

Two-stage Bayesian models – application to ZEDB project

C. Bunea, T. Charitos & R.M. Cooke
 Delft University of Technology, Delft, Netherlands

G. Becker
 RISA, Berlin, Germany

ABSTRACT: A well-known mathematical tool to analyze data for nuclear power facilities is the two-stage Bayesian models. In this paper we review this mathematical model, its underlying assumptions and supporting arguments. Furthermore, we will verify the software implementations of the ZEDB database and compare the results. Lastly, an assessment of the relevance of new developments will take place, while the viability of the two-stage Bayesian approach will be discussed.

1 INTRODUCTION

ZEDB [Becker and Schubert, 1998] is the major German effort to collect data from nuclear facilities. The goal of the project is to create a reliability data base which contains all major plant events: failure events, operational experience, maintenance actions. As a mathematical tool to analyse ZEDB data, a two-stage Bayesian model was chosen. Firstly we identify the standard conditional independence assumptions and derive the general form of the posterior distribution for failure rate λ_0 at plant of interest 0, given failures and observation times at plants 0, 1, ... n . Any departure for the derived mathematical form necessarily entails a departure from the conditional independence assumptions. Vaurio's one stage empirical Bayes model is discussed as an alternative to the two-stage model [Vaurio, 1987]. Hofer [Hofer et al., 1997, Hofer and Peschke, 1999, Hofer, 1999] has criticized the standard two-stage model and proposed an alternative, which is also discussed. Finally, the methods of Pörn and Jeffrey for choosing a non-informative prior distribution are discussed.

2 BAYESIAN TWO-STAGE HIERARCHICAL MODELS

Bayesian two-stage or hierarchical models are widely employed in a number of areas. The common theme of these applications is the assimilation of data from different sources, as illustrated in Figure 1. The data from agent i is characterized by an exposure T_i and a number of events X_i . The exposure T_i is not considered

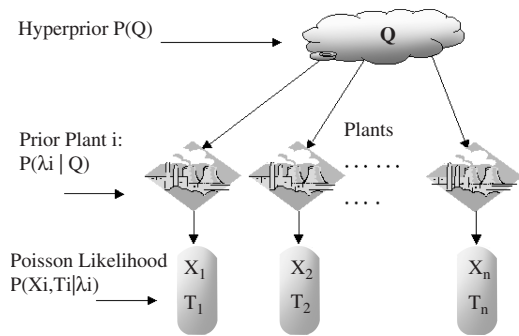


Figure 1. Bayesian two-stage hierarchical model.

stochastic, as it can usually be observed with certainty. The number of events for a given exposure follows a fixed distribution type, in this case Poisson. The parameter(s) of this fixed distribution type are uncertain, and are drawn from a prior distribution. The prior distribution is also of a fixed type, yet with uncertain parameters. In other words, the prior distribution itself is uncertain. This uncertainty is characterized by a hyperprior distribution over the parameter(s) of the prior.

In Figure 1, the hyperprior is a distribution $P(Q)$ over the parameters Q of the prior distribution from which the Poisson intensities $\lambda_1, \dots, \lambda_n$ are drawn. In sum, our model is characterized by a joint distribution:

$$P(X_1, \dots, X_n, \lambda_1, \dots, \lambda_n, Q) \quad (1)$$

To yield tractable models, such models must make two types of assumptions. First, conditional independence

assumptions [Pörn, 1990, Iman and Hora, 1990] are made to factor (1). Second, assumptions must be made regarding the fixed distribution types and the hyperprior distribution $P(Q)$. The conditional independence assumptions may be read from Figure 1, by treating this figure as a “belief net”. In particular, this figure says:

CI.1 Given Q , λ_i is independent of $\{X_j, \lambda_j\}_{j \neq i}$

CI.2 Given λ_i , X_i is independent of $\{Q, \lambda_j, X_j\}_{j \neq i}$

The expression “ X_i is independent of $\{Q, \lambda_j, X_j\}_{j \neq i}$ ” entails that X_i is independent of Q , and X_i is independent of λ_j .

With these assumptions we can derive the conditional probability $P(\lambda_0|X_0, \dots, X_n)$ for the failure rate at plant 0, given X_i failures observed at plant i , $i = 0, \dots, n$. This is sometimes called the posterior probability for λ_0 .

2.1 Derivation of posterior probability for λ_0

We assume throughout that the plant of interest is plant 0. We seek an expression for

$$P(\lambda_0|X_0, \dots, X_n) \quad (2)$$

We step through this derivation, giving the justification for each step. A more detailed exposition is found in [Cooke et al., 1995]. “ \propto ” denotes proportionality, “CI.1” means that conditional independence assumption i is invoked, $i = 1, 2$; “BT” denotes Bayes theorem; and “TP” denotes the law of total probability; and “FB” denotes Fubini theorem. The Fubini theorem (see [Cooke et al., 2002] for a more detailed discussion) authorizes switching the order of integration if the integrals are finite [Royden, 1968, p 269],

$$P(\lambda_0|X_0 \dots X_n) \propto_{BT} P(X_0|\lambda_0, X_1 \dots X_n)P(\lambda_0|X_1 \dots X_n)$$

