

ACCOUNTING FOR PARAMETER UNCERTAINTY IN SIMULATION INPUT MODELING

Faker Zouaoui

Sabre, Inc.
Research Group
Southlake, TX 76092, U.S.A.

James R. Wilson

Department of Industrial Engineering
North Carolina State University
Raleigh, NC 27695-7906, U.S.A.

ABSTRACT

We formulate and evaluate a Bayesian approach to probabilistic input modeling. Taking into account the parameter and stochastic uncertainties inherent in most simulations, this approach yields valid predictive inferences about the output quantities of interest. We use prior information to construct prior distributions on the input-model parameters. Combining this prior information with the likelihood function of sample data observed on the input processes, we compute the posterior parameter distributions using Bayes' rule. This leads to a Bayesian Simulation Replication Algorithm in which: (a) we estimate the parameter uncertainty by sampling from the posterior distribution of the input model's parameters on selected simulation runs; and (b) we estimate the stochastic uncertainty by multiple independent replications of those selected runs. We also formulate some performance evaluation criteria that are reasonable within both the Bayesian and frequentist paradigms. An experimental performance evaluation demonstrates the advantages of the Bayesian approach versus conventional frequentist techniques.

1 INTRODUCTION

Discrete-event simulations, especially those modeling complex systems, are almost all driven by random input processes. A simulation experiment therefore typically requires a number of streams of random variates drawn from specified distributions or input models. The inherent variation in the output of a simulation experiment arising from its dependence on these random inputs is often called *stochastic uncertainty* (Helton 1998). We generally assume that the input models driving the simulation belong to known parametric families. However, uncertainty typically occurs when choosing between different input models. We refer to this second source of variation as *model uncertainty* (Raftery et al. 1995). The parameters on which these models depend are usually assumed to be fixed. In practice,

these parameters are estimated from subjective information (expert opinion) or from real data observed on the input processes. The estimation of unknown parameters gives rise to another source of variation often referred to in the literature as *parameter uncertainty* (Raftery et al. 1995).

Cheng and Holland (1997) consider two methods of assessing how the variation in the simulation output depends on parameter uncertainty and stochastic uncertainty. The first method is based on classical differential analysis, or the δ -method (Stuart and Ord 1994). The main result of Cheng and Holland (1997) shows that under general conditions the total variation in the simulation output is composed of two distinct terms, depending respectively on the parameter uncertainty and stochastic uncertainty. One problem with this method is that certain sensitivity coefficients have to be estimated, and the effort needed to do this increases linearly with the number of unknown input parameters. Moreover, when the number of parameters is large, a problem can occur with spurious variation overinflating the variance estimate. The second method that Cheng and Holland (1997) consider for assessing the variation in the simulation output is the parametric bootstrap (Efron and Tibshirani 1993). Although computationally more expensive, this method does not suffer from the difficulties of the δ -method. It can also be more competitive on the grounds of computational efficiency if the number of unknown parameters is large.

Both of the above methods rely on the assumption that the parameters of the input models are unknown but deterministic quantities. Moreover, the output inferences are implicitly conditional on the selected single input model. The objective of the simulation experiment is therefore to estimate the mean output response as a function of the "true" but unknown parameter values. The most fundamental problem with such approaches to input model selection is that conditional on a single input model and on given values of the parameters for that input model, the output inference underestimates the overall (unconditional) uncertainty in the output quantities of interest, sometimes to a dramatic extent (Kass and Raftery 1995). Moreover, the usual approach to

model selection in the simulation community is commonly guided by a series of goodness-of-fit tests (Law and Kelton 2000). These tests can be highly misleading and very difficult to interpret in a classical statistical framework (Berger and Delampady 1987).

All these difficulties can be avoided, if one adopts a Bayesian approach that incorporates prior information on competing models and their parameters in a rigorous manner. We can compute the posterior probabilities using Bayes' rule for all competing models and their parameters; and then we can make a composite inference that takes account of model and parameter uncertainty in a formally justifiable way (Zouaoui 2001). Even if prior information is not readily available, there are methods to perform a full Bayesian analysis that rely on noninformative priors and thus will give more weight to the observed data, but will still incorporate the model and parameter uncertainties that are due to our lack of knowledge about the input processes driving our simulation experiment. These methods generalize the classical inferences conditional on the choice of a single input model and its parameters, and they work for both small and large sample sizes.

In this paper, we use a Bayesian approach to account for both stochastic uncertainty and parameter uncertainty in simulation; and we estimate the effects of these sources of uncertainty on the output quantity of interest. Hence, we assume that the functional forms of all input models are known, perhaps based on our prior knowledge of the processes driving the simulation model—a situation that sometimes occurs in simulation applications. In Zouaoui and Wilson (2001b), we extend our approach to account for model uncertainty as well as parameter uncertainty and stochastic uncertainty.

The rest of this paper is organized as follows. In Section 2 we define the symbolism used to describe our layout for probabilistic simulation experiments. In Section 3 we survey some recent developments concerning the use of the bootstrap method for classical frequentist simulation input modeling. In Section 4 we detail our Bayesian framework for handling parameter and stochastic uncertainties in discrete-event simulation; and this development leads to our “Bayesian Simulation Replication Algorithm” for designing simulation experiments. To evaluate the performance of both Bayesian and frequentist input modeling techniques on a fair and consistent basis, in Section 5 we formulate appropriate evaluation criteria and illustrate the application of these criteria to a simulation of a single-server queue. For more complete details on the development presented in this paper, see Zouaoui and Wilson (2001a).

2 THE SIMULATION EXPERIMENT

A simulation experiment, in its basic form, consists of making m independent runs of the simulation model and

observing a single output performance measure of interest, y , from each run. For simplicity, we assume that the simulation model is driven by a single sequence $\{X_1, X_2, \dots\}$ of independent and identically distributed input random variables, from which we observe the random sample $\mathbf{x} = (x_1, \dots, x_n)$. Multiple independent random input sequences will be treated in a similar fashion.

