Bayesian Statistics: An Introduction for the Practicing Reliability Engineer

Carsten H. Botts

ABSTRACT

This article introduces and reviews some of the principles and methods used in Bayesian reliability. It specifically discusses methods used in the analysis of success/no-success data and describes a simple Monte Carlo algorithm that can be used to calculate the posterior distribution of a system's reliability. This algorithm is especially useful when a system's reliability is modeled through the reliability of its subcomponents, yet only system-level data are available.

INTRODUCTION

A common way to measure the reliability of a system is to determine the probability that it will pass, or survive, a stress test. This typically requires several system tests, but it may be prohibitively expensive to conduct many tests of a sophisticated system, such as an air-to-air missile¹ or medical device.²

Bayesian statistical methods can help in such a situation, since they enable inclusion of other types of data (such as computer simulation experiments or subject-matter-expert opinions). Bayesian methods may also be necessary because many modern systems do not fail during testing. With no failures, it is difficult for classical statistics to accurately quantify the probability of failure.

This article begins by briefly reviewing Bayesian methods. It then discusses how these methods can be used to learn more about the probability of a system surviving a test, and it concludes by describing a straightforward algorithm for calculating a total system's reliability once it has been tested. This algorithm is simple, produces an exact answer, and is not mentioned in the Bayesian reliability literature. The efficiency of this algorithm is illustrated with an example.

BAYESIAN METHODS

This section introduces the basics of Bayesian principles and Bayesian statistical methodology. The most effective way to introduce this concept is to contrast it with the principles and methodology of classical statistics. The biggest difference between Bayesian and classical statistics is in how probability is defined. In classical statistics, probability is the long-run frequency of an event. So for a fixed (and unknown) parameter such as a population mean, μ ,

$$\mathbb{P}(3.66 \le \mu \le 4.11) = \begin{cases} 1 & \text{if true} \\ 0 & \text{if not} \end{cases}$$
(1)

In words, Eq. 1 states that the fixed parameter μ is either in the stated interval or it is not. Bayesian statisticians think about probability in a different way. In Bayesian statistics, probability is the belief that a statement is true. So if one believes (based on their experience and/or the data that they have seen) that μ is within the stated interval with 95% probability, it would be fair to say that

$$\mathbb{P}(3.66 \le \mu \le 4.11) = .95.$$
(2)

The objective and point of Bayesian statistics is to calculate probabilities like the one in Eq. 2 and to ensure that this calculation is scientifically respected.

To calculate such a probability, a Bayesian statistician begins with a prior distribution. Assuming that the unknown parameter of interest is θ , this prior distribution is typically denoted as $\pi(\theta)$. The prior distribution indicates where the user believes the parameter θ to be before data are observed or collected. Assume, for example, that we purchased a coin at a magic shop. Upon the purchase, the shop owner tells us that the coin will more often turn up heads than tails. In this case, we will let $\theta = \mathbb{P}(H)$ and define the prior distribution $\pi(\theta)$ for all values of θ between 0 and 1. This prior will also be more heavily weighted toward values of 1 to indicate that, a priori, the coin is expected to turn up heads more frequently than tails.

Once the prior is formulated, data are collected. The distribution of the data conditioned on a value of θ is written as $p(\mathbf{x}|\theta)$, i.e.,

$$(X_1, X_2, \ldots, X_n) \sim p(x_1, x_2, x_3, \ldots, x_n | \boldsymbol{\theta}) = p(\mathbf{x} | \boldsymbol{\theta}),$$

where $\mathbf{x} = (x_1, x_2, \dots, x_n)$. The function $p(\mathbf{x}|\theta)$ is also referred to as the likelihood of θ .

With the prior and the likelihood, the posterior distribution, typically denoted as $\pi(\theta | \mathbf{x})$, can be calculated. The posterior is calculated using Bayes's rule.³ This calculation is shown below:

$$\pi(\boldsymbol{\theta}|\mathbf{x}) = \frac{p(\mathbf{x},\boldsymbol{\theta})}{p(\mathbf{x})} = \frac{p(\mathbf{x}|\boldsymbol{\theta})\pi(\boldsymbol{\theta})}{\int_{\Theta} p(\mathbf{x}|\boldsymbol{\theta})\pi(\boldsymbol{\theta})d\boldsymbol{\theta}} \propto p(\mathbf{x}|\boldsymbol{\theta})\pi(\boldsymbol{\theta}), \quad (3)$$

where Θ is the set of all possible values of θ . The formula given in Eq. 3 makes sense: the posterior is proportional to the prior distribution of θ (where we thought θ was before collecting data) times the likelihood (where the data suggests θ to be).

