}$ of prior distributions over λ_k , as introduced in Section 3.

4.1 Bayesian Estimation of Component Scale Parameter with Right-censored Lifetimes

Consider observing a system until t_{now} , where the system has K different types of components, and for each type k there are n_k components in the system. Denoting the number of type k components that have failed by t_{now} by e_k , there are $n_k - e_k$ components still functioning at t_{now} . We denote the corresponding vector of observations by

$$\mathbf{t}_{e_k; n_k}^k = \left(\underbrace{t_1^k, \dots, t_{e_k}^k}_{e_k \text{ failure times}}, \underbrace{t_{\text{now}}^+, \dots, t_{\text{now}}^+}_{n_k - e_k \text{ censored obs.}} \right), \quad (9)$$

where t^+ indicates a right-censored observation.

According to Bayes' rule, multiplying the prior density and the likelihood (which accounts for right-censored observations through the cdf terms) gives a term proportional to the density of the posterior distribution for λ_k :

$$f(\lambda_k | n_k^{(0)}, y_k^{(0)}, \mathbf{t}_{e_k; n_k}^k) \propto f(\lambda_k) [1 - F(t_{\text{now}} | \lambda_k)]^{n_k - e_k} \prod_{i=1}^{e_k} f(t_i^k | \lambda_k) \quad (10)$$

Conjugacy is preserved and we get $\lambda_k | n_k^{(0)}, y_k^{(0)}, \mathbf{t}_{e_k; n_k}^k \sim \text{IG}(n_k^{(n)} + 1, n_k^{(n)} y_k^{(n)})$, where

$$n_k^{(n)} + 1 = n_k^{(0)} + e_k + 1 \quad (11)$$

$$n_k^{(n)} y_k^{(n)} = n_k^{(0)} y_k^{(0)} + (n_k - e_k) (t_{\text{now}})^\beta + \sum_{i=1}^{e_k} (t_i^k)^\beta \quad (12)$$

are the updated parameters of the inverse gamma distribution.

4.2 System Reliability using the Survival Signature

The structure of complex systems can be visualized by reliability block diagrams, an example is given in Figure 3. Components are represented by boxes or nodes, and the system works when a path from the left end to the right exists which passes only through working components. In a system with n components, the state of the components can be expressed by the state vector $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \{0, 1\}^n$, with $x_i = 1$ if the i th component functions and $x_i = 0$ if not. The structure function $\phi : \{0, 1\}^n \rightarrow \{0, 1\}$, defined for all possible \mathbf{x} ,

takes the value 1 if the system functions and 0 if the system does not function for state vector \mathbf{x} (Barlow and Proschan 1975). Most real-life systems are coherent, which means that $\phi(\mathbf{x})$ is non-decreasing in any of the components of \mathbf{x} , so system functioning cannot be improved by worse performance of one or more of its components. Furthermore, one can usually assume that $\phi(0, \dots, 0) = 0$ and $\phi(1, \dots, 1) = 1$.

The survival signature (Coolen and Coolen-Maturi 2012) is a summary of the structure function for systems with K groups of exchangeable components. Denoted by $\Phi(l_1, \dots, l_K)$, with $l_k = 0, 1, \dots, n_k$ for $k = 1, \dots, K$, it is defined as the probability for the event that the system functions given that precisely l_k of its n_k components of type k function, for each $k \in \{1, \dots, K\}$. Essentially, this creates a K -dimensional partition for the event $T_{\text{sys}} > t$, such that $R_{\text{sys}}(t) = P(T_{\text{sys}} > t)$ can be calculated using the law of total probability,

$$\begin{aligned} P(T_{\text{sys}} > t) &= \sum_{l_1=0}^{m_1} \cdots \sum_{l_K=0}^{m_K} P(T_{\text{sys}} > t \mid C_t^1 = l_1, \dots, C_t^K = l_K) P\left(\bigcap_{k=1}^K \{C_t^k = l_k\}\right) \\ &= \sum_{l_1=0}^{n_1} \cdots \sum_{l_K=0}^{n_K} \Phi(l_1, \dots, l_K) P\left(\bigcap_{k=1}^K \{C_t^k = l_k\}\right) \\ &= \sum_{l_1=0}^{n_1} \cdots \sum_{l_K=0}^{n_K} \Phi(l_1, \dots, l_K) \prod_{k=1}^K P(C_t^k = l_k), \end{aligned} \quad (13)$$

where $P(C_t^k = l_k)$ is the (predictive) probability that exactly l_k components of type k function at time t , and the last equality holds as we assume that components of different types are independent. Note that for coherent systems, the survival signature $\Phi(l_1, \dots, l_K)$ is non-decreasing in each l_k .

4.3 Posterior Predictive Distribution

In calculating the system reliability using (13), the component-specific predictive probabilities $P(C_t^k = l_k)$ need to use all information available at time t_{now} , which, in the Bayesian framework, are given by the posterior predictive distribution $P(C_t^k = l_k \mid n_k^{(0)}, y_k^{(0)}, \mathbf{t}_{e_k; n_k}^k)$, $l_k = 0, 1, \dots, n_k - e_k$. (Remember that e_k type k components have failed by t_{now} , such that there can be at most $n_k - e_k$ working components beyond time t_{now} .) This posterior predictive distribution is obtained as

$$\begin{aligned} &P(C_t^k = l_k \mid n_k^{(0)}, y_k^{(0)}, \mathbf{t}_{e_k; n_k}^k) \\ &= \binom{n_k - e_k}{l_k} \int [P(T^k > t \mid T^k > t_{\text{now}}, \lambda_k)]^{l_k} \times \\ &\quad [P(T^k \leq t \mid T^k > t_{\text{now}}, \lambda_k)]^{n_k - e_k - l_k} f(\lambda_k \mid n_k^{(0)}, y_k^{(0)}, \mathbf{t}_{e_k; n_k}^k) d\lambda_k. \end{aligned} \quad (14)$$

Now, by the Weibull assumption (3), one has

$$\begin{aligned} P(T^k \leq t \mid T^k > t_{\text{now}}, \lambda_k) &= \frac{P(t_{\text{now}} < T^k \leq t \mid \lambda_k)}{P(T^k > t_{\text{now}} \mid \lambda_k)} \\ &= \frac{F(t \mid \lambda_k) - F(t_{\text{now}} \mid \lambda_k)}{1 - F(t_{\text{now}} \mid \lambda_k)} = 1 - e^{-\frac{t^{\beta k} - (t_{\text{now}})^{\beta k}}{\lambda_k}}. \end{aligned} \quad (15)$$

