

BAYESIAN IDEAS AND DISCRETE EVENT SIMULATION: WHY, WHAT AND HOW

Stephen E. Chick

INSEAD

Technology and Operations Management Area

Boulevard de Constance

77300 Fontainebleau, FRANCE

ABSTRACT

Bayesian methods are useful in the simulation context for several reasons. They provide a convenient and useful way to represent uncertainty about alternatives (like manufacturing system designs, service operations, or other simulation applications) in a way that quantifies uncertainty about the performance of systems, or about inputs parameters of those systems. They also can be used to improve the efficiency of discrete optimization with simulation and response surface methods. Bayesian methods work well with other decision theoretic tools, and can therefore provide a link from traditional operations-level experiments to higher-level managerial decision-making needs, in addition to improving the efficiency of computer experiments.

1 INTRODUCTION

If simulation is only defined to be a numerical analysis tool that can estimate expectations and distributions of random variables, given a set of simulation inputs, then one might argue that there is no role for Bayesian analysis in stochastic, discrete event simulation. Indeed, most of the foundations of discrete event simulation emerged from the classical, or frequentist, statistical viewpoint. That viewpoint appeals to statistical properties of repeated sampling, assuming that the inputs to the system are known, fixed quantities. That is a natural viewpoint to adopt, given that simulation analysts can fully control the inputs to a simulation.

If simulation is recognized to be a tool that also models existing or proposed systems, however, then things are different. Statistical input parameters are typically unknown, and are estimated from expert judgment and statistical data. Understanding how uncertainty about input parameters influences outputs is an example of *uncertainty analysis*. Glynn (1986) explored Bayesian methods for uncertainty analysis in simulation experiments, by expanding the traditional role of estimating

$$\alpha = h(E[Y])$$

to a broader role of estimating the mean and distribution of the random variable $\alpha(\theta) = h(E[Y | \theta])$. The distribution of $\alpha(\theta)$ depends upon unknown parameters, θ . The unknown parameters, in turn, have a distribution $p(\theta)$ that can be assessed with data from the modeled system.

Chick (1997) reviewed few works to the time of that work that applied Bayesian ideas to simulation, then suggested a broader range of application areas than uncertainty analysis, including *ranking and selection*, *response surface modeling*, and *experimental design*. A central theme of the Bayesian approach is to use probability distributions to describe all uncertainty, not just randomness from stochastic behavior of simulated systems. That includes uncertainty about input parameters or output performance measures. Decision theory methods are often added to inform a modeler about decisions that a manager (which simulated system should be implemented?) or analyst (which simulation inputs to choose for the next run?) must make.

This tutorial discusses some basic building blocks for Bayesian reasoning, and identifies several applications to simulation experiments. Issues that arise in the Bayesian framework include: the need to specify initial uncertainty about unknown parameters by specifying *prior distributions* for unknown quantities (such as unknown outputs, unknown input parameters or unknown metamodel parameters); the specification of *likelihood models* to relate unknown parameters to observable data, if applicable; and *numerical tools* to update beliefs about unknown quantities as data becomes available using *Bayes' rule* to obtain posterior distributions for unknown quantities. But quantifying uncertainty is not enough—one needs to understand how uncertainty affects the potential outcomes of *decisions*. Decisions under uncertainty are guided by the principle of minimizing the *expected loss* associated with a choice of design parameters for a simulated system. Loss functions can also lead to experimental design criteria for structuring simulation experiments.

This tutorial updates a similar WSC Advanced Tutorial (Chick 2004) that was later extended significantly (Chick

2006). For reasons of space, the citations to the literature in this paper are woefully incomplete (apologies to the many articles that deserve mention). Please consult Chick (2006) for a fuller literature review on the development of theoretical techniques for Bayesian methods in simulation experiments; for applications of those tools (to scheduling, insurance, finance, traffic modeling, public health, waterway safety, supply chain and other areas), the relationship of Bayesian methods to deterministic simulations; and to subjective probability and Bayesian statistics in general.

2 MAIN CONCEPTS

We represent a stochastic simulation as a deterministic function of several types of inputs, with

$$Y_r = g(\theta_p, \theta_e, \theta_c; U_r), \quad (1)$$

so that Y_r is the output of the r th replication, θ_p is a vector of statistical input parameters that can be inferred with data from the real system a model is intended to represent; θ_e is a potentially unknown parameter that describes a system's operating environment, but that is *not* inferrable from a stream of potentially available data; θ_c are control parameters (or design variables) for the modeled system; and U_r are jointly independent uniform random variables (e.g., samples from different portions of a uniform random number generator to provide randomness in the output).

A subjective probabilist represents all uncertain quantities with probability distributions, including θ_p and θ_e , not only uniform variates U and nonuniform random variates X . We add subscripts if additional specificity is needed, e.g., to refer to the j -th variate for the i -th source of randomness during the r -th replication, we write X_{rij} . See Figure 1. The r may be dropped to describe data collected from the actual system being simulated. That data would be used to infer the parameters of the statistical distributions to describe the system.

One reason a simulation experiment may be run is to estimate the function g because its exact form is not known. Because of this, it is often useful to study *metamodels* of simulation models to predict the outputs of a simulation model (or the simulated system) when a full simulation takes a long time to run. Examples of metamodels are linear regression models and Gaussian random fields (GRFs). Section 3.4 describes some Bayesian methods to describe uncertainty about metamodel parameters ψ .

The specification of prior distributions and Bayesian inference with data are discussed in Sections 2.1 and 2.2. Asymptotic theorems are presented in Section 2.3, followed by a discussion of loss functions that can be used to design sampling allocations for simulation experiments.

2.1 Exchangeability and Input Parameters

An important simulation design issue is the selection of appropriate input distributions to characterize the stochastic behavior of the modeled system. This section reviews basic ideas and important theorems for inferring input parameters with the Bayesian formalism. The presentation is in the context of selecting a parameter θ for a specific candidate distribution (e.g., *one* of the Bernoulli, exponential, or gamma distributions) for input into a computer simulation. Section 3.2 explores input selection if multiple candidate models for a given source of randomness is proposed.

For a Bayesian, the idea of exchangeability is preferred to the idea of independent and identically distributed (i.i.d.) random variables. Let $X = (X_1, X_2, \dots, X_N)$ be a generic vector of random variables on an outcome space Ω . A probability p on Ω is said to be *exchangeable* if it is invariant with respect to permutations of the coordinates (e.g., $p(x_1, x_2, \dots, x_n) = p(x_{s_1}, x_{s_2}, \dots, x_{s_n})$ for permutations s on $\{1, 2, \dots, n\}$ for arbitrary n). Exchangeability is a weaker assumption than independence, but is important because it motivates conditionally independent sequences that are useful in simulation experiments for parameter inference.