This article focuses on how to use Bayesian methods to learn more about the probability that a system survives a test. The subsection titled The Prior and Posterior of One Subsystem specifically discusses how Bayesian methods are used to learn about the survival probability of one system, and the subsection following that one, The Prior and Posterior of the Entire System, discusses the methods necessary to learn about the survival probability of one system composed of multiple subsystems. The algorithm discussed in the latter subsection is elementary and uses no approximations when calculating its answer.

BAYESIAN METHODS FOR BERNOULLI EXPERIMENTS

The Prior and Posterior of One Subsystem

Let us first assume that we are interested in learning about the probability that a system passes an endurance test of some sort. We will denote this probability as θ , and we will conduct *n* trials/tests on this system and record *X*, the number of times (out of the *n* trials) that the system passes a test.

To do a Bayesian analysis on θ , we begin by specifying a prior distribution for it. The beta distribution is often used as the prior for the probability of success in a sequence of success/failure trials.^{4,5} The beta distribution is specified by two parameters and is especially convenient in cases such as this since it is a conjugate prior (i.e., it produces a posterior distribution of the same form). The beta prior takes the form

$$\pi(\theta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha - 1} (1 - \theta)^{\beta - 1} \qquad 0 \le \theta \le 1,$$

where $\Gamma(\cdot)$ is the gamma function. This prior has mean

Prior Mean =
$$\frac{\alpha}{\alpha + \beta}$$

and variance

Prior Var =
$$\frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}$$
.

The values of α and β (α , $\beta > 0$) are selected to reflect the user's prior belief. This prior belief is often informed in a variety of ways, such as expert opinion, computer simulation, or prior experiments. A user who believes that θ is small (<0.5) would set $\alpha < \beta$ (making the prior mean < 0.5). A user who believes that θ is large (>0.5) would set $\alpha > \beta$. The confidence in these prior beliefs is, of course, reflected in the variance of the prior. If a user wanted to set the prior mean of θ to 0.4, they could set α = 2 and β = 3, making the prior variance 0.04. If a user wanted to elevate the confidence in this statement (that the prior mean of θ is 0.4), they could adjust the values of α and β to α = 20 and β = 30, making the prior variance 0.004. And if a user knew absolutely nothing about θ , they would set $\alpha = \beta = 1$, in which case the prior for θ is uniform over the interval (0, 1).

Johns Hopkins APL Technical Digest, Volume 38, Number 1 (2025), Article 2501649 www.jhuapl.edu/technical-digest There is no consistent selection of α and β in reliability studies. Leoni et al.⁶ set $\alpha = 3$ and $\beta = 1$ in one of their reliability studies, Burke and Harman⁷ set $\alpha = 7.2$ and $\beta = 0.8$ in one of their reliability studies, and Martz et al.⁸ set $\alpha = 27.3$ and $\beta = 0.5$. When priors are informed by simulation results or previous experiments, analysts sometimes set

$$\alpha = n_{\rm pr} \cdot \hat{\theta}^{\rm pr} + 1$$
 and $\beta = n_{\rm pr} \cdot (1 - \hat{\theta}^{\rm pr}) + 1$,

where $\hat{\theta}^{pr}$ is a prior estimate of θ , and n_{pr} is some positive number that represents the confidence the analyst has in the simulation or experiment informing the prior;⁹ think of n_{pr} as the effective sample size that informs the prior. The greater this effective sample size that informs the prior, the more peaked the prior distribution is near the prior estimate of θ . If no confidence exists in the simulation informing the prior, then $n_{pr} = 0$ and the prior would be flat.

