Ann. Inst. Statist. Math.
Vol. 55, No. 1, 7–19 (2003)

BAYESIAN ESTIMATION OF SYSTEM RELIABILITY IN BROWNIAN STRESS-STRENGTH MODELS

SANJIB BASU AND RAMA T. LINGHAM

Division of Statistics, Department of Mathematical Sciences, Northern Illinois University, DeKalb, IL 60115-2854, U.S.A., e-mail: basu@niu.edu; rama@math.niu.edu

(Received November 9, 2000; revised December 3, 2001)

Abstract. A stress-strength system fails as soon as the applied stress, X, is at least as much as the strength, Y, of the system. Stress and strength are time-varying in many real-life systems but typical statistical models for stress-strength systems are static. In this article, the stress and strength processes are dynamically modeled as Brownian motions. The resulting stress-strength system is then governed by a time-homogeneous Markov process with an absorption barrier at 0. Conjugate as well as non-informative priors are developed for the model parameters and Markov chain sampling methods are used for posterior inference of the reliability of the stress-strength system. A generalization of this model is described next where the different stress-strength systems are assumed to be exchangeable. The proposed Bayesian analyses are illustrated in two examples where we obtain posterior estimates as well as perform model checking by cross-validation.

Key words and phrases: Cross-validation, first-passage time, Gibbs sampler, hitting time, non-informative prior, prediction.

1. Introduction

A physical system, whether it is a single component or a large structure, possesses an intrinsic 'strength' Y. The system itself is typically operating subject to some kind of environmental 'stress' X and is rendered out of commission at its time of failure, T say, which is reached as soon as X is at least as much as Y. Typically, statistical inference focuses on reliability of the system, which is given by $R = P(Y > X)$. For example, Weerahandi and Johnson (1992) study a situation where X is the chamber pressure generated by ignition in a solid propellant rocket engine and Y is the rocket chamber burst strength. When $X \geq Y$, the chamber blows up, resulting in a failure.

The stress-strength model was first considered by Birnbaum (1956). Traditionally, stochastic models for stress-strength (SS) systems are *static* in the sense that potential data on X and Y are considered not to involve the time of system operation. In such models, m observations on X and n observations on Y are gathered. The distribution of X and Y are modeled parametrically or nonparametrically. For example, Weerahandi and Johnson (1992) and many others assume that X and Y are normally distributed. Inferences for R, such as point and interval estimation as well as hypothesis testing, are then sought based on the assumed model. Comprehensive reviews of frequentist inference for stress-strength models are given in Basu (1985) and Johnson (1988). For a review of Bayesian work in this area, see Ghosh and Sun (1998).

From a practical point of view, the status of a stress-strength system clearly changes dynamically with time. The *dynamic* approach to modeling SS-systems accommodates a time-dependent stress process, $X(t)$, and a strength process $Y(t)$. Let $Z(t) = Y(t) - X(t)$ represent the difference between strength and stress at time t. In the reliability engineering literature, $Z(t)$ is called the limit-state function. The system fails as soon as stress exceeds strength and hence the failure time, T, is the first-passage functional given by

$$T = \inf\{0 \leq t \leq \infty : Z(t) \leq 0\}. \tag{1.1}$$

Basu and Ebrahimi (1983) have discussed this dynamic model for SS-systems. Reliability of the system is a function of time. For a specific time point $t^* > 0$, the reliability of a SS-system, $R(t^*)$, is the chance that the system will not fail within $(0, t^*]$. It follows from (1.1) that

$$R(t^*) = P(T > t^*) = P\left(\inf_{0<t\leq t^*} Z(t) > 0\right). \tag{1.2}$$

As mentioned earlier, a common assumption in the static case is that X and Y are independently normally distributed. Ebrahimi and Ramallingam (1993) motivate and generalize this to a dynamic model, where $X(t)$ and $Y(t)$ are independent Brownian motion (BM) processes with mean functions $\mu_X t$ and $\mu_Y t$ and variance functions $\sigma_X^2 t$ and $\sigma_Y^2 t$ respectively. However, note that we observe $X(t)$ and $Y(t)$ and hence the difference $Z(t)$ until time T when $Z(t)$ hits zero for the first time and the system fails. Observations on the system cease after its failure. In particular, we actually observe the *stopped processes* $X^*(t)$, $Y^*(t)$, and $Z^*(t) = Y^*(t) - X^*(t)$ where $X^*(t) = X(t)$ if $t < T$ and $= X(T)$ if $t \geq T$, and $Y^*(t)$, $Z^*(t)$ are similarly defined. While the latent $X(t), Y(t)$ and hence $Z(t)$ are Gaussian processes, the observed stopped process $Z^*(t)$ is no longer Gaussian but a time-homogeneous Markov process (see Section 2).

Our inferential goals for this model are multifold. The first goal is Bayesian estimation of model parameters. The second goal is predictive inference. Consider a new system, identical in characteristic with the observed ones, which is beginning to operate with $Z(0) = z > 0$. For a future time t^*, the reliability of this new system is

$$R_z(t^*) = P_z(T > t^*) = P_z\left(\inf_{0\leq t\leq t^*} Z(t) \leq 0\right),$$

where the subscript z in P_z refers to the assumption that $Z(0) = z$. We seek inferences for $R_z(t^*)$, such as point and interval estimates. The third goal is inference about an unfailed observed system. Suppose that $Z^*(t_M) = z > 0$ at the end of the monitoring period $[0, t_M]$. The objective here is to predict the reliability of this system for a future time point $t^* > t_M$. It however follows that

$$R_{z,t_M}(t^*) = P(T > t^* \mid Z^*(t_M) = z) = R_z(t^* - t_M),$$

i.e., the reliability of an equivalent system at a time $t^* - t_M$ where the system is started at $Z(0) = z$. Thus, the second and the third cases are equivalent.

The plan for this paper is as follows. In Section 2, we summarize relevant results about the involved likelihood function. Conjugate and default priors for the model are developed in Section 3. We establish posterior propriety under the default prior in Section 4 and describe how Markov chain sampling methods can be used for posterior analysis. In Section 5, we illustrate the proposed analysis in a data set and compare the performances of the maximum likelihood estimate with the Bayesian estimate. The proposed model is extended to an exchangeable setup in Section 6 and Bayesian inference for this exchangeable model is developed. Section 7 describes another application where we obtain posterior point and interval estimates of reliability and perform model checking from a cross-validated prediction viewpoint.