With this and the posterior (7) substituted into (14), this gives

$$\begin{aligned}
P(C_t^k = l_k | n_k^{(0)}, y_k^{(0)}, \mathbf{t}_{e_k; n_k}^k) &= \binom{n_k - e_k}{l_k} \int \left[e^{-\frac{t^{\beta_k} - (t_{\text{now}})^{\beta_k}}{\lambda_k}} \right]^{l_k} \left[1 - e^{-\frac{t^{\beta_k} - (t_{\text{now}})^{\beta_k}}{\lambda_k}} \right]^{n_k - e_k - l_k} \times \\
&\quad \frac{(n_k^{(n)} y_k^{(n)})^{n_k^{(n)} + 1}}{\Gamma(n_k^{(n)} + 1)} \lambda_k^{-(n_k^{(n)} + 1) - 1} e^{-\frac{n_k^{(n)} y_k^{(n)}}{\lambda_k}} d\lambda_k \\
&= \binom{n_k - e_k}{l_k} \sum_{j=0}^{n_k - e_k - l_k} (-1)^j \binom{n_k - e_k - l_k}{j} \frac{(n_k^{(n)} y_k^{(n)})^{n_k^{(n)} + 1}}{\Gamma(n_k^{(n)} + 1)} \times \\
&\quad \int \lambda_k^{-(n_k^{(n)} + 1) - 1} \exp \left\{ -\frac{(l_k + j)(t^{\beta_k} - (t_{\text{now}})^{\beta_k}) + n_k^{(n)} y_k^{(n)}}{\lambda_k} \right\} d\lambda_k. \quad (16)
\end{aligned}$$

The terms remaining under the integral form the core of an inverse gamma distribution (5) with parameters $n_k^{(n)} + 1$ and $n_k^{(n)} y_k^{(n)} + (l_k + j)(t^{\beta_k} - (t_{\text{now}})^{\beta_k})$, allowing to solve the integral using the corresponding normalization constant. We thus have, for $l_k \in \{0, 1, \dots, n_k - e_k\}$,

$$\begin{aligned}
P(C_t^k = l_k | n_k^{(0)}, y_k^{(0)}, \mathbf{t}_{e_k; n_k}^k) &= \binom{n_k - e_k}{l_k} \sum_{j=0}^{n_k - e_k - l_k} (-1)^j \binom{n_k - e_k - l_k}{j} \left(\frac{n_k^{(n)} y_k^{(n)}}{n_k^{(n)} y_k^{(n)} + (l_k + j)(t^{\beta_k} - (t_{\text{now}})^{\beta_k})} \right)^{n_k^{(n)} + 1} \\
&= \sum_{j=0}^{n_k - e_k - l_k} (-1)^j \frac{(n_k - e_k)!}{l_k! j! (n_k - e_k - l_k - j)!} \left(\frac{n_k^{(n)} y_k^{(n)}}{n_k^{(n)} y_k^{(n)} + (l_k + j)(t^{\beta_k} - (t_{\text{now}})^{\beta_k})} \right)^{n_k^{(n)} + 1} \\
&= \sum_{j=0}^{n_k - e_k - l_k} (-1)^j \frac{(n_k - e_k)!}{l_k! j! (n_k - e_k - l_k - j)!} \times \\
&\quad \left(\frac{n_k^{(0)} y_k^{(0)} + \sum_{i=1}^{e_k} (t_i^k)^{\beta_k} + (n_k - e_k)(t_{\text{now}})^{\beta_k}}{n_k^{(0)} y_k^{(0)} + \sum_{i=1}^{e_k} (t_i^k)^{\beta_k} + (n_k - e_k - l_k - j)(t_{\text{now}})^{\beta_k} + (l_k + j)t^{\beta_k}} \right)^{n_k^{(0)} + e_k + 1}. \quad (17)
\end{aligned}$$

These posterior predictive probabilities can also be expressed as a cumulative probability mass function (cmf)

$$F(l_k | n_k^{(0)}, y_k^{(0)}, \mathbf{t}_{e_k; n_k}^k) = P(C_t^k \leq l_k | n_k^{(0)}, y_k^{(0)}, \mathbf{t}_{e_k; n_k}^k) = \sum_{j=0}^{l_k} P(C_t^k = j | n_k^{(0)}, y_k^{(0)}, \mathbf{t}_{e_k; n_k}^k). \quad (18)$$

4.4 Optimizing over Sets of Parameters

Together with (17), (13) allows to calculate the system reliability $R_{\text{sys}}(t | t > t_{\text{now}})$ for fixed prior parameters $(n_k^{(0)}, y_k^{(0)})$, $k = 1, \dots, K$. In Section 3, we argued for using sets of priors $\mathcal{M}^{(0)}$, which allow for vague and incomplete prior knowledge, and provide prior-data

conflict sensitivity. We will thus use, for each component type, a set of priors $\mathcal{M}_k^{(0)}$ defined by varying $(n_k^{(0)}, y_k^{(0)})$ in a prior parameter set $\Pi_k^{(0)} = [\underline{n}_k^{(0)}, \bar{n}_k^{(0)}] \times [\underline{y}_k^{(0)}, \bar{y}_k^{(0)}]$, and the objective is to obtain the bounds

$$\underline{R}_{\text{sys}}(t | t > t_{\text{now}}) = \min_{\Pi_1^{(0)}, \dots, \Pi_K^{(0)}} R_{\text{sys}}(t | t > t_{\text{now}}, \cup_{k=1}^K \{\Pi_k^{(0)}, \mathbf{t}_{e_k; n_k}^k\}), \quad (19)$$

$$\bar{R}_{\text{sys}}(t | t > t_{\text{now}}) = \max_{\Pi_1^{(0)}, \dots, \Pi_K^{(0)}} R_{\text{sys}}(t | t > t_{\text{now}}, \cup_{k=1}^K \{\Pi_k^{(0)}, \mathbf{t}_{e_k; n_k}^k\}), \quad (20)$$

where we suppress in notation that $\underline{R}_{\text{sys}}(t | t > t_{\text{now}})$ and $\bar{R}_{\text{sys}}(t | t > t_{\text{now}})$ depend on prior parameter sets and data.