Figures 1 and 2 illustrate what these prior distributions look like. The prior in Figure 1 puts large probability on low values of θ and does so by setting the value of α to be significantly less than the value of β . The prior in Figure 2 puts large probability on high values of θ and does so by setting the value of β to be smaller than the value of α . Also observe that the prior is much more peaked for small values of θ in Figure 1 than it is for large values of θ in Figure 2. This is a consequence of the difference in the values between α and β . The difference is larger for the prior in Figure 1 than it is for the prior in Figure 2.

Let us now assume that we observe x successes of the system out of n tests conducted. In this case, the likelihood is binomial,

$$p(x|\theta) = \binom{n}{x} \theta^{x} (1-\theta)^{n-x}$$

making the posterior distribution of θ

$$\pi(\theta|x) = \frac{p(x|\theta)\pi(\theta)}{\int_{\Theta} p(x|\theta)\pi(\theta)d\theta} = \frac{\binom{n}{x}\theta^{x}(1-\theta)^{n-x}\frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}\theta^{\alpha-1}(1-\theta)^{\beta-1}}{\int_{0}^{1}\left[\binom{n}{x}\theta^{x}(1-\theta)^{n-x}\frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)}\theta^{\alpha-1}(1-\theta)^{\beta-1}\right]d\theta}.$$
(4)

There is a simple way to calculate the posterior distribution in Eq. 4 without having to evaluate the integral in the denominator. First observe that the expression in the denominator is not a function of θ ; it is a normalizing constant independent of θ , and for this reason we can write

$$\pi(\theta|x) \propto p(x|\theta) \pi(\theta)$$

Eliminating all multiplicative constants in $p(x|\theta) \pi(\theta)$ that do not depend on θ , we get that

$$\pi(\theta|x) = c \cdot \theta^{x+\alpha-1} (1-\theta)^{n-x+\beta-1}$$





Figure 1. The prior $\pi(\theta)$ with $\alpha = 2$ and $\beta = 10$. With this selection of α and β , the prior is peaked at low values of θ .

Figure 2. The prior $\pi(\theta)$ with $\alpha = 7$ and $\beta = 3$. With this selection of α and β , the prior is peaked at high values of θ .

Johns Hopkins APL Technical Digest, Volume 38, Number 1 (2025), Article 2501649 www.jhuapl.edu/technical-digest C. H. Botts





Figure 3. The posterior distribution with n = 10, x = 1, $\alpha = 2$, and $\beta = 10$.

Figure 4. The posterior distribution with n = 10, x = 2, $\alpha = 7$, and $\beta = 3$.

where *c* is some constant such that

$$\int_0^1 c \cdot \theta^{x+\alpha-1} (1-\theta)^{n-x+\beta-1} d\theta = 1$$

The posterior $\pi(\theta | x)$ takes the form of a beta distribution, making

$$c = \frac{\Gamma(\alpha + \beta + n)}{\Gamma(\alpha + x)\Gamma(\beta + n - x)}$$

The posterior of θ is thus a beta distribution with parameters α^{pst} and β^{pst} where

$$\alpha^{\text{pst}} = \alpha + x$$
, and
 $\beta^{\text{pst}} = \beta + n - x$.

The plots in Figures 3 and 4 show the posteriors corresponding to the priors shown in Figures 1 and 2, respectively. In Figure 3, one success was observed after 10 trials, emphasizing even more that the value of θ is small. Observe how the posterior in this case is more peaked at small values of θ than the prior was. In Figure 4, two successes were observed in 10 trials, indicating that the probability of success was much smaller than the prior anticipated. Observe how, in this case, the peak of the posterior has significantly shifted to smaller values of θ .

The Prior and Posterior of the Entire System

Let us now put this problem in the context of one large system that is composed of several subsystems. If all the subsystems have to work for the entire system to work, how do the posterior distributions of the subsystem reliabilities inform the distribution of the total system's reliability? And how would testing the entire system (as a whole) affect the posterior of the total system reliability? The next two subsections address these questions.

Subsystem Test Sizing

This subsection focuses on how the distributions of the subsystem survival probabilities affect the distribution of the total system's survival probability. If there are S subsystems, and the entire system fails if any one of the subsystems fail (i.e., the subsystems work in series), then the success probability of the total system, $\theta_{Tot Sys}$, is calculated as

$$\begin{aligned}
\theta_{\text{Tot Sys}} &= \mathbb{P}(\text{Success of Subsys 1}) \times \mathbb{P}(\text{Success of Subsys 2}) \times \cdots \\
\times \mathbb{P}(\text{Success of Subsys S}) \\
&= \prod_{j=1}^{S} \theta_j,
\end{aligned}$$
(5)



Figure 5. Flowchart of system composed of five subsystems (Subs), two of which work in parallel.

where θ_i is the success probability of the *j*th subsystem.

