

Bayesian Nonparametric Estimation via Gibbs Sampling for Coherent Systems with Redundancy*

Hani Doss and Fred W. Huffer

Department of Statistics
Florida State University
Tallahassee, Florida 32306

Kevin L. Lawson

Department of Mathematics and Statistics
Air Force Institute of Technology
Wright-Patterson Air Force Base, Ohio 45433

July 1994

Abstract

We consider a coherent system S consisting of m independent components for which we do not know the distributions of the components' lifetimes. If we know the structure function of the system, then we can estimate the distribution of the system lifetime by estimating the distributions of the lifetimes of the individual components. Suppose that we can collect data under the 'autopsy model', wherein a system is run until a failure occurs and then the status (functioning or dead) of each component is obtained. This test is repeated n times. The autopsy statistics consist of the age of the system at the time of breakdown and the set of parts that are dead by the time of breakdown. We develop a nonparametric Bayesian estimate of the distributions of the component lifetimes and then use this to obtain an estimate of the distribution of the lifetime of the system. The procedure is applicable to machine-test settings wherein the machines have redundant designs. A parametric procedure is also given.

Key words and phrases: Gibbs sampling, autopsy model, incomplete data, nonparametric, Dirichlet prior.

*Research supported by Air Force Office of Scientific Research Grant F49620-94-1-0028.

1 Introduction and Summary

Consider a coherent system S consisting of m independent components for which we do not know the distributions of the component lifelengths. Assume that each of the m components occupies one of two states, functioning or failed. We consider the statistical model in which each element of a sample of n replicates of S is observed until it fails. The observed data consist of the set of components that are in a failed state and the failure time of the system. The failure times of the dead components are not directly observed. The set of dead components and the system failure time comprise the “autopsy statistics” of the system. This model is usually called the autopsy model.

Two statistical problems arise in considering the autopsy model—the problems of estimating the distributions of the component lifelengths and the distribution of the entire system’s lifelength. One approach to estimating the distribution of the system lifelength is to use only the observed system failure times. For example, we could use the empirical distribution function. However, such an approach ignores the (partial) information we have about the components of the system. If the structure function of the system is known, the distribution of the lifelength of the system can, in general, be calculated from knowledge of the distributions of the component lifelengths. Hence, an alternative approach would be to estimate the distributions of the lifelengths of the m components of system S and then use the structure function of S to estimate the distribution of the system lifelength.

Clearly, the component information provided by the autopsy model is quite limited. It is reasonable to consider alternative testing procedures which provide more component information. However, the autopsy model is important when alternative testing procedures such as separate testing of components are not possible or practical. For example, it may be difficult to reproduce the conditions which exist in the functioning system when components are tested separately. For such systems, it is important to obtain every bit of information about the components when they are parts of a machine or other system. In these settings, as well as the case of certain biological systems, the autopsy model is natural. Probabilistic aspects of the autopsy model were considered by Meilijson (1981), Nowick (1990), and Antoine, Doss, and Hollander (1993). Inferential aspects of the autopsy model have been considered only by Watelet (1990) and Meilijson (1994).

The U.S. Air Force’s C-17 transport airplane’s Fuel Quantity (FQ) computer is an example of a system in which the “component” is a logical subsystem whose status can be readily determined in the field. The FQ computer is a parallel system of order two, consisting of an “A” bus and a “B” bus. When this system’s data were first examined by the authors, there were approximately 2440 cumulative flying hours spread among six different prototype C-17’s, each of which contained one FQ computer. We present a data analysis using our procedure in Section 3.

Consider the autopsy statistics as described above. Assume that the system we wish to study is a coherent system (see Chapters 1 and 2 of Barlow and Proschan (1975)). We can use the structure function, the set of dead components, and the failure time of the system to say more about the failure times of each component. Specifically, for each of the n replicates of S , we can classify the failure time of each of the m components as follows:

- [C1] Component failure time is greater than system failure time.
- [C2] Component failure time is less than system failure time.
- [C3] Component failure time is equal to system failure time.
- [C4] Component failure time is either less than or exactly equal to system failure time, but we cannot tell which.

The first two categories correspond to what are usually termed right-censored data and left-censored data, respectively. Right-censored data occur when the component is still alive when the system fails. Left-censored data occur when the component is dead and information contained in the structure function, along with information contained in the set of dead components, allow one to deduce that the component's death occurred prior to the system failure time. Similarly, the third category arises when a component is observed in a failed state and we deduce that it caused the system to fail. The last category occurs whenever all the components in a redundant system or a redundant subsystem belong to the set of dead components. In this case, we know there is exactly one component, whose identity we do not know, with a failure time equal to that of the system, while the failure times of all the other components are strictly less than that of the system.

The last category above is problematic, particularly in frequentist settings, and is at the heart of the issue of “identifiability”. (General references on identifiability are given at the end of this section.) For example, consider a parallel system of order two. Let the distributions for components 1 and 2 be F_1 and F_2 , respectively. Both components will always be dead when the system is observed in a failed state. Under the autopsy model, one can only observe the system failure time, which has distribution F_1F_2 . This simple redundant system is not identifiable in the frequentist sense in that it is not possible to determine F_1 and F_2 from a knowledge of the distribution of the observed data.

We use a Bayesian framework because, for the applications we have in mind, we have a small amount of data but have extensive past experience on the components in other systems. For example, in the case of the C-17 FQ computer, similar but not identical FQ computers exist in other aircraft for which there has been more extensive testing. Thus, we have strong reasons to suspect that a certain parametric family provides a good approximation to the true probability model, and we have some knowledge about the unknown parameters. A Bayes procedure makes sense here since we have only limited testing hours on just six FQ systems, each within a C-17, and we wish to estimate the distribution of the lifelength of the FQ system when it is part of the C-17.

We consider a Bayesian framework for estimation of the distributions of component lifelengths in which the prior distributions on each of the F_i 's give most of their mass to “small neighborhoods of a parametric family”. The prior distributions which we use are derived from the Dirichlet process priors discussed by Ferguson (1973, 1974). The Dirichlet process priors are probability measures on \mathcal{P} parameterized by the set of all finite non-null measures on the real line \mathcal{R} , where \mathcal{P} is the space of all probability measures on \mathcal{R} . Let α be a finite non-null measure on the Borel sets of \mathcal{R} . The random distribution function F is said to have a Dirichlet process prior distribution with parameter α , denoted \mathcal{D}_α , if for every measurable partition $\{B_1, \dots, B_\ell\}$ of \mathcal{R} , the random vector $(F(B_1), \dots, F(B_\ell))$ has

the Dirichlet distribution with parameter vector $(\alpha(B_1), \dots, \alpha(B_\ell))$ (here and throughout the rest of the paper, probability measures are identified with their cumulative distribution functions, and the same symbol is used to denote both a measure and its distribution function whenever convenient). When a prior distribution is put on \mathcal{P} , then for every $t \in \mathcal{R}$, the quantity $F(t)$ is a random variable. Write $H = \alpha/\alpha(\mathcal{R})$, so that H is a probability measure on \mathcal{R} . If $F \sim \mathcal{D}_\alpha$, then $EF(t) = H(t)$, while the quantity $\alpha(\mathcal{R})$ indicates the degree of concentration of \mathcal{D}_α around its “center” H . For example, it is well known that as $\alpha(\mathcal{R}) \rightarrow \infty$, \mathcal{D}_α converges to the point mass at H in the weak topology. Ferguson (1973) showed that the Dirichlet priors have the property that the support of \mathcal{D}_α is the set of all probability measures whose support is contained in the support of H . For example, if the support of H is the positive real axis, then the support of \mathcal{D}_α is the set of distributions of all positive random variables. Ferguson also showed that if $F \sim \mathcal{D}_\alpha$, then F is a.s. discrete.

The priors on each of the F_i 's, $i = 1, \dots, m$, that we use are mixtures of Dirichlet priors. To keep the notation less cumbersome, let $m = 2$. For component 1, we consider a parametric family H_θ , $\theta \in \Theta$, and put a mixture of Dirichlets as the prior on F_1 . That is,

$$F_1 \sim \int \mathcal{D}_{\alpha_\theta} \nu(d\theta),$$

where for each $\theta \in \Theta$, $\alpha_\theta = \alpha_\theta(\mathcal{R})H_\theta$, $0 < \alpha_\theta(\mathcal{R}) < \infty$, and ν is a probability measure on Θ . Similarly, for component 2 we consider the parametric family K_ψ , $\psi \in \Psi$, and put as prior on F_2 the mixture $\int \mathcal{D}_{\beta_\psi} \mu(d\psi)$. For each $\psi \in \Psi$, $\beta_\psi = \beta_\psi(\mathcal{R})K_\psi$, $0 < \beta_\psi(\mathcal{R}) < \infty$, and μ is a probability measure on Ψ .

Our nonparametric Bayesian procedure, which uses mixtures of Dirichlets as priors on the F_i 's, has two advantages. First, we protect against the problems associated with using an incorrectly specified parametric model, such as obtaining an inconsistent estimator. Second, we can avoid the loss of efficiency due to ignoring partial information we may have about a parametric model, since we use prior distributions (on each of the F_i 's) that concentrate their mass around the hypothesized parametric family.

Our approach is based on the Gibbs sampling algorithm as discussed in Gelfand and Smith (1990). We now review the algorithm. Let f_{Y_1, \dots, Y_p} be the joint distribution of the (possibly vector-valued) random variables Y_1, \dots, Y_p . We suppose that we do not know the form of f_{Y_1, \dots, Y_p} , but that we know the conditional distributions $f_{Y_i|Y_j, j \neq i}$, $i = 1, \dots, p$ or that at least we are able to generate observations from these conditional distributions. Suppose we want to sample observations from the joint distribution of the random variables Y_1, \dots, Y_p , or simply an observation from one of the p marginals. The algorithm to generate an observation from f_{Y_1, \dots, Y_p} proceeds as follows. We fix arbitrary starting values $Y_1^{(0)}, \dots, Y_p^{(0)}$ and then update these values. Draw $Y_1^{(1)}$ from $f_{Y_1|Y_j, j \neq 1}(\cdot, Y_2^{(0)}, \dots, Y_p^{(0)})$. Next, draw $Y_2^{(1)}$ from $f_{Y_2|Y_j, j \neq 2}(Y_1^{(1)}, \cdot, Y_3^{(0)}, \dots, Y_p^{(0)})$. Continue until we draw $Y_p^{(1)}$ from $f_{Y_p|Y_j, j \neq p}(Y_1^{(1)}, \dots, Y_{p-1}^{(1)}, \cdot)$. We have now completed one iteration of the scheme by visiting each variable. After k iterations, we have the random variables $(Y_1^{(k)}, \dots, Y_p^{(k)})$. The sequence $(Y_1^{(j)}, \dots, Y_p^{(j)})$, $j = 1, 2, \dots$, is a Markov chain and f_{Y_1, \dots, Y_p} is a stationary distribution of the chain. If one can establish that the chain converges in distribution to f_{Y_1, \dots, Y_p} , then (for large k) $(Y_1^{(k)}, \dots, Y_p^{(k)})$ has a distribution which is approximately equal

to f_{Y_1, \dots, Y_p} . Such observations can be used to estimate f_{Y_1, \dots, Y_p} .

Perhaps the most natural way to implement the Gibbs sampler here is to proceed as is normally done in a Bayesian analysis of missing data problems under conjugacy. That is, consider the pair (parameter θ , missing data): In such a setup, if we knew the missing data, we would easily be able to find the conditional distribution of the parameter θ , and if we knew the parameter θ we would be able to generate the missing data (see for instance the linkage example or the Dirichlet sampling process example in Tanner and Wong (1987)). Indeed, this is precisely the approach taken by Doss (1994), who considers the use of Dirichlet priors for the problem of estimating an unknown distribution F in the presence of censoring. He considers random variables $X_1, \dots, X_n \stackrel{\text{iid}}{\sim} F$, but only observes $X_i \in A_i$, where A_i is a singleton if X_i is uncensored and $A_i = (c_i, \infty)$ if X_i is censored on the right by c_i . His approach is based on a Gibbs sampling algorithm of length 2 involving (\mathbf{X}, F) , where $\mathbf{X} = (X_1, \dots, X_n)$; that is, he generates \mathbf{X} from $\mathcal{L}_{\text{data}}(\mathbf{X} | F)$ and F from $\mathcal{L}_{\text{data}}(F | \mathbf{X})$, where **data** consists of the sets A_1, \dots, A_n . (Here, we use the notation $\mathcal{L}_{W_2}(W_1)$ or $\mathcal{L}(W_1 | W_2)$ to denote the conditional distribution of the random variable W_1 given the random variable W_2 .) The details of carrying out the steps of the algorithm rely on a constructive definition of the Dirichlet prior, given in Sethuraman (1994).