(19) and (20) seem to suggest that a full $2K$ -dimensional box-constraint optimization is necessary, but this is not the case. Remember that $\Phi(l_1, \dots, l_K)$ from (13) is non-decreasing in each of its arguments l_1, \dots, l_K , so if there is stochastic dominance in $F(l_k | n_k^{(0)}, y_k^{(0)}, \mathbf{t}_{e_k; n_k}^k)$, then there is, for each component type k , a prior parameter pair in $\Pi_k^{(0)}$ that minimizes system reliability, and a prior parameter pair in $\Pi_k^{(0)}$ that maximizes system reliability, independently of the other component types. Indeed, stochastic dominance in $F(l_k | n_k^{(0)}, y_k^{(0)}, \mathbf{t}_{e_k; n_k}^k)$ is provided for $y_k^{(0)}$. To see this, note that $y_k^{(0)}$ gives the mean expected lifetime for type k components. Thus, higher values for $y_k^{(0)}$ mean higher expected lifetimes for the components, which in turn increases the probability that many components survive until time t , and with it, decreases the probability of few or no components surviving, so in total giving low probability weight for low values of l_k , and high probability weight for high values of l_k . Therefore, for any fixed value of $n_k^{(0)}$, the lower bound of $F(l_k | n_k^{(0)}, y_k^{(0)}, \mathbf{t}_{e_k; n_k}^k)$ for all l_k is obtained with $\bar{y}_k^{(0)}$, and the upper bound of $F(l_k | n_k^{(0)}, y_k^{(0)}, \mathbf{t}_{e_k; n_k}^k)$ for all l_k is obtained with $\underline{y}_k^{(0)}$. There is however no corresponding result for $n_k^{(0)}$, such that different values of $n_k^{(0)}$ may minimize (or maximize) $F(l_k | n_k^{(0)}, y_k^{(0)}, \mathbf{t}_{e_k; n_k}^k)$ at different l_k 's. Therefore, the $n_k^{(0)}$ values for lower and upper system reliability bounds are obtained by numeric optimization.

Writing out (13), one obtains

$$\begin{aligned} \underline{R}_{\text{sys}}(t | t > t_{\text{now}}) &= \min_{\Pi_1^{(0)}, \dots, \Pi_K^{(0)}} R_{\text{sys}}(t | t > t_{\text{now}}, \cup_{k=1}^K \{\Pi_k^{(0)}, \mathbf{t}_{e_k; n_k}^k\}) \\ &= \min_{\substack{n_1^{(0)} \in [\underline{n}_1^{(0)}, \bar{n}_1^{(0)}] \\ \vdots \\ n_K^{(0)} \in [\underline{n}_K^{(0)}, \bar{n}_K^{(0)}]}} \sum_{l_1=0}^{n_1 - e_1} \cdots \sum_{l_K=0}^{n_K - e_K} \Phi(l_1, \dots, l_K) \prod_{k=1}^K P(C_t^k = l_k | n_k^{(0)}, \underline{y}_k^{(0)}, \mathbf{t}_{e_k; n_k}^k), \end{aligned} \quad (21)$$

such that a K -dimensional box-constraint optimization is needed to obtain $\underline{R}_{\text{sys}}(t | t > t_{\text{now}})$. The result for $\bar{R}_{\text{sys}}(t | t > t_{\text{now}})$ is completely analogous. Computing time can furthermore be saved by computing only those summation terms for which $\Phi(l_1, \dots, l_K) > 0$.

We have implemented the method in the statistical computing environment R (R Core Team 2016), using box-constraint optimization via option L-BFGS-B of `optim`, and intend to release a package containing functions and code to reproduce all results and figures for the examples in Section 6 below.

5 ELICITATION OF PRIOR PARAMETER SETS

To represent expert knowledge on component failure times through bounds for $y_k^{(0)}$ and $n_k^{(0)}$, one can refer to the interpretations as given in Section 2: $y_k^{(0)}$ is the prior expected value of λ_k , where λ_k is linked to expected component lifetimes through (4). $n_k^{(0)}$ can be seen as pseudocount, indicating how strong expert knowledge is trusted in comparison to a sample of size n . Crucially, the approach allows the expert to give ranges $[\underline{y}_k^{(0)}, \bar{y}_k^{(0)}]$ and $[\underline{n}_k^{(0)}, \bar{n}_k^{(0)}]$ instead of requiring a precise answer.

It is also possible to directly link $n_k^{(0)}$ and $y_k^{(0)}$ to observed lifetimes using a prior predictive distribution. Dropping the component index k for ease of notation, this is given by

$$\begin{aligned} f(t \mid n^{(0)}, y^{(0)}) &= \int f(t \mid \lambda) f(\lambda \mid n^{(0)}, y^{(0)}) d\lambda \\ &= \beta t^{\beta-1} (n^{(0)} + 1) \frac{(n^{(0)} y^{(0)})^{n^{(0)}+1}}{(n^{(0)} y^{(0)} + t^\beta)^{n^{(0)}+2}}. \end{aligned} \quad (22)$$

Replacing the prior parameters $n^{(0)}$ and $y^{(0)}$ with their posterior counterparts $n^{(n)}$ and $y^{(n)}$ as defined in (6), the effect of virtual observations on (22), or the corresponding reliability function, can be determined. This allows to determine $n^{(0)}$ and $y^{(0)}$ through a number of ‘what-if’ scenarios, by asking the expert to state what (s)he would expect to learn from observing certain virtual data.

This strategy is known as pre-posterior analysis, being first advocated by Good (1965, p. 19). We recommend to check whether the effects of $\Pi_k^{(0)}$ on the inference of interest (this may not always be the full reliability function) reasonably reflect an expert’s beliefs before the data and in case some specific data become available, both data agreeing with initial beliefs and surprising data. Essentially, we advise to do an analysis like in our examples in Section 6 below, using hypothetical data.

To elicit a meaningful prior distribution, or a set of prior distributions, it is important to ask questions which enable experts to stay close to their actual expertise. Coolen (1996) discussed the possibility of generalizing the usual conjugate prior distributions, for parameters of exponential family models, by including pseudo-data which are right-censored. If the real data set contains such values, then such generalized priors do not lead to more computational complexities, while they can have several advantages. In addition to providing slightly more general classes of prior distributions through an additional hyperparameter, they may enable more realistic elicitation of expert judgements, for example if the expert has no experience with certain components past a specific life time. For more details we refer to Coolen (1996), it should be noted that adopting such generalized prior distributions may also provide more flexibility for modelling the effects of prior-data conflict, this is left as a topic for future research.

6 EXAMPLES

As illustrative example, consider a simplified automotive brake system with four types of component. The master brake cylinder (M) activates all four wheel brake cylinders (C1 – C4), which in turn actuate a braking pad assembly each (P1 – P4). The hand brake mechanism (H) goes directly to the brake pad assemblies P3 and P4; the car brakes when at least one brake pad assembly is actuated. The system layout is depicted in Figure 3, with




Figure 3. Reliability block diagram for a simplified automotive brake system, with those components marked that we assume to fail in the three scenarios.

