Proceedings of the 1987 Winter Simulation Conference
A. Thesen, H. Grant, W. David Kelton (eds.)

LIKELIHOOD RATIO GRADIENT ESTIMATION: AN OVERVIEW

Peter W. Glynn
Department of Industrial Engineering
University of Wisconsin
Madison, WI 53706

ABSTRACT

The likelihood ratio method for gradient estimation is briefly surveyed. Two applications settings are described, namely Monte Carlo optimization and statistical analysis of complex stochastic systems. Steady-state gradient estimation is emphasized, and both regenerative and non-regenerative approaches are given. The paper also indicates how these methods apply to general discrete-event simulations; the idea is to view such systems as general state space Markov chains.

1. INTRODUCTION

Consider a single-server queue in which the service rate θ is a decision variable. Given that $\alpha(\theta)$ is the steady-state cost of running the queue at parameter level θ, one is frequently interested in minimizing $\alpha(\theta)$ over a suitable constraint set. Since $\alpha(\cdot)$ is often difficult to evaluate analytically, Monte Carlo optimization is an attractive methodology. By analogy with deterministic mathematical programming, efficient Monte Carlo gradient estimation is typically an important ingredient of simulation based optimization algorithms. As a consequence, gradient estimation has recently attracted considerable attention in the simulation community. It is our goal, in this paper, to describe one such method for estimating gradients in the Monte Carlo setting, namely the likelihood ratio method.

In Section 2, we describe two important problems which motivate our study of efficient gradient estimation algorithms. Section 3 is devoted to the derivation of the likelihood ratio gradient estimate for transient estimation problems in a discrete-time Markov chain setting. Section 4 extends the methodology to steady-state gradient estimation by using regenerative structure; in Section 5, a technique for non-regenerative systems is explored. Section 6 describes the specialization of these techniques to the Markov chains associated with discrete-event simulations, while Section 7 states some conclusions.

2. EFFICIENT GRADIENT ESTIMATION: MOTIVATING APPLICATIONS

As indicated in the Introduction, one motivation for studying Monte Carlo gradient estimation is for the purpose of optimizing complex stochastic systems. More precisely, consider a stochastic system depending on d decision variables $\theta_1, \theta_2, \ldots, \theta_d$. Let $\alpha(\theta)$ ($\theta = (\theta_1, \ldots, \theta_d)$) be the expected "cost" of running the system at parameter choice θ.

A powerful method for computing the value θ^* which minimizes $\alpha(\cdot)$ is the Robbins-Monro algorithm. This technique recognizes that, under suitable regularity on $\alpha(\cdot)$, θ^* must be a θ-root of the equation

$$\nabla\alpha(\theta) = 0, \tag{2.1}$$

where $\nabla\alpha(\theta)$ is the gradient of $\alpha(\cdot)$ evaluated at θ. The idea then is to construct a stochastic recursion which has the root θ^* as its limit point.

This approach is most clearly illustrated when $d = 1$. In this case, such a recursion is given by

$$\theta_{n+1} = \theta_n - \frac{a}{n} V_{n+1} \tag{2.2}$$

($a > 0$) where the V_n's mimic $\alpha'(\cdot)$ in expectation. More precisely, one is required to compute V_n's with the property that

$$E\{V_{n+1} | V_0, \theta_0, \ldots, V_n, \theta_n\} = \alpha'(\theta_n) \text{ a.s.} \tag{2.3}$$

Under appropriate additional hypotheses, it then follows that there exists finite σ such that

$$\begin{aligned} &\theta_n \to \theta^* \text{ a.s. as } n \to \infty \\ &n^{1/2}(\theta_n - \theta^*) \Rightarrow \sigma N(0,1) \end{aligned} \tag{2.4}$$

where $N(0,1)$ is a standard normal r.v. The key result in (2.4) is the central limit theorem which asserts that θ_n converges to θ^* at rate $O_p(n^{-1/2})$. (A stochastic sequence $\{H_n : n \geq 0\}$ is said to be $O_p(a_n)$ if $\{a_n^{-1} H_n : n \geq 0\}$ is tight.) Since a convergence rate $O_p(n^{-1/2})$ is typically the best that one can expect of a Monte Carlo algorithm (because of central limit effects), this suggests that recursive algorithms of the form (2.2) should lead to reasonably efficient procedures for calculating θ^*. Of course, the critical component of such an algorithm is the sequence of gradient estimates (derivative estimates when $d = 1$) $\{V_n : n \geq 0\}$ appearing in (2.3). Thus, efficient stochastic optimization is one setting which requires gradient estimation.

A second problem context which leads naturally to gradient estimation is statistical estimation for complex stochastic systems. As an example, consider a single-server infinite capacity queue in which the inter-arrival distribution F_a and service distribution F_s are unknown. Suppose that one is given data $X_1, X_2, \ldots, X_n$ for the inter-arrival times and observations $Y_1, \ldots, Y_m$ for the service times, with the goal of estimating the steady-state queue-length α. The parameter α may then be regarded as a function of the inter-arrival and service time distributions i.e. $\alpha = \alpha(F_a, F_s)$. If F_a^* and F_s^* are respectively the "true" inter-arrival and service time distributions, our goal here is to estimate $\alpha^* = \alpha(F_a^*, F_s^*)$ from the data.

Assume that F_a^*, F_s^* are elements of one parameter families of distributions $\{F_a(\theta_1)\}, \{F_s(\theta_2)\}$, respectively, such that $F_a^* = F_a(\theta_1^*), F_s^* = F_s(\theta_2^*)$. We can then reduce the problem of estimating α^* to that of determining $\tilde{\alpha}(\theta_1^*, \theta_2^*)$, when $\tilde{\alpha}(\theta_1, \theta_2) = \alpha(F_a(\theta_1), F_s(\theta_2))$. For example, if $F_a(\theta_1)$ and $F_s(\theta_2)$ are both exponential, the resulting system is an M/M/1 queue with ($\tilde{\alpha}$ can be calculated analytically here)

$$\tilde{\alpha}(\theta_1, \theta_2) = \begin{cases} (\theta_1/\theta_2)(1 - (\theta_1/\theta_2))^{-1}, & \theta_1 < \theta_2 \\ \infty, & \theta_1 \geq \theta_2. \end{cases}$$

On the other hand, if $F_a(\cdot)$ and $F_s(\cdot)$ are uniform on $[0, \theta_1]$ and $[0, \theta_2]$ respectively, $\tilde{\alpha}$ is not available in closed form, and Monte Carlo evaluation may be necessary.