During the j th simulation run, a stream of random numbers \mathbf{u}_j is generated internally within the simulation model. This stream is used to generate the input random variates $\tilde{\mathbf{x}}_j$ by some transformation method, from which the output y_j is computed. It is more convenient to think of both streams as finite for a fixed simulation run length.

One possible method for generating the input random variates $\tilde{\mathbf{x}}_j$, which is commonly used by the simulation software packages, is the inverse transform method. If we let u_{ji} denote the i th random number sampled on the j th simulation run, then \tilde{x}_{ji} can be generated using the inverse transform method as

$$\tilde{x}_{ji} = G_M^{-1}(u_{ji}, \theta_M), \tag{1}$$

where $G_M^{-1}(\cdot)$ is the inverse of the distribution function $G_M(\cdot, \theta_M)$ of the simulation input model M , having θ_M as its d_M -dimensional vector of parameters. Given θ_M , we shall assume in this paper that the conditional distributions of \tilde{X}_{ji} and X_i are the same. Hence, the real data observations are assumed to have been drawn from the distribution $G_M(x, \theta_M)$.

The model and parameter uncertainties are represented by the random variables M and θ_M , respectively, both of which are assumed to depend only on the subjective information or data observed on the target input processes; and the stochastic uncertainty depends only on the randomness of \mathbf{u} . Thus the output of interest from the simulation run, y , can be regarded as an unknown complicated function of \mathbf{u} , M , and θ_M ,

$$y = y(\mathbf{u}, M, \theta_M). \tag{2}$$

In this paper, we limit our attention to the effects of parameter uncertainty and stochastic uncertainty on the distribution of y . For simplicity we drop M from our subsequent expressions, recognizing that they are implicitly dependent on the input model M . (In Zouaoui and Wilson (2001b), we will relax the simplifying assumption that the input model M is known.) Thus equation (2) becomes

$$y = y(\mathbf{u}, \theta), \tag{3}$$

and we let

$$\eta(\theta) = \int y(\mathbf{u}, \theta) d\mathbf{u} \tag{4}$$

denote the expected value of y given θ .

The objective of a classical simulation experiment is generally to estimate $\eta(\theta_0)$, where θ_0 is the true but unknown parameter value, estimated separately from the simulation experiment using real data. It is also of interest to compute a measure of the variability of the simulation output, from which a confidence interval for $\eta(\theta_0)$ can be constructed.

3 CONVENTIONAL APPROACHES TO ESTIMATING PARAMETER AND STOCHASTIC UNCERTAINTY

From the structure of the simulation model described in the previous section, we see that the responses or outputs of the simulation runs for a fixed parameter vector $\theta \in \mathbb{R}^d$ can be written as

$$y_j = y(\mathbf{u}_j, \theta) = \eta(\theta) + e_j(\mathbf{u}_j, \theta), \quad j = 1, \dots, m.$$

The error variable e_j is the random difference between the output of the j th simulation run and $\eta(\theta)$. We generally assume that $E(e_j|\theta) = 0$ and $\text{Var}(e_j|\theta) = \tau^2(\theta)$ for $j = 1, \dots, m$, so that $E(y_j|\theta) = \eta(\theta)$. Hence the mean of the simulation outputs,

$$\bar{y} = \frac{1}{m} \sum_{j=1}^m y_j, \quad (5)$$

is an unbiased estimator of $\eta(\theta)$.

If the maximum likelihood estimator $\hat{\theta}$ of θ_0 is used in (3), then the output is $y_j = y(\mathbf{u}_j, \hat{\theta}) = \eta(\hat{\theta}) + e_j(\mathbf{u}_j, \hat{\theta})$ for $j = 1, \dots, m$, where $\hat{\theta}$ and \mathbf{u}_j are independent random vectors. We have the variance decomposition

$$\begin{aligned} \text{Var}(y_j) &= \text{Var}_{\hat{\theta}}\{E[y_j(\mathbf{u}_j, \hat{\theta})|\hat{\theta}]\} \\ &\quad + E_{\hat{\theta}}\{\text{Var}[y_j(\mathbf{u}_j, \hat{\theta})|\hat{\theta}]\} \\ &= V_{\text{par}} + V_{\text{sto}}, \end{aligned} \quad (6)$$

in which V_{par} denotes the parameter variance and V_{sto} denotes the stochastic or simulation variance.

It is generally easy to estimate the stochastic variance V_{sto} in (6). Based on m independent replications, the most commonly used estimator is

$$\hat{V}_{\text{sto}} = \frac{1}{m-1} \sum_{j=1}^m (y_j - \bar{y})^2, \quad (7)$$

where $y_j = y(\mathbf{u}_j, \hat{\theta})$ is the output of the j th replication, and \bar{y} is the average of all y_j 's and the estimate of the mean response. The parameter variance V_{par} in (6) is more

difficult to estimate. Cheng and Holland (1997) developed a finite-difference estimator of V_{par} .

In the rest of this section, we review a bootstrap method (Efron 1979) for estimating V_{par} . Following the notation of Section 2, we substitute the estimator $\hat{\theta}$ into $G(x, \theta)$ to obtain the fitted distribution $G(x, \hat{\theta})$. We then draw a sample $\mathbf{x}_1^* = (x_{11}^*, x_{12}^*, \dots, x_{1n}^*)$ from $G(x, \hat{\theta})$, perhaps via the inverse transform (1). Corresponding to this bootstrap sample is the bootstrap estimate $\hat{\theta}^*(1)$ of θ_0 , computed from \mathbf{x}_1^* in exactly the same way that $\hat{\theta}$ was computed from \mathbf{x} . Repeating this sampling-and-estimation operation independently B times yields the estimates $\{\hat{\theta}^*(i) : i = 1, \dots, B\}$. We can then carry out B bootstrap simulation experiments, one for each $\hat{\theta}^*(i)$, where each run using the input parameter vector $\hat{\theta}^*(i)$ has the same length as in the original experiment; and altogether m' simulation runs are performed using each $\hat{\theta}^*(i)$. For $i = 1, \dots, B$, this yields the i th set of simulation-generated output responses $\{y_{ij}^* : j = 1, \dots, m'\}$ from which we calculate the sample mean $\bar{y}_i^* = \sum_{j=1}^{m'} y_{ij}^*/m'$. Let $\bar{\bar{y}}^* = \sum_{i=1}^B \bar{y}_i^*/B$ denote the grand mean of the bootstrap sample means.