M	H	C	P	Φ	M	H	C	P	Φ
1	0	1	1	0.25	1	0	2	1	0.50
1	0	1	2	0.50	1	0	2	2	0.83
1	0	1	3	0.75	1	0	3	1	0.75
0	1	0	1	0.50	1	1	0	1	0.50
0	1	0	2	0.83	1	1	0	2	0.83
0	1	1	1	0.62	1	1	1	1	0.62
0	1	1	2	0.92	1	1	1	2	0.92
0	1	2	1	0.75	1	1	2	1	0.75
0	1	2	2	0.97	1	1	2	2	0.97
0	1	3	1	0.88	1	1	3	1	0.88

Table 1. Survival signature values $\notin \{0, 1\}$ for the simplified automotive brake system depicted in Figure 3.

those components marked that we assume to fail in each of the three cases studied below. The values for $\Phi \notin \{0, 1\}$ for the complete system are given in Table 1.

A fixed prior setting, described in Section 6.1, will be combined with three different data scenarios, where one observes failure times in accordance with prior expectations in the first case (Section 6.2), surprisingly early failures in the second case (Section 6.3), and surprisingly late failures in the third case (Section 6.4). In each case, it is assumed that C2, C3, P2 and P3 fail, only the failure times are varied, investigating the effect of learning about these failures on the component level. Effects on posterior reliability bounds for the running system are then discussed for all three cases in Section 6.5.

6.1 Prior assumptions

The prior assumptions, which one can imagine to be determined by an expert, or by a combination of expert knowledge and component test data, are given by the prior parameter sets $\Pi_k^{(0)}$, $k = M, H, C, P$, as described in Table 2. There, $\underline{E}[T_i^k]$ and $\bar{E}[T_i^k]$ give the lower and upper bound for expected component lifetimes, respectively, which then have been transformed to bounds for the scale parameter using (4), resulting in $\underline{y}_k^{(0)}$ and $\bar{y}_k^{(0)}$. For example, according to the expert, the mean time to failure for component type M is between

k	β_k	$\underline{E}[T_i^k]$	$\overline{E}[T_i^k]$	$\underline{y}_k^{(0)}$	$\overline{y}_k^{(0)}$	$\underline{n}_k^{(0)}$	$\overline{n}_k^{(0)}$
M	2.5	5	8	75.4	244.1	2	5
H	1.2	2	20	2.5	39.2	1	10
C	2	8	10	81.5	127.3	1	5
P	1.5	3	4	6.1	9.3	1	10

Table 2. Prior parameter sets for the four component types.

5 and 8 time units, leading to $\underline{y}_M^{(0)} = 75.4$ and $\overline{y}_M^{(0)} = 244.1$, and the expert considers his knowledge on these expected lifetime bounds as having the strength of at least 2 and at most 5 observations.

These prior assumptions for the four component types are visualized in Figure 4, showing the sets of reliability functions corresponding to the prior predictive density (22). The figure thus displays the bounds for the probability that a single component, having been put under risk at time 0, will have failed by time t . The top left graph in Figure 8 shows what the prior assumptions on components signify for the system, depicting the prior bounds for the system reliability on a scale of time elapsed since system startup. For example, the prior probability of the system to survive until time 10 is between 0.03% and 6.91%.

6.2 Case 1: failure times as expected

In the first case, we observe $t_1^C = 6$, $t_2^C = 7$, $t_1^P = 3$, $t_2^P = 4$, and observe the running system until $t_{\text{now}} = 8$, i.e., $t_1^M = t_1^H = t_3^C = t_4^C = t_3^P = t_4^P = 8^+$. (Note that component failure times are numbered by order, not by component number in the system layout.) These observations correspond more or less to prior expectations, and the corresponding posterior predictive component distributions are given in Figure 5. In analogue to Figure 4, Figure 5 displays the bounds for the probability that a single component, having been put under risk at time 0, will have failed by time t , after having seen these (partly censored) observations. For easy comparisons, Figure 5 also contains the prior bounds from Figure 4.

We see that the graphs for M and H do not change dramatically, as there is only one component of each in the system to learn from. For C, the bounds have considerably narrowed, showing the effect of having seen the four observations $t_1^C = 6$, $t_2^C = 7$, $t_3^C = t_4^C = 8^+$. The bounds for P have not narrowed as much; this is due to the two right-censored observations $t_3^P = t_4^P = 8^+$; from the viewpoint of the prior, surviving past time 8 is already quite unusual.

6.3 Case 2: surprisingly early failure times

For the second case, with $t_1^C = 1$, $t_2^C = 2$, $t_1^P = 0.25$, $t_2^P = 0.5$ and $t_{\text{now}} = 2$ (so $t_1^M = t_1^H = t_3^C = t_4^C = t_3^P = t_4^P = 2^+$), we assume to observe surprisingly early failures; the corresponding posterior predictive component distributions are given in Figure 6. Having observed the system only until $t = 2$, the data are not very informative, such that prior and posterior predictive reliability bounds are very similar. For C, the effect of the early failures is however still visible, and posterior imprecision, i.e., the range between lower and upper posterior bound, is notably larger as compared to prior imprecision, and substantially larger than posterior imprecision in case 1. The effect for P is less pronounced, mainly because observing $\mathbf{t}^P = (0.25, 0.5, 2^+, 2^+)$ for P is less extreme as observing $\mathbf{t}^C = (1, 2, 2^+, 2^+)$ for C.




Figure 4. Sets of prior predictive reliability functions for the four component types, illustrating the choice of prior parameter sets $\Pi_k^{(0)}$, $k = M, H, C, P$.

6.4 Case 3: surprisingly late failure times

For the third case, we assume to observe surprisingly late failures, namely $t_1^C = 11$, $t_2^C = 12$, $t_1^P = 8$, $t_2^P = 9$, and $t_{\text{now}} = 12$ (so $t_1^M = t_1^H = t_3^C = t_4^C = t_3^P = t_4^P = 12^+$); the corresponding posterior predictive component distributions are given in Figure 7. Having observed the system for a much longer time than in case 2, the data contain now much more information, resulting in considerable differences between prior and posterior bounds. The effect of these surprisingly late failures is most prominent for P, with a set considerably shifted to the right and having very wide posterior bounds. The posterior set for C also indicates that, after having seen these late failures, one expects type C components to fail much later. This effect is also visible for M and H, but is weaker for them as there is only one component of each in the system.