$$\propto_{CI.2} P(X_0|\lambda_0)P(\lambda_0|X_1 \dots X_n)$$

$$\propto_{TP,BT} P(X_0|\lambda_0) \int_{\lambda_1 \dots \lambda_n} \int_q P(\lambda_0|\lambda_1 \dots \lambda_n, q, X_1 \dots X_n) P(q, \lambda_1 \dots \lambda_n | X_1 \dots X_n) dq d\lambda_1 \dots d\lambda_n$$

$$\propto_{CI.1} P(X_0|\lambda_0) \int_{\lambda_1 \dots \lambda_n} \int_q P(\lambda_0|q)$$

$$P(q, \lambda_1 \dots \lambda_n | X_1 \dots X_n) dq d\lambda_1 \dots d\lambda_n$$

$$\propto_{BT} P(X_0|\lambda_0) \int_{\lambda_1 \dots \lambda_n} \int_q P(\lambda_0|q)P(X_1 \dots X_n|q, \lambda_1 \dots \lambda_n) P(q, \lambda_1 \dots \lambda_n) dq d\lambda_1 \dots d\lambda_n$$

$$\propto_{CI.1,2;FT} P(X_0|\lambda_0) \int_q P(\lambda_0|q)$$

$$\int_{\lambda_1 \dots \lambda_n} [\prod_{i=1 \dots n} P(X_i|\lambda_i)P(\lambda_i|q)d\lambda_1 \dots d\lambda_n] P(q) dq$$

$$\propto P(X_0|\lambda_0) \int_q P(\lambda_0|q)$$

$$\prod_{i=1 \dots n} \int [P(X_i|\lambda_i)P(\lambda_i|q)d\lambda_i] P(q) dq$$

Expression (4) is normalized by integrating over all λ_0 .

2.2 Summary of significant features

1. If $Q = q_0$ is known with certainty, then there is no influence from X_1, \dots, X_n on λ_0 . Indeed, in this case the posterior density in λ_0 is simply proportional to $P(X_0|\lambda_0) P(\lambda_0|q)$.
2. As the numbers $X_i, T_i, i = 1 \dots n$ get large, $X_i/T_i \rightarrow \lambda_i$, then the Poisson likelihood $P(X_i|\lambda_i)$ converges to a Dirac measure concentrating mass at the point $X_i = T_i \lambda_i$. In the limit the “hyperposterior”

$$\prod_{i=1 \dots n} [\int P(X_i|\lambda_i)P(\lambda_i|q)d\lambda_i] P(q) \quad (5)$$

becomes

$$[\prod_{i=1 \dots n} P(\lambda_i = X_i/T_i|q)] P(q) \quad (6)$$

(6) corresponds to the situation where $P(q)$ is updated with observations $\lambda_1, \dots, \lambda_n$. Note that as the observation time increases, the number n does not change. If n is only modest (say in the order 10) then the effect of the hyperprior will never be dominated by the effect of observations (see [Cooke et al., 2002] for a more detailed discussion). We say that the hyperprior persists in the posterior distribution $P(\lambda_0|X_0, \dots, X_n)$.

3. It is shown in [Cooke et al., 1995], [Hennings and Meyer, 1999] that improper hyperpriors $P(q)$ do not always become proper when multiplied by $[\prod_{i=1 \dots n} P(\lambda_i = X_i/T_i|q)]P(q)$. In other words, the hyperposterior may well remain improper.

2.3 Selected literature review

Two-stage Bayesian models have been implemented by various authors. [Kaplan, 1983] used a log normal prior with a Poisson likelihood, which of course is not a natural conjugate. This method has been implemented by ZEDB. [Iman and Hora, 1989, 1990] and [Hora and Iman, 1987] proposed a natural conjugate

gamma prior. [Vaurio, 1987] proposed a one-stage empirical Bayes approach, using other plants to determine the prior. The SKI data bank [1987] uses a two-stage model developed by Pörn [1990]. This model was reviewed in [Cooke et al., 1995], and further discussed in [Meyer and Hennings, 1999]. Recently [Hofer et al., 1997], [Hofer and Peschke, 1999] and [Hofer, 1999] have suggested that an incorrect chance mechanism underlies the two-stage models, and have proposed their own model. In this section we briefly review these developments.

In the two-stage Bayesian models considered here a Poisson likelihood is used. The prior is usually gamma or log normal. The second stage places a hyperprior distribution over the parameters of the prior gamma or log normal distribution. We briefly recall the definitions and elementary facts of the Poisson, Gamma, and log normal distributions in Table 1 below.