Our initial approach was simply to extend the technique of Doss (1994) to our setting; however, the approach fails since it turns out that the procedure produces a Markov chain which does not converge to the posterior distribution. We present an entirely different algorithm that produces a Markov chain which we show does converge to the posterior.

We now give a preliminary explanation of our procedure for the case where the priors on each of the F_j 's are single Dirichlets; i.e. $F_j \sim \mathcal{D}_{\alpha_j}$. Let S_i , $i = 1, \dots, n$, be the vector of lifelengths of the m components (for system i) if we could see them all. Let **data** be the set of autopsy statistics for the n systems. (It may be helpful to think of the case of parallel systems. In this case, **data** consists simply of the n system failure times.) The algorithm proceeds as follows. Fix arbitrary starting values $S_1^{(0)}, \dots, S_n^{(0)}$. Generate $S_1^{(1)} \sim \mathcal{L}_{\text{data}}(S_1 | S_2^{(0)}, \dots, S_n^{(0)})$. Next, generate $S_2^{(1)} \sim \mathcal{L}_{\text{data}}(S_2 | S_1^{(1)}, S_3^{(0)}, \dots, S_n^{(0)})$. Continue until we generate $S_n^{(1)} \sim \mathcal{L}_{\text{data}}(S_n | S_1^{(1)}, \dots, S_{n-1}^{(1)})$. We have now completed one iteration of the procedure. We repeat the procedure a large number of times and use the realizations of the chain to estimate $\mathcal{L}_{\text{data}}(S_1, \dots, S_n)$. There are two key points that allow this procedure to produce a Markov chain which converges to the posterior distribution.

- The conditional distribution of the lifelength of the j^{th} component of system i depends only on the autopsy statistics for system i and on the current set of lifelengths for the j^{th} components of $S_\ell, \ell \neq i$.
- The joint unconditional distribution of the n lifelengths of component j can be described in full.

The estimate of $\mathcal{L}_{\text{data}}(S_1, \dots, S_n)$ can be used to obtain an estimate of $\mathcal{L}_{\text{data}}(F_1, \dots, F_m)$.

Note that, in contrast to Doss (1994), we deal only with the random lifelengths (of the components) and bypass entirely the problem of generating the infinite-dimensional F_j 's.

In Section 2, we explore the algorithm further. In Section 2.1 we give a Gibbs sampling algorithm for our problem under parametric assumptions. In Section 2.2, we present

a simple extension of the procedure of Doss (1994) to our setting and show that this leads to a reducible Markov chain. In Section 2.3 we give a detailed description of the algorithm and a rigorous proof of its convergence for the case of parallel systems of order 2. In Section 2.4 we indicate the modifications necessary when dealing with an arbitrary coherent system. In Section 3 we illustrate our nonparametric procedure on data pertaining to the C-17 FQ computer.

The last portion of this section is used to summarize other relevant research in the statistical literature. First, we define the term identifiability. Suppose system S has m components and the component distributions are denoted F_1, \dots, F_m . If F_1, \dots, F_m can be recovered from a knowledge of the true distributions of the autopsy statistics, then we say that F_1, \dots, F_m are identifiable and that S is an identifiable system. For example, any parallel system is nonidentifiable. Meilijson (1981), Nowick (1990), and Antoine, Doss, and Hollander (1993) considered probabilistic aspects of the autopsy model. In these three papers, the authors identify conditions on the structure function of the system and on the distributions of the component lifelengths that guarantee identifiability.

Relevant work on estimating the distribution of the system lifelength under the autopsy model by first estimating the distributions of the component lifelengths and then using the structure function of the system can be found in Watelet (1990) and Meilijson (1994). Watelet (1990) considered two estimators for the autopsy model. He developed nonparametric estimators for the F_j 's. Meilijson (1994) estimated the parameters of these distributions from the empirical estimate of $\mathcal{L}(Z, D)$, where Z is the lifetime of the machine and D is the diagnostic set of parts that had died by time Z , by maximum likelihood from incomplete data via the EM algorithm. Meilijson assumes the F_j 's are drawn from well-behaved parametric families.

Under the assumption that the failure times of the dead components are known, Doss, Freitag, and Proschan (1989) also considered inferential aspects of estimating the distribution of the lifelength of the system by first estimating the distribution of the lifelengths of the system's components. In their model, they have more information than is available under the autopsy model.

2 Development and Convergence of the Algorithm

In order to keep the notation as light as possible Sections 2.1 and 2.3 deal only with parallel systems of order 2 (for which we have n replicates). It will be clear from the remarks in Section 2.4.1 that knowledge of how to deal with a parallel "module" makes it possible to handle an arbitrary coherent system.

2.1 A Gibbs Sampling Algorithm for the Autopsy Model under Parametric Assumptions

For the i^{th} replicate, $i = 1, \dots, n$, let the observed data (system failure time) be m_i , the maximum of the two component lifelengths, X_i and Y_i . Let $\mathbf{Z} = (\mathbf{X}, \mathbf{Y})$, where $\mathbf{X} = (X_1, \dots, X_n)$, $\mathbf{Y} = (Y_1, \dots, Y_n)$. Let $X_i \stackrel{\text{iid}}{\sim} F_\theta$ with $\theta \in \Theta$ and suppose $\theta \sim \nu$, where ν is a conjugate prior distribution on θ . Let $Y_i \sim G_\psi$ with $\psi \in \Psi$ and suppose $\psi \sim \mu$,

where μ is a conjugate prior distribution on ψ . Assume that F_θ and G_θ are absolutely continuous distribution functions. To make the discussion as easy as possible to follow, consider the case where $F_\theta = \mathcal{E}(\theta)$ (the exponential distribution with parameter θ) with $\theta \sim \nu = \mathcal{G}(a_1, b_1)$ (the Gamma distribution with shape parameter a_1 and scale parameter b_1) and $G_\psi = \mathcal{E}(\psi)$ with $\psi \sim \mu = \mathcal{G}(a_2, b_2)$.

Before describing the algorithm, we introduce the following notation. If H is a distribution function, B is a set, and X is distributed according to H , then H_B will denote the conditional distribution function of X given that $X \in B$; that is

$$H_B(A) = H(A \cap B)/H(B) \quad (2.1)$$

when $H(B) > 0$.

The algorithm proceeds as follows.

Give arbitrary initial values to $(X, Y)_i^{(0)}$ such that $X_i^{(0)} \vee Y_i^{(0)} = m_i$.

For $k = 1, \dots, K$:

1. Generate $(\theta, \psi)^{(k)} \sim \mathcal{L}_{\text{data}}((\theta, \psi) \mid (\mathbf{X}, \mathbf{Y})^{(k-1)})$.
2. Generate $(X, Y)_i^{(k)} \sim \mathcal{L}_{\text{data}}((X, Y)_i \mid (\theta, \psi)^{(k)})$, independently for each i , $i = 1, \dots, n$.

We now describe these two steps in more detail.

In step 1 of the algorithm, $\mathcal{L}_{\text{data}}((\theta, \psi) \mid \mathbf{Z}^{(k-1)})$ is the product of two gamma distributions, $\mathcal{G}(a_1 + n, b_1 + \sum_{i=1}^n X_i^{(k-1)})$ and $\mathcal{G}(a_2 + n, b_2 + \sum_{i=1}^n Y_i^{(k-1)})$; i.e. θ and ψ are generated independently. (Note that knowledge of \mathbf{X} and \mathbf{Y} make knowledge of the observed data superfluous.) To carry out step 2, first let f_θ and g_ψ be the densities of F_θ and G_ψ , respectively. Then, for each i , $i = 1, \dots, n$, set

$$p_i^{(k)} = \frac{f_{\theta^{(k)}}(m_i)G_{\psi^{(k)}}(m_i)}{f_{\theta^{(k)}}(m_i)G_{\psi^{(k)}}(m_i) + g_{\psi^{(k)}}(m_i)F_{\theta^{(k)}}(m_i)},$$

where $F_\lambda = G_\lambda = \mathcal{E}(\lambda)$. Now, generate $\mathbf{Z}^{(k)}$ by setting $(X, Y)_i^{(k)}$, $i = 1, \dots, n$, as follows:

$$(X, Y)_i^{(k)} = \begin{cases} (m_i, V), & \text{where } V \sim G_{\psi^{(k)}, [0, m_i]} & \text{with probability } p_i^{(k)}, \\ (V, m_i), & \text{where } V \sim F_{\theta^{(k)}, [0, m_i]} & \text{with probability } (1 - p_i^{(k)}) \end{cases}$$

where the notation used for the conditional distribution functions is given by (2.1).

To estimate $\mathcal{L}_{\text{data}}((\theta, \psi))$, we use the sequence of generated $\mathbf{Z}^{(k)}$'s to approximate the mixture

$$\int \mathcal{L}_{\text{data}}((\theta, \psi) \mid \mathbf{Z}) d\mathcal{L}_{\text{data}}(\mathbf{Z})$$

by, say, $\frac{1}{K} \sum_{j=1}^K \mathcal{L}_{\text{data}}((\theta, \psi) \mid \mathbf{Z}^{(j)})$.

If ν and μ are not conjugate priors for F_θ and G_ψ , respectively, but they are continuous distribution functions with univariate log-concave density functions, then we can still apply the above algorithm by using an efficient rejection scheme of Gilks and Wild (1992).

2.2 A Naïve Extension of the Algorithm of Doss (1994)

Doss (1994) uses a Gibbs sampling algorithm of length 2 involving (\mathbf{X}, F) . The prior he put on F is actually a mixture of Dirichlets, but we shall consider a single Dirichlet prior for our naïve extension, since if the procedure fails for a single Dirichlet, it will fail *a fortiori* for a mixture.

Suppose we have n replicates of a parallel system of order 2. Let \mathbf{Z} , \mathbf{X} , and \mathbf{Y} be defined as in Section 2.1. Suppose $X_i \stackrel{\text{iid}}{\sim} F$ and $Y_i \stackrel{\text{iid}}{\sim} G$, $i = 1, \dots, n$. Suppose $F \sim \mathcal{D}_\alpha$ and $G \sim \mathcal{D}_\beta$, where \mathcal{D}_α and \mathcal{D}_β are Dirichlet priors with parameter measures α and β . We observe $\mathbf{data} = (m_1, \dots, m_n)$, where $m_i = X_i \vee Y_i$, $i = 1, \dots, n$. Our goal is to estimate $\mathcal{L}_{\mathbf{data}}(F, G)$.

We know that the conditional distribution of (F, G) given (\mathbf{X}, \mathbf{Y}) is

$$\mathcal{L}_{\mathbf{data}}((F, G) \mid \mathbf{Z}) = \mathcal{D}_{\alpha + \sum_{i=1}^n \delta_{X_i}} \otimes \mathcal{D}_{\beta + \sum_{i=1}^n \delta_{Y_i}}, \quad (2.2)$$

where \otimes denotes product measure; i.e. F and G are independent. Also, given an updated (F, G) , we can generate a random \mathbf{Z} conditional on the data. To run the Gibbs sampling algorithm, we first set initial values $\mathbf{Z}^{(0)}$ and $(F, G)^{(0)}$. Then, for some large K , execute the following loop for $k = 1, \dots, K$:

1. Draw $(F, G)^{(k)}$ from $\mathcal{L}_{\mathbf{data}}((F, G) \mid \mathbf{Z}^{(k-1)})$.
2. Draw $\mathbf{Z}^{(k)}$ from $\mathcal{L}_{\mathbf{data}}(\mathbf{Z} \mid (F, G)^{(k)})$.