Simulation is very often concerned with conceptually infinite ($\lim N \rightarrow \infty$) exchangeable sequences of random variables (e.g., replications $r = 1, 2, \dots$ or service times X_{rij} for $j = 1, 2, \dots$). A key theorem (de Finetti 1990 or Bernardo and Smith 1994) for conceptually infinite exchangeable sequences of Bernoulli random variables says that outcomes are conditionally independent, given the limiting fraction of heads, $\Theta = \lim_{N \rightarrow \infty} \sum_{i=1}^N X_i / N$, with some mixture distribution $\pi(\theta)$,

$$\lim_{N \rightarrow \infty} p(x_n) = \int \left\{ \prod_{i=1}^n f(x_i | \theta) \right\} d\pi(\theta), \quad (2)$$

where $p(x_i | \theta) = f(x_i | \theta) = \theta^{x_i} (1 - \theta)^{1 - x_i}$ is a conditional probability when considered as a function of x_i and a likelihood when written as a function of θ . A mixture of conceptually infinite sequences of exchangeable random variables written in the form Equation (2) for an arbitrary parameter θ , distribution $\pi(\theta)$ and likelihood model f is called a *de Finetti-type representation*, although judgments stronger than exchangeability, such as invariance to sums, may be required to justify them (Barlow and Mendel 1992).

That notation anticipates the convention of writing a *prior distribution* as $\pi(\cdot)$, representing the *a priori* belief that the parameter takes on a given value. That representation allows for the inference of θ from data $x_n = (x_1, \dots, x_n)$ via *Bayes rule*,

$$p(\theta | x_n) = \frac{\pi(\theta)p(x_n | \theta)}{p(x_n)}. \quad (3)$$

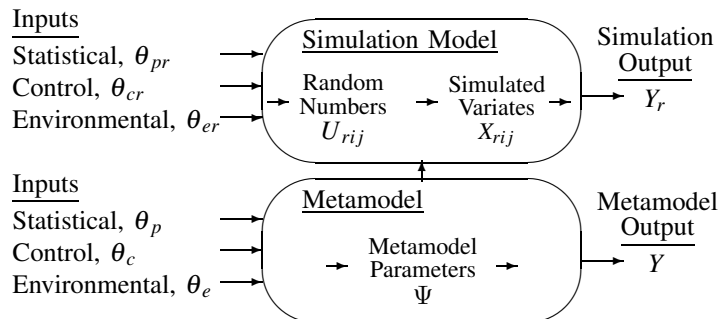


Figure 1: Simulation Takes Multiple Types of Inputs and Meta-models Predict Outputs for Unsimulated Input Values

The *posterior probability* $p(\theta | x_n)$ of θ , given x_n , summarizes uncertainty about θ via the likelihood model and prior distribution.

2.2 Prior Probability Distributions

Bayesian methods require probability distributions to quantify initial uncertainty before data is observed. The selection of a prior distribution is controversial. Bruno de Finetti (1990) argues that a prior distribution is a subjective expression of uncertainty, and that You (yes, You) may justifiably specify a different distribution than I, since we may have different beliefs about the likelihood of a given event. Savage (1972) suggests a process for eliciting a prior distribution from a modeler through the evaluation of “fair bets”. Kahneman et al. (1982) illustrate potential pitfalls with eliciting probability judgments and present techniques to counter them. While this may seem “too subjective” and open to biases (Edwards 1984), the ability to include prior information provides necessary flexibility and can be considered an advantage of the approach. Frequentist methods apply only with data, and problems remain (e.g., Section 3.2).

To avoid the impression of subjectivity, several “automated” mechanisms have nonetheless been proposed to support the selection of a prior distribution. When a lot of data are available, the likelihood function will be the dominant term in Bayes’ rule, rather than the prior distribution, so these methods may be helpful. The first approach is to obtain a prior distribution for a parameter of an infinite exchangeable sequence as a limiting case of an *indifference judgment* for a finite exchangeable sequences. For the finite exchangeable sequence of $\{0, 1\}$ outcomes, if each of the finite set of alternatives $\theta_N \in \{0/N, 1/N, \dots, (N-1)/N, 1\}$ is judged equally likely for each N , then $\lim_{N \rightarrow \infty} p(\theta_N) \xrightarrow{D} \text{uniform}[0, 1]$, the prior probability model used by Laplace (1812) to assess his prior probability that the sun would come up tomorrow. That approach is coordinate dependent (indifference judgments for θ and $\log \theta$ give different results). Jeffreys (1946) suggested $\pi(\theta) \propto |H(\theta)|^{1/2} d\theta$, where H is the expected

information in one observation,

$$H(\theta) = \mathbb{E}_X \left[- \frac{\partial^2 \log p(X | \theta)}{\partial \theta^2} \Big|_{\theta} \right], \quad (4)$$

because it has the attractive property of being invariant with respect to coordinate changes in θ . *Jeffreys’ prior* for Bernoulli sampling is a $\text{beta}(1/2, 1/2)$ distribution. For some models, Jeffreys’ prior is improper (does not integrate to one), but may be useful if the data results in a proper posterior after Bayes’ rule is formally applied.

A third approach that is mathematically convenient is to assume a *conjugate prior* distribution, meaning that the posterior distribution has the same functional form as the prior distribution. For Bernoulli(θ) sampling, the $\text{beta}(\alpha, \beta)$ distribution with probability density function (pdf) $f(\theta) \propto \theta^{\alpha-1} (1-\theta)^{\beta-1}$ is a conjugate prior. If data x_n is observed, with $s_n = \sum_{i=1}^n x_i$, then the posterior pdf is $f(\theta | x_n) \propto \theta^{\alpha+s_n-1} (1-\theta)^{\beta+n-s_n-1}$, a $\text{beta}(\alpha+s_n, \beta+n-s_n)$ distribution. Conjugate prior distributions exist for all members of the regular exponential family (Bernardo and Smith 1994), including the exponential, normal, gamma, lognormal, Wishart, Bernoulli, geometric, and Poisson distributions, as well as linear regression models with normally distributed error. The $\text{uniform}[0, 1]$ distribution is in the conjugate family for Bernoulli sampling—it is a $\text{beta}(1, 1)$ distribution.

Conjugate prior distributions are often selected to be as *noninformative* as possible, meaning that probability is spread “evenly” over the parameter space. Although evenly is subjectively defined, heuristics are available for members of the regular exponential family, whose likelihood function can be written $p(x | \theta) = a(x)h_0(\theta) \exp[\sum_{j=1}^d c_j \phi_j(\theta) h_j(x)]$ for some $a(\cdot), h_0(\cdot), c_j, \phi_j(\cdot), h_j(\cdot)$. The conjugate prior distribution is $p(\theta) = [K(t)]^{-1} [h_0(\theta)]^{n_0} \exp[\sum_{j=1}^d c_j \phi_j(\theta) t_j]$. The posterior distribution given n conditionally independent data points then has parameters $n_0 + n$ and the sum of $t = (t_1, t_2, \dots, t_d)$ and the sufficient statistics (Bernardo and

Smith 1994). The parameter n_0 is therefore interpreted by some to be the “strength” of the prior, measured in terms of the number of samples. In that case, evenly spreading probability can be taken to be selecting n_0 close to 0, while insuring that the prior is still proper. Jaynes (1983) suggests a fourth approach that is common in image and signal processing: maximum entropy methods define “diffuse” prior with respect to a background measure, subject to moment constraints on the parameters. Berger (1994) and Kass and Wasserman (1996) discuss on default prior distributions and sensitivity analysis with respect to them.