If S = 5 and the block diagram of the subsystems is as shown in Figure 5, then the system fails if subsystem 1, 4, 5, or both 2 and 3 fail. In this case, the success probability of the entire system would be calculated as

 $\begin{array}{ll} \theta_{\text{Tot Sys}} &= & \mathbb{P}\left(\text{Success of Subsys 1}\right) \times \mathbb{P}\left(\text{Success of Subsys 2 or 3}\right) \times \mathbb{P}\left(\text{Success of Subsys 4}\right) \times \mathbb{P}\left(\text{Success of Subsys 5}\right) \\ &= & \theta_1\left(\theta_2 + \theta_3 - \theta_2\theta_3\right)\theta_4\theta_5. \end{array}$

The value of $\theta_{\text{Tot Sys}}$ is thus the product and/or sum of beta random variables. The resulting distribution of a random variable such as $\theta_{\text{Tot Sys}}$ has been derived in a number of publications, but this distribution is very complicated and thus difficult to work with analytically.^{1,10,11,12,13} The distribution of $\theta_{\text{Tot Sys}}$ is easy to work with and understand, however, using Monte Carlo methods. Since the posterior distribution of all the components of the system take the form of a beta distribution with known parameters, assuming independence of the subsystems, we can easily simulate n_{Sim} values from the prior distribution of $\theta_{\text{Tot Sys}}$. This requires simulating n_{Sim} S-tuples of (θ_1 , θ_2 , θ_3 ,..., θ_S). With each simulated S-tuple, we can calculate a value of $\theta_{\text{Tot Sys}}$. The algorithm for generating n_{Sim} values of $\theta_{\text{Tot Sys}}$ for a system in series is given in Procedure 1; that for a system as shown in Figure 5 is given in Procedure 2.

Figures 6-12 illustrate how the posterior distributions of subsystem reliability affect the distribution of $\theta_{\text{Tot Sys}}$. In the simulations performed, we assumed that the entire system was composed of three subsystems (S = 3) and

Procedure 1: Simulating n_{Sim} values of $\theta_{\text{Tot Sys}}$ when the subsystems work in series

 $\begin{array}{c} \text{input} : \pi(\theta_1|x_1), \pi(\theta_2|x_2), \dots, \pi(\theta_S|x_S), \text{ and } n_{\text{Sim}}, \text{ where } x_i \text{ is the number of successes of subsystem } i \\ \text{output: } \theta_{\text{Tot Sys}}^{(1)}, \theta_{\text{Tot Sys}}^{(2)}, \dots, \theta_{\text{Tot Sys}}^{(n_{\text{Sim}})} \\ \text{for } i \leftarrow 1 \text{ to } n_{\text{Sim}} \text{ do} \\ \\ \theta_{\text{Tot Sys}}^{(i)} \leftarrow 1 \\ \text{for } j \leftarrow 1 \text{ to } S \text{ do} \\ \\ \\ \theta_{\text{Tot Sys}}^{(i)} \leftarrow \theta_{\text{fot Sys}}^{(i)} \sim \pi(\theta_j|x_j) \\ \\ \theta_{\text{Tot Sys}}^{(i)} \leftarrow \theta_{\text{Tot Sys}}^{(i)} \leftarrow \theta_{j}^{(i)}. \end{array}$

Procedure 2: Simulating n_{Sim} values of $\theta_{\text{Tot Svs}}$ when the five subsystems work as shown in Figure 14 (series and parallel)

C. H. Botts



Figure 6. The prior and posterior of the first subsystem with $\alpha = 5$ and $\beta = 2$.