To carry out the second step in the above loop, we need to compute the probability that each component’s lifelength takes on the observed maximum, since one of the values must. Consider the case where $n = 1$. Let the observed maximum be denoted by m . We assume that $m > 0$ and that α and β are continuous measures. Suppose the initial value of $\mathbf{Z}^{(0)}$ is $(X_1, Y_1)^{(0)} = (m, V)$, for some $V < m$. These initial values give rise to $(F, G)^{(1)}$ via the first step of the algorithm. By (2.2), $F^{(1)} \sim \mathcal{D}_{\alpha + \delta_m}$, but by definition of the Dirichlet prior, this implies that $F^{(1)}(\{m\})$ is distributed as a Beta distribution with parameters $(\alpha + \delta_m)(\{m\})$ and $(\alpha + \delta_m)(\{m\}^c)$, and therefore, $F^{(1)}(\{m\}) > 0$ with probability 1. Next, recall from the previous section that if $F \sim \mathcal{D}_\gamma$, then $EF(A) = \gamma(A)/\gamma(\mathcal{R})$ for any Borel set A . Since $G^{(1)} \sim \mathcal{D}_{\beta + \delta_V}$, we have $EG^{(1)}(\{m\}) = \frac{\beta + \delta_V}{\beta(\mathcal{R}) + 1}(\{m\}) = 0$. Thus, $G^{(1)}(\{m\}) = 0$ with probability 1. In other words, with probability 1, $F^{(1)}$ has an atom at m and $G^{(1)}$ does not, from which it is clear that $P\{X^{(1)} = m \mid X^{(1)} \vee Y^{(1)} = m\} = 1$. As we continue to run the algorithm, $X^{(2)}, X^{(3)}, \dots$ will each take on the value m . The case where $(X_1, Y_1)^{(0)} = (V, m)$ is handled by symmetry. Thus, we see that the starting point does not get “washed out” and therefore, the algorithm produces a Markov chain that cannot converge to the posterior distribution.

2.3 The Algorithm We Propose: Detailed Description and Convergence for Parallel Systems of Order 2

Before introducing the setup and notation needed for the algorithm, we briefly describe the “extended Pólya urn scheme” as given in Blackwell and MacQueen (1973). Define

a sequence $\{T_1, T_2, \dots\}$ of random variables as a *Pólya sequence with parameter α* if for every $B \subset \mathcal{R}$, we have $P(T_1 \in B) = \alpha(B)/\alpha(\mathcal{R})$, and for every n ,

$$P(T_{n+1} \in B \mid T_1, \dots, T_n) = \left(\alpha(B) + \sum_{i=1}^n \delta_{T_i}(B) \right) / \left(\alpha(\mathcal{R}) + n \right). \quad (2.3)$$

Blackwell and MacQueen proved that if $\{T_1, T_2, \dots\}$ is a Pólya sequence with parameter α , then the empirical distribution of $\{T_1, \dots, T_n\}$ converges a.s. to a limiting discrete measure H . Furthermore, $H \sim \mathcal{D}_\alpha$. Also, given H , the random variables T_1, T_2, \dots are iid $\sim H$. In addition, for every n ,

$$T_1, \dots, T_n \text{ are exchangeable.} \quad (2.4)$$

Recall that X_i and Y_i are the lifelengths of components 1 and 2 in system i . For the X_i 's, we consider a parametric family H_θ , $\theta \in \Theta \subset \mathcal{R}^{d_1}$, and put a mixture of Dirichlets as the prior on F . That is,

$$F \sim \int \mathcal{D}_{\alpha_\theta} \nu(d\theta), \quad (2.5)$$

where for each $\theta \in \Theta$, $\alpha_\theta = \alpha_\theta(\mathcal{R})H_\theta$, $0 < \alpha_\theta(\mathcal{R}) < \infty$, and ν is a probability measure on Θ . Similarly, for the Y_i 's, we consider the parametric family K_ψ , $\psi \in \Psi \subset \mathcal{R}^{d_2}$, and put the mixture $\int \mathcal{D}_{\beta_\psi} \mu(d\psi)$ as the prior on G , where for each $\psi \in \Psi$, $\beta_\psi = \beta_\psi(\mathcal{R})K_\psi$, $0 < \beta_\psi(\mathcal{R}) < \infty$, and μ is a probability measure on Ψ . We will assume that for each θ and ψ , H_θ and K_ψ are absolutely continuous, with continuous densities. Given F and G , X_1, X_2, \dots are iid $\sim F$, and Y_1, Y_2, \dots are iid $\sim G$. Also, for every n , X_1, \dots, X_n are exchangeable, as are Y_1, \dots, Y_n , by (2.4). We assume that F is independent of G . It follows that (X_1, \dots, X_n) is independent of (Y_1, \dots, Y_n) . We observe

$$\mathbf{data} = (m_1, \dots, m_n), \quad \text{where } X_i \vee Y_i = m_i, \quad i = 1, \dots, n.$$

Our goal is to estimate $\mathcal{L}_{\mathbf{data}}(F, G)$. As in Section 2.2, knowledge of (\mathbf{X}, \mathbf{Y}) makes knowledge of \mathbf{data} superfluous.

Recall that S_i , $i = 1, \dots, n$, is the vector of lifelengths for system i . For the case of a parallel system of order 2, $S_i = (X_i, Y_i)$. To unify the notation, let $S_0 = (\theta, \psi)$ and let $\mathbf{S} = (S_0, \dots, S_n)$.

The algorithm proceeds as follows. Fix arbitrary starting values $S_0^{(0)}, \dots, S_n^{(0)}$. Then, cycle through the $N = n + 1$ elements of \mathbf{S} in order; i.e. at time t , we update element $K = K(t) = (t - 1) \bmod N$ by generating

$$S_K^{(t)} \sim \mathcal{L}_{\mathbf{data}}(S_K \mid S_j^{(t-1)} \quad \text{for } j \neq K). \quad (2.6)$$

At time jN , each component of \mathbf{S} has been updated j times. Our algorithm generates vectors $\mathbf{S}^{(t)}$, $t = 1, 2, \dots$, where $\mathbf{S}^{(t)}$ is the same as $\mathbf{S}^{(t-1)}$ except for the one element $S_K^{(t)}$ which has been updated at time t according to (2.6). Note that this notation differs slightly from that of Section 1, where $\mathbf{S}^{(t)}$ was formed by updating *all* the elements of $\mathbf{S}^{(t-1)}$.

In our algorithm, the updating of $S_i^{(t)}$, $i = 1, \dots, n$ is accomplished by the use of the following lemma, the proof of which is a calculation.

Lemma 2.1 Suppose A and B are distribution functions on $[0, \infty)$ satisfying $A = A_c + A_d$ and $B = B_c + B_d$, where A_c and B_c are absolutely continuous with continuous derivatives denoted A'_c and B'_c , and A_d and B_d are discrete. Let X^* and Y^* be independent with distributions A and B , respectively. For $m > 0$, we can generate a pair (X, Y) with

$$(X, Y) \sim \mathcal{L}\left((X^*, Y^*) \mid X^* \vee Y^* = m\right) \quad (2.7)$$

by the mixing procedure we now describe.

Define probability distributions on $[0, m)$ by

$$\bar{A}_c(v) = \frac{A_c(v)}{A_c(m)}, \quad \bar{B}_c(v) = \frac{B_c(v)}{B_c(m)}, \quad \bar{A}_d(v) = \frac{A_d(v)}{A_d(m-)}, \quad \bar{B}_d(v) = \frac{B_d(v)}{B_d(m-)}.$$

If the denominator in any of the equations above is 0, then we set the corresponding probability distribution to δ_0 . Let V_1, V_2, V_3, V_4 be random variables with distributions $\bar{A}_c, \bar{B}_c, \bar{A}_d, \bar{B}_d$, respectively. We now take

$$(X, Y) = \begin{cases} (V_1, m) & \text{with probability } p_1, \\ (m, V_2) & \text{with probability } p_2, \\ (V_3, m) & \text{with probability } p_3, \\ (m, V_4) & \text{with probability } p_4, \\ (m, m) & \text{with probability } p_5, \end{cases} \quad (2.8)$$

where the mixing probabilities p_1, \dots, p_5 are given by the following formulas.

If $A_d(\{m\}) + B_d(\{m\}) > 0$, then

$$p_1 = \frac{A_c(m)B_d(\{m\})}{D}, \quad p_2 = \frac{A_d(\{m\})B_c(m)}{D}, \quad p_3 = \frac{A_d(m-)B_d(\{m\})}{D},$$

$$p_4 = \frac{A_d(\{m\})B_d(m-)}{D}, \quad \text{and} \quad p_5 = \frac{A_d(\{m\})B_d(\{m\})}{D},$$

with D such that $\sum_{i=1}^5 p_i = 1$. If $A_d(\{m\}) = B_d(\{m\}) = 0$, then

$$p_1 = \frac{A_c(m)B'_c(m)}{D}, \quad p_2 = \frac{A'_c(m)B_c(m)}{D}, \quad p_3 = \frac{A_d(m)B'_c(m)}{D},$$

$$p_4 = \frac{A'_c(m)B_d(m)}{D}, \quad \text{and} \quad p_5 = 0,$$

with D such that $\sum_{i=1}^5 p_i = 1$.

By (2.3) and (2.4), for the t^{th} step of the algorithm, the conditional distribution of X_i , given the other $(n-1)$ X_ℓ 's and θ , is given by

$$A_i^{(t)} = \frac{\alpha_{\theta^{(t)}} + \sum_{\ell \neq i} \delta_{X_\ell^{(t)}}}{\alpha_{\theta^{(t)}}(\mathcal{R}) + n - 1}. \quad (2.9)$$

Similarly, for the t^{th} step of the algorithm, the conditional distribution of Y_i , given the other $(n-1)$ Y_ℓ 's and ψ , is

$$B_i^{(t)} = \frac{\beta_{\psi^{(t)}} + \sum_{\ell \neq i} \delta_{Y_\ell^{(t)}}}{\beta_{\psi^{(t)}}(\mathcal{R}) + n - 1}. \quad (2.10)$$

Thus we can generate $S_i^{(t)} = (X_i^{(t)}, Y_i^{(t)})$ by using Lemma 2.1 with $A = A_i^{(t)}$, $B = B_i^{(t)}$ and $m = m_i$.

To update $S_0^{(t)}$ in the above algorithm, we will need formulas for updating the ‘‘mixing measures’’ ν and μ . The formula for the conditional distribution of θ given \mathbf{X} is well known. Proposition 2.1 below is a special case of Lemma 1 of Antoniak (1974). The formula for the conditional distribution of ψ given \mathbf{Y} will be evident by symmetry. The notation $\#(\mathbf{v})$ is used to denote the number of distinct values in the vector \mathbf{v} .

Proposition 2.1 *Assume that for each $\theta \in \Theta$, H_θ is absolutely continuous, with a density h_θ that is continuous on \mathcal{R} . If the prior of F is given by (2.5), then the (marginal) posterior distribution of θ given $\mathbf{X} = \mathbf{x}$ is*

$$\nu_{\mathbf{x}}(d\theta) = c(\mathbf{x}) \left(\prod^{\text{dist}} h_\theta(x_i) \right) \frac{(\alpha_\theta(\mathcal{R}))^{\#(\mathbf{x})} \Gamma(\alpha_\theta(\mathcal{R}))}{\Gamma(\alpha_\theta(\mathcal{R}) + n)} \nu(d\theta), \quad (2.11)$$

where the ‘dist’ in the product indicates that the product is taken over distinct values of x_i only, Γ is the gamma function, and $c(\mathbf{x})$ is a normalizing constant.

From (2.11) and the independence of F and G , we can update $S_0^{(t)}$ by independently generating $\theta^{(t)}$ and $\psi^{(t)}$ from $\nu_{\mathbf{x}^{(t)}}$ and $\mu_{\mathbf{y}^{(t)}}$, respectively.