Using a gamma prior with parameters as above, the term $\prod_{i=1}^n \int P(X_i|\lambda_i)P(\lambda_i|q)d\lambda_i$ in (4) becomes after carrying out the integration:

$$\prod_{i=1}^n \frac{\Gamma(X_i + \alpha)}{\Gamma(X_i + 1)\Gamma(\alpha)} \left(\frac{\beta}{\beta + T_i}\right)^\alpha \left(\frac{T_i}{\beta + T_i}\right)^{X_i} \\ = \prod_{i=1}^n \frac{(X_i + \alpha - 1)!}{X_i!(\alpha - 1)!} \left(\frac{\beta}{\beta + T_i}\right)^\alpha \left(\frac{T_i}{\beta + T_i}\right)^{X_i} \quad (7)$$

Further calculation to solve equation (4) must be performed numerically. It is shown in [Cooke et al., 1995] that improper hyperpriors may remain improper after assimilating observations. The asymptotic behavior of the “hyperposterior”

$$P(\alpha, \beta | X_1 \dots X_n, T_1 \dots T_n) \propto P(X_1 \dots X_n, T_1 \dots T_n | \alpha, \beta) P(\alpha, \beta)$$

will essentially be determined by the maximum of $P(X_1 \dots X_n, T_1 \dots T_n | \alpha, \beta)$. The significant fact is that $P(X_1 \dots X_n, T_1 \dots T_n | \alpha, \beta)$ has no maximum; it is asymptotically maximal along a ridge, see (Figure 2).

Table 1.

	Poisson
Density	$P(X T, \lambda) = \frac{(\lambda T)^X}{X!} e^{-\lambda T}, \lambda > 0, T > 0$
Mean	λT
Variance	λT
	Gamma
Density	$f(\lambda \alpha, \beta) = \frac{\lambda^\alpha}{\Gamma(\alpha)} \beta^\alpha e^{-\beta \lambda}, \alpha > 0$
Mean	α/β
Variance	α/β^2
	Lognormal
Density	$f(\lambda \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma\lambda} e^{-\frac{(\ln \lambda - \mu)^2}{2\sigma^2}}, \sigma > 0$
Mean	$e^{\mu + \sigma^2/2}$
Variance	$e^{2\mu + \sigma^2} (e^{\sigma^2} - 1)$

2.4 Vaurio

[Vaurio, 1987] proposed an analytic empirical Bayes approach to the problem of assimilating data from other plants. A simple one-stage Bayesian model for one plant would use a Poisson likelihood with intensity λ , and a $\Gamma(\lambda|a, b)$ prior. Updating the prior with X_i failures in time T_i yields a $\Gamma(\lambda|\alpha + X_i, \beta + T_i)$ posterior. Vaurio proposes to use data from the population of plants to choose the $\Gamma(\lambda|\alpha, \beta)$ prior by moment fitting. Any other two moment prior could be used as well. Data from other plants are not used in updating, hence, this is a one-stage model.

The model is consistent, in the sense that as $X_i, T_i \rightarrow \infty$, with $X_i/T_i \rightarrow \lambda_i$, his model does entail that $E(\lambda_i | X_i, T_i) \rightarrow \lambda_i$. Elegance and simplicity are its main advantages. Disadvantages are that it cannot be applied if all $X_i = 0$, or if the population consists of only 2 plants. Further, numerical results indicate that the model is non-conservative when the empirical failure rate at plant 0 is low and the empirical failure rates at other plants are high. A final criticism, which applies to most empirical Bayes models is that the data for the plant of interest is used twice, once to estimate the prior and once again in the Poisson likelihood.

2.5 Hofer

Hofer has published a number of articles [Hofer et al., 1997], [Hofer, 1999] and [Hofer and Peschke, 1999] in which the two-stage models are faulted for using a “wrong chance mechanism”, and a new model is proposed. He does not explicitly formulate conditional independence assumptions, and does not derive the posterior by conditionalizing the joint as done above. Rather, the model is developed by shifting between the point of view of “observing λ_i ” and “observing (X_i, T_i) ”. [Hofer, 1997] criticizes [Hora and Iman, 1990] for using the wrong order of integration of improper

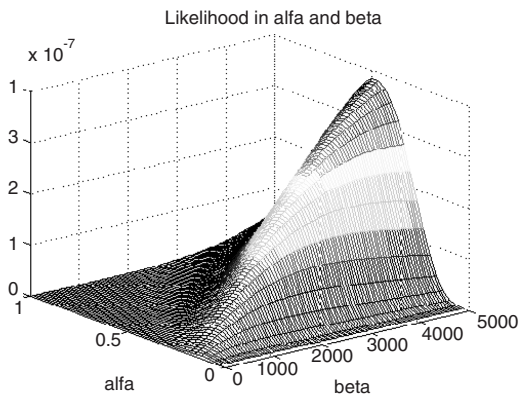


Figure 2.

integrals (see [Cooke et al., 2002] for a more detailed discussion). In later publications, a “deeper” reason is found to reside in the use of a “wrong chance mechanism”. Hofer’s model appears to result in a posterior of the following form:

$$P(\lambda_0|X_0..X_n) \propto P(X_0|\lambda_0) \int_{(q)} P(\lambda_0|q) \int \int_{\sqrt{\lambda_1..n}} \frac{[\prod_{i=1..n} P(\lambda_i|q)]P(q)}{\int_{(q)} [\prod_{i=1..n} P(\lambda_i|q')]P(q')dq'} [\prod_{i=1..n} P(\lambda_i|q)]P(\lambda_1... \lambda_n)d\lambda_1...d\lambda_n dq \quad (8)$$

Notice that this does not appear to have the form of (4).