6.5 Reliability bounds for the running system

Figure 8 depicts the set of prior system reliability functions, together with the sets of posterior system reliability function for the three cases, on a scale of elapsed time since system startup. Due to this time scale, posterior system reliability is 1 at t_{now} , as it is known

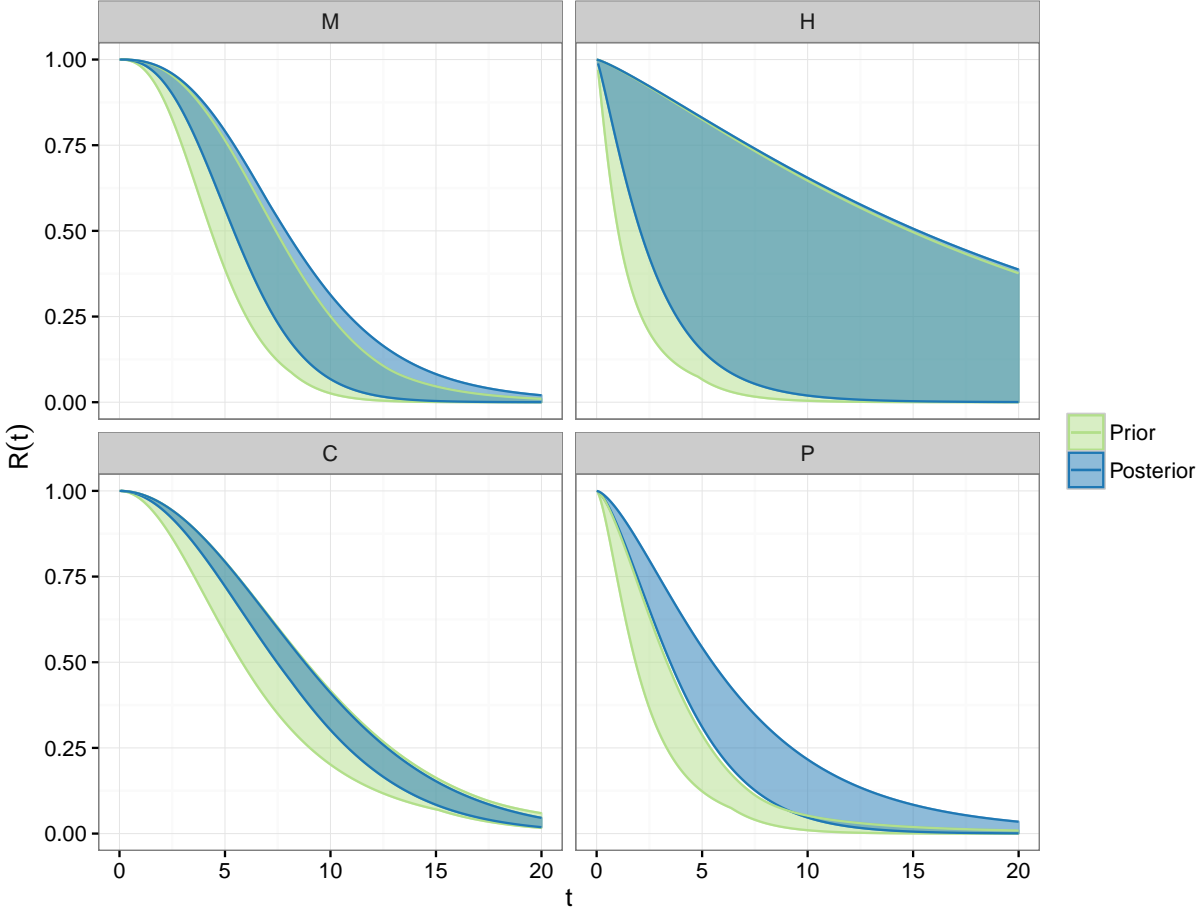


Figure 5. Sets of posterior predictive reliability functions for the four component types for observations in line with prior expectations (case 1).

that the system has survived until $t_{\text{now}} = 8, 2, 12$ in case 1,2,3, respectively. For all three cases, the posterior bounds drop faster after t_{now} than the prior bounds drop after $t = 0$ since the components in the system have aged until t_{now} and so are expected to fail sooner. In the case of surprisingly early failures, posterior bounds are mostly within prior bounds, this is due to t_{now} being close to 0 and weakly informative data in this scenario; the posterior bounds are nevertheless wider than those for case 1; posterior bounds are widest for case 3. This most visible in Figure 9, which shows $\bar{R}_{\text{sys}}(t) - \underline{R}_{\text{sys}}(t)$, the difference between upper and lower bound of prior and posterior system reliability. The left panel shows imprecision on the scale of elapsed time since system startup like in Figure 8; the right panel shows imprecision on the scale of prospective time instead, indicating how far in the future periods are for which estimation of system reliability is most uncertain. Posterior imprecision is indeed considerably lower in case 1, where failure times were more or less like expected. On the prospective timescale, one can see that periods of heightened uncertainty are closer to the present for the posteriors, while uncertainty is considerably reduced for periods further in the future.