The natural estimate for α^* is $\hat{\alpha} = \hat{\tilde{\alpha}}(\hat{\theta}_1, \hat{\theta}_2)$, where $\hat{\theta}_1$ is an estimate for θ_1^* calculated from $X_1, \ldots, X_n$ and $\hat{\theta}_2$ is an estimate for θ_2^* derived from $Y_1, \ldots, Y_m$; $\hat{\tilde{\alpha}}(\cdot)$ is a Monte Carlo estimate of $\tilde{\alpha}(\cdot)$. To calculate the error in $\hat{\alpha}$ as an estimate of α^*, note that

$$\begin{aligned} \hat{\alpha} - \alpha^* &= [\hat{\tilde{\alpha}}(\hat{\theta}_1, \hat{\theta}_2) - \tilde{\alpha}(\hat{\theta}_1, \hat{\theta}_2)] \\ &\quad + [\tilde{\alpha}(\hat{\theta}_1, \hat{\theta}_2) - \tilde{\alpha}(\theta_1^*, \theta_2^*)]. \end{aligned} \tag{2.5}$$

The first term on the right-hand side of (2.5) is error incurred from the Monte Carlo estimation of $\tilde{\alpha}(\hat{\theta}_1, \hat{\theta}_2)$; the second term, which is (conditionally) independent of the first, reflects the intrinsic error in α^* due to uncertainty in the data sets. The error in the first term can be estimated from conventional output analysis procedures. For the second, note that if $\alpha(\cdot)$ is differentiable, then

$$\tilde{\alpha}(\hat{\theta}_1, \hat{\theta}_2) - \tilde{\alpha}(\theta_1^*, \theta_2^*) \approx \nabla\tilde{\alpha}(\theta^*)(\hat{\theta} - \theta^*).$$

Typically, the vector $\hat{\theta} - \theta^*$ will be a mean zero multivariate normal, with a covariance matrix that can be easily estimated from the data sets. (This occurs, for example, if the $\hat{\theta}_i$'s are maximum likelihood estimators for the θ_i^*'s.) To calculate the distribution of the second term, it therefore remains to compute $\nabla\tilde{\alpha}(\theta^*)$ or, more precisely, its estimator $\nabla\tilde{\alpha}(\hat{\theta})$. For analytically intractable models (such as the single-server infinite capacity queue with uniform inter-arrival and service time distributions), this entails calculating a gradient via Monte Carlo simulation.

The situation described above in the single-server queueing context is typical of many statistical problems that arise in the analysis of complex stochastic systems. To fully resolve the statistical error then generally requires Monte Carlo estimation of an appropriate gradient.

3. DERIVATION OF LIKELIHOOD RATIOS FOR MARKOV CHAINS

In this section, we derive likelihood ratio gradient estimators for discrete-time Markov chains. Our view is that discrete-event simulations can be characterized probabilistically as discrete-time Markov chains. In particular, suppose that one views the "state" as incorporating all that information about the discrete-event system which needs to be computationally updated on every transition of the process (e.g. event list, clock times, and physical state). Thus, one can view a computer program for a discrete-event simulation as an implementation of the recursion

$$X_{n+1} = f_n(X_n, \eta_{n+1}) \tag{3.1}$$

where X_n is the "state" of the system at the n'th transition, and η_{n+1} is a vector incorporating all new random variables which need to be computed in order to calculate X_{n+1} from X_n. The mappings f_n are complicated functions which are rarely considered explicitly by the simulator, but which are mathematical representations of the computational algorithm used to obtain X_{n+1} from X_n and η_{n+1}. We will return to this point in Section 6 when we consider generalized semi-Markov processes. In any case, any sequence $\{X_n : n \geq 0\}$ satisfying (3.1) is Markov, since

$$P\{X_{n+1} \epsilon \cdot | X_0, \ldots, X_n\} = P\{f_n(X_n, \eta_{n+1}) \epsilon \cdot | X_n\} = Q_n(X_n, \cdot)$$

where $Q_n(x, \cdot) = P\{f_n(x, \eta_{n+1}) \epsilon \cdot\}$. The above equality follows from the fact that the new r.v.'s which are generated at the $(n+1)$'st transition are independent of everything previously generated. In most discrete-event simulations, the transition mechanism is time-homogeneous, so that $f_n \equiv f$ and $\eta_{n+1} \overset{D}{=} \eta$; the Markov chain

$\{X_i : n \geq 0\}$ is then itself time-homogeneous so that $Q_n \equiv Q$.

Note that for most discrete-event systems, the Markov chain $X = \{X_n : n \geq 0\}$ defined by (3.1) has both a complicated state space and complex transition rule. To simplify our exposition here, we therefore start by considering likelihood ratio gradient estimates for discrete state space Markov chains. For each θ in some open set, suppose that $P(\theta)$ is the transition matrix associated with the choice θ of the parameter value. We further assume that a cost $g(\theta, i_0, \ldots, i_n)$ is incurred when the sample sequence $(X_0, \ldots, X_n)$ takes on the values $(i_0, \ldots, i_n)$. In this case, the expected "cost" of running the chain X at parameter value θ takes the form

$$\alpha(\theta) = E_\theta g(\theta, X_0, \ldots, X_n) \tag{3.2}$$

where $E_\theta(\cdot)$ reflects the fact that the probabilistic dynamics of X are governed by $P(\theta)$.

If $E_\theta(\cdot)$ were independent of θ, our solution to the Monte Carlo gradient estimation problem would be trivial, namely to simulate i.i.d. replicates of the random vector $\nabla g(\theta, X_0, \ldots, X_n)$. The trick is therefore to transform (3.2) into a representation where the expectation operator is independent of θ. To accomplish this, observe that

$$\begin{aligned}\alpha(\theta) &= \sum_{i_0,\ldots,i_n} g(\theta, i_0, \ldots, i_n)\mu(\theta, i_0) \prod_{k=0}^{n-1} P(\theta, i_k, i_{k+1}) \\ &= \sum_{i_0,\ldots,i_n} g(\theta, i_0, \ldots, i_n) L_n(\theta, i_0, \ldots, i_n)\mu(\theta_0, i_0) \prod_{k=0}^{n-1} P(\theta_0, i_k, i_{k+1})\end{aligned} \tag{3.3}$$

where

$$L_n(\theta, i_0, \ldots, i_n) = \frac{\mu(\theta, i_0)}{\mu(\theta_0, i_0)} \prod_{k=0}^{n-1} \frac{P(\theta, i_k, i_{k+1})}{P(\theta_0, i_k, i_{k+1})}. \tag{3.4}$$

We assume here implicitly (as throughout this paper) that appropriate positivity conditions are in force so as to guarantee that no divisions by zero occur in (3.4).