that these subsystems worked in series. The priors of the three subsystems are shown in black in Figures 6–8. The subsystems are then tested with $n_1 = 2$, $n_2 = 5$, and $n_3 = 4$, where n_j is the number of times the *j*th subsystem is tested. The resulting posteriors are shown in red in Figures 6–8, and 10,000 draws from the resulting distribution of $\theta_{\text{Tot Sys}}$ are shown in Figure 9. The subsystems were also tested at $n_1 = 11$, $n_2 = 14$, and $n_3 = 12$, and the corresponding posteriors are shown in blue in Figures 6–8. Observe that these posteriors are more peaked (more informed) than the others since the sample sizes are larger. The distribution of $\theta_{\text{Tot Sys}}$ corresponding to these larger sample sizes is shown in Figure 10. Observe how the variance of this posterior is smaller than that shown in Figure 9; this is because the subsystem sample sizes are larger.



Figure 7. The prior and posterior of the second subsystem with $\alpha = 3$ and $\beta = 2$.



Figure 8. The prior and posterior of the third subsystem with $\alpha = 2$ and $\beta = 2$.

Total System Test Sizing

Let us now investigate how testing the entire system (and not just its individual components) affects the posterior distribution of $\theta_{Tot Sys}$. Updating the prior distribution of $\theta_{Tot Sys}$ given test results on the total system is more challenging than updating the subcomponent values of θ because, in this case, the original distribution of $\theta_{Tot Sys}$ is not a beta distribution. Recall that the prior distribution of $\theta_{Tot Sys}$ was analytically challenging to work with and, as a result, was obtained using Monte Carlo methods. It is not uncommon for practitioners to approximate this prior with another (perhaps beta) distribution^{1,11,14,15,16,17} to make the posterior analysis simpler and more convenient. Others redefine the priors of the independent components entirely just so the prior of the total system's reliability is analytically tractable.



Figure 9. The resulting distribution of $\theta_{\text{Tot Sys}}$ when $n_1 = 2$, $n_2 = 5$, and $n_3 = 4$.

Johns Hopkins APL Technical Digest, Volume 38, Number 1 (2025), Article 2501649 www.jhuapl.edu/technical-digest



Figure 10. The resulting distribution of $\theta_{\text{Tot Sys}}$ when $n_1 = 11$, $n_2 = 14$, and $n_3 = 12$.

Zoh et al.,¹⁸ for example, set the priors of the components to negative log-gamma distributions, yet they still had to execute complicated Markov chain Monte Carlo methods to do posterior analysis of the total system.

These workarounds and approximations to make posterior analysis easier are not necessary. This article describes a simple and quick Monte Carlo method that updates the prior distribution of $\theta_{\text{Tot Sys}}$ after testing the entire system. This method is a simple application of the algorithm proposed by Rubin.¹⁹ Rubin observed that a sample from the posterior distribution of a parameter can be obtained by first generating values from its prior and then generating data conditioned on these sampled values. Those values of the parameter for which the generated data match the observed data follow the posterior distribution.



Figure 11. The resulting distribution of $\theta_{\text{Tot Sys}}$ when $n_1 = 2$, $n_2 = 5$, $n_3 = 4$, and $n_{\text{TS}} = 4$.



Figure 12. The resulting distribution of $\theta_{\text{Tot Sys}}$ when $n_1 = 11$, $n_2 = 14$, $n_3 = 12$, and $n_{\text{Tot}} = 7$.

To apply this algorithm in our case, we begin by writing the posterior for $\theta_{Tot\;Svs}$ as

$$\pi \left(\theta_{\text{Tot Sys}} | x_{\text{TS}} \right) \propto p \left(x_{\text{TS}} | \theta_{\text{Tot Sys}} \right) \pi \left(\theta_{\text{Tot Sys}} \right)$$

where

$$p\left(x_{\text{TS}}|\boldsymbol{\theta}_{\text{Tot Sys}}\right) = \binom{n_{\text{TS}}}{x_{\text{TS}}} \boldsymbol{\theta}_{\text{Tot Sys}}^{x_{\text{TS}}} \left(1 - \boldsymbol{\theta}_{\text{Tot Sys}}\right)^{n_{\text{TS}}-x_{\text{TS}}}, (6)$$