2. Sampling model

We assume that we have independent observations from n identical stress-strength systems. Observations on the i-th system are taken periodically at pre-determined time grids $0 = t_{i0} < t_{i1} < \cdots < t_{iM_i}$ at which point monitoring stops (time truncation). We use $\underset{\sim}{t}_i$ to denote the sequence of time points $(t_1, \ldots, t_M)$. The periods $t_{ij} - t_{ij-1}$, $j = 1, \ldots, M_i$ are not assumed to be equal. Finally, observation stops if the system fails during the monitoring period. In particular, we observe the stopped processes $Z_i^*(t)$, $i = 1, \ldots, n$ on discrete time-grids.

If the i-th system does not fail during the monitoring period then t_{iM_i} represents the last monitored time point. For a failed system, the failure time T_i may occur in between two time grids $t_{i,j-1}$ and $t_{i,j}$. However, since failure is a catastrophic event, we assume that the failure time T_i is exactly recorded even when it occurs within a time-grid interval. In such a case, for notational convenience, we relabel the point on the grid immediately following the failure as t_{iM_i} and observation stops after failure.

If we had observation from a Brownian motion (BM) on a discrete time-grid, it would have been easy to write down the likelihood function utilizing the independent normal increments of the BM. We instead observe the stopped process $Z^*(t)$. The complication in the distributional property of the process $Z^*(t)$ stems from the restriction that $Z^*(t)$ is known to be positive for $0 \le t < T$ until it is absorbed at 0 and we stop observing the process.

The description of the transition kernel for the $Z^*(t)$ process thus has two parts. Given that the process is currently at $Z^*(0) = a > 0$, (i) it may move to a point $z > 0$ at a future time-point t, or (ii) it may get absorbed at 0 by time t. Ebrahimi and Ramallingam (1993) showed that $\{Z^*(t), t \ge 0\}$ is a time-homogeneous Markov process with the following transition kernel describing the cases (i) and (ii) respectively.

(i) $P(Z^*(t) > z \mid Z^*(0) = a, \mu, \sigma^2) = \int_{b=z}^{\infty} \psi_t(a, b, \mu, \sigma^2) db$ for $a > 0$, $z \ge 0$, $t > 0$, and

(ii) $P(Z^*(t) = 0 \mid Z^*(0) = a, \mu, \sigma^2) = 1 - H_a(t, \mu, \sigma^2)$ for $a > 0$, $t > 0$,
where $\mu = \mu_Y - \mu_X, \sigma^2 = \sigma_X^2 + \sigma_Y^2$, and

$$\psi_t(a, b; \mu, \sigma^2) = \frac{1}{\sigma\sqrt{t}} \left[\phi\left(\frac{-b + a + \mu t}{\sigma\sqrt{t}}\right) - \phi\left(\frac{-b - a + \mu t}{\sigma\sqrt{t}}\right) \exp\left(\frac{-2\mu a}{\sigma^2}\right) \right], \tag{2.1}$$

$$H_a(t; \mu, \sigma^2) = \Phi\left(\frac{a + \mu t}{\sigma\sqrt{t}}\right) - \left\{ \Phi\left(\frac{-a + \mu t}{\sigma\sqrt{t}}\right) \exp\left(\frac{-2\mu a}{\sigma^2}\right) \right\}. \tag{2.2}$$

We can now proceed to write down the form of the likelihood function. For the i-th system, we observe snapshots of the trajectory at discrete time-points $\{z_i^*(t_{ij}), j = 1, \ldots, M_i\}$. All n systems start at time $t_{i0} = 0$ with a known positive value. At any intermediate t_{ij}, $z_i^*(t_{ij})$ is clearly > 0 and hence the likelihood is determined by (i) above. If the system failed at the last time point $z_i^*(t_{i,M_i}) = 0$ then the likelihood contribution is determined from (ii) above. Let

$$h_a(t;\mu,\sigma^2) = -\frac{d}{dt}H_a(t;\mu,\sigma^2) = \frac{a}{\sigma t^{3/2}\sqrt{2\pi}}\exp\left[-\frac{(a+\mu t)^2}{2t\sigma^2}\right]. \tag{2.3}$$

If we now let $\Delta_{ij} = t_{ij} - t_{i,j-1}$ denote the grid widths, the likelihood function is then

$$L(\mu,\sigma^2) = \prod_{i=1}^{n}\prod_{j=1}^{M_i} g_{\Delta_{ij}}(z_i^*(t_{i,j-1}), z_i^*(t_{ij}), \mu; \sigma^2) \tag{2.4}$$

where $g_t(a,b;\mu,\sigma^2)$ combines cases (i) and (ii) in a concise form as

$$g_t(a,b;\mu,\sigma^2) = \begin{cases} \psi_t(a,b;\mu,\sigma^2), & a>0,\ b>0 \\ h_a(t;\mu,\sigma^2), & a>0,\ b=0 \\ 1, & a=0,\ b=0. \end{cases} \tag{2.5}$$

One of our inferential goals is the reliability of such a stress-strength system at time t^* when initially $Z(0) = z > 0$. From (1.2), $R_z(t^*;\mu,\sigma^2) = P(T > t^* \mid \mu,\sigma^2)$. If the mean parameter $\mu \le 0$, then, based on a well-known Brownian motion result, the hitting time T has density $h_z(t,\underset{\sim}{\theta})$, which is an inverse Gaussian distribution. If $\mu > 0$, then as is expected, the process may not reach 0 in finite time and $P(T < \infty \mid Z(0) = z, \mu, \sigma^2) = \exp(-\frac{2\mu z}{\sigma^2})$. From the above distributional results, it now follows that the reliability at a future time t^* is

$$R_z(t^*;\mu,\sigma^2) = H_z(t^*;\mu,\sigma^2). \tag{2.6}$$

3. Priors

For Bayesian analysis, one needs to specify prior distribution for the unknown parameters μ and σ^2 of the model. A conjugate type prior for the transformed parameters $\theta_1 = \mu/\sigma$ and $\theta_2 = 1/\sigma$ can be hierarchically specified as follows:

- θ_1 has a $N(\theta_0, 1/\tau_0)$ prior and θ_2^2 has an independent Gamma(α_0,β_0) prior with mean α_0/β_0.
- The hyperparameters have a prior $p(\theta_0,\tau_0,\alpha_0,\beta_0)$. For example, $\theta_0 \sim N(\nu, 1/\lambda)$, $\tau_0 \sim$ Gamma(a,b), $\alpha_0 \sim$ Gamma(c,d) and $\beta_0 \sim$ Gamma(e,f).

For later reference, we will denote this conjugate specification as $\pi_c(\theta_1,\theta_2)$.