Cheng and Holland (1997) show how the variance of \bar{y}_i^* depends on both the stochastic variance and the parameter variance. When each bootstrap is an exact replica of the original experiment (i.e. $m' = m$), Cheng and Holland suggest using the sample variance of the $\{\bar{y}_i^*\}$,

$$S_B^2 = \frac{1}{B-1} \sum_{i=1}^B (\bar{y}_i^* - \bar{\bar{y}}^*)^2, \quad (8)$$

as an estimate for the variance of the original sample mean \bar{y} specified in (5). Using (7) as an estimate for V_{sto} and (8), we estimate the parameter variance by

$$\hat{V}_{\text{par}} = S_B^2 - \hat{V}_{\text{sto}}. \quad (9)$$

An alternative bootstrap-based approach to estimating V_{par} is to compute $\hat{V}_{\text{sto}}^*(i) = \sum_{j=1}^{m'} (y_{ij}^* - \bar{y}_i^*)^2 / (m' - 1)$ as an additional estimate of $\tau^2(\theta_0)$ for $i = 1, \dots, B$; and then we can use $\hat{V}_{\text{sto}}^* = [\hat{V}_{\text{sto}} + \sum_{i=1}^B \hat{V}_{\text{sto}}^*(i)] / (B + 1)$ in place of \hat{V}_{sto} in (9).

Finally, we discuss how we can construct a bootstrap confidence interval for $\eta(\theta_0)$. One possibility is to construct a classical t -type confidence interval using the bootstrap variance estimates given in this section. However, this interval generally performs poorly in practice. Efron and Tibshirani (1993) explain this phenomenon and suggest methods to overcome such a drawback. A confidence interval that generally behaves better for such cases is the bootstrap

percentile-type confidence interval. From the order statistics $\bar{y}_{(1)}^* \leq \bar{y}_{(2)}^* \leq \dots \leq \bar{y}_{(B)}^*$ of the $\{\bar{y}_i^* : i = 1, \dots, B\}$, we obtain the $100(1 - \alpha)\%$ percentile-type confidence interval

$$[\beta_L^*, \beta_U^*] \approx [\bar{y}_{(\lceil B\alpha/2 \rceil)}^*, \bar{y}_{(\lceil B(1-\alpha/2) \rceil)}^*]. \quad (10)$$

4 BAYESIAN APPROACH

We describe in this section the applicability of the Bayesian approach to our basic structure (3) of the discrete-event simulation model. We provide methods to estimate the posterior mean response and to construct a confidence interval for that quantity.

4.1 Estimating Mean Response

We observe a random sample $\mathbf{x} = (x_1, \dots, x_n)$ from our selected input model. Let θ be its d -dimensional vector of parameters with prior distribution $p(\theta)$. We assume that the hyperparameters of the prior distribution are either known or estimated using moment matching or some other empirical Bayes method (Carlin and Louis 1996). We are not being fully Bayesian here by stopping the hierarchy at the second stage. However, we argue here that in most simulation applications in operations research and industrial engineering (as opposed, for example to applications in econometrics), prior information will be generally vague if it exists at more than one level down the hierarchy.

For our basic simulation model structure, we derive the posterior mean response given \mathbf{x} .

Theorem 1 *If the simulation response y has the form (3), then*

$$E(y|\mathbf{x}) = \int \eta(\theta) p(\theta|\mathbf{x}) d\theta.$$

Zouaoui (2001) and Zouaoui and Wilson (2001a) contain proofs of all results presented in this paper.

Our first objective is to develop an approach that accounts fully for parameter and stochastic uncertainty and that can be extended easily to account for model uncertainty. This approach should also have a theoretical appeal, and most importantly it should perform well in practice. Figure 1 summarizes the Bayesian Simulation Replication Algorithm that we have developed to implement a Bayesian approach to simulation input modeling. The net effect of the algorithm is to account fully for the uncertainty in the parameters of the input model as well as to account fully for the usual stochastic uncertainty. The algorithm can be seen as an uncertainty decomposition algorithm. The inner loop will be used to generate estimates of the stochastic uncertainty, whereas the outer loop will estimate the parameter uncertainty. The next subsection gives a detailed explanation of how to estimate the stochastic and param-

```

for  $r = 1, \dots, R$ 
  generate the parameter vector  $\theta^r$  from  $p(\theta|\mathbf{x})$ 
  set the parameter vector  $\theta \leftarrow \theta^r$ 
  for  $j = 1, \dots, m$ 
    set the random-number input  $\mathbf{u} \leftarrow \mathbf{u}_j$ 
    perform the simulation run using  $\mathbf{u}$  and  $\theta$ 
    calculate the output response  $y_{rj} = y(\mathbf{u}, \theta)$ 
  end loop
  compute  $\bar{y}_r = \sum_{j=1}^m y_{rj}/m$ 
end loop
compute the grand mean  $\bar{\bar{y}} = \sum_{r=1}^R \bar{y}_r/R$ 
as an estimate for  $E(y|\mathbf{x})$ 
    
```

Figure 1: Bayesian Simulation Replication Algorithm

eter variances and construct a confidence interval on the posterior mean response.

4.2 Assessing Output Variability

We try to assess the variability of the simulation output based on simple response surface models, given the objective of estimating the mean response in our simulation study. The following analysis can be extended to more complicated models, but such an extension is beyond the scope of this paper. We develop two methods of estimating the parameter and stochastic variances in the simulation output.