Let π be the posterior distribution of \mathbf{S} given **data** on the space $(\mathcal{R}^{2n} \times \Theta \times \Psi, \mathcal{B})$, where \mathcal{B} is the collection of Borel sets on $\mathcal{R}^{2n} \times \Theta \times \Psi$. This π is also a stationary distribution for the chain $\{\mathbf{S}^{(jN)}\}_{j=0}^\infty$. We now show that the distribution of the Markov chain $\{\mathbf{S}^{(jN)}\}_j$ converges to π at a rate that is geometric and independent of the starting point. Before stating our theorem, we introduce some notation. Define

$$D = \{ \mathbf{s} : x_i, y_i \geq 0 \text{ and } x_i \vee y_i = m_i \text{ for } i = 1, \dots, n, \theta \in \Theta, \psi \in \Psi \},$$

and let $P^j(\mathbf{s}, C)$ be the j -step transition probabilities for the chain; i.e.

$$P^j(\mathbf{s}, C) = P(\mathbf{S}^{(jN)} \in C \mid \mathbf{S}^{(0)} = \mathbf{s}).$$

Note that if ν is absolutely continuous with respect to Lebesgue measure, then so is $\nu_{\mathbf{x}}$. This implies $\nu_{\mathbf{x}}$ has a density with respect to Lebesgue measure, which we shall denote $\nu'_{\mathbf{x}}$. In the proof of Theorem 1, we shall need to find bounds on $\nu'_{\mathbf{x}}(\theta)$ for fixed θ as \mathbf{x} varies over the compact set $[0, m_{(n)}]^n$. Here, $m_{(n)} = \max(m_1, \dots, m_n)$. This would be straightforward if $\nu'_{\mathbf{x}}(\theta)$ were continuous in \mathbf{x} , but inspection of (2.11) clearly shows this is not the case. For this reason, we introduce functions $f_i(\mathbf{x}, \theta)$ on $\mathcal{R}^i \times \Theta$, defined by

$$f_i(\mathbf{x}, \theta) = c_i(\mathbf{x}) \left(\prod_{j=1}^i h_\theta(x_j) \right) \frac{(\alpha_\theta(\mathcal{R}))^i \Gamma(\alpha_\theta(\mathcal{R}))}{\Gamma(\alpha_\theta(\mathcal{R}) + n)} \nu'(\theta), \quad \text{for } i = 1, \dots, n,$$

where the values $c_i(\mathbf{x})$ are constants such that, for each \mathbf{x} , $f_i(\mathbf{x}, \theta)$ is a density in θ . Note that

$$\nu'_{\mathbf{x}}(\theta) = f_{\#(\mathbf{x})}(\mathbf{x}^{\text{dist}}, \theta),$$

where \mathbf{x}^{dist} is comprised of the distinct values of \mathbf{x} arranged in any order.

Theorem 1 Consider the Markov chain $\{\mathbf{S}^{(jN)}\}_j$ resulting from applying the above algorithm to a parallel system of order 2. Suppose that

1. The observed maximum values m_1, \dots, m_n are distinct and positive.
2. There exists a compact set $\Gamma_1 \subset \Theta$ with $\nu(\Gamma_1) > 0$ such that $\alpha'_\theta(x)$ exists, is positive, and is continuous in both θ and x for $(\theta, x) \in \Gamma_1 \times [0, m_{(n)}]$; similarly there exists a compact set $\Gamma_2 \subset \Psi$ corresponding to β'_ψ .
3. The prior ν is absolutely continuous with respect to Lebesgue measure and for each $i, i = 1, \dots, n$, $f_i(\cdot, \cdot)$ is positive and continuous on $[0, m_{(n)}]^i \times \Gamma_1$; an analogous condition holds for the prior and posterior of ψ .

Then, there exists a value $\lambda > 0$ such that

$$\sup_{C \in \mathcal{B}, \mathbf{s} \in D} |P^j(\mathbf{s}, C) - \pi(C)| < e^{-\lambda j} \quad \text{for all } j \geq 2. \quad (2.12)$$

Recall that if $\{Z_k\}_{k=0}^\infty$ is a Markov chain on $(\mathcal{Z}, \mathcal{F})$ with n -step transition probabilities $P^n(z, C)$ which satisfy the *Doebelin condition*, i.e. for some probability measure ρ on $(\mathcal{Z}, \mathcal{F})$, some positive integer n_0 , and some $\epsilon > 0$,

$$P^{n_0}(z, C) \geq \epsilon \rho(C) \quad \text{for all } z \in \mathcal{Z} \text{ and all } C \in \mathcal{F}, \quad (2.13)$$

then

$$\sup_{C \in \mathcal{F}} |P^n(z, C) - \pi(C)| \leq (1 - \epsilon)^{\lfloor n/n_0 \rfloor} \quad \text{for all } z \in \mathcal{Z}. \quad (2.14)$$

The proof of this fact involves a coupling argument which we review below. We may write

$$P^{n_0}(z, \cdot) = (1 - \epsilon)\eta(z, \cdot) + \epsilon\rho(\cdot) \quad \text{for all } z \in \mathcal{Z}, \quad (2.15)$$

where $\eta(z, \cdot) = (P^{n_0}(z, \cdot) - \epsilon\rho(\cdot))/(1 - \epsilon)$. By the Doebelin condition (2.13), $\eta(z, \cdot)$ is a probability measure. Thus, $P^{n_0}(z, \cdot)$ is a convex combination of two measures, *the second of which does not involve z* . Therefore, the representation in (2.15) allows us to view each step in the evolution of the Markov chain $\{Z_{\ell n_0}\}_\ell$ as a coin-tossing experiment followed by a draw from either $\eta(z, \cdot)$ or $\rho(\cdot)$. This is the key that makes possible the coupling argument.

Now run two chains $\{W_\ell\}_\ell$ and $\{Y_\ell\}_\ell$ as follows. Let $W_0 = w_0$ and $Y_0 \sim \pi$. If the current state of the two chains is (w_{n-1}, y_{n-1}) , generate W_n and Y_n by first tossing a coin with probability of heads equal to ϵ . If the toss results in a head, select $V \sim \rho$, and set *both* W_n and Y_n equal to V . If the toss results in tails, independently select W_n from

$\eta(w_{n-1}, \cdot)$ and Y_n from $\eta(y_{n-1}, \cdot)$. It is clear that $\{W_\ell\}_\ell$ and $\{Y_\ell\}_\ell$ are Markov chains with transition probabilities $P^{n_0\ell}(z, C)$ and that $Y_\ell \sim \pi$ for all ℓ . Moreover,

$$P\{W_\ell \neq Y_\ell \text{ for some } \ell \geq k\} \leq (1 - \epsilon)^k.$$

This argument shows that the Markov chain $\{Z_{\ell n_0}\}_\ell$ satisfies

$$\sup_{C \in \mathcal{F}} |P\{Z_{n_0 k} \in C\} - \pi(C)| \leq (1 - \epsilon)^k,$$

from which (2.14) follows, since $\sup_{C \in \mathcal{F}} |P^n(z, C) - \pi(C)|$ is nonincreasing in n .

Results of the form (2.12) are generally proved by verifying the Doeblin condition (2.13). However, we shall prove our theorem by dealing with the underlying coupling argument directly and explicitly because our arguments are then easier to follow.

Proof of Theorem 1 Consider starting the Gibbs sampler from two different initial states $\mathbf{S}^{(0)}$ and $\tilde{\mathbf{S}}^{(0)}$ and producing two sequences $\{\mathbf{S}^{(t)}\}$ and $\{\tilde{\mathbf{S}}^{(t)}\}$. We shall “couple” these sequences by defining them on the same probability space in such a way that

$$P\{\mathbf{S}^{(2N)} = \tilde{\mathbf{S}}^{(2N)}\} > \epsilon, \tag{2.16}$$

where ϵ is positive and can be chosen independently of the starting states $\mathbf{S}^{(0)}$ and $\tilde{\mathbf{S}}^{(0)}$. This clearly suffices to prove the theorem.

Condition (2.16) indicates that our proof requires two passes of the algorithm (or equivalently, two “cycles” of the Gibbs sampler) to couple the two sequences. In the first pass, we show that there is a positive lower bound not depending on the starting states, for the probability that for each $i, i = 1, \dots, n$, the minima $X_i \wedge Y_i$ and $\tilde{X}_i \wedge \tilde{Y}_i$ are not equal to any of the observed maxima m_1, \dots, m_n . Then in the second pass, we show that there is also a positive lower bound for the probability that for each $i, i = 1, \dots, n$, $X_i = \tilde{X}_i$ and $Y_i = \tilde{Y}_i$; i.e. $S_i = \tilde{S}_i$. The probability that $S_0 = \tilde{S}_0$ is handled in a manner similar to the proof of (2.14).

A simple example may help to identify the issues involved. Suppose that $n = 2$ and $m_1 < m_2$ with starting states $\mathbf{S}^{(0)} = (S_0^{(0)}, S_1^{(0)}, S_2^{(0)}) = ((\theta, \psi), (\ell_1, m_1), (m_1, m_2))$ and $\tilde{\mathbf{S}}^{(0)} = ((\tilde{\theta}, \tilde{\psi}), (m_1, \ell_2), (m_2, m_1))$, where $\ell_1, \ell_2 < m_1$. At time $t = 1$, we update θ, ψ and $\tilde{\theta}, \tilde{\psi}$, but this is not important. At time $t = 2$, we update (X_1, Y_1) and $(\tilde{X}_1, \tilde{Y}_1)$, and this is where the problem arises. For the sequence $\{\mathbf{S}^{(t)}\}$, $A_1^{(2)}$ has an atom at m_1 , but $B_1^{(2)}$ does not, so $P(X_1^{(2)} = m_1, Y_1^{(2)} < m_1) = 1$. For the sequence $\{\tilde{\mathbf{S}}^{(t)}\}$, $\tilde{B}_1^{(2)}$ has an atom at m_1 , but $\tilde{A}_1^{(2)}$ does not, so $P(\tilde{X}_1^{(2)} < m_1, \tilde{Y}_1^{(2)} = m_1) = 1$. Thus, with probability 1, the two sequences are *not* coupled during the first pass of the algorithm. In the proof, we will show that after one pass of the algorithm, the atoms at m_1 and m_2 are “out of the way” and thus, there is a positive probability (that is independent of the starting states) that $\mathbf{S}^{(6)} = \tilde{\mathbf{S}}^{(6)}$, making (2.16) true.

The following will involve a specific implementation of the algorithm described by (2.6) through (2.11), strictly for use in our coupling argument. Let $\{U_{jk}; j = 1, \dots, \infty; k = 1, 2, 3, 4\}$ be an array of independent $\mathcal{U}(0, 1)$ (the uniform distribution on $(0, 1)$) random variables. We shall generate $\theta^{(t)}, \psi^{(t)}, X_i^{(t)}$, and $Y_i^{(t)}$ in terms of these uniform random

variables using certain functions ρ_1, ρ_2 , and ρ_3 . These functions and their properties are described in detail later in this section. Specifically, we define the functions and show that using them produces the desired conditional distribution of $\mathbf{S}^{(t)}$ given $\mathbf{S}^{(t-1)}$.

For t with $(t-1) \bmod N = 0$, define $S_0^{(t)} = (\theta^{(t)}, \psi^{(t)})$ by

$$\theta^{(t)} = \rho_1(\mathbf{X}^{(t-1)}, U_{t1}, U_{t2}) \quad \text{and} \quad \psi^{(t)} = \rho_2(\mathbf{Y}^{(t-1)}, U_{t3}, U_{t4}), \quad (2.17)$$

and for t with $(t-1) \bmod N = K \neq 0$, define $S_K^{(t)} = (X_K^{(t)}, Y_K^{(t)})$ by

$$(X_K^{(t)}, Y_K^{(t)}) = \rho_3(\theta^{(t-1)}, \psi^{(t-1)}, m_K, \mathbf{X}_{(-K)}^{(t-1)}, \mathbf{Y}_{(-K)}^{(t-1)}, U_{t1}, U_{t2}, U_{t3}), \quad (2.18)$$

where $\mathbf{W}_{(-K)}$ is used to denote the $(n-1)$ -tuple obtained by deleting the K^{th} element of the n -tuple \mathbf{W} .

For the process $\{\tilde{\mathbf{S}}^{(t)}\}$, the quantities $\tilde{\theta}^{(t)}, \tilde{\psi}^{(t)}, \tilde{X}^{(t)}$, and $\tilde{Y}^{(t)}$ are defined in the same way except that θ, ψ, \mathbf{X} , and \mathbf{Y} are everywhere replaced by $\tilde{\theta}, \tilde{\psi}, \tilde{\mathbf{X}}$, and $\tilde{\mathbf{Y}}$. The *same* uniform variates are used in generating both $\mathbf{S}^{(t)}$ and $\tilde{\mathbf{S}}^{(t)}$.