On the other hand, the available prior information may be weak for a variety of reasons including automation of the reliability assessment without the benefit of careful subjective elicitations and one may want to specify an objective prior. Several strategies to formalize a (typically improper) default prior to facilitate the Bayesian analysis are given in the excellent review paper of Kass and Wasserman (1996). In this section, we primarily focus on two such priors, the Jeffreys prior and the reference prior. Many authors have derived default priors for the static stress-strength model. Thompson and Basu (1993) derive reference priors when the stress and strength are both exponentially distributed. Ghosh and Sun (1998) provide reference as well as probability matching priors for different stress-strength models.

In the sampling scheme of our dynamic model, the observation process is truncated at failure and hence, for the failed systems, the stopping times T_i's are part of the data. The resulting log-likelihood for these stopped processes is quite complicated as shown in (2.4) in the previous section. As a result, the information matrix is difficult to compute and the formal derivation of the Jeffreys and the reference priors for (μ, σ^2) is a formidable task. We take a pragmatic approach to the prior construction and *pretend* that the data, $\underset{\sim}{Z}_i^*$ are obtained from the *unstopped* BM $\{Z_i(t), t \geq 0\}$, $i = 1, 2, \ldots, n$. It should be remarked that a special case (i.e. $n = 1$) of such data from the (unstopped) Brownian motion process has been considered by Sivaganesan and Lingham (2000). Using the independent normal increment structure of the BM, it follows that the information matrix is

$$I(\mu, \sigma^2) = \text{diag}\left[\frac{\sum_{i=1}^n t_{iM_i}}{\sigma^2}, \frac{\sum_{i=1}^n M_i}{2\sigma^4}\right].$$

The most frequently used default prior is Jeffreys prior which is proportional to the positive square root of the determinant of the information matrix. Note, further that the information matrix is block-diagonal. For such information matrices, Datta and Ghosh (1995) have derived the reference prior explicitly. Based on their result, we obtain default priors as below:

$$\pi_q(\mu, \sigma^2) = (\sigma^2)^{-q/2} I(-\infty < \mu < \infty, \sigma^2 > 0) \quad \text{for} \quad q \geq 2, \text{ or, alternatively}$$
$$\pi_q(\mu, \sigma) = \sigma^{1-q} I(-\infty < \mu < \infty, \sigma > 0) \quad \text{for} \quad q \geq 2.$$

If we plug-in $q = 2$ and $q = 3$, we respectively obtain the "reference" prior and the "Jeffreys" prior. In terms of the reparametrization $\underset{\sim}{\theta} = (\theta_1, \theta_2)$ with $\theta_1 = \mu/\sigma$ and $\theta_2 = 1/\sigma$ the default priors are given by

$$\pi_q(\underset{\sim}{\theta}) = \theta_2^{q-4} I(-\infty < \theta_1 < \infty, \theta_2 > 0). \tag{3.1}$$

4. Posterior analysis

4.1 *Propriety of the posterior*

The likelihood and the priors for our stress-strength model were described in Sections 2 and 3 respectively. Here, we pursue posterior analysis of the model. Note that the default prior on $\underset{\sim}{\theta}$ described in (3.1) is improper. A very important question in such a case is the propriety of the posterior. This is especially a crucial issue in Markov chain Monte Carlo (MCMC) analysis where an improper posterior may not be immediately obvious in the full conditional distributions or in the transition kernel of the Markov chain but may lead to completely erroneous results.

Result 1. Let $N = \sum_{i=1}^n M_i$. Assume that $N > 1$ and the improper default priors are given by $\pi_q(\underset{\sim}{\theta}) = \theta_2^{q-4}$, where $q \geq 2$. Let D denote the observed data. Then,

(i) The posterior $\pi_q(\underset{\sim}{\theta} \mid D)$ is proper.

(ii) In addition, θ_1, θ_2, μ and σ have finite posterior means.

PROOF. We will only prove (i) as the proof of (ii) is similar. Moreover, for simplicity of exposition, we assume that each of the n observed systems is 'alive' (has not reached 0) at the end of the monitoring time. The case when some of the systems fail during monitoring can be dealt with in an analogous manner.

Fix $t > 0$, $a > 0$, $b > 0$, $\underset{\sim}{\theta} \in \Theta$. Then, one can show that

$$\psi_t(a, b; \underset{\sim}{\theta}) < \theta_2 \cdot \frac{1}{\sqrt{t}} \phi\left[\frac{\theta_1 t - (b - a)\theta_2}{\sqrt{t}}\right].$$

The (i,j)-th term in the likelihood, $L(\underset{\sim}{\theta})$ is $\psi_{\Delta_{ij}}(a_{ij}, b_{ij})$ (see (2.4)), where $a_{ij} = z_i^*(t_{i,j-1}), b_{ij} = z_i^*(t_{i,j})$. Then, using the above inequality, for a positive constant c_1, we get,

$$L(\underset{\sim}{\theta}) < c_1 \theta_2^N e^{-1/2[A-(C^2/B)]\theta_2^2} \times e^{-B/2[(\theta_1-(\theta_2 C/B))]^2}$$

where

$$A = \sum_{i=1}^{n}\sum_{j=1}^{M_i} \frac{(b_{ij}-a_{ij})^2}{\Delta_{ij}}, \qquad B = \sum_{i=1}^{n}\sum_{j=1}^{M_i} \Delta_{ij}, \quad \text{and} \quad C = \sum_{i=1}^{n}\sum_{j=1}^{M_i} (b_{ij}-a_{ij}).$$

Since the prior is $\pi_q(\underset{\sim}{\theta}) = \theta_2^{q-4}$ and the posterior is $\pi_q(\underset{\sim}{\theta} \mid D) = c_2 \cdot \pi_q(\underset{\sim}{\theta}) \times L(\underset{\sim}{\theta})$ for normalizing constant c_2, we get,

$$\int_{\theta_1=-\infty}^{\infty} \pi_q(\underset{\sim}{\theta} \mid D) d\theta_1 < c_1 c_2 \cdot (\sqrt{2\pi/B}) \theta_2^{(N+q-4)} e^{-1/2[A-(C^2/B)]\theta_2^2}.$$

Since $A > C^2/B$ by Cauchy-Schwarz inequality, it now follows that

$$\int_{\theta_2=0}^{\infty}\int_{\theta_1=-\infty}^{\infty} \pi_q(\underset{\sim}{\theta} \mid D) d\underset{\sim}{\theta} < \infty$$

which establishes the propriety of the posterior. □

Remark. The condition that $N > 1$ is needed for the validity of the Cauchy-Schwarz inequality and simply means that we should plan to get at least two observations from the n systems that are being monitored. As there are two unknown parameters that govern the Brownian systems, this is a natural condition to impose to insure that the posterior is proper.