4.2.1 Classical Output Analysis

Following our basic simulation structure of Section 2, we see that the output responses from the simulation runs performed by the Bayesian Simulation Replication Algorithm of Figure 1 are given by

$$y_{rj} = y(\mathbf{u}_j, \theta^r) = \eta(\theta^r) + e_j(\mathbf{u}_j, \theta^r), \quad (11)$$

for $r = 1, \dots, R$ and $j = 1, \dots, m$. We generally assume that

$$E(e_j|\theta^r) = 0 \quad \text{and} \quad \text{Var}(e_j|\theta^r) = \tau^2, \quad (12)$$

where τ^2 does not depend on θ^r . Given that our main objective is to estimate the overall mean response, we further assume that

$$\eta(\theta^r) = \beta + \delta_r(\theta^r), \quad (13)$$

where $\beta = \beta(\mathbf{x}) = E_{\theta^r}[\eta(\theta^r)] = \int \eta(\theta) p(\theta|\mathbf{x}) d\theta = E(y|\mathbf{x})$ from Theorem 1; and

$$E_{\theta^r}(\delta_r) = 0 \quad \text{and} \quad \text{Var}_{\theta^r}(\delta_r) = \sigma^2. \quad (14)$$

Based on these assumptions, the posterior variance can be written as the sum of two variances measuring the stochastic and parameter uncertainty, respectively.

Theorem 2 *If (11)–(14) hold, then*

$$\text{Var}(y|\mathbf{x}) = \tau^2 + \sigma^2.$$

The response surface model given by (11)–(14) is known in the statistical literature as the classical random-effects model (Rao 1997), where one estimates β , τ^2 , and σ^2 using the simulation-generated statistics specified in Figure 1 as follows: the posterior mean response is estimated by

$$\hat{\beta} = \bar{\bar{y}}; \quad (15)$$

the variance due to stochastic uncertainty is estimated by

$$\hat{\tau}^2 = \frac{\sum_{r=1}^R \sum_{j=1}^m (y_{rj} - \bar{y}_r)^2}{R(m-1)}; \quad (16)$$

and finally the variance due to parameter uncertainty is estimated by

$$\hat{\sigma}^2 = \frac{\sum_{r=1}^R (\bar{y}_r - \bar{\bar{y}})^2}{(R-1)} - \frac{\hat{\tau}^2}{m}. \quad (17)$$

In many applications, especially the ones where we are totally ignorant about the output performance measure, the above formulation delivers reasonable point estimates. However, there are two important drawbacks of using such an approach. The first drawback concerns the estimate $\hat{\sigma}^2$, which can be negative. The problem of a negative estimate for a variance component can be avoided by setting it equal to zero, but this creates new issues (Rao 1997). The main drawback comes from the fact that substituting point estimates for the above parameters ignores our real uncertainty about them. We suggest in §4.2.2 a full Bayesian treatment of the same model with noninformative priors based on normally distributed simulation output observations.

In addition to point estimates, we can also construct a confidence interval for β from the output of the Bayesian Simulation Replication Algorithm given in Figure 1. Similar to the bootstrap approach, we can construct an approximate $100(1 - \alpha)\%$ Bayesian percentile confidence interval for β

as

$$[\beta_L, \beta_U] \approx [\bar{y}_{(\lceil R\alpha/2 \rceil)}, \bar{y}_{(\lceil R(1-\alpha/2) \rceil)}], \quad (18)$$

where the quantities $\bar{y}_{(1)} \leq \bar{y}_{(2)} \leq \dots \leq \bar{y}_{(R)}$ denote the order statistics of the $\{\bar{y}_r : r = 1, \dots, R\}$ specified in Figure 1. These intervals generally perform better in practice than t -type confidence intervals (Efron and Tibshirani 1993).

4.2.2 Bayesian Output Analysis

Under the hierarchical normal model, we assume that the output data $\{y_{rj} : r = 1, \dots, R \text{ and } j = 1, \dots, m\}$ specified in Figure 1 are independently normally distributed, that is

$$y_{rj} | \mu_r, \tau \sim N(\mu_r, \tau^2). \quad (19)$$

For $r = 1, \dots, R$, we also assume that the parameter μ_r , which corresponds to $\eta(\theta^r)$ in equation (11), is also normally distributed:

$$\mu_r | \beta, \sigma \sim N(\beta, \sigma^2), \quad r = 1, \dots, R. \quad (20)$$

Note that although assumption (19) generally holds in practice because of the simulation output usually being an average of a large number of output random variables, assumption (20) may not be consistent with the form of the posterior distribution $p(\theta|\mathbf{x})$.

To complete the Bayesian formulation, we assume a noninformative prior distribution for (β, τ, σ) , with $\sigma > 0$ and $\tau > 0$; specifically we take $p(\beta, \tau, \sigma) \propto \tau^{-1}$. It follows from this assumption that the joint posterior density of all parameters satisfies the relation

$$p(\boldsymbol{\mu}, \beta, \tau, \sigma | \mathbf{y}) \propto \frac{1}{\tau} \prod_{r=1}^R p(\mu_r | \beta, \sigma) \prod_{r=1}^R \prod_{j=1}^m p(y_{rj} | \mu_r, \tau), \quad (21)$$

where $\boldsymbol{\mu} = (\mu_1, \dots, \mu_R)$ and $\mathbf{y} = \{y_{rj} : r = 1, \dots, R; j = 1, \dots, m\}$.

Many numerical methods such as conditional maximization can be used to obtain posterior point estimates for the parameters of interest. However, we are also interested to compute posterior confidence intervals for these parameters; and Markov Chain Monte Carlo (MCMC) methods are appropriate for such inferences. In fact, we can even obtain a large sample from the posterior distribution of each parameter from which we can estimate its density.