The essence of the proof consists of showing that there is a positive probability (which is independent of the starting states) of coupling the two processes if we implement the algorithm using ρ_1, ρ_2 , and ρ_3 (i.e. (2.16) holds). To this end, we introduce a sequence of events A_1, A_2, \dots, A_{2N} defined as follows. Let $\mathcal{M} = \{m_1, m_2, \dots, m_n\}$. For $t = 1$ and $t = N + 1$, we define

$$A_t = \{\theta^{(t)} = \tilde{\theta}^{(t)}, \psi^{(t)} = \tilde{\psi}^{(t)}, \theta^{(t)} \in \Gamma_1, \psi^{(t)} \in \Gamma_2\}.$$

(The sets Γ_1 and Γ_2 are the compact sets given by assumption 2 of Theorem 1.) For $2 \leq t \leq N$ and $K = (t-1) \bmod N$ we define

$$A_t = \{X_K^{(t)} \wedge Y_K^{(t)} \notin \mathcal{M}, \tilde{X}_K^{(t)} \wedge \tilde{Y}_K^{(t)} \notin \mathcal{M}\}.$$

For $N + 2 \leq t \leq 2N$ and $K = (t-1) \bmod N$ we define

$$A_t = \{X_K^{(t)} = \tilde{X}_K^{(t)}, Y_K^{(t)} = \tilde{Y}_K^{(t)}, X_K^{(t)} \wedge Y_K^{(t)} \notin \mathcal{M}\}.$$

Note that $\bigcap_{t=1}^{2N} A_t$ implies that $\mathbf{S}^{(2N)} = \tilde{\mathbf{S}}^{(2N)}$. We can write

$$P\left(\bigcap_{t=1}^{2N} A_t\right) = \prod_{t=1}^{2N} \zeta_t,$$

where $\zeta_1 = P(A_1)$ and $\zeta_t = P(A_t \mid \bigcap_{j < t} A_j)$ for $t > 1$. We shall show that each of the values ζ_t are bounded away from 0 by quantities which do not depend on $\mathbf{S}^{(0)}$ and $\tilde{\mathbf{S}}^{(0)}$. In the last part of this section, we establish these bounds in Facts 1–6. Specifically, Fact 1 is used to bound ζ_1 and ζ_{N+1} ; and Facts 2 and 3 are used to bound ζ_2, \dots, ζ_N . Now consider $N + 2 \leq t \leq 2N$ and $K = (t-1) \bmod N$. Conditional on $\bigcap_{j < t} A_j$, the vectors $\mathbf{X}_{(-K)}^{(t-1)}, \tilde{\mathbf{X}}_{(-K)}^{(t-1)}, \mathbf{Y}_{(-K)}^{(t-1)}$, and $\tilde{\mathbf{Y}}_{(-K)}^{(t-1)}$ contain no entries equal to m_K . (Here we are using assumption 1 of Theorem 1; i.e. the values m_1, \dots, m_n are distinct.) This will allow us

to use Fact 6 to bound $\zeta_{N+2}, \zeta_{N+3}, \dots, \zeta_{2N}$, which will establish (2.16) and complete the proof of Theorem 1.

Details of Proof of Theorem 1 We shall now define in detail the functions ρ_1, ρ_2 , and ρ_3 and demonstrate the properties of these functions that are used in the proof of our theorem. In this discussion we will use the following notation. If H is a distribution function, then H^\dagger will denote a function with the property that $H^\dagger(U) \sim H$, when U is a $\mathcal{U}(0, 1)$ random variable. Such an H^\dagger always exists. In the case of a distribution function on \mathcal{R}^1 , the function H^\dagger may be taken to be $H^\dagger(y) = \inf\{x : H(x) > y\}$. For economy of notation, if h is a density, then h^\dagger is used if the distribution function associated with h has not been introduced. Finally, let U_1, U_2, U_3, U_4 be independent $\mathcal{U}(0, 1)$ random variables.

By (2.17), we generate θ and ψ from uniform random variables using functions ρ_1 and ρ_2 defined below. Define

$$\phi_1(u) = \inf_{\mathbf{x} \in E} \nu'_{\mathbf{x}}(u),$$

where $E = \{\mathbf{x} : 0 \leq x_i \leq m_i, 1 \leq i \leq n\}$,

$$\pi_1 = \int \phi_1(u) du,$$

and densities

$$\bar{\phi}_1(u) = \frac{\phi_1(u)}{\pi_1} \quad \text{and} \quad \bar{\nu}_{\mathbf{x}}(u) = \frac{\nu'_{\mathbf{x}}(u) - \phi_1(u)}{1 - \pi_1}.$$

Note that $\pi_1 > 0$. This is because by assumption 3 of Theorem 1, for each i , the function f_i is bounded away from 0 over $[0, m_{(n)}]^i \times \Gamma_1$. Therefore, for each $u \in \Gamma_1$, $\phi_1(u)$ is positive. It follows that $\int_{\Theta} \phi_1(u) du \geq \int_{\Gamma_1} \phi_1(u) du > 0$.

Note that in the special case of the ‘‘exponential/gamma’’ setup used in our data analysis in Section 3, we have

$$\phi_1(u) = \inf_{\mathbf{x} \in E} \mathcal{G}(u \mid a(\mathbf{x}), b(\mathbf{x}))$$

where $a(\mathbf{x}) = a_1 + \#(\mathbf{x})$, and $b(\mathbf{x}) = b_1 + \sum^{\text{dist}} x_i$. For \mathbf{x} with $0 \leq x_i \leq m_i$, $i = 1, \dots, n$, $(a(\mathbf{x}), b(\mathbf{x}))$ ranges over a compact set which excludes $(0, 0)$. Thus, it is clear that $\phi_1(u) > 0$ for all $u > 0$.

Now we define the function ρ_1 . For given values \mathbf{x}, u_1, u_2 , we let

$$\rho_1(\mathbf{x}, u_1, u_2) = \begin{cases} \bar{\phi}_1^\dagger(u_2) & \text{if } u_1 \leq \pi_1, \\ \bar{\nu}_{\mathbf{x}}^\dagger(u_2) & \text{otherwise.} \end{cases}$$

It is easy to verify that $\rho_1(\mathbf{x}, U_1, U_2) \sim \nu_{\mathbf{x}}$.

For generating ψ we define a function ρ_2 in a similar manner. Define

$$\phi_2(u) = \inf_{\mathbf{y} \in E} \mu'_{\mathbf{y}}(u),$$

$$\pi_2 = \int \phi_2(u) du,$$

and densities

$$\bar{\phi}_2(u) = \frac{\phi_2(u)}{\pi_2} \quad \text{and} \quad \bar{\mu}_{\mathbf{y}}(u) = \frac{\mu'_{\mathbf{y}}(u) - \phi_2(u)}{1 - \pi_2}.$$

Note that $\pi_2 > 0$. Now we define ρ_2 . For given values \mathbf{y}, u_1, u_2 , we let

$$\rho_2(\mathbf{y}, u_1, u_2) = \begin{cases} \bar{\phi}_2^\dagger(u_2) & \text{if } u_1 \leq \pi_2, \\ \bar{\mu}_\mathbf{y}^\dagger(u_2) & \text{otherwise.} \end{cases}$$

It is easy to verify that $\rho_2(\mathbf{y}, U_3, U_4) \sim \mu_\mathbf{y}$.

In the proof of our theorem, we need the following fact regarding ρ_1 and ρ_2 . Let $\mathbf{x}, \tilde{\mathbf{x}}, \mathbf{y}, \tilde{\mathbf{y}}$ be arbitrary fixed values in E . Define $\theta = \rho_1(\mathbf{x}, U_1, U_2)$, $\tilde{\theta} = \rho_1(\tilde{\mathbf{x}}, U_1, U_2)$, $\psi = \rho_2(\mathbf{y}, U_3, U_4)$, and $\tilde{\psi} = \rho_2(\tilde{\mathbf{y}}, U_3, U_4)$. From the definitions of ρ_1 and ρ_2 it is clear that

$$P\{\theta = \tilde{\theta}, \theta \in \Gamma_1\} \geq \pi_1 \int_{\Gamma_1} \bar{\phi}_1(u) du = \int_{\Gamma_1} \phi_1(u) du > 0,$$

and similarly,

$$P\{\psi = \tilde{\psi}, \psi \in \Gamma_2\} \geq \pi_2 \int_{\Gamma_2} \bar{\phi}_2(u) du = \int_{\Gamma_2} \phi_2(u) du > 0,$$

so that by independence we obtain the following.

Fact 1

$$P\{\theta = \tilde{\theta}, \psi = \tilde{\psi}, \theta \in \Gamma_1, \psi \in \Gamma_2\} \geq \int_{\Gamma_1} \phi_1(u) du \int_{\Gamma_2} \phi_2(u) du.$$

Note that this bound does not involve $\mathbf{x}, \tilde{\mathbf{x}}, \mathbf{y}$, and $\tilde{\mathbf{y}}$.

By equation (2.18), the generation of $S_K = (X_K, Y_K)$ will be done using a function $\rho_3(\theta, \psi, m, \mathbf{x}, \mathbf{y}, u_1, u_2, u_3)$, where $m > 0$, $\mathbf{x} = (x_1, \dots, x_{n-1})$ and $\mathbf{y} = (y_1, \dots, y_{n-1})$ are $(n-1)$ -tuples with $x_i, y_i \geq 0$ for all i . The function ρ_3 will be defined in the following discussion.

In Lemma 2.1, take

$$A = \frac{\alpha_\theta + N_\mathbf{x}}{\alpha_\theta(\mathcal{R}) + n - 1} \quad \text{and} \quad B = \frac{\beta_\psi + N_\mathbf{y}}{\beta_\psi(\mathcal{R}) + n - 1},$$

where $N_\mathbf{x} = \sum_{i=1}^{n-1} \delta_{x_i}$ and $N_\mathbf{y} = \sum_{i=1}^{n-1} \delta_{y_i}$. Define p_1, \dots, p_5 and the distributions $\bar{A}_c, \bar{B}_c, \bar{A}_d, \bar{B}_d$ as in Lemma 2.1. These quantities are all (implicitly) functions of θ, ψ, \mathbf{x} , and \mathbf{y} .

Now define ρ_3 by

$$\rho_3(\theta, \psi, m, \mathbf{x}, \mathbf{y}, u_1, u_2, u_3) = \begin{cases} (\bar{A}_c^\dagger(u_3), m) & \text{if } 0 < u_1 \leq p_1 + p_2 \text{ and} \\ & 0 < u_2 \leq p_1/(p_1 + p_2), \\ (m, \bar{B}_c^\dagger(u_3)) & \text{if } 0 < u_1 \leq p_1 + p_2 \text{ and} \\ & p_1/(p_1 + p_2) < u_2 \leq 1, \\ (\bar{A}_d^\dagger(u_3), m) & \text{if } p_1 + p_2 < u_1 \leq p_1 + p_2 + p_3, \\ (m, \bar{B}_d^\dagger(u_3)) & \text{if } p_1 + p_2 + p_3 < u_1 \leq p_1 + p_2 + p_3 + p_4, \\ (m, m) & \text{if } p_1 + p_2 + p_3 + p_4 < u_1 \leq 1. \end{cases}$$

It is clear that the pair $(X, Y) = \rho_3(\theta, \psi, m, \mathbf{x}, \mathbf{y}, U_1, U_2, U_3)$ has the distribution given in (2.7) since the function ρ_3 merely describes a particular way to carry out the mixing procedure in Lemma 2.1.

The function ρ_3 has two properties which we use in the proof of our theorem. Pick arbitrary fixed values $\theta \in \Theta$, $\psi \in \Psi$, $m > 0$, and $\mathbf{x}, \tilde{\mathbf{x}}, \mathbf{y}, \tilde{\mathbf{y}}$ in \mathcal{R}^{n-1} . Define

$$\begin{aligned}(X, Y) &= \rho_3(\theta, \psi, m, \mathbf{x}, \mathbf{y}, U_1, U_2, U_3) \quad \text{and} \\ (\tilde{X}, \tilde{Y}) &= \rho_3(\theta, \psi, m, \tilde{\mathbf{x}}, \tilde{\mathbf{y}}, U_1, U_2, U_3).\end{aligned}$$

First note that $X \wedge Y$ is generated from a continuous distribution when $U_1 \leq p_1 + p_2$. Therefore we have the following fact.