Probability modeling is inherently subjective—even so-called “objective” methods require the subjective specification of a likelihood model. One standard Bayesian practice is to use a slightly informative conjugate distribution for the unknown mean, by choosing it to be proper but diffuse (Gilks et al. 1996). For example, the conjugate prior for an unknown mean of a normal distribution is also a normal distribution. A diffuse prior would be $\text{Normal}(0, \sigma_{big}^2)$ for some large σ_{big}^2 . Conjugate prior distributions are mathematically convenient, but care is still required with their use, as with any statistical analysis, Bayesian or otherwise.

2.3 Asymptotic Theorems

Classical asymptotic theorems (laws of large numbers, LLN; central limit theorem, CLT; e.g., Billingsley 1986) have Bayesian interpretations when considered to be conditional on the mean and standard deviation of an infinite exchangeable sequence.

A Bayesian extension of the LLN allows for sample averages to converge to random variables rather than to “true” means.

Theorem 1 (Bayesian LLN) *If \bar{X}_n and \bar{Y}_m are respectively the averages of n and m exchangeable random quantities X_i (the two averages may or may not have some terms in common), the probability that*

$$|\bar{X}_n - \bar{Y}_m| > \varepsilon$$

may be made arbitrarily small by taking n and m sufficiently large (de Finetti 1990, p. 216 assumes a finite variance).

Even though the modes of Bayesian posterior distributions may not be the true mean, an asymptotic normality property holds for posterior distributions of parameters.

Theorem 2 (Posterior Normality) *For each n , let $p_n(\cdot)$ be the posterior pdf of the d -dimensional parameter θ_n given $x_n = (x_1, \dots, x_n)$, let $\tilde{\theta}_n$ be its mode, and define the $d \times d$ Bayesian observed information matrix Σ_n^{-1} by*

$$\Sigma_n^{-1} = - \left. \frac{\partial^2 \log p_n(\theta | x_n)}{\partial \theta^2} \right|_{\tilde{\theta}_n}. \quad (5)$$

Then $\phi_n = \Sigma_n^{-1/2}(\theta_n - \tilde{\theta}_n)$ converges in distribution to a standard (multivariate) normal random variable (Bernardo and Smith 1994, Prop 5.14 needs regularity conditions).

Theorem 2 asserts that uncertainty about the value of the unknown parameter value can be approximated asymptotically with a normal distribution. The Bayesian observed information Σ_n^{-1} is a measure of precision of the posterior distribution of θ , and behaves asymptotically like the frequentist observed information (which ignores the prior distribution) under rather general conditions, but the interpretation differs somewhat. The classical analog of Theorem 2 asserts that the MLE is asymptotically normally distributed about a “true” parameter θ_0 (Law and Kelton 2000), rather than describing uncertainty about θ . The mode $\tilde{\theta}_n$ is often called a MAP (maximum a posteriori probability) estimator.

2.4 Expected Value of Information

The fact that input uncertainty is described by probability distributions allows the modeler to assess the *expected value of information* of additional data. The expected value of information is useful in experimental design. It measures the value of resolving uncertainty with respect to a *loss function* $\mathcal{L}(d, \omega)$ that describes the loss when a *decision* d is chosen when the state of nature is ω . The expected improvement in the loss given by the information in an experiment is a Bayesian experimental design criterion.

The value of information idea directly leads to the selection procedures in Section 3.3. A simplified version of that problem adapted from de Groot (1970) illustrates the key concepts. Suppose we must decide whether or not the unknown mean W of a normal distribution (known sampling variance σ^2) is smaller (decision $d = 1$) or larger ($d = 2$) than w_0 . Conditionally independent samples $X_n = (X_1, X_2, \dots, X_n)$, with $p(X_i) \sim \text{Normal}(w, \sigma^2)$ given $W = w$, can be used to infer the value of the mean. The decision maker designs an experiment (chooses n) to balance the cost of sampling, cn , and the expected penalty if the wrong answer is chosen. Here the penalty for incorrect selection is the opportunity cost $\mathcal{L}(1, w)$, the difference between the actual value of w and w_0 when the wrong answer is selected, and 0 if the right answer is selected.

$$\mathcal{L}(1, w) = \begin{cases} 0 & \text{if } w \leq w_0 \\ w - w_0 & \text{if } w > w_0, \end{cases}$$

$$\mathcal{L}(2, w) = \begin{cases} w_0 - w & \text{if } w \leq w_0 \\ 0 & \text{if } w > w_0. \end{cases}$$

Since the mean is not known exactly, there is a potential penalty for incorrectly specifying whether W is smaller or larger than w_0 . We model uncertainty about W with a $\text{Normal}(\mu, 1/\tau)$ prior distribution, which is conjugate for normal sampling with an unknown mean and known variance (Bernardo and Smith 1994). Here, τ is the *precision* in our

uncertainty about W . Observing $X_n = x_n$ would reduce the uncertainty and result in the posterior distribution

$$p(w | x_n) \sim \text{Normal}(z, \tau_n^{-1}),$$

where

$$z = \text{posterior mean of } W = E[W | x_n] = \frac{\tau\mu + \frac{n}{\sigma^2}\bar{x}_n}{\tau + \frac{n}{\sigma^2}},$$

and $\tau_n = \text{posterior precision of } W = \tau + n/\sigma^2$.

The variance τ_n^{-1} equals the posterior variance approximation Σ_n in Equation (5) because Σ_n is based on a normal distribution approximation.

The posterior mean z influences the decision, but it depends upon n , which must be selected before X_n is observed. We need the *predictive distribution* $p(z)$ of the posterior mean $Z = E[W | X_n] = (\tau\mu + n\bar{X}_n)/\tau_n$ to see how n samples influences the decision d . The conditional distribution of \bar{X}_n given w is $\text{Normal}(w, \sigma^2/n)$. Mixing over the prior distribution of W implies that the predictive distribution for Z is $\text{Normal}(\mu, \tau_z^{-1})$, where

$$\tau_z = \tau(\tau + n/\sigma^2)/(n/\sigma^2). \quad (6)$$

The variance τ_z^{-1} of Z is 0 when $n \rightarrow 0$ (no new information). If $n \rightarrow \infty$ (perfect information about w), then $\text{Var}[Z] \rightarrow \tau^{-1}$, the prior variance for W .