The idea behind MCMC (Gilks et al. 1996) is to simulate a random walk in the space of $(\boldsymbol{\mu}, \beta, \tau, \sigma)$ which converges to a stationary distribution that is the joint posterior distribution $p(\boldsymbol{\mu}, \beta, \tau, \sigma | \mathbf{y})$. The most widely used MCMC method is the *Gibbs Sampler* algorithm (Casella and George 1992). Figure 2 summarizes briefly the steps of the algorithm. For

our problem this algorithm can be easily implemented in any statistical software package, given its simplicity and the fact that we can generate variates easily from the following required conditional distributions (Gelman et al. 1995):

$$\mu_r | \beta, \tau, \sigma, \mathbf{y} \sim N\left(\frac{\frac{\beta}{\sigma^2} + \frac{m\bar{y}_r}{\tau^2}}{\frac{1}{\sigma^2} + \frac{m}{\tau^2}}, \frac{1}{\frac{1}{\sigma^2} + \frac{m}{\tau^2}}\right), \quad (22)$$

for $r = 1, \dots, R$;

$$\beta | \boldsymbol{\mu}, \tau, \sigma, \mathbf{y} \sim N\left(\frac{1}{R} \sum_{r=1}^R \mu_r, \frac{\sigma^2}{R}\right); \quad (23)$$

$$\tau^2 | \boldsymbol{\mu}, \beta, \sigma, \mathbf{y} \sim \text{IG}\left(\frac{mR}{2}, \frac{1}{2} \sum_{r=1}^R \sum_{j=1}^m (y_{rj} - \mu_r)^2\right); \quad (24)$$

and

$$\sigma^2 | \boldsymbol{\mu}, \beta, \tau, \mathbf{y} \sim \text{IG}\left(\frac{R-1}{2}, \frac{1}{2} \sum_{r=1}^R (\mu_r - \beta)^2\right), \quad (25)$$

where the density function of a random variable Z having an Inverse Gamma distribution $\text{IG}(\nu, \phi)$ with shape parameter ν and scale parameter ϕ is defined as

$$p(z|\nu, \phi) = \frac{\phi^\nu e^{-\phi/z}}{\Gamma(\nu)z^{\nu+1}}, \quad z > 0, \nu > 0, \phi > 0,$$

and where $\Gamma(\cdot)$ denotes the gamma function. The classical estimates of β , τ^2 , and σ^2 given by equations (15), (16), and (17), respectively, are good initial estimates with which to start the Gibbs sampler algorithm. If the initial estimate $\hat{\sigma}^2$ is negative, then we can set it to a small value, say 0.001. The output inference from the Gibbs sampler algorithm is based on $T = 50000$ iterations, with a warm-up period of $T^* = 5000$ iterations. These relatively large numbers of iterations are generally required for convergence because of the high autocorrelation between successively sampled values of some parameters (Spiegelhalter et al. 1996). The point estimates of β , τ^2 , and σ^2 are given in Figure 2. A $100(1 - \alpha)\%$ credible interval for β can be constructed from the output of the Gibbs sampler algorithm as

$$[\beta_L, \beta_U] \approx [\beta_{(\lceil(T-T^*)\alpha/2\rceil)}, \beta_{(\lceil(T-T^*)(1-\alpha/2)\rceil)}], \quad (26)$$

where the quantities $\beta_{(1)} \leq \beta_{(2)} \leq \dots \leq \beta_{(T-T^*)}$ denote the order statistics of the $\{\beta_t : t = T^* + 1, \dots, T\}$ generated by the Gibbs sampler algorithm of Figure 2. Credible intervals similar to (26) can also be constructed for τ^2 and σ^2 .

Set β_0 , τ_0^2 , and σ_0^2 to their classical estimates (15), (16), and (17), respectively.

For $t = 1, \dots, T^*, T^* + 1, \dots, T$, generate:

$\mu_{rt} \sim p(\mu_r | \beta_{t-1}, \tau_{t-1}, \sigma_{t-1}, \mathbf{y})$ as in (22),
for $r = 1, \dots, R$, and

take $\boldsymbol{\mu}_t = (\mu_{1t}, \dots, \mu_{Rt})$;

$\beta_t \sim p(\beta | \boldsymbol{\mu}_t, \tau_{t-1}, \sigma_{t-1}, \mathbf{y})$ as in (23);

$\tau_t^2 \sim p(\tau^2 | \boldsymbol{\mu}_t, \beta_t, \sigma_{t-1}, \mathbf{y})$ as in (24); and

$\sigma_t^2 \sim p(\sigma^2 | \boldsymbol{\mu}_t, \beta_t, \tau_t, \mathbf{y})$ as in (25).

end loop

Compute estimates for the parameters of interest:

$$\hat{\beta} = \sum_{t=T^*+1}^T \beta_t / (T - T^*)$$

$$\hat{\tau}^2 = \sum_{t=T^*+1}^T \tau_t^2 / (T - T^*)$$

$$\hat{\sigma}^2 = \sum_{t=T^*+1}^T \sigma_t^2 / (T - T^*)$$

Figure 2: Gibbs Sampler Algorithm

5 PERFORMANCE EVALUATION

In this section we propose some criteria for comparing Bayesian and frequentist approaches within the discrete-event simulation framework. We will test our ideas empirically by applying them to a queueing system. In the following discussion, we take the view that a frequentist is an analyst who seeks to estimate an unknown parameter based only on the model that has been adopted for the observable data, whereas a Bayesian is one who seeks to estimate the parameter by appropriately combining his or her prior intuition with the information content in the data.

5.1 Evaluation Criteria

5.1.1 Point Estimation

Suppose that a sample $\mathbf{x} = (x_1, x_2, \dots, x_n)$ of independent random variables is to be drawn; and given the value of a scalar parameter θ_0 , \mathbf{x} has the conditional probability density $g_{\theta_0}(\mathbf{x})$. We assume that the unknown parameter θ_0 is a random variable and refer to the distribution π_0 of θ_0 as the “true prior” distribution. Here it is important to distinguish between (a) the “true” parameter vector θ_0 of the inputs \mathbf{x} to the real system with response y_0 versus (b) the parameter θ of the input model for the simulation with response $y = y(\mathbf{u}, \theta)$, even when θ is sampled from the posterior distribution $p(\theta|\mathbf{x})$ as in the Simulation Replication Algorithm of Figure 1.

For a realization θ_0 obtained from π_0 (which will be unknown to the Bayesian and frequentist analysts), we as-

sume that we can compute the true output performance measure of the real system $\beta_0(\theta_0) = E(y_0|\theta_0)$ (which will also be unknown to both analysts). The analysts will then attempt to estimate $\beta = \beta(\mathbf{x}) = E(y|\mathbf{x})$ using their respective simulation procedures based solely on observing the experimental outcomes $\{x_1, x_2, \dots, x_n\}$ drawn from g_{θ_0} .