Fact 2 *If $U_1 \leq (p_1 + p_2) \wedge (\tilde{p}_1 + \tilde{p}_2)$, then both $X \wedge Y$ and $\tilde{X} \wedge \tilde{Y}$ are generated from continuous distributions.*

In conjunction with this fact, we note that $p_1 + p_2 \geq f/(f + 1)$, where

$$f = \frac{\alpha_\theta(m) \wedge \beta_\psi(m)}{n - 1}.$$

To see this, we go back to the expressions for p_1, \dots, p_5 given in Lemma 2.1. We find that

$$A_d(m)p_1 = A_c(m)(p_3 + p_5) \quad \text{and} \quad B_d(m)p_2 = B_c(m)(p_4 + p_5).$$

Since

$$\begin{aligned}A_d(m) &\leq \frac{n - 1}{\alpha_\theta(\mathcal{R}) + n - 1}, \quad B_d(m) \leq \frac{n - 1}{\beta_\psi(\mathcal{R}) + n - 1}, \\ A_c(m) &= \frac{\alpha_\theta(m)}{\alpha_\theta(\mathcal{R}) + n - 1}, \quad \text{and} \quad B_c(m) = \frac{\beta_\psi(m)}{\beta_\psi(\mathcal{R}) + n - 1},\end{aligned}$$

these relations imply

$$(n - 1)p_1 \geq \alpha_\theta(m)(p_3 + p_5) \quad \text{and} \quad (n - 1)p_2 \geq \beta_\psi(m)(p_4 + p_5).$$

Thus

$$p_1 + p_2 \geq f \cdot (p_3 + p_4 + p_5) = f \cdot (1 - p_1 - p_2),$$

so that $p_1 + p_2 \geq f/(1 + f)$ as desired. From assumption 2 of Theorem 1, we know that f is bounded below for $\theta \in \Gamma_1$ and $\psi \in \Gamma_2$. Therefore, we can state the following fact.

Fact 3 *For $\theta \in \Gamma_1$ and $\psi \in \Gamma_2$, the value $(p_1 + p_2)$ is bounded away from 0 by some quantity which does not depend on θ, ψ, \mathbf{x} , or \mathbf{y} .*

Now observe that A and \tilde{A} have the same continuous part; that is $A_c = \tilde{A}_c$. This implies that $\bar{A}_c = \tilde{\bar{A}}_c$ and thus $(\bar{A}_c)^\dagger = (\tilde{\bar{A}}_c)^\dagger$. Similarly, $(\bar{B}_c)^\dagger = (\tilde{\bar{B}}_c)^\dagger$. This leads to our next property.

Fact 4 *If $U_1 \leq (p_1 + p_2) \wedge (\tilde{p}_1 + \tilde{p}_2)$ and*

$$U_2 \leq \frac{p_1}{p_1 + p_2} \wedge \frac{\tilde{p}_1}{\tilde{p}_1 + \tilde{p}_2} \quad \text{or} \quad U_2 > \frac{p_1}{p_1 + p_2} \vee \frac{\tilde{p}_1}{\tilde{p}_1 + \tilde{p}_2},$$

then $X = \tilde{X}$, $Y = \tilde{Y}$, and $X \wedge Y$ is generated from a continuous distribution.

Also note that when $N_{\mathbf{x}}(\{m\}) = N_{\mathbf{y}}(\{m\}) = 0$, we have

$$\frac{p_1}{p_1 + p_2} = \frac{\alpha_\theta(m)\beta'_\psi(m)}{\alpha_\theta(m)\beta'_\psi(m) + \alpha'_\theta(m)\beta_\psi(m)}$$

so that $p_1/(p_1 + p_2)$ does not depend on \mathbf{x} and \mathbf{y} in this situation. Now, using properties in assumption 2 of Theorem 1, we obtain Fact 5.

Fact 5 *If $N_{\mathbf{x}}(\{m\}) = N_{\mathbf{y}}(\{m\}) = 0$, then for $\theta \in \Gamma_1$ and $\psi \in \Gamma_2$, the value $p_1/(p_1 + p_2)$ can be bounded away from 0 and 1 by quantities which do not depend on θ, ψ, \mathbf{x} , and \mathbf{y} .*

Combining Facts 3, 4, and 5 leads to our last property.

Fact 6 *If $N_{\mathbf{x}}(\{m\}) = N_{\mathbf{y}}(\{m\}) = N_{\tilde{\mathbf{x}}}(\{m\}) = N_{\tilde{\mathbf{y}}}(\{m\}) = 0$, $\theta \in \Gamma_1$, and $\psi \in \Gamma_2$, then the probability of the event*

$$\{X = \tilde{X}, Y = \tilde{Y}, \text{ and } X \wedge Y \text{ is generated from a continuous distribution}\}$$

is bounded away from 0 by a quantity which does not depend on $\theta, \psi, \mathbf{x}, \mathbf{y}, \tilde{\mathbf{x}}$, or $\tilde{\mathbf{y}}$.

This concludes the details of the proof of Theorem 1.

Note that

$$\mathcal{L}_{\text{data}}(F, G) = \int (\mathcal{D}_{\alpha_\theta + \sum_{i=1}^n \delta_{X_i}} \otimes \mathcal{D}_{\beta_\psi + \sum_{i=1}^n \delta_{Y_i}}) \mathcal{L}_{\text{data}}(d\theta, d\psi, d\mathbf{X}, d\mathbf{Y}).$$

(In particular, the marginal posterior distribution of F is a mixture of Dirichlets, and a similar statement holds for G .) Let

$$(F, G)^{(jN)} \sim \mathcal{D}_{\alpha_\theta(jN) + \sum_{i=1}^n \delta_{X_i^{(jN)}}} \otimes \mathcal{D}_{\beta_\psi(jN) + \sum_{i=1}^n \delta_{Y_i^{(jN)}}}.$$

Corollary 2.1 *Let $Q_{\mathbf{s}}^{(jN)}$ denote the distribution of $(F, G)^{(jN)}$ when the chain $\{\mathbf{S}^{(\ell N)}\}_\ell$ is started at \mathbf{s} . Then under the conditions of Theorem 1, we have*

$$\sup_{A \in \mathcal{A}} |Q_{\mathbf{s}}^{(jN)}\{(F, G)^{(jN)} \in A\} - P_{\text{data}}\{(F, G) \in A\}| \leq e^{-\lambda j} \quad \text{for all } \mathbf{s} \quad (2.19)$$

where λ is the same λ that appears in the statement of Theorem 1.

(In (2.19), $\mathcal{A} = \mathcal{A}_0 \times \mathcal{A}_0$, where \mathcal{A}_0 is the smallest σ -field on the set of probability measures on \mathcal{R} such that the map $P \mapsto P(B)$ is measurable for each Borel set $B \subset \mathcal{R}$.)

Proof of Corollary 2.1 Let $\pi_{\mathbf{s}}^{(jN)}$ be the distribution of $\mathbf{S}^{(jN)}$ when the Markov chain $\{\mathbf{S}^{(jN)}\}_j$ is started at \mathbf{s} . Fix $A \in \mathcal{A}$ and let

$$f_A(\mathbf{s}) = \left(\mathcal{D}_{\alpha_\theta + \sum_{i=1}^n \delta_{X_i}} \otimes \mathcal{D}_{\beta_\psi + \sum_{i=1}^n \delta_{Y_i}} \right)(A).$$

We know that

$$\begin{aligned}
|Q_{\mathbf{s}^{(0)}}^{(jN)}(A) - P_{\text{data}}\{(F, G) \in A\}| &= \left| \int f_A(\mathbf{s}) d\pi_{\mathbf{s}^{(0)}}^{(jN)}(\mathbf{s}) - \int f_A(\mathbf{s}) d\pi(\mathbf{s}) \right| \\
&= \left| \int f_A(\mathbf{s}) d(\pi_{\mathbf{s}^{(0)}}^{(jN)} - \pi)(\mathbf{s}) \right| \tag{2.20} \\
&= \left| \int f_A(\mathbf{s}) d(\pi_{\mathbf{s}^{(0)}}^{(jN)} - \pi)^+(\mathbf{s}) - \int f_A(\mathbf{s}) d(\pi_{\mathbf{s}^{(0)}}^{(jN)} - \pi)^-(\mathbf{s}) \right|.
\end{aligned}$$

(For a signed measure λ , the representation $\lambda = \lambda^+ - \lambda^-$ is the standard Jordan decomposition of λ .) Since $f_A(\cdot)$ is a measurable function of \mathbf{s} that satisfies $0 \leq f_A(\mathbf{s}) \leq 1$ for all \mathbf{s} , we see that each of the two integrals in the last line of (2.20) is bounded by $\exp(-\lambda j)$, and this proves Corollary 2.1.

Thus, if for example we want to estimate, for fixed v , the density of $\mathcal{L}_{\text{data}}(F(v))$, we would use the mixture

$$\frac{1}{L} \sum_{j=1}^L \text{beta} \left(\left(\alpha_{\theta^{(jN)}} + \sum_{i=1}^n \delta_{X_i^{(jN)}} \right) (0, v], \left(\alpha_{\theta^{(jN)}} + \sum_{i=1}^n \delta_{X_i^{(jN)}} \right) (v, \infty) \right)$$

where $\text{beta}(a, b)(\cdot)$ is the Beta density.

2.4 The Algorithm for Arbitrary Systems

Here we discuss the implementation and convergence of the algorithm in the general case. A discussion of the issues of identifiability and frequentist consistency appears in Lawson (1994).

2.4.1 Implementation of the Algorithm

Let the autopsy statistics for system i , $i = 1, \dots, n$ be (T_i, D_i) , where T_i is the death time of the system and D_i is the set of components that are dead at time T_i . Recall that after examining (T_i, D_i) , each component in system i is put into exactly one of the categories C1, C2, C3, or C4 described in Section 1. For a component in Category C1, one generates an observation according to the distribution $(A_i^{(t)})_{(T_i, \infty)}$, where $A_i^{(t)}$ is defined in (2.9) (i.e. the distribution $A_i^{(t)}$ restricted to (T_i, ∞) and renormalized to be a probability measure). Similarly, for components in Category C2, we generate an observation from $(A_i^{(t)})_{[0, T_i]}$. For a component in Category C3, nothing needs to be done.

Suppose there are k components that fall in Category C4. We then use an extension of Lemma 2.1 describing the conditional distribution of k independent random variables ($k \geq 2$), whose distributions have both absolutely continuous and discrete components, given the value of their maximum. The necessary formulas are easy to derive but require elaborate notation to write down explicitly, and so are not given here. We note however, that the needed computer algorithm is relatively easy to implement.

We remark that the case of an arbitrary coherent system is no more difficult than that for a general parallel system if we note that the set of components in Category C4 changes from system to system.

2.4.2 Convergence of the Algorithm

When considering the case of an arbitrary system, it is helpful to first look at the situation when the prior distribution on each F_j is a single Dirichlet, i.e. there is no mixing. In this case, the updating of the lifelength of component j in system i is based on (2.9), where $\alpha_{\theta(t)}$ is replaced simply by α . When updating a component in Category C1, the probability of drawing from the fixed probability measure proportional to $\alpha(\cdot \cap (T_i, \infty))$ is bounded below by $\alpha((T_i, \infty))/(\alpha((T_i, \infty)) + n - 1)$, independently of the current state of the chain. A similar statement holds for the lifelengths of components in Category C2. We have already explained how to deal with the lifelengths for components in Category C4 in Section 2.3. Thus a coupling argument along the lines of the proof of Theorem 1 gives convergence at a uniform geometric rate.

When the priors on the F_j 's are mixtures of Dirichlets, a difficulty arises in that the distributions $\alpha_{\theta(t)}$ in general need not have a uniform lower bound. For parallel systems we were able to find a uniform lower bound for the posterior distribution of θ given the lifelengths \mathbf{X} only because \mathbf{X} is known to lie in a compact set. Since the lifelengths of components in Category C1 do not lie in a compact set, this argument no longer applies. For general systems convergence of the Markov chain can be established using the lower bounds established in Theorem 1 in conjunction with Theorem 1 of Athreya, Doss, and Sethuraman (1992), which gives simple ergodicity (i.e. convergence, but not at a geometric rate).

3 Analysis of U.S. Air Force C-17 Fuel Quantity Computer Data

We illustrate our algorithm on data involving survival times of the Fuel Quantity (FQ) Computer system of the C-17 transport aircraft. The test program will eventually involve six aircraft being flown for approximately 10000 cumulative hours. Our data set is taken relatively early in the test program, since only 2440 flight hours had been accumulated at the time of this writing. The data, listed below in Table 1, fall into one of three categories (we denote the failure times of the “A-bus” and “B-bus” as X and Y , respectively):

- The FQ computer fails (both buses are dead) and the maximum survival time, say t_0 , of the two buses is observed; i.e. we have the usual autopsy statistics (system failure time and set of dead components). This type of observation has the form “ $X \vee Y = t_0$ ”.
- The two components are checked at time t_1 and time t_2 . Both buses are alive at t_1 , but one of the buses, say “B”, is in a failed state at t_2 . Even though the “A-bus” is alive at t_2 , the FQ computer is replaced. This situation generates two observations, which have the form $Y \in (t_1, t_2]$ and $X \in (t_2, \infty)$.
- Both components of the FQ computer are alive when the data are taken, but the aircraft had flown for t_3 hours. The failure times for both buses lie in the interval $[t_3, \infty)$. Thus, two observations are generated: $X \in (t_3, \infty)$ and $Y \in (t_3, \infty)$.