The experimental design that minimizes risk (the cost of sampling plus expected losses due to a potentially incorrect decision) is the n that minimizes a nested expectation, an inner expectation corresponds to the expected loss after X_n is observed, an outer expectation averages over X_n ,

$$\rho(n) = cn + E_{X_n}[E_W[\mathcal{L}(d(X_n), W) | X_n]]. \quad (7)$$

One technique to analyze $E_W[\mathcal{L}(d(X_n), W) | X_n]$ is to obtain an auxiliary loss function \mathcal{L}^* that has the same optimal decision, but simplifies the loss function by making the loss of one of the decisions equal to 0. Adding a function of w does not change the optimal decision (de Groot 1970). Set $\mathcal{L}^*(d, w) = \mathcal{L}(d, w) - \mathcal{L}(1, w)$, which is 0 if $d = 1$ and is $w_0 - w$ if $d = 2$. Then

$$E_W[\mathcal{L}^*(d(X_n), W) | X_n] = \begin{cases} 0 & \text{if } d = 1 \\ w_0 - Z & \text{if } d = 2. \end{cases} \quad (8)$$

The decision that minimizes the loss in Equation (8) is to assert $d(X_n) = 2$ (“bigger”) if the posterior mean exceeds the threshold, $Z > w_0$, and to assert $d(X_n) = 1$ (“smaller”) if $Z \leq w_0$.

The expectation over the outcomes X_n can be determined with well-known tables because the decision depends upon X_n only through Z , and Z has a normal distribution. The

expected loss can be determined by the standard normal loss $\Psi[s] = \int_s^\infty (t-s)\phi(t)dt = \phi(s) - s(1 - \Phi(s))$ for expected lost sales in the newsvendor problem if demand is normally distributed (e.g., Porteus 2002).

$$\begin{aligned} E[\mathcal{L}^*(d(X_n), W)] &= E_{X_n}[E_W[\mathcal{L}^*(d(X_n), W) | X_n]] \\ &= - \int_{w_0}^\infty (z - w_0)p(z | X_n)dz \\ &= -\tau_z^{-\frac{1}{2}} \Psi[\tau_z^{\frac{1}{2}}(w_0 - \mu)]. \end{aligned}$$

The expected loss of the original loss function is recovered by adding back $E[\mathcal{L}(1, W)]$, using the prior distribution of W for the expectation.

$$\begin{aligned} E[\mathcal{L}(d(X_n), W)] &= \tau^{-\frac{1}{2}} \Psi[\tau^{\frac{1}{2}}(w_0 - \mu)] \\ &\quad - \tau_z^{-\frac{1}{2}} \Psi[\tau_z^{\frac{1}{2}}(w_0 - \mu)]. \end{aligned} \quad (9)$$

The *expected value of information* for m samples is the difference between Equation (9) when $n = 0$ and when $n = m$ (τ_z depends on n). Combine Equation (9) with Equations (6)-(7), note that $d\Psi/ds = \Phi(s) - 1$ and $d\tau_z/dn = -\tau^2\sigma^2/n^2$, and take the derivative with respect to n (relaxing the integer assumption) to obtain an optimality condition for the sample size.

$$\frac{\partial \rho}{\partial n} = \frac{1}{2} \tau_z^{-\frac{3}{2}} \phi[\tau_z^{\frac{1}{2}}(w_0 - \mu)] \cdot \frac{-\tau^2\sigma^2}{n^2} + c = 0.$$

For diminishing costs $c \rightarrow 0$, the sample size is large. Since $\tau_z \rightarrow \tau$ as $n \rightarrow \infty$, the optimal sample size n is (asymptotically) approximately

$$n^* = \left(\frac{\tau^{\frac{1}{2}}\sigma^2\phi[\tau^{\frac{1}{2}}(w_0 - \mu)]}{2c} \right)^{1/2}. \quad (10)$$

This argument illustrates the basic ideas of loss functions, and the use of predictive distributions for future samples to infer the expected value of information of sampling. *The technique of adding functions of the unknowns can be useful to simplify the derivation of the optimal solution. Asymptotic approximations are a further tool to identify criteria-based sampling plans.* Extensions of this basic argument justify the expected value of information based selection procedures that are summarized in Section 3.3.

An alternate mechanism to approximate the effect of information on parameter uncertainty is based on a thought experiment for the posterior probabilities of parameters. For members of the regular exponential family, the asymptotic variance approximation Σ_n in Equation (5) simplifies to the form $H^{-1}(\theta)/(n_0 + n)$, where H is the expected information from one observation (Equation (4)), when a canonical conjugate prior distribution is used (Bernardo and Smith

1994). To approximate the effect of collecting m additional samples on the parameter uncertainty, given that n samples have already been collected, one could presume that the posterior distribution changes from $\text{Normal}(\tilde{\theta}_n, \Sigma_n)$ to

$$\text{Normal}\left(\tilde{\theta}_n, \Sigma_n \frac{n_0 + n}{n_0 + n + m}\right). \quad (11)$$

This transformation reflects an appropriate scaling of the posterior precision, and the idea is used in a frequentist context for estimating how many replications are required to achieve a confidence interval of a given size (Law and Kelton 2000). Chen (1996) uses this type of approximation for the Bayesian posterior distribution of the unknown means of several simulated systems in order to motivate a class of ranking and selection procedures called the *OCBA*. Ng and Chick (2006) use the idea to plan experiments to reduce input uncertainty in a way that reduces output uncertainty.

2.5 Entropy and Kullback-Leibler Discrepancy

Kullback-Leibler discrepancy is a useful measure of the difference between two distributions. For discrete distributions \tilde{p} and p , the discrepancy is $\delta(p \parallel \tilde{p}) = \sum \tilde{p}_i \log(\tilde{p}_i/q_i)$. Two continuous distributions for a random variable X with densities \tilde{f} and $f_\theta = f(x \mid \theta)$ have discrepancy

$$\delta(f_\theta \parallel \tilde{f}) = \int \tilde{f}(x) \log \frac{\tilde{f}(x)}{f(x \mid \theta)} dx. \quad (12)$$

One application for discrepancy is as a loss function for a decision maker that must specify a probability distribution. If the decision-maker believes that the distribution is \tilde{f} , and loses $\delta(f \parallel \tilde{f})$ if he/she provides a distribution f , then the decision-maker should honestly report \tilde{f} to minimize the expected loss (Bernardo and Smith 1994).

A second application of discrepancy is as a design criterion for experiments designed to assist with parameter estimation (Bernardo 1979, Smith and Verdinelli 1980). In the simulation context, this could mean selecting a design matrix \mathbf{d}_Θ of r vectors of inputs $(\theta_{pi}, \theta_{ei}, \theta_{ci})$ for $i = 1, 2, \dots, r$ with output \mathbf{Y} in order to best differentiate the posterior distribution of the response parameters $\boldsymbol{\psi}$ from the prior distribution for $\boldsymbol{\psi}$. In other words, the goal is to select the \mathbf{d}_Θ from a set of possible design matrices to maximize the expected divergence of posterior and prior distributions for $\boldsymbol{\psi}$, the expectation over the predictive distribution of the output \mathbf{Y} .

$$\int p(\mathbf{Y} \mid \mathbf{d}_\Theta) \left(\int p(\boldsymbol{\psi} \mid \mathbf{Y}) \log \frac{p(\boldsymbol{\psi} \mid \mathbf{Y})}{p(\boldsymbol{\psi})} d\boldsymbol{\psi} \right) d\mathbf{Y} \quad (13)$$

The predictive distribution $p(\mathbf{Y} \mid \mathbf{d}_\Theta)$ of future outputs depends upon the current uncertainty about $\boldsymbol{\psi}$ and the design

matrix \mathbf{d}_Θ . This approach is essentially an expected value of information design criterion, as in Section 2.4, except now the loss function is the Kullback-Leibler discrepancy rather than the opportunity cost.