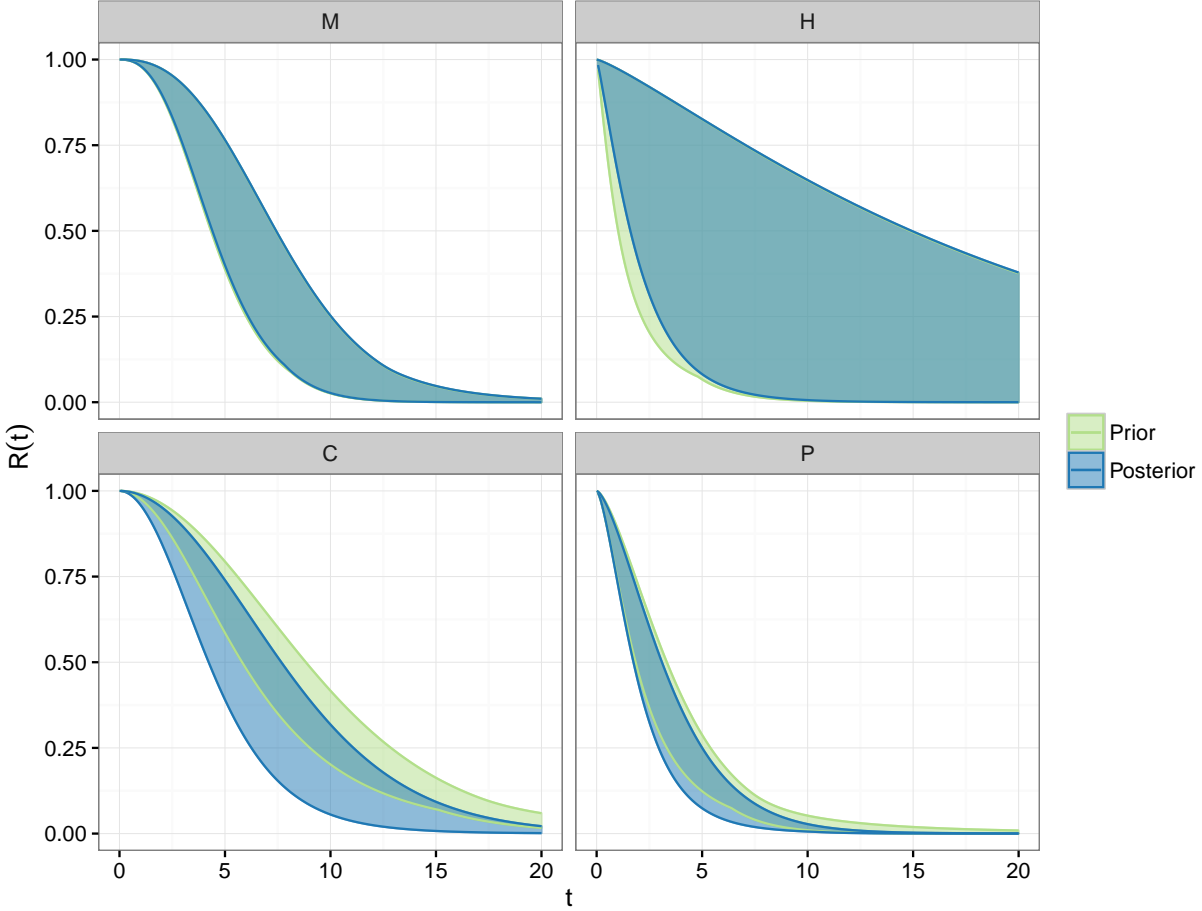


Figure 6. Sets of posterior predictive reliability functions for the four component types for surprisingly early failures (case 2).

7 CONCLUDING REMARKS

In this paper we presented a robust Bayesian approach to reliability estimation for systems of arbitrary layout, and showed how the use of sets of prior distributions results in increased imprecision, i.e., more cautious probability statements, in case of prior-data conflict (cases 2 and 3 in Section 6), while giving more precise reliability bounds when prior and data are in agreement (case 1 in Section 6). The parameters through which prior information is encoded have a clear interpretation and are thus easily elicited, either directly or with help of the prior predictive (22). Calculation of lower and upper predictive system reliability bounds is tractable, requiring only a simple K -dimensional box-constrained optimization in Equation (21).

We think that increased imprecision is an appropriate tool for mirroring prior-data conflict when considering sets of priors as is done in both the robust Bayesian and imprecise probability framework. We want to emphasize, however, that this tool may be useful already for just highlighting ‘conflict’ between multiple information sources, and that we do not think that the resulting set of posteriors, although it can form a meaningful basis, must necessarily be used for all consequential inferences, as a strict Bayesian would posit. We be-

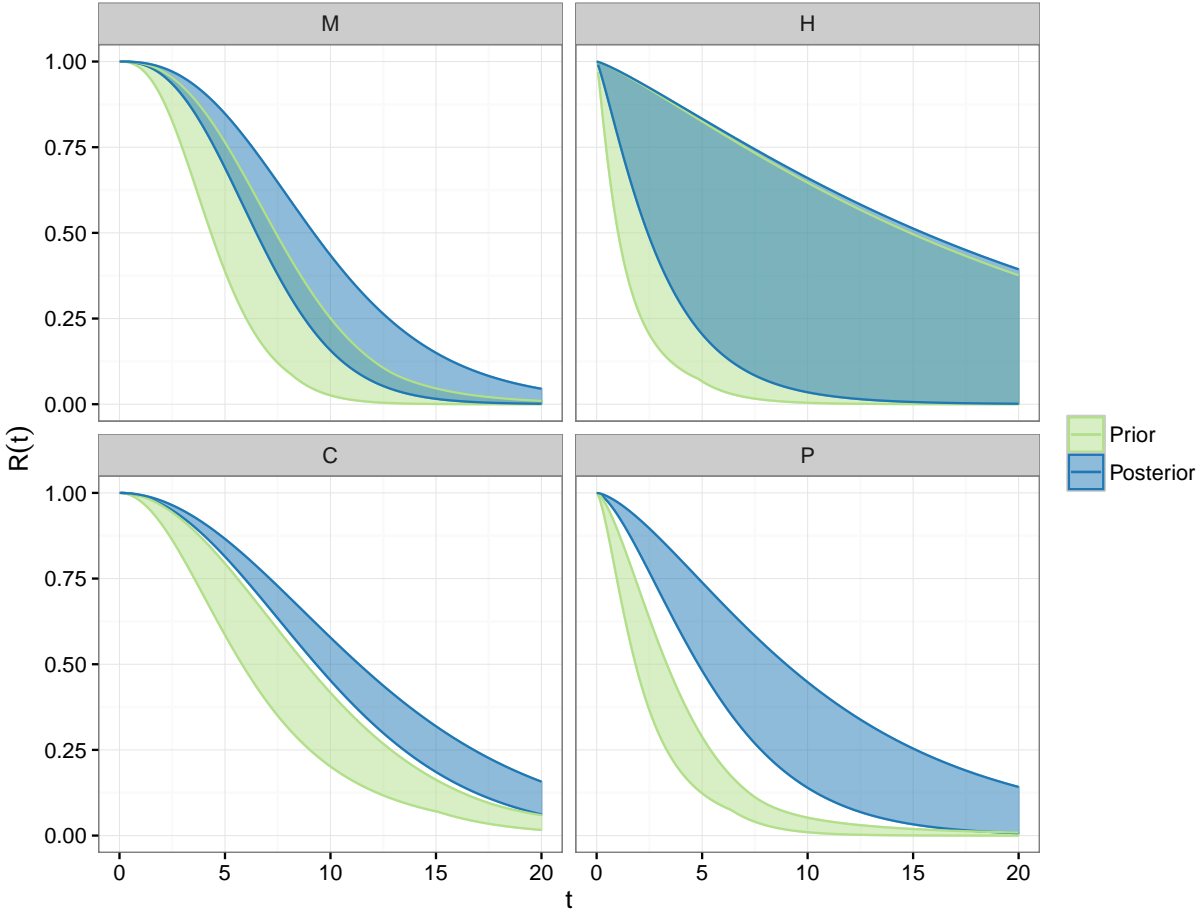


Figure 7. Sets of posterior predictive reliability functions for the four component types for surprisingly late failures (case 3).

lieve an analyst is free to reconsider any aspect of a model (of which the choice of prior can be seen to form a part) after seeing the data, and so may use our method only for becoming aware of a conflict between prior and data.