<i>Aircraft Id</i>	<i>Event</i>	<i>hours</i>	<i>Observation</i>
P-1	FQ Computer failed	43.4	$X_1 \vee Y_1 = 43.4$
P-1	FQ Computer failed	236.8	$X_2 \vee Y_2 = 236.8$
P-2	FQ Computer failed	244.0	$X_3 \vee Y_3 = 244.0$
T-1	A-bus and B-bus alive	11.9	$Y_4 \in (11.9, 15.4]$ $X_4 \in (15.4, \infty)$
	B-bus dead	15.4	
	FQ Computer replaced	15.4	
P-4	A-bus and B-bus alive	174.4	$Y_5 \in (174.4, 181.8]$ $X_5 \in (181.8, \infty)$
	B-bus dead	181.8	
	FQ Computer replaced	181.8	
T-1	A-bus and B-bus alive	819.6	$X_6 \in (819.6, \infty)$
	no more data available		$Y_6 \in (819.6, \infty)$
P-1	A-bus and B-bus alive	85.0	$X_7 \in (85.0, \infty)$
	no more data available		$Y_7 \in (85.0, \infty)$
P-2	A-bus and B-bus alive	476.4	$X_8 \in (476.4, \infty)$
	no more data available		$Y_8 \in (476.4, \infty)$
P-3	A-bus and B-bus alive	24.5	$X_9 \in (24.5, \infty)$
	independent software failure		$Y_9 \in (24.5, \infty)$
P-3	A-bus and B-bus alive	71.7	$X_{10} \in (71.7, \infty)$
	no more data available		$Y_{10} \in (71.7, \infty)$
P-4	A-bus and B-bus alive	68.4	$X_{11} \in (68.4, \infty)$
	no more data available		$Y_{11} \in (68.4, \infty)$
P-5	A-bus and B-bus alive	173.4	$X_{12} \in (173.4, \infty)$
	no more data available		$Y_{12} \in (173.4, \infty)$

Table 1. C-17 Fuel Quantity Computer Data. The first three lines contain observed FQ Computer failure times. The next two categories (four observations) occurred when the FQ computer was replaced due to a failure in one bus. The remaining observations occurred because no more data is available.

Note that this data structure is a bit more complex than the data structure in the autopsy model; however, the required modifications to the algorithm involve no real difficulties.

The reader may wonder why there is a need for a computer, as opposed to a simple analogue gauge, to deal with fuel quantity. Indeed, this is not a frivolous issue. There is actually a need for a computer even during level flight, since the aircraft maintains its desired center of gravity via fuel transfer from one wing to another. This task is further complicated as the aircraft flies at different angles or possibly under turbulence. The FQ computer receives the current angle of flight from another computer and uses this information, along with readings from a series of probes in each fuel tank, to make accurate fuel quantity calculations. Also, the Mission Computer requires input from the FQ computer to make range calculations.

We analyzed the C-17 data using our proposed algorithm. We took our prior on both F and G to be (2.5), where H_θ is the exponential distribution with parameter θ (the mean

is $1/\theta$). We assumed $\alpha_\theta(\mathcal{R})$ to be constant in θ and considered three cases: $\alpha_\theta(\mathcal{R}) = 1$, $\alpha_\theta(\mathcal{R}) = 10$, and $\alpha_\theta(\mathcal{R}) = 100$. We took $\nu = \mathcal{G}(a, b)$ (the Gamma distribution with shape parameter a and scale parameter b), since we wished to center the prior around the family of exponential distributions, and the Gamma is conjugate for this family. From (2.11), $\nu_{\mathbf{X}} = \mathcal{G}(a + \#(\mathbf{X}), b + \sum^{\text{dist}} X_i)$, and $\mu_{\mathbf{Y}}$ is similarly defined.

We elicited the prior of a computer systems engineer from the C-17 Special Programs Office by asking his opinion about the FQ computer Mean Time Before Failure (MTBF) with respect to two reports. The first report, supplied by the C-17 manufacturer, provides target numbers for each “Logical Replaceable Unit” (LRU), including the FQ computer. The manufacturer guarantees that the MTBF for each LRU, computed at the end of the acceptance testing period, will exceed that LRU’s target number. The C-17 engineer thought it was highly likely (probability of .9) that the FQ Computer MTBF would exceed the target number (which was 1300 hours). The second report, supplied to the C-17 Special Programs Office by the manufacturer’s design group, contains a list of “mature” MTBF numbers for each LRU being evaluated. These numbers represent an average of MTBF’s, by LRU, across many different aircraft which have similar LRUs. The data come from maintenance data accumulated following the acceptance testing periods for each aircraft (hence the word “mature”). Since these numbers come from mature aircraft, the C-17 engineer thought it was quite unlikely (probability of .1) that the FQ Computer MTBF for the acceptance testing period would exceed the mature MTBF number (which was 3167 hours). Thus, we took 1300 hours and 3167 hours to be the .1 and .9 quantiles of the distribution of the MTBF for the FQ Computer.

If $X \sim \mathcal{E}(\theta)$ (the exponential distribution with parameter θ) and $Y \sim \mathcal{E}(\psi)$, then $\text{MTBF} = E(X \vee Y) = (1/(\theta + \psi))[1 + \theta/\psi + \psi/\theta]$. If θ, ψ are iid $\sim \mathcal{G}(a, b)$, then the .1 and .9 quantiles of the distribution of the MTBF are equal to 1300 and 3167, respectively, when $a = 6.04424$ and $b = 6835.32$. Note that if the conditional distribution of X given θ is $\mathcal{E}(\theta)$ and θ is distributed as $\mathcal{G}(a, b)$, then the unconditional cumulative distribution function of X is $F(t) = 1 - (b/(b+t))^a$, which is a “shifted Pareto” distribution with parameters a and b . The prior distributions of X and Y are each shifted Pareto distributions with parameters 6.04424 and 6835.32. The prior distribution of $X \vee Y$ is the product of two such distributions.

Of particular interest to the C-17 engineers is the question of how the lifelength of a future FQ computer, as well as the lifelengths of a future A-bus and B-bus, would be distributed. The Bayes approach is especially well-suited to answer such a question. Figures 1, 2, and 3 give the prior density and a representation of the posterior distribution for the future lifelengths of the C-17 FQ computer (maximum lifelength of the A-bus and B-bus), the A-bus, and the B-bus, for the cases $\alpha_\theta(\mathcal{R}) = 1$, $\alpha_\theta(\mathcal{R}) = 10$, and $\alpha_\theta(\mathcal{R}) = 100$. Note that for each posterior distribution, the prior distribution is a shifted Pareto distribution with parameters 6.04424 and 6835.32 (or the product of two such distributions in the case of Figure 1). Recall from Table 1 that three system failures were observed, at times 43.4, 236.8, 244.0. This causes all the posterior distributions to have atoms at these three failure times. The masses at these three failure times have been removed and plotted as distinct spikes, with their masses labeled separately. It is interesting to note that the distribution of the lifelength of the B-bus has much more mass to the left of 236.8 hours than the other two distributions, making it necessary to use a

larger scale for the vertical axis in Figure 3. This can be attributed to the two pairs of observations (X_4, Y_4) and (X_5, Y_5) (see Table 1), which were generated due to a failure in the B-bus while the A-bus was still functioning. For low values of $\alpha_\theta(\mathcal{R})$ (the cases $\alpha_\theta(\mathcal{R}) = 1$ and $\alpha_\theta(\mathcal{R}) = 10$), these two observations cause the algorithm to assign a fairly high value to the conditional probability that the lifelength of the A-bus is set to the observed maximum and the lifelength of the B-bus is set to a value less than the observed maximum. The masses at the three observed maximum values can be explicitly seen in the plots for the cumulative distribution functions, given in Lawson (1994). Figures showing the posterior density estimates for the future lifelengths of the maximum, A-bus, and B-bus, with the spikes smoothed, also appear in Lawson (1994). From Figures 1, 2 and 3 it can be seen that the masses at the three observed FQ computer failure times account for much of the mass in the posterior distributions for the cases $\alpha_\theta(\mathcal{R}) = 1$ and $\alpha_\theta(\mathcal{R}) = 10$.

The means of the posterior distributions of the lifelengths of the FQ computer, for the cases $\alpha_\theta(\mathcal{R}) = 1$, $\alpha_\theta(\mathcal{R}) = 10$, and $\alpha_\theta(\mathcal{R}) = 100$, are 1254, 1230, and 1230 hours, respectively. At this relatively early stage of the study, we conclude that the performance of the FQ computer is not as good as the Air Force would like to see, as the means are just below the minimum acceptable MTBF. However, we caution that the study will continue for an additional two years beyond the close of our data set, and that some of the early failures experienced can be directly attributed to ongoing design changes. It will be interesting to rerun the algorithm on the updated data set at the close of the acceptance testing period. Additional details on our analysis of the C-17 FQ computer data appear in Lawson (1994).

Acknowledgements


We are very grateful to Lieutenant Colonel Larry Wolf and Captain Brian Kearns of the Air Force Operational Test and Evaluation Center, and to Captain Ron Cleaves, Bob Armstrong, and Bob Fooks of the C-17 Special Programs Office for their help in obtaining data for the C-17 Fuel Quantity Computer.

References


- Antoine, R., Doss, H. and Hollander, M. (1993). On the identifiability in the autopsy model of reliability theory. *J. Appl. Prob.* **30** 913–930.
- Antoniak, C. (1974). Mixtures of Dirichlet processes with applications to Bayesian non-parametric problems. *Ann. Statist.* **2** 1152–74.
- Barlow, R. E. and Proschan, F. (1981). Statistical theory of reliability and life testing. To Begin With, Silverspring, Maryland.
- Blackwell, D. and MacQueen J. B. (1973). Ferguson distributions via Pólya urn schemes. *Ann. Statist.* **1** 353–355.
- Doss, H. (1994). Bayesian nonparametric estimation for incomplete data via successive substitution sampling. *Ann. Statist.*, to appear.

- Doss, H., Freitag, S. and Proschan, F. (1989). Estimating Jointly System and Component Reliabilities Using a Mutual Censorship Approach. *Ann. Statist.* **17** 764–782.
- Ferguson, T. S. (1973). A Bayesian analysis of some nonparametric problems. *Ann. Statist.* **1** 209–230.
- Ferguson, T. S. (1974). Prior distributions on spaces of probability measures. *Ann. Statist.* **2** 615–629.
- Gelfand, A. and Smith, A. F. M. (1990). Sampling-based approaches to calculating marginal densities. *J. Amer. Statist. Assoc.* **85** 398–409.
- Geyer, C. J. (1992). Practical Markov Chain Monte Carlo (with discussion). *Statist. Sci.* **7** 473–511.
- Gilks, W. R. and Wild, P. (1992). Adaptive rejection sampling for Gibbs sampling. *Appl. Statist.* **41** 337–348.
- Lawson, K. L. (1994). Bayesian Nonparametric Estimation Via Gibbs Sampling for Coherent Systems With Redundancy. Ph.D. Dissertation, Florida State University.
- Meilijson, I. (1981). Estimation of the lifetime distribution of the parts from the autopsy statistics of the machine. *J. Appl. Prob.* **18** 829–838.
- Meilijson, I. (1994). Competing risks on coherent reliability systems. Estimation in the parametric case. *J. Amer. Statist. Assoc.*, to appear.
- Nowick, S. (1990). Identifiability problems in coherent systems. *J. Appl. Prob.* **28** 862–872.
- Sethuraman, J. (1994). A constructive definition of Dirichlet priors. *Statistica Sinica*, to appear.
- Tanner, M. A. and Wong, W. (1987). The calculation of posterior distribution by data augmentation. *J. Amer. Statist. Assoc.* **82** 528–550.
- Watelet, L. (1990). Nonparametric Estimation of Component Life Distribution in Meilijson's Competing Risks Model. Ph.D. Dissertation, University of Washington.


Prior Density for Lifelength of Maximum of A-bus & B-bus



Posterior Density for Lifelength of a Future Maximum, large atoms removed, $\alpha=1$



Posterior Density for Lifelength of a Future Maximum, large atoms removed, $\alpha=10$



Posterior Density for Lifelength of a Future Maximum, large atoms removed, $\alpha=100$




Figure 1. Prior and Posterior Density Estimates for the lifelength of the FQ Computer (Maximum of A-bus and B-bus), $\alpha(R) = 1$, $\alpha(R) = 10$, and $\alpha(R) = 100$.