Returning to (3.4), we can easily verify that

$$\alpha(\theta) = E_{\theta_0} g(\theta, X_0, \ldots, X_n) L_n(\theta) \equiv E_{\theta_0} \tilde{g}(\theta) \tag{3.5}$$

where

$$L_n(\theta) = L_n(\theta, X_0, \ldots, X_n).$$

The crucial point in (3.5) is that the expectation operator appearing on the right-hand side is independent of θ. One therefore obtains that $\nabla\alpha(\theta) = E_{\theta_0} \nabla\tilde{g}(\theta)$. Specifically, one has the relation

$$\frac{\partial\alpha}{\partial\theta_i}(\theta) = E_{\theta_0} \frac{\partial\tilde{g}}{\partial\theta_i}(\theta)$$

where

$$\frac{\partial\tilde{g}}{\partial\theta_i}(\theta) = \frac{\partial g}{\partial\theta_i}(\theta, X_0, \ldots, X_n) L_n(\theta) + g(\theta, X_0, \ldots, X_n)\frac{\partial}{\partial\theta_i} L_n(\theta)$$

and

$$\frac{\partial L_n}{\partial\theta_i}(\theta) = \frac{\partial\mu}{\partial\theta_i}(\theta, X_0) \cdot \frac{L_n(\theta)}{\mu(\theta, X_0)} + \sum_{j=0}^{n-1} \frac{\partial P(\theta, X_j, X_{j+1})}{\partial\theta_i} \cdot \frac{L_n(\theta)}{P(\theta, X_j, X_{j+1})}$$

Thus, by simulating $X_0, \ldots, X_n$ under initial distribution $\mu(\theta_0)$ and transition matrix $P(\theta_0)$, we can calculate $\partial\tilde{g}(\theta)/\partial\theta_i$ and thereby estimate $\nabla\alpha(\theta)$. Observe that the estimator $\nabla\tilde{g}(\theta)/\partial\theta_i$ contains the product terms

$$L_n(\theta) = \prod_{k=0}^{n-1} \frac{P(\theta, X_k, X_{k+1})}{P(\theta_0, X_l, X_{k+1})} \cdot \frac{\mu(\theta, X_0)}{\mu(\theta_0, X_0)}.$$

We claim that the choice $\theta_0 = \theta'$ is particularly convenient for evaluating $\nabla\alpha(\theta')$. In this case, $L_n(\theta') = 1$ so that the computation involved in calculating the estimator $\nabla\tilde{g}(\theta')$ is reduced. Furthermore, for large n, this choice substantially reduces the variance of $\nabla\tilde{g}(\theta')$. To see this, note that $E_{\theta_0} L_n(\theta') = 1$, so that

$$\mathrm{var}_{\theta_0} L_n(\theta') = E_{\theta_0} L_n^2(\theta') - 1.$$

Assuming that $P(\theta_0)$ is positive recurrent with stationary probabilities $\pi(\theta_0)$,

$$\begin{aligned}\frac{1}{n}\log L_n^2(\theta') &= \frac{1}{n}\sum_{k=0}^{n-1} \log \frac{P^2(\theta', X_k, X_{k+1})}{P^2(\theta_0, X_k, X_{k+1})} + \frac{1}{n}\log\frac{\mu^2(\theta', X_0)}{\mu^2(\theta_0, X_0)} \\ &\sim \sum_{i,j} \pi(\theta_0, i) P(\theta_0, i, j) \log\left\{\frac{P^2(\theta', i, j)}{P^2(\theta_0, i, j)}\right\} \equiv \varphi\end{aligned}$$

Hence,

$$\frac{1}{n} E_{\theta_0} \log L_n^2(\theta') \sim \varphi$$

so that, by Jensen's inequality,

$$E_{\theta_0} L_n^2(\theta') \geq \exp(n\varphi/2)$$

for n sufficiently large. Since $E_{\theta_0} L_n^2(\theta') \geq (E_{\theta_0} L_n(\theta'))^2 = 1$, it follows that $\varphi \geq 0$ (with strict inequality holding when $L_n(\theta')$ is non-deterministic). We conclude that if $\theta_0 \neq \theta'$, the variance of $L_n(\theta')$ generally grows exponentially fast in n. One would expect this exponential growth to significantly impact the variance of $\nabla \tilde{g}(\theta')$ for large n.

We turn now to the generalization of this approach to Markov chains having a general state space; this generalization is necessary in order to apply this methodology to Markov chains of the type arising in discrete-event simulation. The analog of the initial distribution vector $\mu(\theta)$ is an initial probability distribution

$$\mu(\theta, A) \equiv P_\theta\{X_0 \epsilon A\}$$

whereas the transition matrix $P(\theta)$ is replaced by the transition kernel

$$P(\theta, x, A) \equiv P_\theta\{X_{n+1} \epsilon A | X_n = x\}.$$

We require that $\mu(\theta), P(\theta)$ have densities, in the sense that

$$\begin{aligned} \mu(\theta, A) &= \int_A u(\theta, y)\mu(dy) \\ P(\theta, x, A) &= \int_A p(\theta, x, y) P(x, dy), \end{aligned} \tag{3.6}$$

for some (σ-finite) measures $\mu, P(x, \cdot)$. It is easily verified that

$$\begin{aligned} \alpha(\theta) &= E_\theta g(\theta, X_0, \ldots, X_n) \\ &= E_{\theta_0} g(\theta, X_0, \ldots, X_n) L_n(\theta) \end{aligned}$$

where

$$L_n(\theta) = \frac{u(\theta, X_0)}{u(\theta_0, X_0)} \prod_{k=0}^{n-1} \frac{p(\theta, X_k, X_{k+1})}{p(\theta_0, X_k, X_{k+1})},$$

and $E_\theta(\cdot)$ is the expectation operator corresponding to the probability measure $P_\theta\{X_0 \epsilon dx_0, \ldots, X_n \epsilon dx_n\} = \mu(\theta, dx_0) \cdot \prod_{k=0}^{n-1} P(\theta, x_k, dx_{k+1})$. As in the discrete state space case, choosing $\theta_0 = \theta'$ makes sense in evaluating $\nabla\alpha(\theta')$ via Monte Carlo simulation. In this case, we obtain $\nabla\alpha(\theta') = E_{\theta_0} \nabla \tilde{g}(\theta')$, where