The employed robust Bayesian setting provides many further modelling opportunities beyond the explicit reaction to prior-data conflict. These opportunities have not yet been explored and provide a wide field for further research. An example is our currently ongoing investigation into extending the present model to allow also for an appropriate reflection of very strong agreement between prior and data (Walter and Coolen 2016).

There are many aspects to further develop in our analysis and modeling. The general approaches used in this paper, namely the use of sets of conjugate priors for component lifetime models and the survival signature to calculate the system reliability, can be used with other parametric component lifetime distributions that form a canonical exponential family, since for such distributions, a canonical conjugate prior using the same canonical parameters $n^{(0)}$ and $y^{(0)}$ can be constructed, for details see, e.g., Bernardo and Smith (2000, p. 202 and 272f), or Walter (2013, p. 8). Likewise straightforward to implement is, e.g., the analysis of the effect of replacing failed components in the system on system reliability

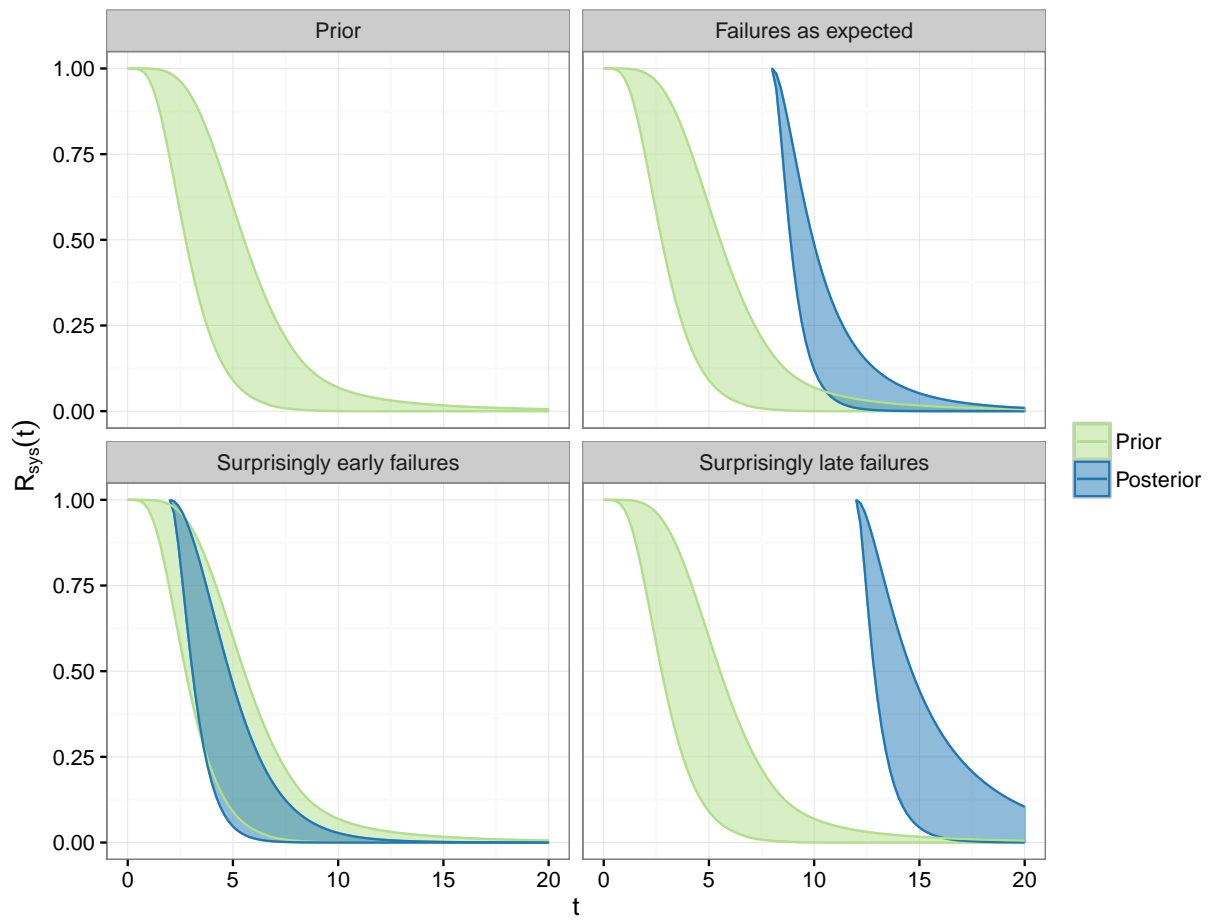


Figure 8. Sets of prior and posterior system reliability functions for the three cases on a time showing time elapsed since system startup.

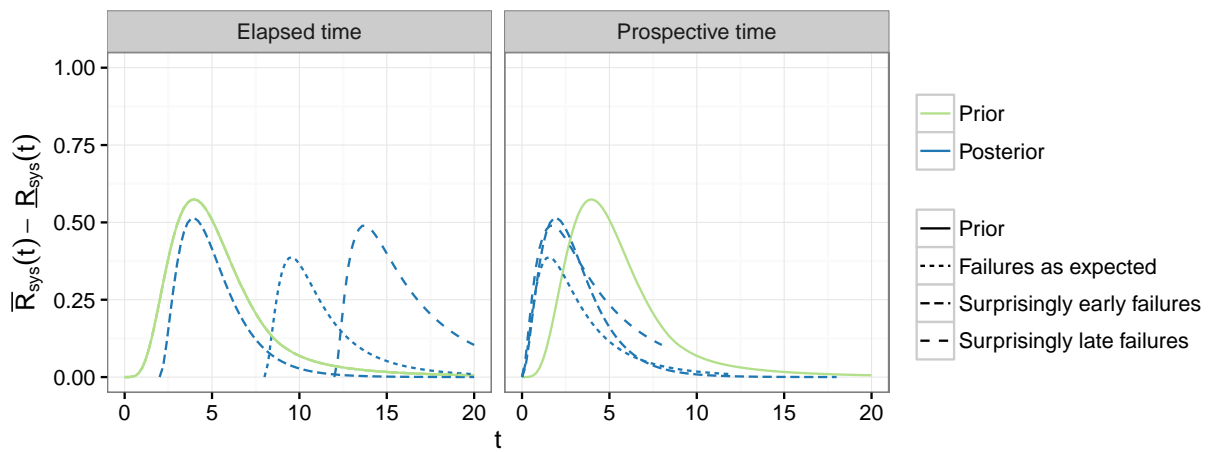


Figure 9. Imprecision of prior and posterior system reliability sets.

bounds. Criteria for the trade-off between the cost of replacement and the gain in reliability would have to be adapted for the interval output of our model, leading to very interesting research questions.