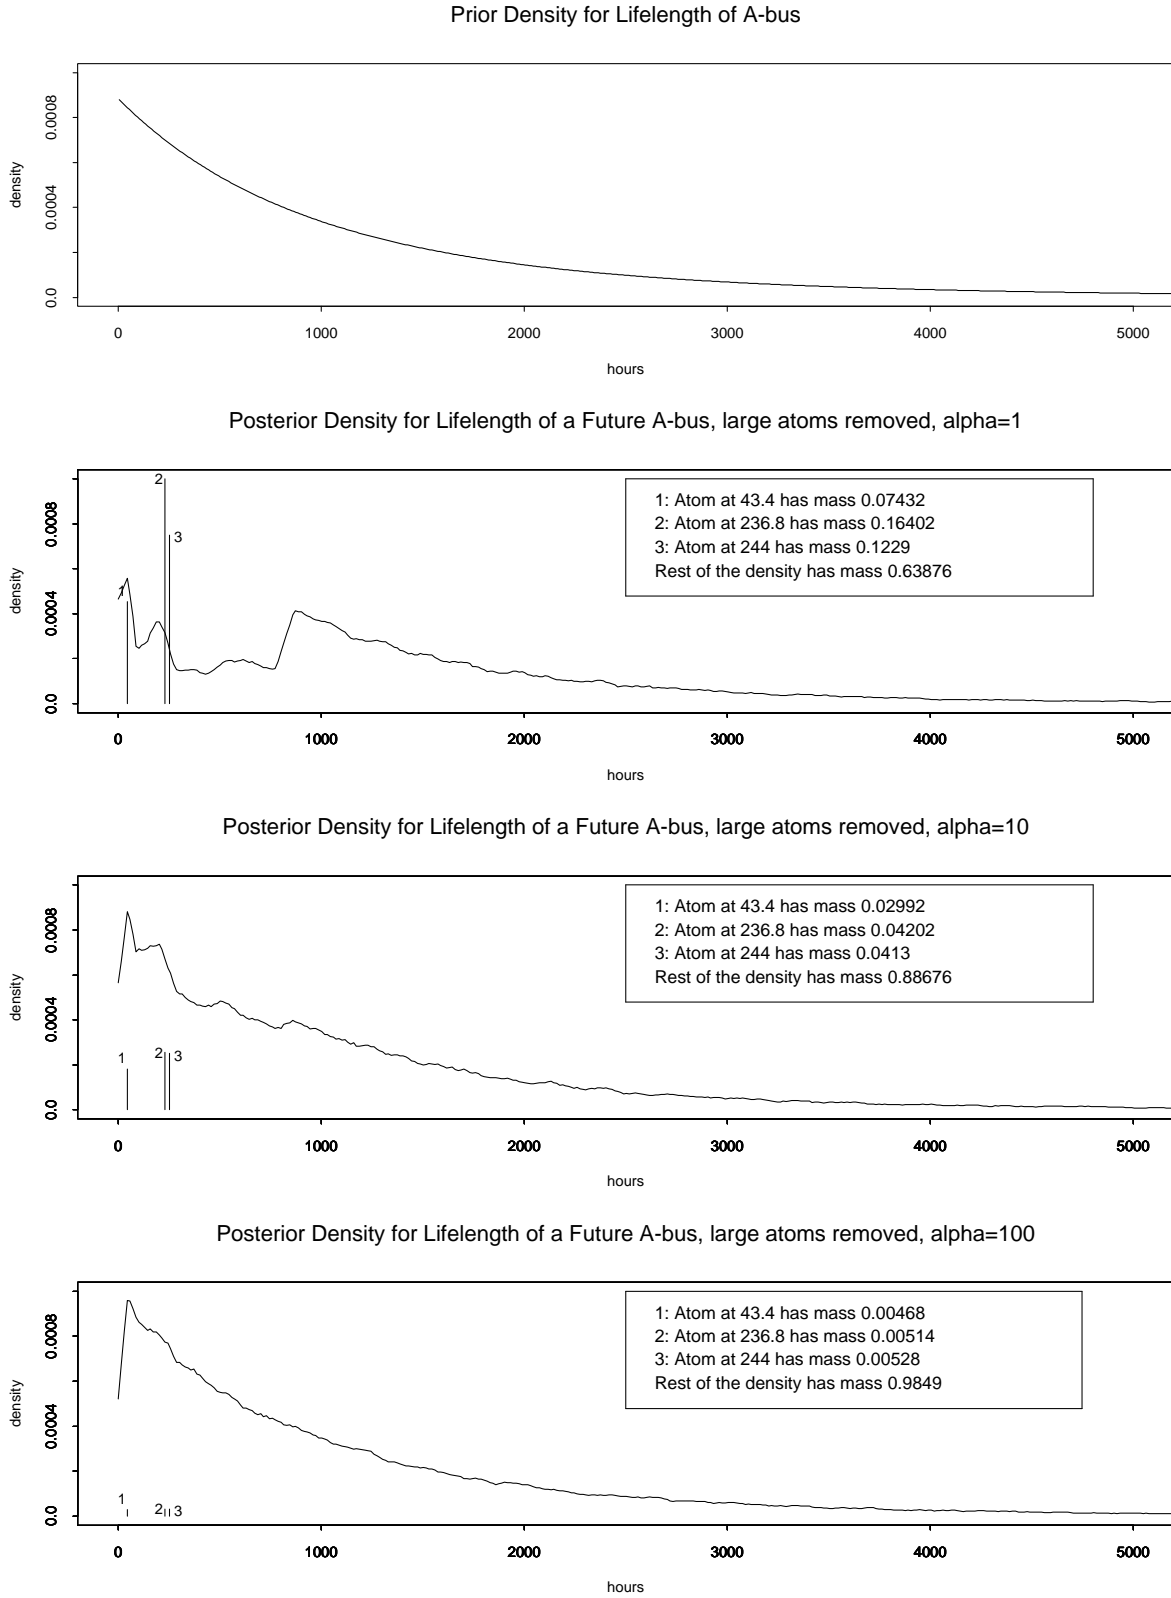


Figure 2. Prior and Posterior Density Estimates for the lifelength of the FQ Computer's A-bus, $\alpha(R) = 1$, $\alpha(R) = 10$, and $\alpha(R) = 100$.

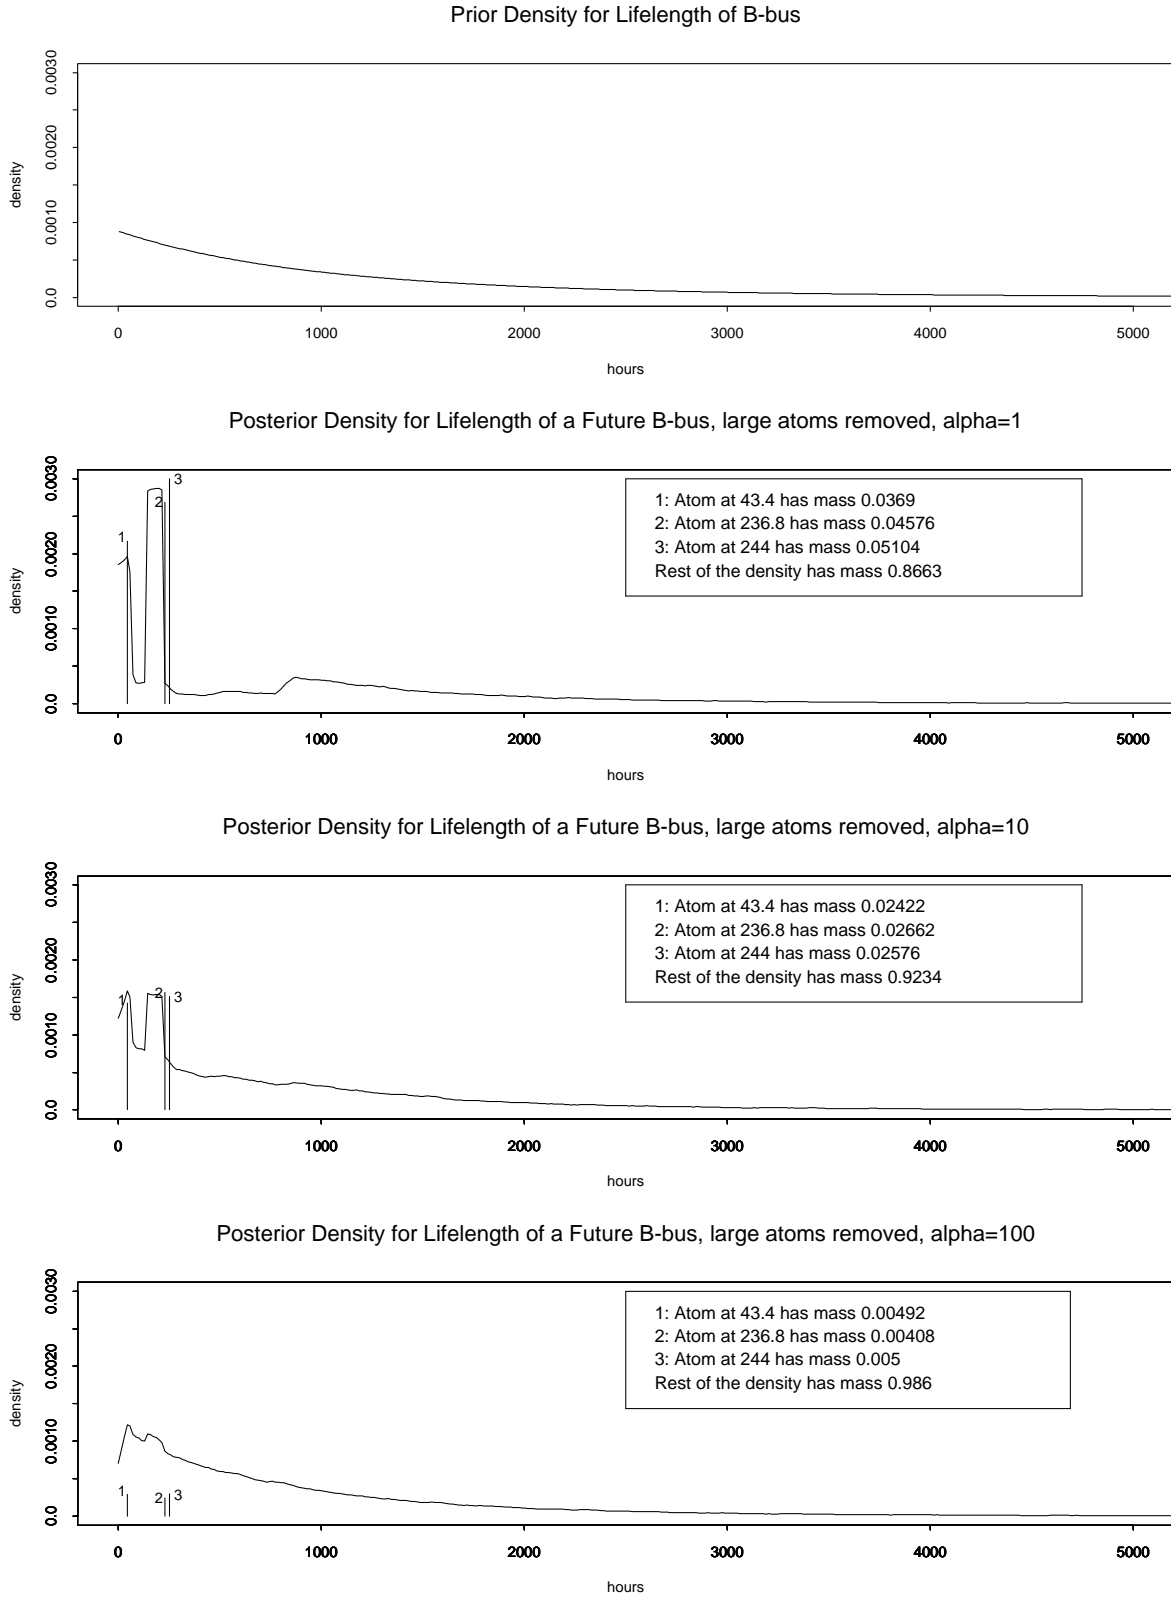


Figure 3. Prior and Posterior Density Estimates for the lifelength of the FQ Computer's B-bus, $\alpha(R) = 1$, $\alpha(R) = 10$, and $\alpha(R) = 100$.