Although Hofer does not explicitly formulate his conditional independence assumptions, he does use them. E.g. he uses CI.1 to derive the expression in the denominator (see equation (4) of [Becker and Hofer, 2001]). If CI.1 holds, then necessarily

$$\int_{(q)} [\prod_{i=1..n} P(\lambda_i|q')]P(q')dq' = \int_{(q)} P(\lambda_1... \lambda_n|q)P(q)dq = P(\lambda_1... \lambda_n) \quad (9)$$

and (8) reduces to (4). If CI.1 does not hold, then the origin of the product $\prod P(\lambda_i|q)$ is unclear. Hofer says that $P(\lambda_1... \lambda_n) = \prod r(\lambda_i)$, where $r(\lambda_i)$ is a noninformative prior, which he takes to be constant. This entails that the λ_i are unconditionally independent. It is not difficult to show that if λ_i are unconditionally independent, and independent given q , that then λ_i is independent of q . Indeed, independence implies that for all λ_i , $P(\lambda_1, \dots, \lambda_n) = \prod P(\lambda_i) = \prod \int P(\lambda_i|q) P(q)dq$. By conditional independence; $P(\lambda_1, \dots, \lambda_n) = \int P(\lambda_1, \dots, \lambda_n|q)P(q)dq = \int \prod P(\lambda_i|q)P(q)dq$. The λ_i are identically distributed given q ; take $\lambda_i = \lambda$; $i = 1, \dots, n$. These two statements imply $[\int P(\lambda|q)^n P(q)dq]^{1/n} = \int P(\lambda|q)P(q)dq$. Since the integrand is non-negative, this implies that $P(\lambda|q) = constant = P(\lambda)$ [Hardy, Littlewood and Polya, 1983, p 143]. This would make the entire two-stage model quite senseless. If $P(\lambda_1... \lambda_n) = \prod r(\lambda_i) = constant$ in the numerator of (8), but not in the denominator, then (8) is not equivalent to (4), but rests on conflicting assumptions.

In any event, if (8) does not reduce to (4) then the assumptions CI.1, CI.2 do not both hold. Hofer does not say which assumptions are used to derive (8), in fact (8) is not derived mathematically, but is “woven together” from shifting points of view. The danger of such an approach is that conflicting assumptions may

be inadvertently introduced. This appears to be the case, as the λ_i are at one point assumed to be independent, and at another point are assumed to be conditionally independent given q .

2.6 Hyperpriors

Pörn (1990) introduces a two-stage model with a gamma prior for λ , similar to (Hora and Iman, 1990). He provides an argument for choosing the following non-informative (improper) densities for the parameters v, μ' : $g(v) = 1/v, v > 0$, and $k(\mu') = 1/\sqrt{\mu'(1 + \mu')}$, $\mu' > 0$, where $v = 1/\alpha$ is the coefficient of variation and $\mu' = T\alpha/\beta$ is the expected number of failures at time T , given α and β . Assuming independence between these parameters and transforming back to the hyperparameters α and β , $\alpha = 1/v^2 \beta = T/v^2 \mu$, a joint (improper) hyperprior density for α and β is obtained proportional to: $1/\beta \sqrt{\alpha(\alpha + \beta/T)}$.

Another frequently used principle, called Jeffrey’s rule, is to choose the non-informative prior $P(q)$ for a set of parameters q proportional to the square root of the determinant of the information matrix $\Phi_n(q)$:

$$P(q) \propto \sqrt{|\Phi_n(q)|}, \Phi_n(q) = E\left\{ \frac{\partial^2 L}{\partial q_i \partial q_j} \right\} \quad (10)$$

and L is the log-likelihood function for the set of the parameters.

Hora and Iman (1990) apply this rule to the two-dimensional parameter vector (α, β) for the Gamma distribution of the failure rate λ . They get the (approximate) improper hyperprior: $1/\alpha^{1/2} \beta, \alpha, \beta > 0$.

3 ZEBD SOFTWARE VERIFICATION

Three data sets are used to check the concordance with the results from [Becker and Hofer, 2001]. Differences between our results and those of ZEDB reflect differences that may arise from an independent implementation based on public information. Although ZEDB recommends the lognormal model, both the lognormal and gamma models are supported, and both are benchmarked here. The data sets are shown in Tables 2–4.

3.1 Gamma model

The computation may be broken into three steps: Firstly, truncate the range of (α, β) to a finite rectangle. Then identify a range for λ_0 which contains all the “plausible” values. Finally, for every “plausible” value of λ_0 , evaluate numerically the integrals over α and β , and interpolate to find the 5%, 50% and 95% quantiles.

Table 2. Data set 1 (4 in [Becker and Hofer, 2001]).

Nr. failures	Obs. time	Nr. failures	Obs. time
7	24000	0	24000
1	24000	0	24000
3	24000	0	24000
<u>2</u>	<u>24000</u>	2	24000
1	24000	0	24000
2	24000	0	24000

Table 3. Data set 2 (2 in [Becker and Hofer, 2001]).