Other applications of discrepancy include the maximum entropy prior distribution mentioned above (Jaynes 1983), and for input distribution selection, as in Section 3.2 below.

3 APPLICATIONS

3.1 Uncertainty Analysis

A sensitivity analysis tests how the mean simulation output depends upon one or more input parameters as that parameter is varied (estimating $E[g(\theta) \mid \mathcal{E}]$ as a function of θ , given all information \mathcal{E}). Uncertainty analysis entails propagating input parameter uncertainty about Θ through to uncertainty about outputs Y . Even if a simulation has no random number stream, a distribution on unknown inputs means that the output is random.

An unbiased estimator of the mean output $E[Y \mid \mathcal{E}]$ with both stochastic (from u) and systemic (or parameter) uncertainty accounted for is obtained from the Bayesian model average (BMA) in Figure 2, which averages over random inputs sampling according to the distribution $p(\theta \mid \mathcal{E})$ (Draper 1995, Chick 2001). Zouaoui and Wilson (2003) explore the relative magnitude of stochastic and systemic uncertainty with variations on the BMA, and discuss how to update the estimate should new data become available (so the algorithm need not be rerun from scratch). Importance sampling techniques can reweigh estimates accordingly (with likelihood ratio determined as the ratio of the “new” posterior divided by the “old” distribution). Andradóttir and Glynn (2004) examine the estimation of $E[Y(\Theta) \mid \mathcal{E}]$ when there may be bias in the estimates of Y given θ , when quasi-random sequences are used in place of the pseudo-random sequences assumed by Figure 2, or when numerical techniques like Simpson’s rule are employed to select values of θ .

Other issues for sensitivity analysis include estimation of the distribution of the conditional expectation $E[Y \mid \Theta, \mathcal{E}]$. When Y is a deterministic function of Θ , then naive Monte Carlo simulation can be used with traditional kernel estimation techniques to assess the distribution of $Y(\Theta)$. When

```

for  $r = 1, \dots, R$  replications
  sample parameter  $\theta_r$  from  $p(\theta \mid \mathcal{E})$ 
  for  $i = 1, 2, \dots, n$ 
    generate simulation output  $y_{ri}$  given input  $\theta_r$ 
  end loop
end loop
Estimate  $E[Y \mid \mathcal{E}]$  with  $\bar{y} = \frac{1}{R} \sum_{r=1}^R \sum_{i=1}^n y_{ri}/n$ .

```

Figure 2: Bayesian Model Average (BMA)

the simulation is stochastic (depends on the random number stream u), then $E[Y | \theta, \mathcal{E}]$ is imperfectly estimated for any given θ . Given several technical conditions (e.g., univariate continuous-valued θ , monotonic mean response), Steckley and Henderson (2003) derive asymptotically optimal ways of selecting by cleverly selecting r and m in Figure 2 to produce a kernel density estimator based on the output. Their work builds upon earlier work by Lee and Glynn (1999) to estimate the distribution function of $E[Y | \Theta, \mathcal{E}]$ for the case of discrete-valued θ .

3.2 Selecting from Multiple Candidate Distributions

Input selection in simulation practice may consider $q > 1$ candidate distributions for potential input to a model. The usual approach (Law and Kelton 2000) is to find the MLE for each candidate distribution, assess the goodness of fit for each model, then select a model that “fits well”. While this is practically appealing, there are known problems with the approach (Lindley 1957, Berger and Pericchi 1996). *In the simulation context, input uncertainty can make standard confidence intervals for the mean output almost meaningless if the classical approach is used* (Barton and Schruben 2001; Chick 2001, 2006), because an excellent simulation estimate based on point estimates for inputs misses the boat if there is uncertainty about input values, which is typical for simulations of real systems.

A Bayesian approach with model uncertainty is basically the same as for parameter uncertainty alone, as in Section 2.1 above, except now that a prior probability distribution needs to be placed on the model/parameter combination, $\pi(m, \theta_m)$, a mixed discrete-continuous model, where $m \in \{1, 2, \dots, q\}$ indexes the set of candidate distributions, and θ_m is the parameter for the m -th candidate distribution. As data x_n becomes available, the BMA then requires sampling from the joint posterior $p(m, \theta_m | x_n)$. This can be accomplished by composition, sampling the input model then the parameter with $p(m | x_n)p(\theta_m | m, x_n)$.

Chick (2001) first illustrated Bayesian input modeling in a stochastic simulation context, and suggested a method-of-moments approach for assessing prior distributions for the unknown parameters of each candidate model. Zouaoui and Wilson (2004) noted a decoupling of stochastic uncertainty from two types of structural uncertainty (that due to uncertainty about the candidate model, plus that due to uncertainty about the parameters of the candidate models) under special conditions, provided a variance reduction for the BMA and numerical analysis.

Selecting models according to $p(m, \theta_m | \mathcal{E})$ is consistent in that if one of the entertained models is the true model, then the true model is selected if enough data is observed and some regularity conditions hold. When the true model is not among those being considered, Bayesian model selection

chooses the model that is closest to the true model in terms of Kullback-Leibler divergence (Berk 1966).

3.3 Ranking and Selection

Procedures to select the best of a finite set of simulated systems using Bayesian approaches have been derived from three very different angles.

Chick and Inoue (2001a) used expected value of information ideas that extend Equations (9)-(10) to obtain two-stage value of information procedures (VIPs) with independent replications that identify the best of a finite set of simulated systems. Equation (9) was generalized to allow for unknown and potentially different variances for multiple simulated systems. The loss for $k > 2$ systems was approximated by the sum of losses for $k - 1$ pairwise comparisons between the system selected in the absence of additional replications and each of the other systems. Additional replications for the second stage are allocated so minimize that loss function using an asymptotic (in the number of replications) approximation like that used for Equation (10). A variation improves the (Bayesian posterior) probability of correct selection with the 0-1 loss function (loss of 1 if the wrong system is selected, 0 if the best is selected), and sequential variations exist (Chick and Inoue 2001a, Chick and Inoue 2002). The procedures, named $\mathcal{LL}(\mathcal{B}), \mathcal{LL}(\mathcal{S}), 0-1(\mathcal{B}), 0-1(\mathcal{S})$ depending upon whether an opportunity cost (aka linear loss) or 0-1 loss function is used, and whether a budget constraint or sequential sampling apply, are empirically quite effective for identifying the best system with respect to several figures of merit. The idea has also been extended to handle common random numbers with screening to improve efficiency. The analysis to justify them also requires missing data techniques (Chick and Inoue 2001b).