and $n_{\rm TS}$ and $x_{\rm TS}$ are the number of tests (and successes) of the total system. Given that $x_{\rm TS}^*$ successes have been observed from $n_{\rm TS}$ trials of the total system, we sample from the posterior $\pi(\theta_{\rm Tot Sys}|x_{\rm TS} = x_{\rm TS}^*)$ by first simulating from the prior of $\theta_{\rm Tot Sys}$, $\pi(\theta_{\rm Tot Sys})$. (An algorithm similar to the ones shown in Procedure 1 or 2 could be used to sample from the prior, $\pi(\theta_{\rm Tot Sys})$.) We then condition on these sampled values of $\theta_{\rm Tot Sys}$ to generate candidate values of $x_{\rm TS}$ from the likelihood shown in Eq. 6. The simulated values of $\theta_{\rm Tot Sys}$ for which the likelihood generates $x_{\rm TS} = x_{\rm TS}^*$ are then considered to be an exact sample from the posterior. The details of this algorithm (assuming the subsystems work in series; minor changes to the first "for" loop of the algorithm would be necessary if the subsystems did not work in series) are given in Procedure 3.

The plots in Figures 11 and 12 show how the distribution of $\theta_{\text{Tot Sys}}$ changes when tests on the entire system are executed. The plot in Figure 11 shows how the distribution of $\theta_{\text{Tot Sys}}$ changes from the distribution in Figure 9 when four successes are observed out of four tests on the entire system. Observe that with this extra evidence of success, the distribution of $\theta_{\text{Tot Sys}}$ shifts to the right. The same story is told in Figure 12. It shows how the distribution of $\theta_{\text{Tot Sys}}$ changes from the distribution in Figure 10 when five successes are observed out of seven tests on the entire system. This distribution moves to the right and is also more peaked.

Johns Hopkins APL Technical Digest, Volume 38, Number 1 (2025), Article 2501649 www.jhuapl.edu/technical-digest

Procedure 3: Simulating n_{Sim} values of $\theta_{\text{Tot Sys}}$ from $\pi \left(\theta_{\text{Tot Sys}} | x_{\text{TS}} = x_{\text{TS}}^* \right)$

 $\begin{array}{l} \text{input} : \pi(\theta_1|x_1), \pi(\theta_2|x_2), \dots, \pi(\theta_S|x_S), n_{\text{Sim}}, n_{\text{TS}}, x_{\text{TS}}^* \\ \text{output:} \ \theta_{\text{Tot Sys}}^{(1)}, \theta_{\text{Tot Sys}}^{(2)}, \dots, \theta_{\text{Tot Sys}}^{(n_{\text{Sim}})} \\ i \leftarrow 1 \\ \text{while } i < n_{\text{Sim}} \text{ do} \\ \left| \begin{array}{c} \theta_{\text{Tot Sys}}^{(i), \text{ cand}} \leftarrow 1 \\ \text{for } j \leftarrow 1 \text{ to } S \text{ do} \\ \\ \left| \begin{array}{c} \text{Generate } \theta_j^{(i)} \sim \pi(\theta_j|x_j) \\ \theta_{\text{Tot Sys}}^{(i), \text{ cand}} \leftarrow \theta_{\text{Tot Sys}}^{(i), \text{ cand}} + \theta_{\text{Tot Sys}}^{(i)} \\ \end{array} \right| \\ \left| \begin{array}{c} \text{Generate } x_{\text{TS}} \sim \text{Binomial} \left(n_{\text{TS}}, \theta_{\text{Tot Sys}}^{(i), \text{ cand}} \right) \\ \text{if } x_{\text{TS}} = x_{\text{TS}}^* \text{ then} \\ \\ \left| \begin{array}{c} \theta_{\text{Tot Sys}}^{(i)} \leftarrow \theta_{\text{Tot Sys}}^{(i), \text{ cand}} \\ \text{i} \leftarrow i + 1 \end{array} \right| \\ \end{array} \right| \\ \end{array} \right|$

It is critical to understand the distinction between the two examples discussed above. In the first case (with the resulting distribution of $\theta_{\rm Tot~Sys}$ shown in Figure 11), the total system was tested four times and four successes were observed. Because the system works in series, a successful test of the entire system implies a successful test of each component. The posterior distribution of $\theta_{\rm Tot\;Sys}$ can thus easily be calculated by simply updating the posteriors of the system's three components and then applying Procedure 1. This is not true for the second example (with the resulting distribution of $\theta_{\text{Tot Sys}}$ shown in Figure 12). Recall that in the second example, the entire system was tested seven times, but only five successes were observed. Since it is not clear which component(s) failed (causing the failure of the entire system), the posteriors of the components cannot be updated, and Procedure 1 cannot be applied. In this case, Procedure 3 is necessary in calculating the posterior of $\theta_{\text{Tot Sys}}$.