4.2 *Sampling method*

The posterior of the proposed model is analytically intractable due to the complicated form of the likelihood in (2.4). Instead, we estimate relevant posterior quantities based on samples drawn from the posterior via Markov chain Monte Carlo (MCMC). MCMC methods construct a Markov chain whose invariant distribution is the posterior and sample path averages of this Markov chain are used to estimate posterior quantities. An easy way to construct this chain is Gibbs sampling whose transition kernel is the product of the full conditional distributions. In our case, we only have two parameters θ_1 and θ_2. However, the full conditional distribution of each parameter is rather complex. An alternative is to use Metropolis-Hastings algorithm but it raises several other issues such as the choice of the proposal distributions. Due to the complex functional form of the likelihood, it is not clear what proposal distribution(s) would provide an efficient sampler. Here, the following result is useful.

Result 2. The joint posterior $\pi_q(\theta_1, \theta_2 \mid D)$ based on the default prior of (3.1) is a log-concave function of θ_1. Same is true for θ_2.

PROOF. Note that $\log \pi_q(\theta_1, \theta_2 \mid D) = \log c + \log \pi_q(\underset{\sim}{\theta}) + \log L(\underset{\sim}{\theta})$ where c is the normalizing constant. We first show log-concavity in θ_1. Since the prior $\pi_q(\underset{\sim}{\theta}) = \theta_2^{q-4}$ does not depend on θ_1, we need to show that the log-likelihood is concave in θ_1. From

(2.4), $\log L(\underset{\sim}{\theta}) = \sum_{i=1}^{n} \sum_{j=1}^{M_i} \log g_{\Delta_{ij}}(z_i^*(t_{i,j-1}), z_i^*(t_{ij}))$. Based on the form of $g(\ )$ in (2.5), it suffices to prove that $\psi_t(\ )$ and $h_a(\ )$ are log-concave in θ_1. Ebrahimi and Ramallingam (1993) showed that $\frac{d^2 \log \psi_t}{d\theta_1^2} = -t$ and $\frac{d^2 \log h_a(t)}{d\theta_1^2} = -t$. It follows that $\psi_t(\ )$, $h_a(\ )$ and hence the posterior $\pi_q(\theta_1, \theta_2 \mid D)$ are log-concave in terms of θ_1. The proof of log-concavity in terms of θ_2 is more involved since the prior $\pi_q(\underset{\sim}{\theta}) = \theta_2^{q-4}$ is not log-concave in θ_2 by itself. We consider $\log \pi_q(\underset{\sim}{\theta}) + \log L(\underset{\sim}{\theta})$ and the proof follows along similar lines by showing that the second partial $\frac{d^2}{d\theta_2^2}$ is non-positive. □

Result 2 establishes that the full conditional distributions $\pi_q(\theta_1 \mid D, \theta_2)$ and $\pi_q(\theta_2 \mid D, \theta_1)$ are log-concave and one can draw samples directly from these full conditional distributions by the adaptive rejection sampling (ARS) method of Gilks and Wild (1992).

If one instead wants to use the conjugate type prior $\pi_c(\underset{\sim}{\theta})$ described in Section 3, Markov chain sampling from the posterior proceeds along similar lines. The full conditional distributions of θ_1 and θ_2 are still log-concave; the former because both the likelihood (Result 2) and the prior are log-concave in θ_1 and the proof of the latter is similar to Result 2. The full conditional distributions of the hyperparameters θ_0, τ_0 and β_0 are respectively Normal, Gamma and Gamma due to the conjugate specification. Finally, the full conditional of α_0 can be shown to be log-concave.

To summarize, we use Gibbs sampling to construct a Markov chain with the posterior as its invariant distribution. We repeatedly and alternately draw samples from each full conditional distribution, either directly or by ARS. Finally, relevant posterior quantities are estimated as sample path averages of this Markov chain.

5. Application

We illustrate our proposed analysis in a data set with $n = 10$ systems. The data are simulated from a dynamic stress-strength (SS) model where the difference (= strength − stress) process $Z(t)$ is a BM with $\mu_{\text{true}} = -0.2$ and $\sigma_{\text{true}}^2 = 1$. The data are shown in Fig. 1. Each SS system started at $Z(0) = 5$ and is observed at periods of 0.5. Monitoring stopped at $t = 25$. Six out of the ten systems failed during this monitoring period.

Ebrahimi and Ramallingam (1993) showed that for this dynamic SS model, the maximum likelihood estimates (MLEs) of μ and σ are unique. We obtained $\widehat{\mu}_{\text{mle}}$ and $\widehat{\sigma}_{\text{mle}}$ using numerical optimization. The results are shown in Table 1.

We obtained posterior inferences for the proposed Bayesian model with the default prior of (3.1) with $q = 2$. The adaptive rejection based Gibbs sampler is run for 15,000 iterations after a burn-in of 5,000. The estimated posterior mean, median and 95% credible interval for μ and σ are shown in Table 1. Both the mle and the posterior

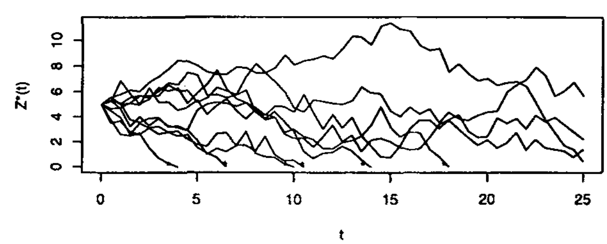

Fig. 1. Trajectories of 10 $Z^*(t)$ systems when $\mu = -0.2$, $\sigma = 1$.

Table 1. MLE and posterior inference for μ and σ.

	True Value	MLE	Posterior		
			Mean	Median	95% Credible interval
μ	−0.2	−0.25	−0.2507	−0.2503	(−0.414, −0.08597)
σ	1.0	2.35	1.035	1.033	(0.9631, 1.116)

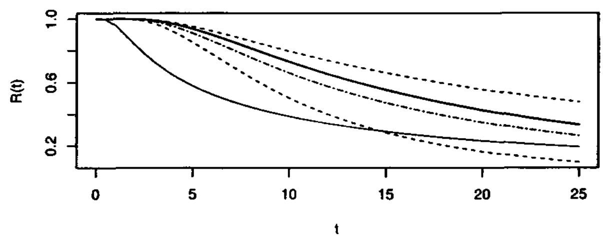

Fig. 2. Reliability estimation. The thick solid line is the "true" reliability. The thin solid line is the mle. The middle dashed line is the posterior mean and the outer dashed lines represent a pointwise 90% credible band.

estimates of μ are close to the true value. However, $\widehat{\sigma}_{\text{mle}}$ clearly overestimates while the posterior estimates of σ are close to σ_{true}.