Our criterion for assessing the performance of a given estimator $\hat{\beta}$ of $\beta_0(\theta_0)$ will be the average risk \mathcal{R} of $\hat{\beta}$ relative to the true prior distribution π_0 of θ_0 ,

$$\begin{aligned} \mathcal{R}(\pi_0, \hat{\beta}) &= E_{\pi_0} \left(E_{g_{\theta_0}} \left\{ [\hat{\beta}(\mathbf{x}) - \beta_0(\theta_0)]^2 \right\} \right) \quad (27) \\ &= \int \int [\hat{\beta}(\mathbf{x}) - \beta_0(\theta_0)]^2 g_{\theta_0}(\mathbf{x}) \pi_0(\theta_0) d\mathbf{x} d\theta_0, \end{aligned}$$

where we take the simulation-based estimator $\hat{\beta}(\mathbf{x}) = \bar{y}$ when we adopt the Bayesian Simulation Replication Algorithm of Figure 1. When we adopt the frequentist approach, we compute the maximum likelihood estimator $\hat{\theta}^{\text{mle}} = \hat{\theta}^{\text{mle}}(\mathbf{x})$ based on the original input data, and then in (27) we take

$$\hat{\beta}(\mathbf{x}) = \frac{1}{m} \sum_{j=1}^m y(\mathbf{u}_j, \hat{\theta}^{\text{mle}}(\mathbf{x})) \quad (28)$$

based on a set of m independent simulation runs.

In the Bayesian approach we will use the ‘‘operational’’ prior π , chosen by the Bayesian analyst, to compute the posterior distribution $p(\theta|\mathbf{x})$ of θ . The operational prior will generally be noninformative or minimally informative. It is also important to stress that the true prior π_0 is entirely unknown to the Bayesian and frequentist analysts. The real contest will be to find the estimator that minimizes $\mathcal{R}(\pi_0, \hat{\beta})$.

The criterion we propose for assessing performance makes sense and is reasonable within the Bayesian and frequentist paradigms. Although the notion of a ‘‘true prior’’ seemingly conflicts with the ‘‘degree of belief’’ interpretation of prior distributions espoused by many Bayesians, it should be noted that the Bayesians’ degree of belief about θ is reflected in π , not π_0 , and that the idea of a true prior distribution makes perfect sense in a computer-assisted experiment where θ_0 is generated at random from π_0 , whose form is unknown to both analysts. If the true θ_0 is constant and π_0 is taken as degenerate at that constant, then the average risk reduces to the mean squared error (MSE), which is widely used by frequentists to judge estimators. In our framework, the average risk is precisely the squared error in estimating $\beta_0(\theta_0)$ by $\hat{\beta}$, averaged over all the randomness in the problem.

5.1.2 Interval Estimation

The Bayesian approach can be also used to develop interval estimation procedures (§4.2.2) to assess its performance compared to the classical and bootstrap methods. As in the previous subsection, our strategy will be to use noninformative priors. The main properties used to assess the performance of a Bayesian confidence interval for $\beta_0(\theta_0)$ are the average interval length and variance of the interval length as well as its coverage probability. We now examine the coverage of a Bayesian confidence interval. Adopting the notation

$$\mathcal{B}(\mathbf{x}) \equiv [\hat{\beta}_L(\mathbf{x}), \hat{\beta}_U(\mathbf{x})], \quad (29)$$

we define the coverage probability of the Bayesian confidence interval (29) as

$$\begin{aligned} \mathcal{C}(\beta_0(\theta_0)) &= E_{\pi_0} \left\{ E_{g_{\theta_0}} [I_{\{\beta_0(\theta_0) \in \mathcal{B}(\mathbf{x})\}}] \right\} \quad (30) \\ &= \int \int I_{\{\beta_0(\theta_0) \in \mathcal{B}(\mathbf{x})\}} g_{\theta_0}(\mathbf{x}) \pi_0(\theta_0) d\mathbf{x} d\theta_0, \end{aligned}$$

where $I_{\{\beta_0(\theta_0) \in \mathcal{B}(\mathbf{x})\}}$ denotes the indicator function of the event $\{\beta_0(\theta_0) \in \mathcal{B}(\mathbf{x})\}$ that the Bayesian confidence interval (29) contains $\beta_0(\theta_0)$.

The average length of the confidence interval (29) is

$$\begin{aligned} \mathcal{L} &= E_{\pi_0} \left\{ E_{g_{\theta_0}} [\hat{\beta}_U(\mathbf{x}) - \hat{\beta}_L(\mathbf{x})] \right\} \quad (31) \\ &= \int \int [\hat{\beta}_U(\mathbf{x}) - \hat{\beta}_L(\mathbf{x})] g_{\theta_0}(\mathbf{x}) \pi_0(\theta_0) d\mathbf{x} d\theta_0, \end{aligned}$$

and the variance of the confidence interval length is given by

$$\begin{aligned} \mathcal{V} &= E_{\pi_0} \left(E_{g_{\theta_0}} \left\{ [\hat{\beta}_U(\mathbf{x}) - \hat{\beta}_L(\mathbf{x}) - \mathcal{L}]^2 \right\} \right) \quad (32) \\ &= \int \int [\hat{\beta}_U(\mathbf{x}) - \hat{\beta}_L(\mathbf{x}) - \mathcal{L}]^2 g_{\theta_0}(\mathbf{x}) \pi_0(\theta_0) d\mathbf{x} d\theta_0. \end{aligned}$$

5.2 Experimental Design

We developed in the previous subsections point and interval estimation criteria to assess the performance of the frequentist and Bayesian approaches in a discrete-event simulation framework. However, criteria (27), (30), (31), and (32) will be generally hard to compute analytically in most of the applications and should be estimated from the output realizations. The performance evaluation of the Bayesian and frequentist approaches is generally conducted by Monte Carlo computer-assisted simulation experiments. The output of these experiments can be used to estimate our performance criteria. Figure 3 summarizes the protocol we used to

conduct N Monte Carlo experiments and deliver estimates for the average risk of each estimator and its standard error, the mean and the variance of the interval lengths, and their coverage probabilities.