EXAMPLE

This example illustrates the efficiency of Rubin's algorithm when calculating total system reliability. Consider two types of systems/fault trees, each with m different types of components. The first fault tree works in series, and in the second, every other component operates in series. Figures 13 and 14 illustrate these fault trees.

The efficiency of Rubin's algorithm is demonstrated with a Monte Carlo study. We initially place Beta(α,β) priors on all the components with $\alpha = 999$ and $\beta = 0.5$. (Such priors assume a prior reliability mean of 0.9995, which is not uncommon for highly reliable systems.) For each system/fault tree, and for a specific value of m, n_j (the number of trials for subsystem j) and x_j (the number of successful trials for subsystem j) can then be simulated. The number of trials is simulated from a Poisson($\lambda = 5$) distribution, and the number of successful trials is simulated from a binomial distribution with parameters n_i









Table 1. Ave	erage time	(in second	s) to	generate	10,000
values from	the posteri	or of the tota	ıl syst	em's reliab	ility

Fault Tree 1	Fault Tree 2
0.16	0.16
0.30	0.31
0.45	0.46
0.62	0.62
0.76	0.77
	Fault Tree 1 0.16 0.30 0.45 0.62 0.76

and $p_j = \alpha/(\alpha + \beta) = 999/(999 + 0.5)$. With these simulated test sizes, and assuming that the entire system is tested twice (with both tests being a success), the time it took to generate 10,000 draws from the posterior of the total system's reliability can be calculated. Table 1 shows the results of doing this 100 times and calculating the average time it took to generate these posterior values for m = 10, 20, 30, 40, and 50.

It is clear from Table 1 that this algorithm efficiently calculates the posterior of the total system's reliability.

CONCLUSION

This article reviews some of the methodologies related to Bayesian reliability. It initially focuses on success/failure data of systems and their subsystems. The article specifically addresses how the number of subsystem (or total system) tests affects the reliability of the entire system. It also presents a simple and efficient Monte Carlo method that can be employed to update the prior of a total system's reliability when only data from the total system are available.

REFERENCES

- ¹ H. F. Martz, R. A. Waller, and E. T. Fickas, "Bayesian reliability analysis of series systems of binomial subsystems and components," *Technometrics*, vol. 30, no. 2, pp. 143–154, 1988, https://doi.org/10.1080/0040 1706.1988.10488361.
- ² H. Cao, C. Yao, and Y. Yuan, "Bayesian approach for design and analysis of medical device trials in the era of modern clinical studies," *Med. Rev.*, vol. 3, no. 5, 2023, pp. 408–424, https://doi.org/10.1515/ mr-2023-0026.
- ³ T. R. Bayes, "An essay towards solving a problem in the doctrine of chances," *Philos. Trans. Roy. Soc. London*, vol. 53, pp. 370–418, 1763, https://doi.org/10.1098/rstl.1763.0053.
- ⁴ A. Gelman, J. Carlin, H. Stern, D. Dunson, A. Vehtari, and D. Rubin, Bayesian Data Analysis, 3rd ed. New York: Chapman and Hall/CRC, 2013.
- ⁵ B. Carlin and T. Louis, Bayes and Empirical Bayes Methods for Data Analysis, 3rd ed. New York: Chapman and Hall, 2008.
- ⁶ L. Leoni, F. B. Toroody, S. Khalaj, F. De Carlo, A. B. Toroody, and M. Abaei, "Bayesian estimation for reliability engineering: Addressing the influence of prior choice," *Int. J. Environ. Res. Public Health*, vol. 18, no. 7, art. 3349, 2021, https://doi.org/10.3390/ijerph18073349.