$$\frac{\partial \tilde{g}}{\partial \theta_i}(\theta) = \frac{\partial g}{\partial \theta_i}(\theta', X_0, \ldots, X_n) + g(\theta', X_0, \ldots, X_n) \frac{\partial}{\partial \theta_i} L_n(\theta') \tag{3.7}$$

and

$$\frac{\partial L_n}{\partial \theta_i}(\theta') = \frac{\partial u(\theta', X_0)/\partial \theta_i}{u(\theta', X_0)} + \sum_{j=0}^{n-1} \frac{\partial p(\theta', X_j, X_{j+1})/\partial \theta_i}{p(\theta', X_j, X_{j+1})}. \tag{3.8}$$

Thus, under the density hypothesis (3.6), it is straightforward to calculate an unbiased estimator for $\nabla\alpha(\theta')$:

1. Generate $X_0, \ldots, X_n$ under $\mu(\theta')$ and $P(\theta')$.
2. Calculate the r.v.'s $\partial\tilde{g}(\theta)/\partial\theta_i$ and $\partial L_n(\theta')/\partial\theta_i$ from (3.6) and (3.7) and the sample path $X_0, \ldots, X_n$ generated in 1.

By replicating steps 1 and 2, one can easily construct an estimator (just use the sample mean) which converges to $\nabla\alpha(\theta')$ at rate $O_p(t^{-1/2})$ (use the multivariate central limit theorem) in the computational effort t. We have therefore obtained a gradient estimator which converges at the best possible Monte Carlo convergence rate, namely $O_p(t^{-1/2})$.

Variants of the gradient estimator algorithm described above have been analyzed in Glynn (1986a), Reiman and Weiss (1986), and Rubinstein (1986).

4. STEADY-STATE GRADIENT ESTIMATORS: REGENERATIVE ANALYSIS

The method outlined in Section 3 was valid for cost functionals $g(\theta, X_0, \ldots, X_n)$ which depend on the chain X up to a deterministic finite time horizon n. In fact, the method is equally valid for functionals $g(\theta, X_0, \ldots, X_T)$ depending on the chain up to a stopping time T. To be precise, suppose that

$$\alpha(\theta) = E_\theta g(\theta, X_0, \ldots, X_T),$$

where $E_\theta(\cdot)$ is the expectation on the path-space of $X = \{X_n : n \geq 0\}$ corresponding to initial distribution $\mu(\theta)$ and transition kernel $P(\theta)$. Then, $\nabla\alpha(\theta') = E_{\theta'} \nabla\tilde{g}(\theta')$ where

$$\begin{aligned} \frac{\partial}{\partial \theta_i}\tilde{g}(\theta') &= \frac{\partial g}{\partial \theta_i}(\theta', X_0, \ldots, X_T) \\ &+ g(\theta', X_0, \ldots, X_T) \left\{ \frac{\partial u(\theta', X_0)/\partial \theta_i}{u(\theta', X_0)} + \sum_{j=0}^{T-1} \frac{\partial p(\theta', X_j, X_{j+1})/\partial \theta_i}{p(\theta', X_j, X_{j+1})} \right\} \end{aligned} \tag{4.1}$$

An alternative estimator can be developed when $g(\theta, X_0, \ldots, X_T)$ is an additive functional of the form

$$g(\theta, X_0, \ldots, X_T) = \sum_{j=0}^{T} h(\theta, X_j). \tag{4.2}$$

In this case, we can use the fact that

$$E_{\theta'}\left\{\frac{p(\theta, X_j, X_{j+1})}{p(\theta', X_j, X_{j+1})}|X_0,\ldots,X_j\right\} = 1$$

so that

$$E_{\theta'}\left\{\frac{\partial p(\theta', X_j, X_{j+1})/\partial\theta_i}{p(\theta', X_j, X_{j+1})}|X_0,\ldots,X_j\right\} = 0. \tag{4.3}$$

By conditioning in (4.1) and using relations (4.2) and (4.3), we find that

$$E_{\theta'}\frac{\partial}{\partial\theta_i}\tilde{g}(\theta') = E_{\theta'}\Lambda(h,\theta') \tag{4.4}$$

where

$$\begin{aligned}\Lambda(h,\theta') &= \sum_{j=0}^{T}\frac{\partial}{\partial\theta_i}h(\theta', X_j) + \sum_{j=0}^{T}h(\theta', X_j)\\ &\quad\left\{\sum_{\ell=0}^{j-1}\frac{\partial p(\theta', X_\ell, X_{\ell+1})/\partial\theta_i}{p(\theta', X_\ell, X_{\ell+1})} + \frac{\partial u(\theta', X_0)/\partial\theta_i}{u(\theta', X_0)}\right\}\\ &= \sum_{j=0}^{T}\frac{\partial}{\partial\theta_i}h(\theta', X_j) + \sum_{j=0}^{T-1}\frac{\partial p(\theta', X_j, X_{j+1})/\partial\theta_i}{p(\theta', X_j, X_{j+1})}\sum_{\ell=j+1}^{T}h(\theta', X_\ell)\\ &\quad + \frac{\partial u(\theta', X_0)/\partial\theta_i}{u(\theta', X_0)}\sum_{j=0}^{T}h(\theta', X_j).\end{aligned} \tag{4.5}$$

Relations (4.3) and (4.5) will prove useful in our regenerative analysis of steady-state gradient estimation.