5.3 Application to a Single Server Queue

To provide a practical illustration of the above performance evaluation framework, we consider a simple single server queueing system. For this example, we will just compare the performance of the Bayesian and bootstrap approaches since the computational effort is approximately the same. Moreover, the average risk for the classical and bootstrap method is the same since they both use the same estimator for the mean response. In Zouaoui and Wilson (2001a),

we consider a larger scale example with a degenerate true prior distribution, where the main objective is to compare the interval estimation performance of the Bayesian method with the classical and bootstrap results of Cheng and Holland (1997).

The single server system has customers enter the queue according to a Poisson process with rate λ , having a “true” Gamma prior distribution π_0 with shape parameter $\nu_1 = 50$ and scale parameter ϕ_1 that will take a range of values corresponding to increasing levels of prior traffic intensity. The customers receive an exponential service time with rate μ , having a “true” Gamma prior distribution with shape parameter $\nu_2 = 50$ and scale parameter $\phi_2 = 100$. The objective is to estimate the average time in the system. Monte Carlo experiments are based on data samples of size $n = 1000$, and are repeated $N = 500$ times.

For the Bayesian approach, we chose the operational prior distribution for the arrival rate λ to be noninformative having a density proportional to λ^{-1} . After observing a sample $\mathbf{x} = \{x_1, x_2, \dots, x_n\}$ of interarrival times, we compute the posterior density of λ using Bayes’ rule to obtain a Gamma density with shape parameter n and scale parameter $1/(\sum_{j=1}^n x_j)$. Similarly, choosing the same functional form for the prior distribution of the service rate μ produces a proper posterior Gamma density having a shape parameter n and a scale parameter $1/(\sum_{j=1}^n z_j)$, where $\mathbf{z} = \{z_1, z_2, \dots, z_n\}$ is the observed sample of service times.

Table 1 shows the results for the average risk in estimating the posterior mean response. It appears that the Bayesian point estimator averages a slightly smaller risk in all cases. As shown in Table 2, the main conclusion concerns the excellent performance of the Bayesian method in terms of interval estimation. The coverage probability is very close to the nominal coverage with shorter confidence intervals than their bootstrap counterparts. Although the system considered here is simple and these results may not be generalized for all systems, there is strong empirical evidence in favor of the Bayesian approach. The other examples detailed in Zouaoui and Wilson (2001a, b) further support this conclusion.

6 CONCLUSIONS

In this paper we have proposed a Bayesian approach that uses prior information and data observations for inferences on the distribution of a simulation-generated output response. We have derived an expression for the posterior mean, and we have developed a Bayesian Simulation Replication Algorithm to compute point and confidence interval estimators for the posterior mean response. We have also proposed two response surface models to decompose the posterior variance into two components measuring the parameter and stochastic uncertainty, respectively.

Choose the true prior distribution π_0
 For $i = 1, \dots, N$
 Sample $\theta_{0,i}$ from π_0
 Compute the true output measure $\beta_{0,i} = \beta_0(\theta_{0,i})$
 Sample \mathbf{x}_i from $g_{\theta_{0,i}}(\mathbf{x})$
 Compute the Bayes estimate $\hat{\beta}_i^{\text{Bayes}} = \bar{y}$
 using an operational prior π
 Compute the frequentist estimate $\hat{\beta}_i^{\text{Freq}}$
 as specified in (28) using $\hat{\theta}^{\text{mle}}$
 Compute a CI for $\beta_{0,i}$ using each approach
 end loop
 For each simulation-based approach, compute:
 the estimated average risk as

$$\hat{\mathcal{R}} = \frac{1}{N} \sum_{i=1}^N (\hat{\beta}_i - \beta_{0,i})^2$$
 with Monte Carlo standard error estimate

$$\widehat{\text{SE}}(\hat{\mathcal{R}}) = \left\{ \frac{1}{N(N-1)} \sum_{i=1}^N [(\hat{\beta}_i - \beta_{0,i})^2 - \hat{\mathcal{R}}]^2 \right\}^{\frac{1}{2}}$$
 the estimated coverage probability as

$$\hat{\mathcal{C}} = \frac{1}{N} \sum_{i=1}^N I_{\{\beta_{0,i} \in \mathcal{B}_i(\mathbf{x}_i)\}}$$
 the estimated mean of the CI length as

$$\hat{\mathcal{L}} = \frac{1}{N} \sum_{i=1}^N [\hat{\beta}_U(\mathbf{x}_i) - \hat{\beta}_L(\mathbf{x}_i)]$$
 the estimated variance of the CI length as

$$\hat{\mathcal{V}} = \frac{1}{N-1} \sum_{i=1}^N [\hat{\beta}_U(\mathbf{x}_i) - \hat{\beta}_L(\mathbf{x}_i) - \hat{\mathcal{L}}]^2$$

Figure 3: Monte Carlo Experimental Design

Given the difference in frequentist and Bayesian views, we have developed some criteria for comparing the performance of both approaches in a simulation framework. The basic criteria are: the Bayes risk for point estimation, and the expected length of the confidence interval and its coverage probability for interval estimation. We have conducted Monte Carlo experiments on a queueing application that show clear empirical evidence in favor of the Bayesian approach mainly in terms of coverage probability.

In Zouaoui and Wilson (2001b) we will extend the Bayesian framework presented in this paper to account for model uncertainty as well as parameter and stochastic uncertainty. It is not clear how the δ and bootstrap methods could be extended to account for model uncertainty.