Our main goal here is inference for reliability of the system. We consider a new system, say the $(n+1)$-th one, which will also start at $Z(0) = 5$ and estimate its reliability $R_5(t^*; \mu, \sigma^2)$ at a future time t^* using (2.6). Figure 2 shows the "true" reliability $R(t^*; \mu_{\text{true}}, \sigma^2_{\text{true}})$ which we obtain for $0 \le t^* \le 25$ at grids of 0.5. By invariance, the mle of $R_{z_0}(t^*; \mu, \sigma^2)$ is $R_{z_0}(t^*; \widehat{\mu}_{\text{mle}}, \widehat{\sigma^2}_{\text{mle}})$ which is also shown in Fig. 2. Finally, at each t^* grid, we estimate the posterior mean and a 90% credible interval for $R_{z_0}(t^*; \mu, \sigma^2)$. For example, the posterior mean at a fixed t^* is obtained as the average of $\{R_{z_0}(t^*; \mu_{(k)}, \sigma^2_{(k)}), k = 1, \ldots, K\}$ where $(\mu_{(k)}, \sigma_{(k)})$ are the values drawn at the k-th iteration of the Gibbs sampler and K is the total number of post-burn in iterations. The posterior estimate of reliability is clearly superior to the mle and the (pointwise) posterior credible band contains the true reliability curve.

6. An exchangeable model

We have so far modeled the strength $X_i(t)$ and stress $Y_i(t)$ processes for the $i = 1, \ldots, n$ systems as independent copies of $\text{BM}(\mu_X, \sigma_X^2)$ and $\text{BM}(\mu_Y, \sigma_Y^2)$ respectively. However, often each SS system has its own characteristics which are reflected in its observed path. An attractive way to model this diversity among systems is through a random effects type model where we assume that the strength $X_i(t)$ and stress $Y_i(t)$ processes for the i-th system are $\text{BM}(\mu_{Xi}, \sigma^2_{Xi})$ and $\text{BM}(\mu_{Yi}, \sigma^2_{Yi})$ respectively, $i = 1, \ldots, n$. Let $\mu_i = \mu_{Yi} - \mu_{Xi}, \sigma_i^2 = \sigma^2_{Xi} + \sigma^2_{Yi}$ and let $\theta_{1i} = \mu_i/\sigma_i$ and $\theta_{2i} = 1/\sigma_i$, $i = 1, \ldots, n$ be the transformed parameters. Also, let $Z_i(t) = Y_i(t) - X_i(t)$ be the latent difference process which we do not observe but instead observe the stopped process $Z_i^*(t)$ and observation stops at failure time T_i as soon as $Z_i(t) \le 0$ for the first time. The exchangeability among the n systems is modeled by assuming that $\underset{\sim}{\theta}_i$, $i = 1, \ldots, n$ are i.i.d. from $\pi(\underset{\sim}{\theta}_i \mid \gamma)$ as follows:

• $\theta_{1i} = \mu_i/\sigma_i$, $i = 1, \ldots, n$ are i.i.d. $\sim N(\theta_0, 1/\tau_0)$. Independently, $\theta_{2i}^2 = 1/\sigma_i^2$, $i = 1, \ldots, n$ are i.i.d. Gamma (α_0, β_0) with scale $= 1/\beta_0$.

• The hyperparameter $\gamma = (\theta_0, \tau_0, \alpha_0, \beta_0)$ has a prior $\pi(\gamma)$.

We shall posit that

$$\pi(\gamma) = \pi_1(\theta_0, \tau_0) \times \pi_2(\alpha_0, \beta_0). \tag{6.1}$$

No standard approaches to deriving the reference prior under hierarchical set-ups are available in the literature. Yet, one may consider the following "default" priors for π_1 and π_2 in (6.1).

$$\pi_1(\theta_0, \tau_0) = 1/\tau_0 \quad \text{and} \quad \pi_2(\alpha_0, \beta_0) = \sqrt{(\alpha_0 PG(1, \alpha_0) - 1)/\beta_0^2\alpha_0}, \tag{6.2}$$

where $PG(1, x) = \sum_{i=0}^{\infty}(x+i)^{-2}$ is the PolyGamma function. Then, in view of the propriety results in Yang and Berger (1998), having integrated out the second level parameter γ, it is clear that the marginal distribution of $(\theta_{11}, \ldots, \theta_{1n}, \theta_{21}, \ldots, \theta_{2n})$ is proper. It follows that the marginal distribution of the data is proper. Furthermore, the full conditional distributions of θ_0, τ_0 are available in Bernardo and Smith ((1994), p. 440) and the full conditional distributions of α_0, β_0 can be easily obtained from Sun and Ye (1996).

Alternatively, one may specify a conjugate prior (γ) as before as: $\theta_0 \sim N(\nu, 1/\lambda)$, $\tau_0 \sim$ Gamma(a, b), $\alpha_0 \sim$ Gamma(α_1, α_2), $\beta_0 \sim$ Gamma(β_1, β_2) and θ_0, τ_0, α_0, β_0 are a priori independent. The posterior analysis of the exchangeable model under these conjugate priors structure can be performed once again via Gibbs sampling. The full conditional distribution of each of the hyperparameters is immediate. In fact, they are Normal, Gamma and Gamma once again for θ_0, τ_0 and β_0 respectively due to the conjugate structure. The full conditional of α_0 is log-concave and one can draw samples by adaptive rejection (Gilks and Wild (1992)).

Either with the default or with the conjugate prior model for the hyperparameters, the main issue in the MCMC implementation here is how to draw samples of $\theta_i = (\theta_{1i}, \theta_{2i})$ from their full conditional distributions at the first stage of the model.

Result 3. The full conditional distribution of θ_{1i} and θ_{2i} are log-concave, $i = 1, \ldots, n$.

PROOF. The log-likelihood from (2.4) is given by $\log L(\theta_{1i}, \theta_{2i}) = \sum_{j=1}^{M_i} \log g_{\Delta_{ij}}$ $(z_i^*(t_{i,j-1}), z_i^*(t_{ij}))$. The remainder of the proof is similar to the proof of Result 2 and mostly rests on the fact that $\psi_t(\ )$ and $h_a(\ )$ are log-concave functions of θ_{1i} and θ_{2i}. The proof here is even easier since the Gamma(α_0, β_0) prior for θ_{2i} is log-concave when $\alpha_0 \geq 1$. The case of $\alpha_0 < 1$ is similar to the default prior case as in Result 2. □

The implementation of the Gibbs sampler is now obvious. The parameters θ_{1i}, θ_{2i}, $i = 1, \ldots, n$ are sampled from their full conditional distributions by adaptive rejection whereas the second level parameters μ_0, τ_0, α_0, β_0 are sampled by either adaptive rejection or conjugate sampling.