Consider a family of transition kernels $P(\theta)$ having unique stationary distributions $\pi(\theta)$, and suppose that we wish to calculate the gradient of

$$\alpha(\theta) = \int \pi(\theta, dx)h(\theta, x).$$

Of course, $\alpha(\theta)$ may generally be regarded as the expected cost, under $P_\theta(\cdot)$, of the functional

$$g(\theta, X_0,\ldots) = \liminf_{n\to\infty}\frac{1}{n}\sum_{k=0}^{n-1}h(\theta, X_k).$$

Note that

$$\alpha(\theta) \approx \liminf_{n\to\infty} E_\theta g_n(\theta, X_0,\ldots,X_{n-1})$$

when

$$g_n(\theta, X_0,\ldots,X_{n-1}) = \frac{1}{n}\sum_{k=0}^{n-1}h(\theta, X_k).$$

We can try to approximate $\nabla\alpha(\theta)$ via $E_{\theta'}\nabla\tilde{g}_n(\theta')$, where

$$\begin{aligned}\frac{\partial\tilde{g}_n(\theta')}{\partial\theta_i} &= \frac{1}{n}\sum_{k=0}^{n-1}\partial h(\theta', X_k)/\partial\theta_i\\ &+ \frac{1}{n}\sum_{k=0}^{n-1}h(\theta', X_k)\cdot\left\{\frac{\partial u(\theta', X_0)/\partial\theta_i}{u(\theta', X_0)} + \sum_{j=0}^{n-1}\frac{\partial p(\theta', X_j, X_{j+1})/\partial\theta_i}{p(\theta', X_j, X_{j+1})}\right\}\end{aligned} \tag{4.6}$$

The first sum in (4.6) satisfies a strong law, and therefore converges. The second quantity on the right-hand side is a product of two factors, the first of which satisfies a strong law with limit $\alpha(\theta')$. The second factor, which involves a sum of terms of the form $\frac{\partial p}{\partial\theta_i}(\theta', X_j, X_{j+1})/p(\theta', X_j, X_{j+1})$, can be analyzed via the central limit theorem (use (4.3)), yielding

$$n^{-1/2}\sum_{j=0}^{n-1}\frac{\partial p(\theta', X_j, X_{j+1})/\partial\theta_i}{p(\theta', X_j, X_{j+1})} \Rightarrow \sigma N(0,1) \tag{4.7}$$

for some constant σ. By squaring both sides of (4.7) and taking expectations, we find that

$$\text{var}_{\theta'}\left\{\sum_{j=0}^{n-1}\frac{\partial p(\theta', X_j, X_{j+1})/\partial\theta_i}{p(\theta', X_j, X_{j+1})}\right\} \sim \sigma^2 n,$$

This suggests that

$$\text{var}_{\theta'}\frac{\partial\tilde{g}_n(\theta')}{\partial\theta_i} \sim \sigma^2\cdot\alpha(\theta')^2\cdot n$$

as $n\to\infty$. We conclude that we can expect the variance of $\partial\tilde{g}_n(\theta')/\partial\theta_i$ to increase linearly with n. Thus, in trying to approximate a steady-state gradient, the approximants become increasingly less stable. This conclusion, which was previously observed by Reiman and Weiss (1986), leads one to look for alternative approaches.

One way to do this is to assume that the sequence $X = \{X_n : n \geq 0\}$ possesses readily identifiable regenerative structure. In this case, assuming that μ is a regenerative initial condition with T its associated regeneration time, the ratio formula of regenerative analysis shows that

$$\alpha(\theta) = \frac{E_\theta\sum_{k=0}^{T-1}h(\theta, X_k)}{E_\theta T} \equiv \frac{u(\theta)}{\ell(\theta)}.$$

Then

$$\begin{aligned}\frac{\partial\alpha(\theta')}{\partial\theta_i} &= \frac{1}{\ell(\theta')^2}\left\{\frac{\partial u}{\partial\theta_i}(\theta')\ell(\theta') - \frac{\partial\ell}{\partial\theta_i}(\theta')u(\theta')\right\}\\ &= \frac{1}{\ell(\theta')}\left\{\frac{\partial u}{\partial\theta_i}(\theta') - \alpha(\theta')\frac{\partial\ell}{\partial\theta_i}(\theta')\right\}\\ &= \frac{\partial}{\partial\theta_i}E_\theta\sum_{k=0}^{T-1}w(\theta, X_k)/E_\theta T\end{aligned}$$

where $w(\theta, x) = h(\theta, x) - \alpha(\theta')$. It remains to evaluate the above partial derivative in terms of a quantity amenable to Monte Carlo estimation.

By applying (4.1), we find that

$$\frac{\partial}{\partial\theta_i} E_\theta \sum_{k=0}^{T-1} w(\theta, X_k) = \sum_{j=0}^{T-1} \frac{\partial}{\partial\theta_i} h(\theta', X_j) + \sum_{j=0}^{T-1} \frac{\partial p(\theta', X_j, X_{j+1})/\partial\theta_i}{p(\theta', X_j, X_{j+1})} \sum_{\ell=0}^{T-1} \{h(\theta', X_\ell) - \alpha(\theta')\}$$

Hence, $\nabla\alpha(\theta')$ can be estimated by using the following algorithm.

ALGORITHM A:

1. Choose a sample size $n \geq 1$, where n corresponds to the number of regenerative cycles to be simulated.
2. Generate a sample path $X_0, \ldots, X_{T-1}$ (i.e. generate X over a regenerative cycle) under $P(\theta')$.
3. Calculate the quantities:

$$Q_{11} = T$$
$$Q_{12} = \sum_{j=0}^{T-1} h(\theta', X_j)$$
$$Q_{13} = \sum_{j=0}^{T-1} \frac{\partial}{\partial\theta_i} h(\theta', X_j)$$
$$Q_{14} = \sum_{j=0}^{T-1} \frac{\partial p(\theta', X_j, X_{j+1})/\partial\theta_i}{p(\theta', X_j, X_{j+1})} \cdot T$$
$$Q_{15} = \sum_{j=0}^{T-1} \frac{\partial p(\theta', X_j, X_{j+1})/\partial\theta_i}{p(\theta', X_j, X_{j+1})} \cdot \sum_{\ell=0}^{T-1} h(\theta', X_\ell)$$

4. Replicate steps 2 and 3 n times, thereby obtaining $Q_{ij}, 1 \leq i \leq n, 1 \leq j \leq 5$.
5. Calculate

$$\partial_i \hat{\alpha}_1(\theta') = \frac{\bar{Q}_3(n)}{\bar{Q}_1(n)} + \frac{\bar{Q}_5(n)}{\bar{Q}_1(n)} - \frac{\bar{Q}_2(n)}{\bar{Q}_1(n)} \cdot \frac{\bar{Q}_4(n)}{\bar{Q}_1(n)},$$

where $\bar{Q}_i(n) = \sum_{j=1}^n Q_{ji}/n (1 \leq i \leq 5)$; this estimator converges to $\partial\alpha(\theta')/\partial\theta_i$ as $n \to \infty$.