Chen (1996) and Chen et al. (2005) use a different tack—the thought experiment in Equation (11) that supposes that additional replications won’t change the estimate of the means of each system, but can reduce the variance associated with the estimate. They provide empirically effective procedures (called *OCBA*) with only a few parameters to tune for identifying the best system with high probability. Although *OCBA* is typically derived by assuming that samples are normally distributed, the results of (Bernardo and Smith 1994) can be used to asymptotically justify the idea of the *OCBA* when both n and the number of samples per stage are large. He et al. (2006) provide an *OCBA* variation that reduces the expected opportunity cost of potentially incorrect selections.

See Branke et al. (2005) for a thorough review of the VIP and *OCBA* approaches, with: a comparison of how they compare with the (frequentist) indifference zone approach; a derivation of new procedures; and the identification of highly effective procedures as determined by battery of

numerical tests (innovations for some specific VIP and OCBA procedures).

A third, very new Bayesian approach to selection procedures was proposed by Chick and Gans (2005). That approach takes an economic viewpoint to simulation selection that accounts for both the costs of replications and the discounting costs that occurs due to delays from a simulation analysis. This reframes the simulation selection problem from one of statistic significance to one of expected net present values (the language of managers that make high-level decisions that may be informed by simulation). They show that, given certain conditions (including independent samples from system to system that represent the economic value of implementing a simulated system), the problem of sequentially selecting which system to simulate or implement resembles a variation of a Bayesian bandit problem. The optimal procedure for deciding which system to further simulate or to implement uses a Gittins index and a special continuation region. They show that the Gittins index turns out to be proportional to the value of what might be called an American option on a regular (not geometric) Brownian motion whose drift is unknown but is inferred with Bayes rule and observations from that Brownian motion. That paper also shows how to approximate the Gittins index when samples are normally distributed with a known variance.

3.4 Metamodels

Metamodels are a model of how simulated responses are believed to behave as a function of input parameters, even for values of parameters not yet input to the simulation. This is particularly useful when the simulation model requires extensive computation. In addition to describing the response, metamodels can be used to predict how reductions in input parameter uncertainty can reduce uncertainty about the mean system performance. Since the metamodel is unknown, uncertainty about it can be modeled from a Bayesian perspective. Here we discuss the normal linear model. For information about a very interesting alternative, called Gaussian random function (GRF) metamodels, see also Kennedy and O'Hagan (2001), Santner et al. (2003) and van Beers and Kleijnen (2003, 2004).

The normal linear model is

$$Y = \sum_{\ell=1}^p g_{\ell}(\theta)\beta_{\ell} + Z(\theta; U) = g^T(\theta)\beta + Z(\theta; U), \quad (14)$$

for *known* regression functions g_1, \dots, g_p , potentially unknown regression coefficients β , and a zero-mean random noise $Z(\cdot)$. The conjugate prior is an inverted gamma distribution for the unknown variance σ^2 and a conditionally normal distribution for β given σ^2 , if all factors are active (Bernardo and Smith 1994). The metamodel parameters are then $\psi = (\beta, \sigma^2)$. Raftery, Madigan, and Hoeting (1997)

describe a relatively uninformative prior distribution for ψ for this normal linear model. A vector of simulation output $\mathbf{Y} = \mathbf{y}$ obtained with design matrix \mathbf{d}_{Θ} with inputs $(\theta_{pi}, \theta_{ei}, \theta_{ci})$ for runs $i = 1, \dots, r$ can be used to update the posterior distribution of unknown response parameters, $p(\psi | \mathbf{y}, \mathbf{d}_{\Theta})$, using Bayes' rule.

Identifying important factors (factors with nonzero β_i) can be formulated as a Bayesian model selection problem, selecting from 2^p different response models that are distinguished by the presence or absence of each factor. George and McCulloch (1996) and Cheng (1999) discuss techniques for estimating which factors are active with what probability. Ng and Chick (2004) describe an entropy-based experimental design criterion (cf. Section 2.5) to identify both which factors are active and reduce parameter uncertainty simultaneously.

When the g_i represent the individual dimensions of the unknown parameters (θ_p, θ_e) , the β_{ℓ} are gradients with respect to the inputs. If the model has only statistical input parameters θ_p for which data can be collected (but not parameters θ_e for which no data is available), Ng and Chick (2001) and Zouaoui and Wilson (2003) indicate that output uncertainty can be decoupled asymptotically or under special conditions.

$$\begin{aligned} \text{Var}[\bar{Y} | \mathcal{E}] &\approx \text{stochastic} + \text{parameter uncertainty} \\ &\approx \frac{\hat{\sigma}_0^2}{m} + \frac{\beta \hat{H}_{\theta_p}^{-1} \beta}{n}, \end{aligned}$$

where $\hat{\sigma}_0^2$ is the estimate of the variance from m replications, the MLE $\hat{\theta}_p$ and estimate $\hat{H}_{\theta_p}^{-1}$ of the information in one observation are based on n data points, and some technical conditions hold. This adapts a frequentist result of Cheng and Holland (1997).

Ng and Chick (2001) use that decoupling, applied to uncertainty due to multiple input parameters,

$$\sum_{i=1}^{n_p} \beta_i \hat{H}_{\theta_{pi}}^{-1} \beta_i / n_i,$$

to provide sampling plans to collect further data to reduce input parameter uncertainty to optimally reduce output uncertainty in some sense, assuming that different numbers of data points can be collected to infer the parameters of different sources of randomness (e.g., arrival rates versus service time distributions). Ng and Chick (2006) extend that analysis by accounting for uncertainty about β ; suggest allocations of resources for either running more replications or collecting more data points to reduce the asymptotic output variance approximation; and give a numerical analysis.

To date, much simulation research seeks analytical results for stochastic models, or mechanisms to reduce the variance of estimators due to stochastic noise. Those re-

sults need to be complemented with an *understanding of how performance depends on input uncertainty, and methods to reduce input uncertainty to effectively reduce output uncertainty*. The Bayesian approach is a tool that can help.

4 IMPLEMENTATION

Three basic computational issues for implementing a Bayesian analysis are maximization (e.g., find the MLE $\hat{\theta}$ or MAP $\tilde{\theta}$ estimator for a posterior distribution); integration, either to find a marginal distribution (find $p(\theta_1 | x_n)$ from $p(\theta_1, \theta_2 | x_n)$) or constant of proportionality for a posterior distribution (find $c^{-1} = \int f(x_n | \theta) d\pi(\theta)$, or $p(m | x_n)$ for Section 3.2); and simulation (sample from $p(\theta | x_n)$ in order to estimate $E[g(\theta) | x_n]$). These techniques are described in a variety of sources (Evans and Swartz 1995, Tanner 1996, Gilks et al. 1996, Devroye 2006).

For maximization, a number of methods are available including gradient-based methods (e.g., Newton-Raphson), gradient-free methods (e.g., Nelder-Mead), and simulation-based methods. The expectation-maximization (EM) algorithm is a technique for finding the MAP or MLE when there is missing data or nuisance parameters are to be integrated out (e.g., the MAP of $p(\theta_1 | x_n)$ when it is “messy” but $p(\theta_1, \theta_2 | x_n)$ is easier to manipulate).