Nr. failures	Obs. time	Nr. failures	Obs. time
1	20000	0	6000
0	2000	1	10000
0	4000	<u>2</u>	<u>12000</u>

Table 4. Data set 3 (3 in [Becker and Hofer, 2001]).

Nr. failures	0	0	<u>1</u>
Obs. Time	12000	2000	<u>3000</u>

Remark: The underlined field is the plant of interest.

The likelihood $P(X_1, \dots, X_n, T_1, \dots, T_n | \alpha, \beta)$ as a function of α and β (2.7) is presented in Figure 2. Values for $(X_1, \dots, X_n, T_1, \dots, T_n)$ are taken from data set 1. For uniform hyperpriors, this likelihood is proportional to the hyperposterior distribution $P(\alpha, \beta | X, T)$. Note that $P(\alpha, \beta | X, T)$ does not peak but “ridges”. This means that a “natural” truncation for α and β cannot be defined; that is, we cannot define a finite rectangle for α and β which contains most of the hyperposterior mass. In our simulations, these ranges were chosen in a manner similar to [Cooke et al., 1995], using Pörn’s heuristic. The inability to localize the hyperposterior mass for (α, β) means that we cannot localize the posterior mass

$$P(\lambda_0 | X_0, \dots, X_n) \propto P(X_0 | \lambda_0) \int \int_{\alpha, \beta} P(\lambda_0 | \alpha, \beta)$$

$$\prod_{i=1..n} \left[\int P(X_i | \lambda_i) P(\lambda_i | \alpha, \beta) d\lambda_i \right] P(\alpha, \beta) d\alpha d\beta$$

For each finite rectangle for α, β , the mass in λ_0 will be localized, but other choices for α, β could significantly shift the region in which λ_0 is localized. This means, of course, that the method of truncation in step 1 will influence the plausible values in step 2, and can have a significant effect on the results.

Figures 3 and 4 represent the hyperposterior distribution for Pörn’s approach and Jeffrey’s hyperpriors; also with $(X_1, \dots, X_n, T_1, \dots, T_n)$ from data set 1 below.

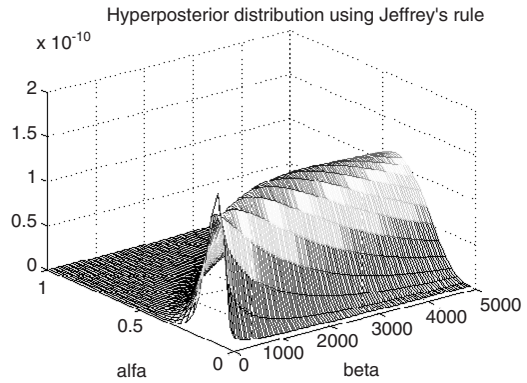


Figure 3.

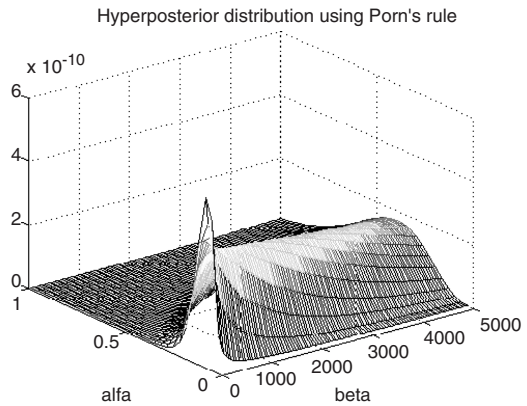


Figure 4.

Table 5. The 5%, 50% and 95% quantiles of the posterior distribution of λ_0 for data set 1.

	Uniform	Pörn	Jeffrey
5%	2.3971 E-5	2.8429 E-5	2.8665 E-5
50%	8.0511 E-5	8.4990 E-5	8.5678 E-5
95%	2.0012 E-4	2.0598 E-4	2.0670 E-4
Ranges		TUD	ZEDB
α :		0.033...1	0.002...351
β :		50...5000	582...232,219

Tables 5–7 compare our results for the uniform, Pörn and Jeffrey prior, and give the integration ranges for α and β for our computation and for the ZEDB results. Table 8 compares the TUD and ZEDB results. Note that the 5% quantile for dataset 3 is a more than a factor three lower in the TUD results. In dataset 1 the agreement is better, as there are more plants, more operational hours and more failures. These differences

Table 6. The 5%, 50% and 95% quantiles of the posterior distribution of λ_0 for data set 2.

	Uniform	Pörn	Jeffrey
5%	4.2371 E-5	3.9306 E-5	4.3845 E-5
50%	1.4603 E-4	1.4278 E-4	1.4908 E-4
95%	2.0012 E-4	2.0598 E-4	2.0670 E-4
Ranges		TUD	ZEDB
α :		0.03...1.2	0.01...1926
β :		50...5000	789...350,098

Table 7. The 5%, 50% and 95% quantiles of the posterior distribution of λ_0 for data set 3.