A second regenerative estimator for $\nabla\alpha(\theta')$ uses (4.3) and (4.5). It is easily shown that

$$\frac{\partial}{\partial\theta_i}\alpha(\theta') = \frac{1}{E_{\theta'}T} \cdot E_{\theta'} \left[\sum_{j=0}^{T-2} \frac{\partial p(\theta', X_j, X_{j+1})/\partial\theta_i}{p(\theta', X_j, X_{j+1})} \cdot \sum_{\ell=j+1}^{T-1} \{h(\theta', X_\ell) - \alpha(\theta')\} \right] + \frac{1}{E_{\theta'}T} E_{\theta'} \sum_{j=0}^{T-1} \frac{\partial}{\partial\theta_i} h(\theta', X_j). \tag{4.8}$$

This gives rise to a second algorithm for estimating $\nabla\alpha(\theta')$.

ALGORITHM B:

1. Choose a sample size $n \geq 1$, where n corresponds to the number of regenerative cycles to be simulated.
2. Generate a sample path $X_0, \ldots, X_{T-1}$ (i.e. generate X over a regenerative cycle) under $P(\theta')$.
3. Calculate the quantities:

$$R_{11} = T$$
$$R_{12} = \sum_{j=0}^{T-1} h(\theta', X_j)$$
$$R_{13} = \sum_{j=0}^{T-1} \frac{\partial}{\partial\theta_i} h(\theta', X_j)$$
$$R_{14} = \sum_{j=0}^{T-2} \frac{\partial p(\theta', X_j, X_{j+1})/\partial\theta_i}{p(\theta', X_j, X_{j+1})} \cdot (T - 1 - j)$$
$$R_{15} = \sum_{j=0}^{T-1} \frac{\partial p(\theta', X_j, X_{j+1})/\partial\theta_i}{p(\theta', X_j, X_{j+1})} \cdot \sum_{\ell=j+1}^{T-1} h(\theta', X_\ell)$$

4. Replicate steps 2 and 3 n times, thereby obtaining $R_{ij}, 1 \leq i \leq n, 1 \leq j \leq 5$.
5. Calculate

$$\partial_i \hat{\alpha}_2(\theta') = \frac{\bar{R}_3(n)}{\bar{R}_1(n)} + \frac{\bar{R}_5(n)}{\bar{R}_1(n)} - \frac{\bar{R}_2(n)}{\bar{R}_1(n)} \cdot \frac{\bar{R}_4(n)}{\bar{R}_1(n)}$$

where $\bar{R}_i(n) = \sum_{j=1}^n R_{ji}/n (1 \leq i \leq 5)$; this estimator converges to $\partial\alpha(\theta')/\partial\theta_i$ as $n \to \infty$.

It is easily verified, via standard arguments, that the estimators described in Algorithms A and B converge at rate $O_p(t^{-1/2})$ in the computational effort t.

5. NON-REGENERATIVE STEADY-STATE GRADIENT ESTIMATORS

We turn now to the case where the sequence $X = \{X_n : n \geq 0\}$ exhibits no obvious regenerative structure. The regenerative results of Section 4 actually provide the key to the analysis.

Turning to (4.8), we note that the second sum can be expressed as a steady-state expectation i.e.

$$\frac{1}{E_{\theta'}T} E_{\theta'} \sum_{j=0}^{T-1} \frac{\partial}{\partial\theta_i} h(\theta', X_j) = \tilde{E}_{\theta'} \frac{\partial}{\partial\theta_i} h(\theta', X_0)$$

where $\tilde{E}_{\theta'}$ is the expectation on the path-space of X associated with transition kernel $P(\theta')$ and initial distribution $\pi(\theta')$. For the first term, a more intricate analysis is necessary.

Let $\tau^k X = (X_k, X_{k+1}, \ldots)$. For a function f defined on the infinite product space, an easy extension of the regenerative ratio formula proves that

$$\frac{1}{E_{\theta'}T}E_{\theta'}\sum_{j=0}^{T-1} f(\tau^j X) = \tilde{E}_{\theta'} f(X)$$

Applying this formula to

$$f(X) = \frac{\partial}{\partial\theta_i}p(\theta', X_0, X_1) \cdot \frac{I(T \geq 1)}{p(\theta', X_0, X_1)} \sum_{j=1}^{T-1}\{h(\theta', X_j) - \alpha(\theta')\},$$

we obtain the relation

$$\begin{aligned} &\frac{\partial}{\partial\theta_i}\alpha(\theta') \\ &= \tilde{E}_{\theta'}\frac{\partial}{\partial\theta_i}h(\theta', X_0) \\ &+ \tilde{E}_{\theta'}\frac{\partial}{\partial\theta_i}p(\theta', X_0, X_1)\frac{I(T \geq 1)}{p(\theta, X_0, X_1)}\sum_{j=1}^{T-1}\{h(\theta', X_j) - \alpha(\theta')\}. \end{aligned}$$

Let $T_1, T_2, \ldots$ be the successive regeneration times for X. By the ratio formula for regenerative processes

$$\alpha(\theta') = \frac{\tilde{E}_{\theta'}\sum_{j=T_k}^{T_{k+1}-1} h(\theta', X_j)}{\tilde{E}_{\theta'}(T_{k+1} - T_k)}$$

and hence

$$\tilde{E}_{\theta'}\sum_{j=T_k}^{T_{k+1}-1}\{h(\theta', X_j) - \alpha(\theta')\} = 0.$$

By the independence of regenerative cycles, we get

$$\begin{aligned} \frac{\partial}{\partial\theta_i}\alpha(\theta') &= \tilde{E}_{\theta'}\frac{\partial}{\partial\theta_i}h(\theta', X_0) \\ &+ \tilde{E}_{\theta'}\frac{\partial}{\partial\theta_i}p(\theta', X_0, X_1) \cdot \frac{1}{p(\theta', X_0, X_1)}\sum_{j=1}^{T_n-1}\{h(\theta', X_j) - \alpha(\theta')\} \\ &= \tilde{E}_{\theta'}\frac{\partial}{\partial\theta_i}h(\theta', X_0) \\ &+ \sum_{j=1}^{\infty}\tilde{E}_{\theta'}\frac{\partial}{\partial\theta_i}p(\theta', X_0, X_1) \cdot \frac{1}{p(\theta', X_0, X_1)}\{h(\theta', X_j) - \alpha(\theta')\} \cdot I(T_n > j) \end{aligned}$$

Let $n \to \infty$ and we obtain

$$\begin{aligned} \frac{\partial}{\partial\theta_i}\alpha(\theta') &= \tilde{E}_{\theta'}\frac{\partial}{\partial\theta_i}h(\theta', X_0) + \sum_{j=1}^{\infty}\tilde{E}_{\theta'}\frac{\partial}{\partial\theta_i}p(\theta', X_0, X_1) \cdot \frac{(h(\theta', X_j) - \alpha(\theta'))}{p(\theta', X_0, X_1)} \\ &= \tilde{E}_{\theta'}\frac{\partial}{\partial\theta_i}h(\theta', X_0) + \sum_{j=1}^{\infty}\tilde{E}_{\theta'}\frac{\partial}{\partial\theta_i}p(\theta', X_0, X_1)\frac{h(\theta', X_j)}{p(\theta', X_0, X_1)} \end{aligned} \tag{5.1}$$

((4.3) was used in the last equality). The important point is that expression (5.1), while derived from a regenerative argument, is <u>independent</u> of regenerative structure.