- ⁷ S. Burke and M. Harman, "Bayesian reliability for complex systems," Scientific Test and Analysis Techniques Center of Excellence (STAT COE) Report, STAT Center of Excellence, Wright-Patterson Air Force Base, OH Aug. 31 2019, https://www.afit.edu/stat/statcoe_files/ Bayesian%20Reliability%20for%20Complex%20Systems1.pdf.
- ⁸ H. F. Martz, R. A. Wailer, and E. T. Fickas, "Bayesian reliability analysis of series systems of binomial subsystems and components," *Technometrics*, vol. 30, no. 22, pp. 143–154, 1988, https://doi.org/10.1080/004 01706.1988.10488361.
- ⁹ V. E. Johnson, T. Graves, M. S. Hamada, and C. S. Reese, "A hierarchical model for estimating the reliability of complex systems," in *Bayesian Statistics 7: Proc. Seventh Valencia Int. Meeting*, eds. J. Bernardo, M. Bayarri, J. Berger, A. David, D. Heckermann, A. Smith, and M. West, 2003, pp. 199–213.
- ¹⁰D. Fan, "The distribution of the product of independent beta variables," Commun. Stat. Theory Methods, vol. 20, no. 12, pp. 4043–4052, 1991, https://doi.org/10.1080/03610929108830755.
- ¹¹H. F. Martz, R. A. Waller, and E. T. Fickas, "Bayesian reliability analysis of complex series/parallel systems of binomial subsystems and components," *Technometrics*, vol. 32, no. 4, pp. 407–416, 1990, https://doi. org/10.1080/00401706.1990.10484727.
- ¹²M. D. Springer and W. E. Thompson, "The distribution of products of beta, gamma, and Gaussian random variables," SIAM J. Appl. Math., vol. 18, no. 4, pp. 721–737, 1970, https://doi.org/10.1137/0118065.
- ¹³J. Tang and A. K. Gupta, "On the distribution of the product of independent beta random variables," *Statist. Probability Lett.*, vol. 2, no. 3, pp. 165–168, 1984, https://doi.org/10.1016/0167-7152(84)90008-7.
- pp. 165–168, 1984, https://doi.org/10.1016/0167-7152(84)90008-7.
 ¹⁴A. R. Abdel-Wahid and A. Winterbottom, "The approximation of system reliability posterior distributions," *J. Statist. Planning Inference*, vol. 16, pp. 267–275, 1987, https://doi.org/10.1016/0378-3758(87)90079-6.
- ¹⁵C. Coelho, "The generalized near-integer gamma distribution: A basis for 'near-exact' approximations to the distribution of statistics which are the product of an odd number of independent beta random variables," *J. Multivariate Anal.*, vol. 89, no. 2, pp. 191–218, 2004, https:// doi.org/10.1016/j.jmva.2003.12.001.
- ¹⁶F. Marques, I. Ghosh, J. Ferreira, and A. Bekker, "A note on the product of independent beta random variables," in Advances in Statistics— Theory and Applications. Emerging Topics in Statistics and Biostatistics, eds. I. Ghosh, N. Balakrishnan, H. Ng, 2021, Springer, pp. 69–83.
- ¹⁷J. Tukey and S. Wilks, "Approximation of the distribution of the product of beta variables by a single beta variable," Ann. Math. Statist., vol. 17, no. 3, pp. 318–324. 1946, https://doi.org/10.1214/aoms/1177730944.
- ¹⁸R. Zoh, A. Wilson, S. Vander Wiel, and E. Lawrence, "The negative log-gamma prior distribution for Bayesian assessment of system reliability," in *Proc. Institution Mech. Eng., Part O: J. Risk Rel.*, vol. 232, no. 3, pp. 308–319, 2018, https://doi.org/10.1177/1748006X17692154.
- ¹⁹D. B. Rubin, "Bayesianly justifiable and relevant frequency calculations for the applied statistician," Ann. Statist., vol. 12, no. 4, pp. 1151–1172, 1984, https://doi.org/10.1214/aos/1176346785.



Carsten H. Botts, Air and Missile Defense Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Carsten H. Botts is an analyst and statistician in APL's Air and Missile Defense Sector. He holds a BS in chemistry from Georgetown Univer-

sity and a PhD, in statistics from Iowa State University. His work centers on analyzing and modeling weapon accuracy using modern computational and statistical methods. His email address is carsten.botts@jhuapl.edu.