To estimate the shape parameters β_k together with the scale parameters λ_k , one could follow standard Bayesian approaches and use a finite discrete distribution for β_k . Developing this together with suitable sets of priors for λ_k , in particular to show the effect of prior-data conflict, is another interesting challenge for future research.

Another further important aspect in system reliability we have not accounted for yet is the possibility of common-cause failures, i.e., failure events where several components fail at the same time due to a shared or common root cause. This could be done by combining the common-cause failure model approaches of Troffaes et al. (2014) and Coolen and Coolen-Maturi (2015).

On a more abstract level, the choice for the set of prior parameters (generating the set of prior component failure distributions) as $[\underline{n}_k^{(0)}, \bar{n}_k^{(0)}] \times [\underline{y}_k^{(0)}, \bar{y}_k^{(0)}]$ has the advantage of allowing for easy elicitation and tractable inferences, but it may not be suitable to reflect certain kinds of prior knowledge. Also, as studied in (Walter et al. 2011) and (Walter 2013, §3.1), the shape of the prior parameter set has a crucial influence on model behaviour like the severity of prior-data conflict reaction. As noted in (Walter 2013, pp. 66f), more general prior parameter set shapes are possible in principle, but may be more difficult to elicit and make calculations more complex.

ACKNOWLEDGEMENTS

Gero Walter was supported by the DIALOG project CAMPI (“Coordinated Advanced Maintenance and Logistics Planning for the Process Industries”).

Appendix I. References

- Ahmadzadeh, F. and Lundberg, J. (2014). “Remaining useful life estimation: review.” *International Journal of System Assurance Engineering and Management*, 5, 461–474.
- Augustin, T., Coolen, F. P. A., de Cooman, G., and Troffaes, M. C. M. (2014). *Introduction to Imprecise Probabilities*. Wiley, New York.
- Barlow, R. and Proschan, F. (1975). *Statistical Theory of Reliability and Life Testing*. Holt, Rinehart and Winston, Inc., New York.
- Berger, J., Moreno, E., Pericchi, L., Bayarri, M., Bernardo, J. M., Cano, J., De la Horra, J., Martín, J., Ríos Insúa, D., Betrò, B., Dasgupta, A., Gustafson, P., Wasserman, L., Kadane, J., Srinivasan, C., Lavine, M., O’Hagan, A., Polasek, W., Robert, C., Goutis, C., Ruggeri, F., Salinetti, G., and Sivaganesan, S. (1994). “An overview of robust Bayesian analysis.” *TEST*, 3, 5–124.
- Bernardo, J. and Smith, A. (2000). *Bayesian Theory*. Wiley, Chichester.
- Coolen, F. P. A. (1996). “On Bayesian reliability analysis with informative priors and censoring.” *Reliability Engineering & System Safety*, 53, 91–98.
- Coolen, F. P. A. and Coolen-Maturi, T. (2012). “Generalizing the signature to systems with multiple types of components.” *Complex Systems and Dependability*, W. Zamojski, J. Mazurkiewicz, J. Sugier, T. Walkowiak, and J. Kacprzyk, eds., Vol. 170 of *Advances in Intelligent and Soft Computing*, Springer, 115–130.
- Coolen, F. P. A. and Coolen-Maturi, T. (2015). “Predictive inference for system reliability after common-cause component failures.” *Reliability Engineering & System Safety*, 135, 27–33.
- Evans, M. and Moshonov, H. (2006). “Checking for prior-data conflict.” *Bayesian Analysis*, 1, 893–914.
- Good, I. J. (1965). *The estimation of probabilities*. MIT Press, Cambridge (MA).
- R Core Team (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, <<https://www.R-project.org/>>.
- Ríos Insua, D. and Ruggeri, F. (2000). *Robust Bayesian Analysis*. Springer.
- Troffaes, M. C. M., Walter, G., and Kelly, D. L. (2014). “A robust Bayesian approach to modelling epistemic uncertainty in common-cause failure models.” *Reliability Engineering & System Safety*, 125, 13–21.
- Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London.
- Walter, G. (2013). “Generalized Bayesian inference under prior-data conflict.” Ph.D. thesis, Department of Statistics, LMU Munich, Department of Statistics, LMU Munich, <<http://edoc.ub.uni-muenchen.de/17059/>>.
- Walter, G. and Augustin, T. (2009). “Imprecision and prior-data conflict in generalized Bayesian inference.” *Journal of Statistical Theory and Practice*, 3, 255–271.
- Walter, G., Augustin, T., and Coolen, F. P. (2011). “On prior-data conflict in predictive Bernoulli inferences.” *ISIPTA ’11: Proceedings of the Seventh International Symposium on Imprecise Probabilities: Theories and Applications*, F. P. Coolen, G. deCooman, T. Fetz, and M. Oberguggenberger, eds., SIPTA, 391–400, <<http://www.sipta.org/isipta11/proceedings/046.html>>.
- Walter, G. and Coolen, F. P. A. (2016). “Sets of priors reflecting prior-data conflict and

- agreement.” *Information Processing and Management of Uncertainty in Knowledge-Based Systems: 16th International Conference, IPMU 2016, Eindhoven, The Netherlands, June 20-24, 2016, Proceedings, Part I*, P. J. Carvalho, M.-J. Lesot, U. Kaymak, S. Vieira, B. Bouchon-Meunier, and R. R. Yager, eds., Cham, Springer International Publishing, 153–164, <http://dx.doi.org/10.1007/978-3-319-40596-4_14>.
- Walter, G., Graham, A., and Coolen, F. P. A. (2015). “Robust Bayesian estimation of system reliability for scarce and surprising data.” *Safety and Reliability of Complex Engineered Systems: ESREL 2015*, L. Podofillini, B. Sudret, B. Stojadinović, E. Zio, and W. Kröger, eds., CRC Press, 1991–1998.