The same expression can be found via a totally different argument. Recall that the stationary distribution $\pi(\theta)$ satisfies

$$\pi(\theta, \cdot) = \int_S P(\theta, x, \cdot)\pi(\theta, dx), \tag{5.2}$$

where S is the state space of X. Then, it is reasonable to assume that $P(\cdot)$ can be expanded as

$$P(\theta' + he_i) = P(\theta') + hQ_i(\theta') + o(h) \tag{5.3}$$

where e_i is the i'th unit vector. If $\pi(\theta' + he_i)$ is (formally) differentiable at $h = 0$, there exists a measure $\eta_i(\theta')$ such that

$$\pi(\theta' + he_i) = \pi(\theta') + h\eta_i(\theta') + o(h). \tag{5.4}$$

Plugging (5.4) and (5.3) into the stationarity equation (5.2) and collecting terms in h yields

$$\begin{aligned} \eta_i(\theta', \cdot) - \int_S \eta_i(\theta', dx)P(\theta', x, \cdot) \\ = \int_S \pi(\theta', dx)Q_i(\theta', x, \cdot). \end{aligned}$$

In operator form, this can be written as

$$\eta_i(\theta')(I - P(\theta')) = \pi(\theta')Q_i(\theta'). \tag{5.5}$$

We wish to solve for $\eta_i(\theta')$. Let $\Pi(\theta', x, \cdot) = \pi(\theta', \cdot)$ for all $x \epsilon S$. Observe that (5.4) implies that $\eta_i(\theta', S) = 0$ (just divide by h and let $h \to 0$), and hence

$$\int_S \eta_i(\theta', dx)\Pi(\theta', x, \cdot) = 0. \tag{5.6}$$

Substituting (5.6) into (5.5) provides

$$\eta_i(\theta')(I - P(\theta') + \Pi(\theta')) = \pi(\theta')Q_i(\theta').$$

Now, for many Markov chains (in particular, aperiodic positive recurrent Harris chains), $P^k(\theta', x, \cdot) \to \pi(\theta', \cdot)$ for all $x \epsilon S$, and it therefore makes sense to assume that

$$(I - P(\theta') + \Pi(\theta'))^{-1} = I + \sum_{k=1}^{\infty}(P^k(\theta') - \Pi(\theta'))$$

exists. (Just use the identity $P(\theta')\Pi(\theta') = \Pi(\theta') = \Pi(\theta')P(\theta')$.) Thus,

$$\eta_i(\theta') = \sum_{k=0}^{\infty}\pi(\theta')Q_i(\theta')P^k(\theta') \tag{5.7}$$

(use (5.6) again). Recall that

$$h\frac{\partial}{\partial\theta_i}\alpha(\theta') \approx \pi(\theta'+he_i)h(\theta'+he_i) - \pi(\theta')h(\theta') \approx h\eta_i(\theta')h(\theta') + \pi(\theta')\frac{\partial}{\partial\theta_i}h(\theta')h \tag{5.8}$$

(Expand and collect terms in h again.) Substituting (5.7) into (5.8), we get

$$\begin{aligned}&\frac{\partial}{\partial\theta_i}\alpha(\theta')\\&= \int_S \pi(\theta',dx)\frac{\partial}{\partial\theta_i}h(\theta',x)\\&+\sum_{k=0}^{\infty}\int_S \pi(\theta',dx)\int_S Q_i(\theta',x,dy)\int_S P^k(\theta',y,dz)h(\theta',z).\end{aligned} \tag{5.9}$$

We now identify $Q_i(\theta',x,dy)$ in our current framework. Note that (see(3.5))

$$P(\theta'+he_i,x,A) = P(\theta',x,A) + h\int_A \frac{\partial}{\partial\theta_i}p(\theta',x,y)\frac{P(\theta',x,dy)}{p(\theta',x,y)}$$

so

$$Q_i(\theta',x,dy) = \frac{\partial p(\theta',x,y)/\partial\theta_i}{p(\theta',x,y)}P(\theta',x,dy). \tag{5.10}$$

Substituting into (5.9) yields (5.1). Formula (5.1) is the fundamental relation for non-regenerative stochastic systems. Notice that the first term on the right-hand side of (5.1) can be consistently estimated via

$$\frac{1}{n}\sum_{k=0}^{n-1}\frac{\partial}{\partial\theta_i}h(\theta',X_k) \tag{5.11}$$

whereas the j'th term in the infinite sum appearing there can be estimated by using

$$\frac{1}{n}\sum_{k=0}^{n-1}\frac{\partial}{\partial\theta_i}p(\theta',X_k,X_{k+1})\frac{h(\theta',X_{j+k})}{p(\theta',X_k,X_{k+1})} \tag{5.12}$$

A standard device for estimating the entire infinite sum is to use an estimator which combines (5.11) and (5.12), namely to use

$$\sum_{j=1}^{\ell(n)}\frac{1}{n-\ell(n)}\sum_{k=0}^{n-\ell(n)-1}\frac{\partial}{\partial\theta_i}p(\theta',X_k,X_{k+1})\frac{h(\theta',X_{j+k})}{p(\theta',X_k,X_{k+1})} + \frac{1}{n+1}\sum_{j=0}^{n}\frac{\partial}{\partial\theta_i}h(\theta',X_j)$$

where $\ell(n)$ is keyed to the sample size n in such a way that $\ell(n)\to\infty$ with $\ell(n)/n \to 0$. The particular choice of $\ell(n)$ effects a compromise between bias and variance effects in estimating the infinite sum (5.1). This approach appeared previously in Glynn (1986b).