For integration, five general techniques apply when analytical results are not available: asymptotic methods, Markov chain methods, importance sampling, adaptive importance sampling, and multiple quadrature. Quadrature is useful when the number of dimensions is not too large. The Laplace method is an interesting approximation for integrals $\int g(\theta) f(\theta | x_n) d\pi(\theta)$ based on asymptotic normality results (like Equation (5)), and it applies even if $f(\theta | x_n)$ is only proportional to the posterior distribution. It also applies for integrating out nuisance parameters if regularity conditions hold. Another effective technique for approximating $p(\theta_1 | x_n)$ (not just the MLE or MAP) when it is “messy” but $p(\theta_1 | \theta_2, x_n)$ and $p(\theta_2 | \theta_1, x_n)$ are easy to manipulate is data augmentation, often called the IP algorithm (for imputation, posterior algorithm). *Importance sampling (IS) remains a powerful methods for efficient integration.*

For simulation of variates, classical methods for generating independent variates from posterior distributions may apply. Posterior distributions are often known only up to a constant of proportionality (the denominator in Bayes rule may be hard to compute). It is therefore important to have a method to simulate variates for arbitrary functions proportional to posterior distributions. *Markov Chain Monte Carlo (MCMC) is the most important of those methods at present.* MCMC constructs a Markov chain whose stationary distribution is the desired posterior distribution $p(\theta | \mathcal{E})$ (Gilks et al. 1996). The state θ_t of the Markov chain can then be sampled to obtain (somewhat correlated) samples from $p(\theta | \mathcal{E})$.

A number of tools implement Bayesian inference. The BUGS and WinBUGS packages implement MCMC tools, and are available on the WWW (Spiegelhalter et al. 1996). BOA, for Bayesian output analysis (Smith 2004), is a set of MCMC diagnostic tools for convergence and data analysis that functions with the R or S-PLUS statistical packages. Gauss and Matlab are also used to implement MCMC methods.

At present, it is possible to input randomized input parameters to some commercial discrete-event simulation packages to implement the BMA algorithm of Figure 2, but interfaces are not yet fully user friendly. A user-friendly tool to implement the BMA and other uncertainty analysis needs in commercial discrete-event simulation packages would be helpful. Uncertainty analysis for other Monte Carlo applications have been available as a spreadsheet tool for some time (e.g., Winston 2000).

5 CONCLUSIONS

A variety of applications of Bayesian methods apply to simulation experiments, including uncertainty analysis, ranking and selection, input distribution modeling, response surface modeling, and experimental designs. One main theme is to represent all uncertainty with probability distributions, to update probability using Bayes’ rule, and to use the expected value of information as a technique to make sampling decisions (e.g., the opportunity cost and 0-1 loss functions for selection procedures, or the Kullback-Leibler divergence for parameter estimation for linear response models). Another theme is to use simulation to efficiently estimate quantities of interest for a Bayesian analysis. Asymptotic approximations are often helpful when exact optimal solutions are difficult to obtain.

An advantage of the Bayesian and decision theoretic tools is that they can easily model the economics of managerial decision-making. Exploiting the economics of decision making is an interesting area for future work, and can help simulation, as a tool, become more useful in broader managerial contexts.

Other research opportunities include further work in: input modeling and uncertainty analysis (kernel estimation of outputs as a result of input uncertainty; the effect of different candidate distributions on uncertainty; tools to help elicit prior distributions for simulation inputs); response modeling (further extending the Gaussian random field work in the world of stochastic simulation; sampling plans for input parameter inference to optimally reduce output uncertainty, including nonasymptotic results, to help understand what data is most important to collect to infer the value of inputs for simulations; theory and improved implementation for the calibration/inverse problem, i.e., how to “invert” probability distributions for output results backwards to obtain a reasonable probability distribution on unknown

“inputs”, a problem that arises regularly in public policy and health contexts); Bayesian methods for experimental designs (estimating quantiles or other non-expected value goals; CRN for unknown input parameters for ranking and selection; non-Gaussian output for ranking and selection and GRFs); improved computational tools (in MCMC and other sampling methods for posterior distributions).

REFERENCES

- Andradóttir, S., and P. W. Glynn. 2004. Computing Bayesian means using simulation. Georgia Tech, ISYE Technical Report.
- Barlow, R. E., and M. B. Mendel. 1992. De Finetti-type representations for life distributions. *Journal of the American Statistical Association* 87:1116–1122.
- Barton, R. R., and L. W. Schruben. 2001. Simulating real systems. In submission.
- Berger, J. O. 1994. An overview of robust Bayesian analysis. *TEST* 3:5–124.
- Berger, J. O., and L. R. Pericchi. 1996. The intrinsic Bayes factor for model selection and prediction. *Journal of the American Statistical Association* 91:109–122.
- Berk, R. 1966. Limiting behaviour of posterior distributions when the model is incorrect. *Annals of Mathematical Statistics* 37 (1): 51–58.
- Bernardo, J. 1979. Expected information as expected utility. *Annals of Statistics* 7:686–690.
- Bernardo, J. M., and A. F. M. Smith. 1994. *Bayesian theory*. Chichester, UK: Wiley.
- Billingsley, P. 1986. *Probability and Measure*. 2nd ed. New York: John Wiley & Sons, Inc.
- Branke, J., S. E. Chick, and C. Schmidt. 2005. Selecting a selection procedure.
- Chen, C.-H. 1996. A lower bound for the correct subset-selection probability and its application to discrete event simulations. *IEEE Transactions on Automatic Control* 41 (8): 1227–1231.
- Chen, C.-H., E. Yücesan, L. Dai, and H. Chen. 2005. Efficient computation of optimal budget allocation for discrete event simulation experiment. *IIE Transactions*: to appear.
- Cheng, R. C. H. 1999. Regression metamodelling in simulation using Bayesian methods. In *Proceedings of the 1999 Winter Simulation Conference*, ed. P. A. Farrow, H. B. Nembhard, D. Sturrock, and G. Evans, 330–335.
- Cheng, R. C. H., and W. Holland. 1997. Sensitivity of computer simulation experiments to errors in input data. *Journal of Statistical Computing and Simulation* 57:219–241.
- Chick, S. E. 1997. Bayesian analysis for simulation input and output. In *Proceedings of the 1997 Winter Simulation Conference*, ed. S. Andradóttir, K. Healy, D. Withers, and B. Nelson, 253–260.
- Chick, S. E. 2001. Input distribution selection for simulation experiments: Accounting for input uncertainty. *Operations Research* 49 (5): 744–758.
- Chick, S. E. 2004. Bayesian methods for discrete event simulation. In *Proceedings of the 2004 Winter Simulation Conference*, ed. R. Ingalls, M. Rossetti, J. Smith, and B. Peters, 89–100.
- Chick, S. E. 2006. Subjective probability and Bayesian methodology. In *Handbook in Operations Research and Management Science: Simulation*, ed. S. Henderson and B. Nelson. Elsevier.
- Chick, S. E., and N. Gans. 2005. The economics of simulation selection procedures. Technical report, INSEAD, Technology and Operations Management Area working paper.
- Chick, S. E., and K. Inoue. 2001a. New two-stage and sequential procedures for selecting the best simulated system. *Operations Research* 49 (5): 732–743.
- Chick, S. E., and K. Inoue. 2001b. New procedures for identifying the best simulated system using common random numbers. *Management Science* 47 (8): 1133–1149.
- Chick, S. E., and K. Inoue. 2002. Corrigendum: New selection procedures. *Operations Research* 50 (3): 566.
- de Finetti, B. 1990. *Theory of Probability*, v. 2. New York: John Wiley & Sons, Inc.
- de Groot, M. H. 1970. *Optimal statistical decisions*. New York: McGraw-Hill.
- Devroye, L. 2006. Nonuniform random variate generation. In *Handbook in Operations Research and Management Science: Simulation*, ed. S. G. Henderson and B. L. Nelson. Elsevier.
- Draper, D. 1995. Assessment and propagation of model uncertainty (with discussion). *Journal of the Royal Statistical Society, Series B* 57 (1): 45–97.
- Edwards, A. W. F. 1984. *Likelihood*. Cambridge University Press.
- Evans, M., and T. Swartz. 1995. Methods for approximating integrals in statistics with special emphasis on Bayesian integration problems. *Statistical Science* 10 (3): 254–272.
- George, E. I., and R. E. McCulloch. 1996. Stochastic search variable selection. In *Markov Chain Monte Carlo in Practice*, ed. W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, 203–214. London: Chapman and Hall.
- Gilks, W. R., S. Richardson, and D. J. Spiegelhalter. 1996. *Markov chain Monte Carlo in practice*. London: Chapman and Hall.
- Glynn, P. 1986. Problems in Bayesian analysis of stochastic simulation. In *Proceedings of the 1986 Winter Simulation Conference*, ed. J. R. Wilson, J. O. Henriksen, and S. D. Roberts, 376–383.