	Uniform	Pörn	Jeffrey
5%	3.3976 E-5	3.6503 E-5	3.6394 E-5
50%	1.7514 E-4	2.1247 E-4	2.0711 E-4
95%	5.8914 E-4	6.6524 E-4	6.5640 E-4
Ranges		TUD	ZEDB
α :		0.0154...0.3846	0.01...503
β :		50...5000	126...91385

Table 8. Comparison TUD (Pörn) and ZEDB for gamma model.

	Dataset 1	Dataset 2	Dataset 3
	Gamma	Model TUD	(Pörn)
5%	2.8429 E-5	3.9306 E-5	3.6503 E-5
50%	8.4990 E-5	1.4278 E-4	2.1247 E-4
95%	2.0598 E-4	3.3217 E-4	6.6524 E-4
	Gamma	Model	ZEDB
5%	3.2518 E-05	1.2220 E-04	1.2141 E-04
50%	6.9926 E-05	1.7247 E-04	2.5766 E-04
95%	1.3044 E-04	3.4473 E-04	7.4076 E-04

are consistent with the results reported in [Cooke et al., 1995], where “stress-testing” the gamma model by exploring the range of plausible choices for α , β resulted in differences up to a factor 5.

3.2 Lognormal model

ZEDB adopted the lognormal distribution as a prior, based on the maximum entropy principle invoked by [Jaynes, 1968]. The uncertainty over parameters μ and σ is expressed by hyperpriors. [Becker, 2001] takes into account four types of hyperprior distribution based on Jeffrey’s rule. [Becker, 2001] proposes four different implementations of Jeffrey’s rule. We caution against the multivariate implementation and version of the Jeffrey’s rule when parameters of different kind e.g. location and scale parameters, are considered. In this case, as [Box and Tiao, 1974] suggested, it is wiser to choose parameters, which can be assumed independent

and then apply the one parameter version of the rule. This is done only in the first and the fourth case below (see [Cooke et al., 2002] for a more detailed discussion); case 1 is used by ZEDB.

- 1st case: Jeffrey’s rule is applied to the parameters μ and σ : In this case the hyperprior has the well-known form $f(\mu, \sigma) \propto 1/\sigma^2$. The same result is obtained, if μ and σ are assumed to be independent, and Jeffrey’s rule is applied twice.
- 2nd case: Jeffrey’s rule is applied to the parameters $\alpha = E(X)$ and $CF = \sqrt{(VAR(X)/E(X))}$ (coefficient of variation). The resulting hyperprior in terms of μ and σ has the form

$$f(\mu, \sigma) \propto \sqrt{\frac{2e^{-2\mu-3\sigma^2}(e^{\sigma^2} - 1)}{\sigma^6}}.$$

- 3rd case: Jeffrey’s rule is applied to the parameters $\alpha = E(X)$ and $b = \sqrt{(VAR(X))}$. We recall that $VAR(X) = CF^2\alpha^2$. The resulting hyperprior in terms of μ and σ has the form

$$f(\mu, \sigma) \propto \sqrt{\frac{2e^{-4(\mu+\sigma^2)}(e^{\sigma^2} - 1)}{\sigma^6}}.$$

- 4th case: Jeffrey’s rule is applied to the parameters $\alpha = E(X)$ and σ , assuming independence between them. In this case has we have:

$$f(\mu, \sigma) = \sqrt{\frac{e^{-2\mu-\sigma^2}(2 + 4\sigma^2)}{\sigma^4}}.$$

For each case the steps in the calculation are similar to those for the gamma model, except that in step 1, truncation is applied to the parameters of the lognormal density.

The likelihood $P(X_1, \dots, X_n, T_1, \dots, T_n | \mu, \sigma)$ as a function of μ and σ (7) is presented in Figure 5, with values for $(X_1, \dots, X_n, T_1, \dots, T_n)$ taken from case 1. For uniform hyperpriors, this likelihood is proportional to the hyperposterior distribution $P(\mu, \sigma | X, T)$. Note that, in contrast to Figure 2, $P(\mu, \sigma | X, T)$ does peak. This means that a “natural” truncation for μ and σ can be defined as any rectangle containing the peak. The choice which such rectangle will have negligible influence on the results.

Figure 6 shows the hyperposterior distribution. Again, the contrast with Figures 3 and 4 is striking. The mass is captured within the μ, σ rectangle containing the peak in Figure 5.

The results are presented in Tables 9–11 and for each hyperprior distribution discussed. For case 1 we include the effect of omitting the square root in the

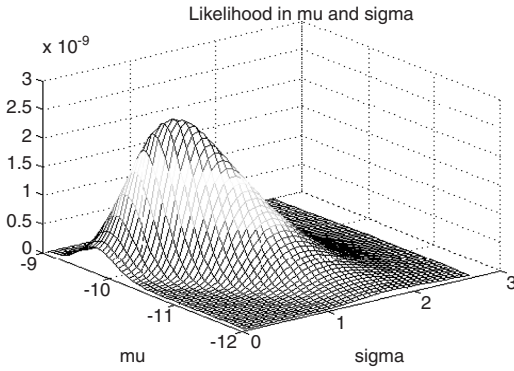


Figure 5.