7. Application 2

We illustrate inference for the proposed exchangeable model for $n = 10$ SS systems shown in Fig. 3. The trajectories in this figure represent 10 stopped $Z_i^*(t)$ processes where the unstopped $Z_i(t)$ processes are $(\mu_i^{(0)}, (\sigma_i^{(0)})^2)$ BMs and observation is stopped as soon as $Z_i(t) \leq 0$. Also, $\theta_{1i}^{(0)} = \mu_i^{(0)}/\sigma_i^{(0)}$, and $(\theta_{2i}^{(0)})^2 = 1/(\sigma_i^{(0)})^2$, $i = 1, \ldots, n$ are, in turn, obtained as random draws from the conjugate priors $N(\theta_0, 1/\tau_0)$ and Gamma(α_0, β_0) respectively (as described in the previous section).

We obtain posterior inference for the reliability $R_i(t^*; \underset{\sim}{\theta}_i)$ of each system, $i = 1, \ldots, 10$ for $0 \leq t^* \leq 25$ at grids of 0.5. Here $R_i(t^*; \underset{\sim}{\theta}_i) = H(t^*; \mu_i, \sigma_i^2)$ as defined in (2.6). This inference is obtained under the following conjugate prior model: $\theta_0 \sim N(0, 10000)$, $\tau_0 \sim$ Gamma(.01, .01), $\alpha_0 \sim$ Gamma(2, 1), $\beta_0 \sim$ Gamma(2, 1). The estimated posterior means as well as pointwise 90% credible bands of $R_i(t^*, \underset{\sim}{\theta}_i)$, $i = 1, \ldots, 10$ are shown in Fig. 4. We compare these point and interval estimates with the "true" reliability $R_i(t^*; \underset{\sim}{\theta}_i^{(0)})$ where the latter is obtained by evaluating the reliability at $\underset{\sim}{\theta}_i^{(0)} = (\theta_{1i}^{(0)}, \theta_{2i}^{(0)})$, i.e., the $\underset{\sim}{\theta}_i$ value used to simulate the observed $Z_i^*(t)$ process. As is expected, the exchangeable model brings in more variation in the estimation process and hence the credible bands are wider than those in Section 5.

The "true" $\theta_{i1}^{(0)}$ values for the fourth and ninth systems are positive and hence their "true" reliability $R_i(t^*, \underset{\sim}{\theta}_i^{(0)})$ stays high throughout the duration of the monitoring period. This results in the true reliability of the ninth system falling outside the 90% credible band. A similar phenomenon for the fourth system is seen in the prediction band shown in Fig. 5.

We now take a critical look at our proposed model and perform model checking from a predictive cross-validation viewpoint as discussed in Gelfand (1996) and many others. We predict the reliability of the i-th system $\widehat{R_i}(t^* \mid Z_{-i}^*)$ at time t^* based on observations on all the other $Z_1^*(t), \ldots, Z_{i-1}^*(t), Z_{i+1}^*(t), \ldots, Z_n^*(t)$ systems. Note that in this prediction, we do not include observations $Z_i^*(t)$ on the i-th system.

We now describe the prediction of $\widehat{R_i}(t^* \mid Z_{-i}^*)$. We have

$$R_i(t^* \mid Z_{-i}^*) = \iint R_i(t^*; \underset{\sim}{\theta}_i)\pi(\underset{\sim}{\theta}_i \mid \gamma, Z_{-i}^*)\pi(\gamma \mid Z_{-i}^*)d\underset{\sim}{\theta}_i d\gamma. \tag{7.1}$$

Moreover, the conditional distribution $\pi(\underset{\sim}{\theta}_i \mid \gamma, Z_{-i}^*)$ is independent of Z_{-i}^* and is simply the prior distribution of $\underset{\sim}{\theta}_i$. In the estimate $\widehat{R_i}(t^* \mid Z_{-i}^*)$, we estimate the integral over γ by Monte Carlo average over $\{\gamma^{(k)}, k = 1, \ldots, K\}$ where these are (approximate) draws from the posterior $\pi(\gamma \mid Z_{-i}^*)$ obtained by Markov chain sampling and K is the number of post-burn-in iterations.

Figure 5 shows the result of this cross-validated prediction for systems $1, 2, \ldots, 10$. We remind the reader that this reliability prediction, for example for system 1, only uses the observed data from systems $2, 3, \ldots, 10$ and does not use observations on system 1.

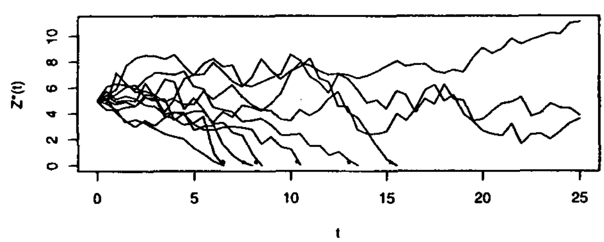

Fig. 3. Trajectories of 10 $Z^*(t)$ systems in the exchangeable model.