Table 1: Average Risk of Classical Bootstrap and Proposed Bayesian Estimators for Average Sojourn Time in a Single Server Queue

ϕ_1	Method	$\widehat{\mathcal{R}}$	$\widehat{SE}(\widehat{\mathcal{R}})$
10	Bootstrap	9.3×10^{-11}	6.0×10^{-12}
	Bayesian	9.1×10^{-11}	5.9×10^{-12}
30	Bootstrap	1.4×10^{-10}	1.0×10^{-11}
	Bayesian	1.3×10^{-10}	9.8×10^{-12}
50	Bootstrap	2.7×10^{-9}	2.0×10^{-10}
	Bayesian	2.4×10^{-9}	1.8×10^{-10}
70	Bootstrap	8.0×10^{-9}	1.1×10^{-9}
	Bayesian	5.9×10^{-9}	6.8×10^{-10}
90	Bootstrap	2.5×10^{-4}	1.2×10^{-4}
	Bayesian	1.1×10^{-5}	4.5×10^{-6}

Table 2: Performance of Nominal 90% Confidence Intervals for Average Sojourn Time in a Single Server Queue

ϕ_1	Method	$\widehat{\mathcal{L}}$	$\widehat{V}^{1/2}/\widehat{\mathcal{L}}$	\widehat{C}
10	Bootstrap	4.9×10^{-5}	9.5×10^{-2}	98.2
	Bayesian	3.0×10^{-5}	9.6×10^{-2}	88.8
30	Bootstrap	6.3×10^{-5}	1.0×10^{-1}	99.0
	Bayesian	3.9×10^{-5}	1.1×10^{-1}	89.6
50	Bootstrap	2.6×10^{-4}	1.7×10^{-1}	99.4
	Bayesian	1.5×10^{-4}	1.5×10^{-1}	90.0
70	Bootstrap	4.0×10^{-4}	1.8×10^{-1}	98.8
	Bayesian	2.3×10^{-4}	1.8×10^{-1}	90.4
90	Bootstrap	7.5×10^{-3}	9.3×10^{-1}	99.4
	Bayesian	4.3×10^{-3}	6.0×10^{-1}	91.6

ACKNOWLEDGMENTS

This research was supported by the National Science Foundation under grant number DMI-9900164. The authors thank Bibhuti B. Bhattacharyya, Sujit K. Ghosh, and Stephen D. Roberts for many enlightening discussions on Bayesian techniques and simulation input modeling.

REFERENCES

Berger, J. O., and M. Delampady. 1987. Testing precise hypotheses. *Statistical Science* 2 (3): 317–352.

Carlin, B. P., and T. A. Louis. 1996. *Bayes and empirical Bayes methods for data analysis*. London: Chapman & Hall.

Casella, G., and E. I. George. 1992. Explaining the Gibbs sampler. *American Statistician* 46:167–174.

Cheng, R. C. H., and W. Holland. 1997. Sensitivity of computer simulation experiments to errors in input data. *Journal of Statistical Computation and Simulation* 57:219–241.

Efron, B. 1979. Bootstrap methods: Another look at the jackknife. *The Annals of Statistics* 7:1–26.

Efron, B., and R. J. Tibshirani. 1993. *An introduction to the bootstrap*. New York: Chapman & Hall.

Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin. 1995. *Bayesian data analysis*. London: Chapman & Hall.

Gilks, W. R., S. Richardson, and D. J. Spiegelhalter. 1996. *Markov chain Monte Carlo in practice*. London: Chapman & Hall.

Helton, J. C. 1998. Uncertainty and sensitivity analysis in the presence of stochastic and subjective uncertainty. *Journal of Statistical Computation and Simulation* 57:3–76.

Kass, R. E., and A. E. Raftery. 1995. Bayes factors. *Journal of the American Statistical Association* 90 (430): 773–795.

Law, A. M., and W. D. Kelton. 2000. *Simulation modeling & analysis*. 3rd ed. New York: McGraw-Hill.

Raftery, A. E., D. Madigan, and C. T. Volinsky. 1995. Accounting for model uncertainty in survival analysis improves predictive performance (with discussion). In *Bayesian statistics 5: Proceedings of the fifth Valencia international meeting, June 5–9, 1994*, ed. J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith, 323–349. Oxford: Oxford University Press.

Rao, P. 1997. *Variance components estimation: Mixed models, methodologies and applications*. London: Chapman & Hall.

Spiegelhalter, D., A. Thomas, N. Best, and W. Gilks. 1996. *BUGS 0.5: Bayesian inference using Gibbs sampling manual (version ii)*. MRC Biostatistics Unit, Institute of Public Health, Robinson Way, Cambridge, UK.

- Stuart, A., and J. K. Ord. 1994. *Kendall's advanced theory of statistics, volume 1: Distribution theory*. 6th ed. London: Edward Arnold.
- Zouaoui, F. 2001. Accounting for input uncertainty in discrete-event simulation. Doctoral dissertation, Graduate Program in Operations Research, North Carolina State University, Raleigh, NC. Available on-line via <http://www.lib.ncsu.edu/etd/public/etd-3949188410121271/etd.pdf> [accessed June 16, 2001].
- Zouaoui, F., and J. R. Wilson. 2001a. Accounting for parameter uncertainty in simulation input modeling. Technical Report, Department of Industrial Engineering, North Carolina State University, Raleigh, North Carolina. Available on-line via <ftp://ftp.ncsu.edu/pub/eos/pub/jwilson/msc1.pdf> [accessed June 16, 2001].
- Zouaoui, F., and J. R. Wilson. 2001b. Accounting for input model and parameter uncertainty in simulation. In *Proceedings of the 2001 Winter Simulation Conference*, ed. B. A. Peters, J. S. Smith, M. W. Rohrer, and D. J. Medeiros, to appear. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers. Available on-line with a more detailed development via <ftp://ftp.ncsu.edu/pub/eos/pub/jwilson/msc2.pdf> [accessed June 16, 2001].

AUTHOR BIOGRAPHIES

FAKER ZOUAOU is an operations research consultant in the Research Group at Sabre, Inc. He received his bachelor's and master's degrees in industrial engineering from Bilkent University, and he received his Ph.D. degree in operations research from North Carolina State University. He is a member of INFORMS. His e-mail address is <faker.zouaoui@sabre.com>.

JAMES R. WILSON is Professor and Head of the Department of Industrial Engineering at North Carolina State University. He currently serves as the corepresentative of INFORMS-College on Simulation to the Board of Directors of the Winter Simulation Conference. He is a member of ASA, ACM, IIE, and INFORMS. His e-mail address is <jjwilson@eos.ncsu.edu>, and his web address is <www.ie.ncsu.edu/jwilson>.