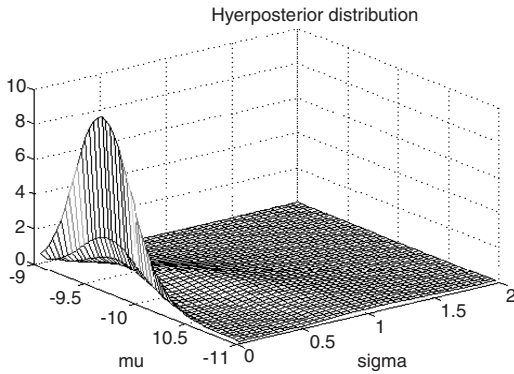


Figure 6.

Table 9. The 5%, 50% and 95% quantiles of the posterior distribution of λ_0 for data set 1.

Quantiles	5%	50%	95%
1st case	2.367 E-5	6.133 E-5	1.217 E-4
2nd case	2.803 E-5	5.805 E-5	1.067 E-4
3rd case	2.548 E-5	5.45 E-5	1.022 E-4
4th case	2.274 E-5	5.751 E-5	1.222 E-4
1st case ZEDB	2.01 E-5	5.91 E-5	1.44 E-4

Table 10. The 5%, 50% and 95% quantiles of the posterior distribution of λ_0 for data set 2.

Quantiles	5%	50%	95%
1st case	2.603 E-5	7.714 E-5	1.381 E-4
2nd case	7.838 E-6	4.817 E-5	1.246 E-4
3rd case	3.016 E-6	1.967 E-5	8.281 E-5
4th case	8.541 E-6	4.852 E-5	1.263 E-4
1st case ZEDB	2.539 E-5	7.710 E-5	2.058 E-4

Table 11. The 5%, 50% and 95% quantiles of the posterior distribution of λ_0 for data set 3.

Quantiles	5%	50%	95%
1st case	2.414 E-5	6.935 E-5	2.501 E-4
2nd case	2.179 E-5	3.794 E-5	1.357 E-4
3rd case	2.126 E-5	3.261 E-5	7.895 E-5
4th case	2.186 E-5	3.860 E-5	1.468 E-4
1st case ZEDB	1.12 E-5	6.228 E-5	3.398 E-4

Jeffrey prior. The corresponding ZEDB results are shown in each table. The differences are smaller than with the gamma model the differences noted above.

4 TRUNCATION

Using a gamma prior, the method of truncation seems to have a large influence on the posterior distribution of λ . It has been shown in section 2.3 that the likelihood in α and β has no maximum, but it is asymptotically maximal along a ridge. [Cooke et al., 1995] showed that different choices of truncation ranges can affect the median and the 95% quantile by a factor 5. In (4), the term $\prod_{i=1, \dots, n} \int [P(X_i|\lambda_i)P(\lambda_i|q)d\lambda_i]$ cannot be calculated analytically when we have a lognormal distribution as prior for λ . Hence, we cannot study the asymptotic behavior of the “hyperposterior” $P(\mu, \sigma | X_1, \dots, X_n, T_1, \dots, T_n) \propto P(X_1, \dots, X_n, T_1, \dots, T_n | \mu, \sigma)P(\mu, \sigma)$ analytically. Performing numerical integration (Figure 9), one can see that a maximum occurs in the likelihood in μ and σ . Hence, if the parameters of the lognormal distribution μ and σ are truncated in a way that includes the bulk of mass around the maximum, then how they are truncated will not make a significant difference. To save time in the computation process, truncation is performed around the significant values of likelihood in μ and σ (Figure 5). Figure 7 shows this same likelihood, but a larger integration rectangle for μ and σ ; integration over the larger rectangle produces effectively the same result. The intervals for integration over λ , were determined as in [Becker and Schubert, 1998]. One remark can be made: if the domains of integration are not large enough the posterior cumulative distribution of λ will not go to one. Using an iterative loop in software implementation, the natural interval of integration can be found.

The possibility of truncating the domains of integration so as to include the bulk of mass around the maximum of the prior is a significant argument in favor of the lognormal prior over the gamma prior. From Table 1 we see that if the variance of a gamma distribution is proportional to the mean, hence they go to zero at the same rate. For the lognormal the variance is proportional to the square of the mean, and hence the variance goes to zero faster than the mean. We can

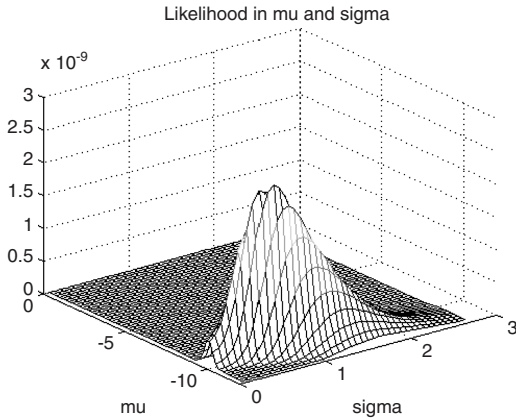


Figure 7. Likelihood μ, σ dataset 1; $\mu = -17.5 \dots -3$, $\sigma = 0.1 \dots 4$, $\lambda = 2 * 10^{-6} \dots 3 * 10^{-4}$.