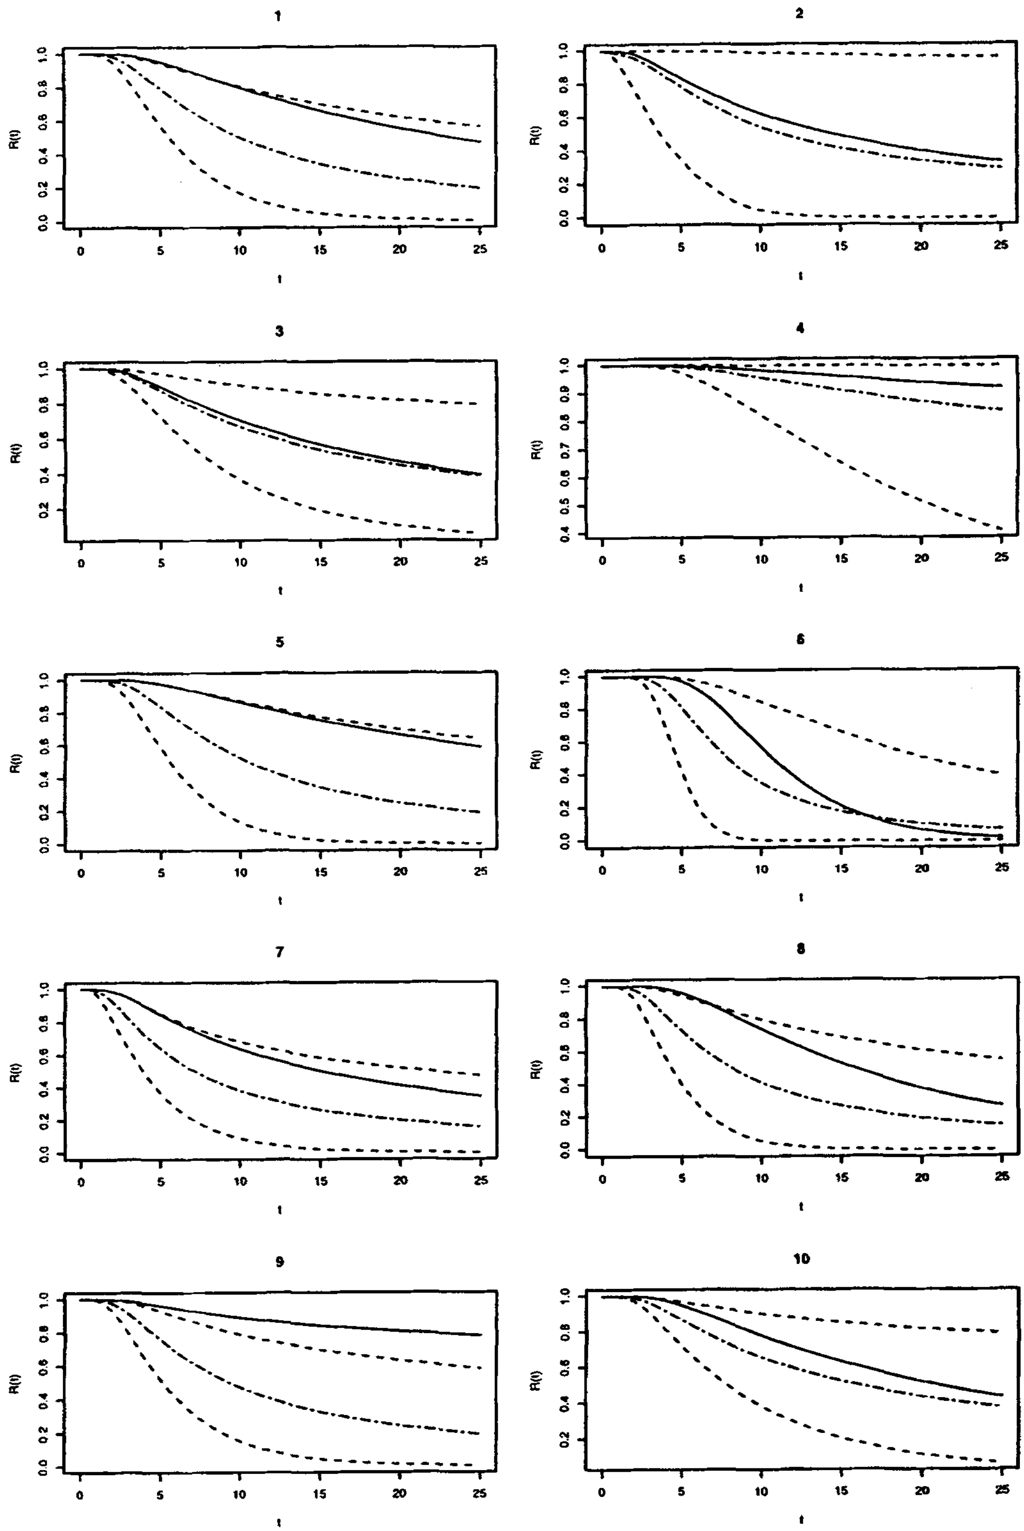

Fig. 4. Reliability estimation in the exchangeable model. The solid line is the "true" reliability. The middle dashed line is the posterior mean and the outer dashed lines represent a pointwise 90% credible band.