- He, D., S. E. Chick, and C.-H. Chen. 2006. The opportunity cost and OCBA selection procedures in ordinal optimization. Submitted.
- Jaynes, E. T. 1983. *E. T. Jaynes: Papers on probability, statistics and statistical physics*. Dordrecht, Holland: D. Reidel Publishing Co.
- Jeffreys, H. 1946. An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London Set A* 186:453–461.
- Kahneman, D., P. Slovic, and A. Tversky. 1982. *Judgment under uncertainty: Heuristics and biases*. Cambridge: Cambridge University Press.
- Kass, R. E., and L. Wasserman. 1996. The selection of prior distributions by formal rules. *Journal of the American Statistical Association* 91 (435): 1343–1370.
- Kennedy, M., and A. O’Hagan. 2001. Bayesian calibration of computer models (with discussion). *Journal of the Royal Statistical Society B* 63:425–464.
- Laplace, P. S. 1812. *Theorie analytique des probabilités*. Paris: Courcier.
- Law, A. M., and W. D. Kelton. 2000. *Simulation modeling & analysis*. 3rd ed. New York: McGraw-Hill, Inc.
- Lee, S.-H., and P. Glynn. 1999. Computing the distribution function of a conditional expectation via Monte Carlo: Discrete conditioning spaces. In *Proceedings of the 1999 Winter Simulation Conference*, ed. P. A. Farrington, H. B. Nembhard, D. Sturrock, and G. Evans, 1654–1663.
- Lindley, D. 1957. A statistical paradox. *Biometrika* 44:187–192.
- Ng, S.-H., and S. E. Chick. 2001. Reducing input distribution uncertainty for simulations. In *Proceedings of the 2001 Winter Simulation Conference*, ed. B. Peters, J. Smith, M. Rohrer, and D. Madeiros, 364–371.
- Ng, S.-H., and S. E. Chick. 2004. Design of follow-up experiments for improving model discrimination and parameter estimation. *Naval Research Logistics* 51 (8): 1129–1148.
- Ng, S.-H., and S. E. Chick. 2006. Reducing parameter uncertainty for stochastic systems. *ACM Transactions on Modeling and Computer Simulation* 15 (1): in press.
- Porteus, E. L. 2002. *Stochastic inventory theory*. Stanford University Press.
- Raftery, A. E., D. Madigan, and J. A. Hoeting. 1997. Bayesian model averaging for linear regression models. *Journal of the American Statistical Association* 92 (437): 179–191.
- Santner, T. J., B. J. Williams, and W. I. Notz. 2003. *The design and analysis of computer experiments*. New York: Springer.
- Savage, L. J. 1972. *The foundations of statistics*. New York: Dover Publications, Inc.
- Smith, A. F. M., and I. Verdinelli. 1980. A note on Bayesian design for inference using hierarchical linear model. *Biometrika* 67 (3): 613–619.
- Smith, B. J. 2004. Bayesian output analysis program (BOA) v. 1.1 user’s manual. Technical report, Univ. Iowa, <<http://www.public-health.uiowa.edu/boa/>>.
- Spiegelhalter, D. J., N. G. Best, W. R. Gilks, and H. Inskip. 1996. Hepatitis B: a case study in MCMC methods. In *Markov Chain Monte Carlo in Practice*, ed. W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, 21–44. London: Chapman and Hall.
- Steckley, S. G., and S. G. Henderson. 2003. A kernel approach to estimating the density of a conditional expectation. In *Proceedings of the 2003 Winter Simulation Conference*, ed. S. Chick, P. Sánchez, D. Ferrin, and D. Morrice, 383–391.
- Tanner, M. A. 1996. *Tools for statistical inference: Methods for exploration of posterior distributions and likelihood functions*. 3rd ed. Springer.
- van Beers, W., and J. P. C. Kleijnen. 2003. Kriging for interpolation in random simulation. *Journal of the Operational Research Society* 54:255–262.
- van Beers, W. C., and J. P. C. Kleijnen. 2004. Kriging interpolation in simulation: A survey. In *Proceedings of the 2004 Winter Simulation Conference*, ed. R. Ingalls, M. Rossetti, J. Smith, and B. Peters.
- Winston, W. 2000. *Financial models using simulation and optimization*. 2nd ed. Palisade Corp.
- Zouaoui, F., and J. R. Wilson. 2003. Accounting for parameter uncertainty in simulation input modeling. *IIE Transactions* 35:781–792.
- Zouaoui, F., and J. R. Wilson. 2004. Accounting for input-model and input-parameter uncertainties in simulation. *IIE Transactions* 36 (11): 1135–1151.

AUTHOR BIOGRAPHY

STEPHEN E. CHICK is Associate Professor of Technology and Operations Management, and is based at the Fontainebleau, France campus of INSEAD. In addition to stochastic simulation, his research focuses on Bayesian statistics, operations management, public health, health economics, decision analysis and computational methods in statistics. He is President of the INFORMS Simulation Society, and is part of the editorial teams of *Operations Research* and *ACM TOMACS*. His e-mail is <stephen.chick@insead.edu>.