Figure 8. Likelihood μ, σ dataset 3; $\mu = -15.5 \dots -6$, $\sigma = -0.1 \dots 4$, $\lambda = 2 * 10^{-5} \dots 3 * 10^{-3}$.

anticipate concentration of mass near $\sigma = 0$ in such cases. In any event, with sufficient observation times failures will be observed and the lognormal prior will peak away from zero and hence admit a good truncation heuristic. For the gamma this is not the case. Figure 8 shows the likelihood of μ, σ for dataset 3. Here a peak is not visible at all, but mass seems to concentrate at $\sigma = 0$. These figures also give the integration ranges for μ, σ , and λ .

5 CONCLUSIONS

Two-stage models provide a valid method for assimilating data from other plants. The conditional independence assumptions are reasonable and yield a

tractable and mathematically valid form for the failure rate a plant of interest, given failures and operational times at other plants in the population. However the choice of hyperprior must be defensible since improper hyperpriors do not always become proper after observations. The lognormal model enjoys a significant advantage over the gamma model in that, as observation time increases, a natural truncation of the hyperparameters μ, σ is possible. In the context of a literature survey, Vaurio's one-stage empirical Bayes model is elegant and simple but will not work with zero observed failures or with a population of two plants. Moreover, Hofer's model appears to rest on shifting viewpoints involving conflicting assumptions. Consistent application of the standard conditional assumptions collapses his model into the form (4), which he criticizes as a "wrong chance model". Further discussion should wait until the conditional independence assumptions and mathematical derivation are clarified.

REFERENCES

- Becker G. 2001. Personal communication.
- Becker G. & Hofer E. 2001. *Schätzung von Zuverlässigkeits-Kennwerten aus der Evidenz verschiedener Anlagen*, GRS-A-2918.
- Becker G. & Schubert B. 1998. ZEDB – A practical application of a 2-stage Bayesian analysis code, ESREL '98 conference, June 16–19, Trondheim, Norway, *Safety and Reliability*, Lydersen, Hansen and Sandtorv (eds), Balkema, Rotterdam, vol. 1, pg. 487–494.
- Box G. & Tiao G. 1974. *Bayesian Inference in Statistical Analysis*, Addison Wesley, New York.
- Cooke R.M., Bunea C., Charitos T. & Mazzuchi T.A. 2002. *Mathematical Review of ZEDB Two-Stage Bayesian Models*, Report 02–45, Department of Mathematics, Delft University of Technology.
- Cooke R.M., Dorrepaal J. & Bedford T.J. 1995. *Review of SKI Data Processing Methodology* SKI Report.
- Hofer E. 1999. On two-stage Bayesian modeling of initiating event frequencies and failure rates, *Reliability Engineering and System Safety*, vol. 66, pg. 97–99.
- Hofer E., Hora S., Iman R. & Peschke J. 1997. On the solution approach for Bayesian modeling of initiating event frequencies and failure rates, *Risk Analysis*, vol. 17, no. 2, pg. 249–252.
- Hofer E. & Peschke J. 1999. Bayesian modeling of failure rates and initiating event frequencies, ESREL 1999, Munich, Germany, *Safety and Reliability*, Schueller and Kafka (eds), Balkema, Rotterdam, vol. 2, pg. 881–886.
- Hora S. & Iman R. 1987. Bayesian analysis of learning in risk analysis, *Technometrics*, vol. 29, no. 2, pg. 221–228.
- Hardy G.H., Littlewood J.E. & Polya G. 1983. *Inequalities*, Cambridge University Press, Cambridge.
- Iman R. & Hora S. 1989. Bayesian methods of modeling recovery times with application to the loss of off-site power at nuclear power plants, *Risk Analysis*, vol. 9, no. 1, pg. 25–36.

- Iman R. & Hora S. 1990. Bayesian modeling of initiating event frequencies at nuclear power plants, *Risk Analysis*, vol. 10, no.1, pg. 103–109.
- Jaynes E. 1968. Prior probabilities, *IEEE Transactions On System Science and Cybernetic*, vol. 4, no. 3, September, pg. 227–241.
- Kaplan S. 1985. Two-stage Poission-type problem in probabilistic risk analysis, *Risk Analysis*, vol. 5, no. 3, pg. 227–230.
- Meyer W. & Hennings W. 1999. Prior distributions in two-stage Bayesian estimation of failure rates, ESREL 1999, Munich, Germany, *Safety and Reliability*, Schueller and Kafka (eds), Balkema, Rotterdam, vol. 2, pg. 893–898.
- Pörn K. 1990. *On Empirical Bayesian Inference Applied to Poisson Probability Models*, Linköping Studies in Science and Technology, Dissertation No. 234, Linköping.
- Royden H.L. 1968. *Real Analysis*, Macmillan, Toronto.
- T-Book Reliability Data of Components in Nordic Nuclear Power Plants*, 1987, ATV Office, Vattenfall, AB, S-162, Vallingby Sweden.
- Vaurio J.K. 1987. On analytic empirical Bayes Estimation of failure rates, *Risk Analysis*, vol. 17, no.3, pg. 329–338.