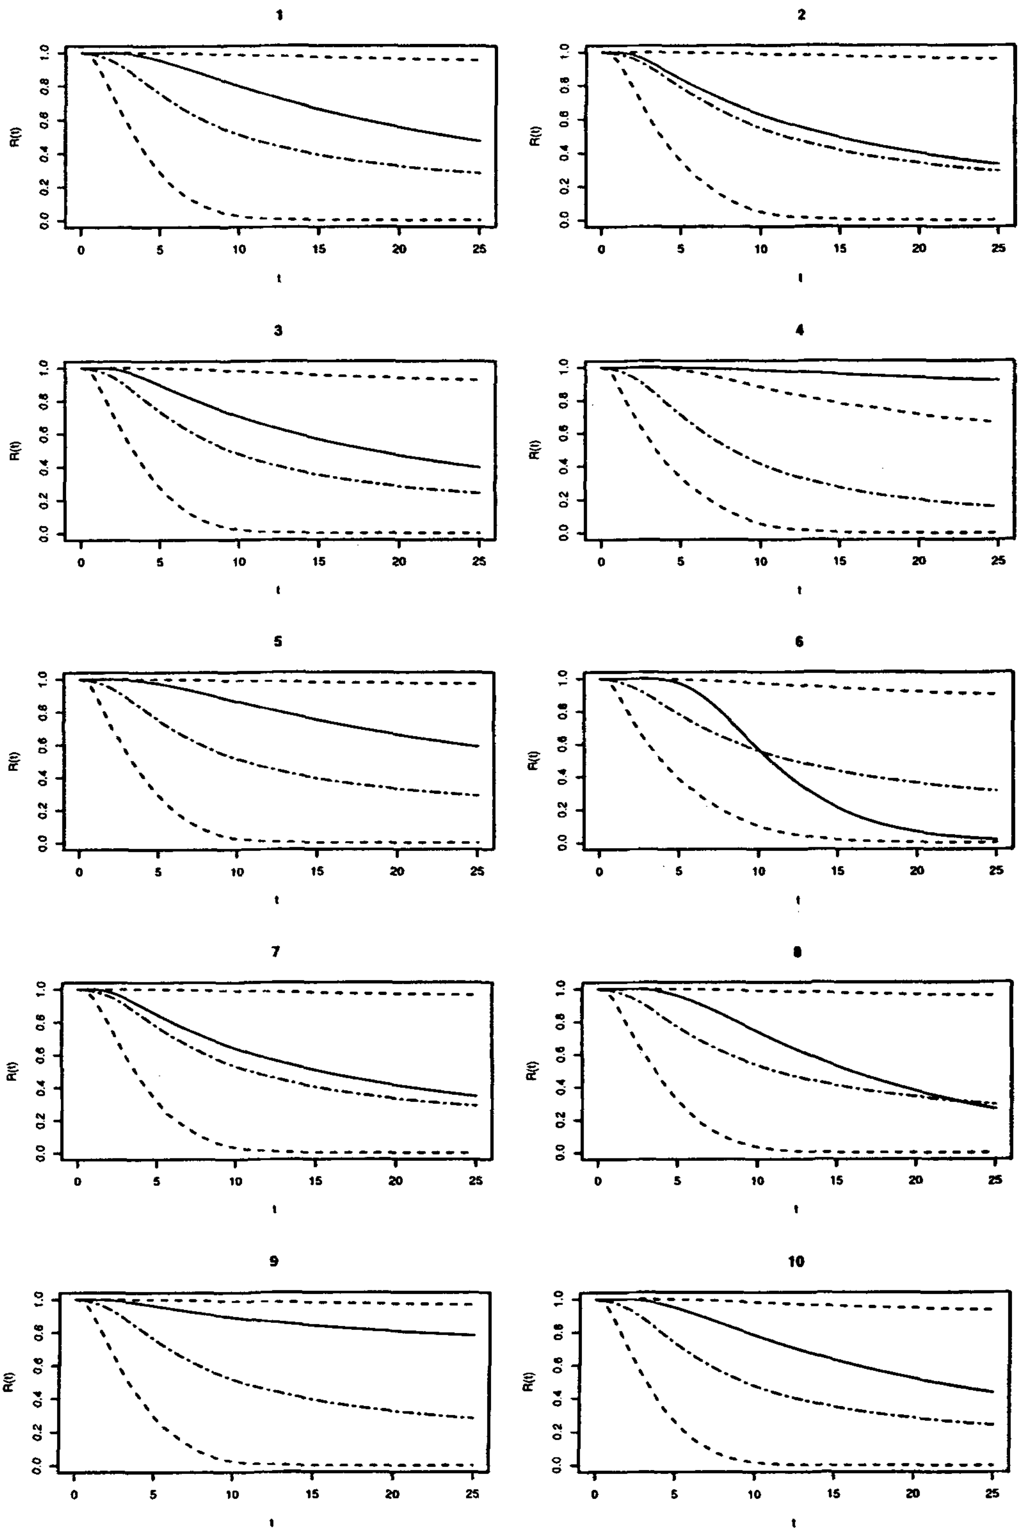

Fig. 5. Cross validated reliability prediction in the exchangeable model. The solid line is the "true" reliability. The middle dashed line is the predictive mean and the outer dashed lines represent a pointwise 90% prediction band.

We obtain the predictive mean (a point prediction) as well as a 90% pointwise prediction band for the reliability. Prediction clearly involves much more variation since here no learning from data is obtained at the first level (as $\pi(\underset{\sim}{\theta}_i \mid \gamma, Z^*_{-i}) = \pi(\underset{\sim}{\theta}_i \mid \gamma)$) and the only updating comes from the second level parameter γ. This is reflected in the wider prediction bands (compared to the credible bands in Fig. 4). We also note that the "true" reliability $R_4(t^*, \underset{\sim}{\theta}_i^{(0)})$ for the 4th system (which, as we noted above, has a positive θ_{1i} value) falls outside the prediction band. Thus the "inclusion" proportions for the 90% credible bands in Fig. 4 and the 90% prediction bands in Fig. 5 are both 9/10 which are what we expected.

References

Basu, A. P. (1985). Estimation of the reliability of complex system—a survey, *The Frontiers of Modern Statistical Inference Procedures* (ed. E. J. Dudewicz), 271–287, American Science Press, Columbus.

Basu, A. P. and Ebrahimi, N. (1983). On the reliability of stochastic systems, *Statist. Probab. Lett.*, **1**, 265–267.

Bernardo, J. M. and Smith, A. F. M. (1994). *Bayesian Theory*, Wiley, New York.

Birnbaum, Z. M. (1956). On a use of the Mann-Whitney statistic, *Proc. Third Berkeley Symp. on Math. Statist. Prob., Vol. 1, Contributions to the Theory of Statistics and Probability*, 13–17, University of California Press, Berkeley.

Datta, G. S. and Ghosh, M. (1995). Some remarks on noninformative priors, *J. Amer. Statist. Assoc.*, **90**, 1357–1363.

Ebrahimi, N. and Ramallingam, T. (1993). Estimation of system reliability in Brownian stress-strength models based on sample paths, *Ann. Inst. Statist. Math.*, **45**, 9–19.

Gelfand, A. E. (1996). Model determination using sampling-based methods, *Markov Chain Monte Carlo in Practice* (eds. W. R. Gilks, S. Richardson and D. J. Spiegelhalter), Chapman and Hall, London.

Ghosh, M. and Sun, D. (1998). Recent developments of Bayesian inference for stress-strength models, *Frontiers in Reliability* (eds. A. P. Basu, S. K. Basu and S. Mukhopadhyay), 143–158, World Scientific, Singapore.

Gilks, W. R. and Wild, P. (1992). Adaptive rejection sampling for Gibbs sampling, *Appl. Statist.*, **41**, 337–348.

Johnson, R. (1988). Stress-strength models for reliability, *Handbook of Statist.*, Vol. **7** (eds. P. R. Krishnaiah and C. R. Rao), 27–54, Elsevier, Amsterdam.

Kass, R. and Wasserman, L. (1996). Selection of prior distributions by formal rules, *J. Amer. Statist. Assoc.*, **91**, 1343–1370.

Sivaganesan, S. and Lingham, R. T. (2000). Bayes factors for a test about the drift of the Brownian motion under non-informative priors, *Statist. Probab. Lett.*, **48**, 163–171.

Sun, D. C. and Ye, K. Y. (1996). Frequentist validity of posterior quantiles for a two-parameter exponential family, *Biometrika*, **83**, 55–65.

Thompson, R. D. and Basu, A. P. (1993). Bayesian reliability of stress-strength systems, *Advances in Reliability*, (ed. A. P. Basu) 411–421, Elsevier, New York.

Weerahandi, S. and Johnson, R. A. (1992). Testing reliability in a stress-strength model when X and Y are normally distributed, *Technometrics*, **34**, 83–91.

Yang, R. and Berger, James O. (1998). A catalog of non-informative priors, Tech. Report, No. 97–42, Duke University, Durham, North Carolilna.