6. LIKELIHOOD RATIO GRADIENT ESTIMATION FOR DISCRETE-EVENT SIMULATIONS

As indicated in Section 3, discrete-event simulations can be viewed as Markov chains living on a general state space. To be precise, discrete-event simulations can be viewed mathematically as a "generalized semi-Markov process' (GSMP). Such a process is characterized by:

> S: a "physical" state space which is countable (e.g. S might be the set of all possible queue-length vectors for a queueing simulation).
>
> E: a set of events to be scheduled (e.g. for each station in a closed queueing network, one needs to schedule an "end of service" event).
>
> $p(s',s,e)$: the probability of jumping from s to s', given that event e triggers the transition from s (e.g. e might correspond to station i completing service, in which case $p(s';s,e)$ might represent the probability of sending a customer from station i to station j; here $s' = s - e_i + e_j$).
>
> r_{se}: the rate at which clock e runs down to zero in state s (e.g. in a queueing network, r_{se} might be unity except for events e which are "interrupted" in state s, in which case $r_{se}=0$).
>
> $F(\cdot;s',e',s,e)$: the probability distribution which schedules event e' in state s', given that the previous state was s and the transition was triggered by e (e.g. these might be service time distributions in a closed queueing network).

In calculating gradients, we allow $p(s';s,e)$ and $F(\cdot;s',e',s,e)$ to depend on the decision parameter θ; the likelihood ratio method is generally inapplicable to problems in which S, E, or r_{se} depends on θ. We further require that $F(\theta;\cdot;s',e',s,e)$ have a density for which the support is independent of θ, so that

$$F(\theta,dx;s',e',s,e) = f(\theta,x;s',e',s,e)\mu(dx)$$

where $\{x: f(\theta,x;s',e',s,e)>0\}$ is independent of θ. (This is the analogue of the positivity condition discussed in Section 3.) This density hypothesis rules out point mass distributions in which θ controls the location of the points; the independent support condition does not permit uniform distributions with support on $[0,\theta]$.

To make a discrete-event simulation Markov, we consider the state of the simulation at transition epochs. Specifically, set $X_n = (S_n, C_n)$, where S_n is the "physical" state occupied at transition n and C_n is the state of the "clocks" on the event scheduling list at the n'th transition. Then, $\{X_n : n \geq 0\}$ is a Markov chain with a complicated state space (inclusion of the clocks makes the state space uncountable). To study the ergodic behavior of a GSMP $\{Y(t) : t \geq 0\}$, note that

$$\frac{1}{t}\int_0^t a(Y(s))ds \approx \frac{\sum_{k=0}^{N(t)} a(S_k)C_k^*}{\sum_{k=0}^{N(t)} C_k^*}$$

for t large, where C_k^* is the time spent in the k'th state visited, and $N(t)$ is the number of transitions by time t. (Note that C_k^* is a simple function of C_k, namely the minimal value of $C_{ke}/r_{S_k,e}$ taken over all clocks e.) If the GSMP is well behaved, we can expect that

$$\frac{1}{t}\int_0^t a(Y(s))ds \rightarrow \tilde{E}_{\theta'}a(S_0)C_0^*/\tilde{E}_{\theta'}C_0^* \; P_{\theta'} a.s.$$

as $t \rightarrow \infty$. The objective of calculating steady-state gradients for GSMP's therefore reduces to estimating the gradients of $\tilde{E}_{\theta'}a(S_0)C_0^*$ and $\tilde{E}_{\theta'}C_0^*$. This can be done by the methods of Sections 4 and 5 (apply to $h(X_k) = a(S_k)C_k^*$ and $h(X_k) = C_k^*$). It remains only to identify the analogue of

$$\frac{\partial}{\partial\theta}p(\theta', X_j, X_{j+1})/p(\theta', X_j, X_{j+1}) \tag{6.1}$$

for this particular class of Markov chains in which $X_j = (S_j, C_j)$. Note that under parameter θ, $X_{j+1} = (S_{j+1}, C_{j+1})$ is obtained from $X_j = (S_j, C_j)$ by:

a.) making a state transition from S_j to S_{j+1} with a probability $p(\theta, S_{j+1}; S_j, e_j^*)$, where e_j^* is the event that triggered the transition from S_j.

b.) certain clocks belonging to the (random) set O_{j+1} continue to be scheduled in S_{j+1} and run down deterministically there.

c.) the remaining events $e \epsilon N_{j+1}$ active in S_{j+1} are scheduled according to the distributions $F(\theta, \cdot; S_{j+1}, e, S_j, e_j^*)$, and set to new values $C_{j+1,e}$.

The analogue of (6.1) can be easily verified to be

$$\frac{\partial}{\partial\theta}p(\theta', S_{j+1}; S_j, e_j^*) \cdot \frac{1}{p(\theta', S_{j+1}; S_j, e_j^*)} + \sum_{e \epsilon N_{j+1}} \frac{\partial}{\partial\theta}f(\theta', C_{j+1,e}; S_{j+1}, e, S_j, e_j^*) \cdot \frac{1}{f(\theta', C_{j+1,e}; S_{j+1}, e, S_j, e_j^*)} \tag{6.2}$$

The algorithm discussed in Sections 4 and 5 can then be applied to general discrete-event simulations, by substituting (6.2) appropriately.

7. CONCLUSION

We have shown that gradient estimation plays an important role in the optimization of stochastic systems, as well as in their statistical analysis. The likelihood ratio method described here is easily applied to discrete-event simulations of arbitrary complexity (see Section 6), and does not require case-by-case analysis for implementation. On the other hand, this method is inapplicable to problems in which the settings of deterministic event times are decision variables. (See the density conditions in Section 6.) Such problems frequently arise in a manufacturing context. Nevertheless, we believe that the methods described here form a promising avenue for future research.

ACKNOWLEDGEMENTS

The author gratefully acknowledges the support of the United States Army under Contract No. DAAG29-84-K-0030 and by the National Science Foundation under NSF Grant ECS-840-4809.

REFERENCES

Glynn, P. W. (1986a). Stochastic approximation for Monte Carlo optimization. In: *Proceedings of the 1986 Winter Simulation Conference*, p. 356–365.

Glynn, P. W. (1986b). Sensitivity analysis for stationary probabilities of a Markov chain. In: *Proceedings of the Fourth Army Conference on Applied Mathematics and Computing.*

Reiman, M. I. and Weiss, A. (1986). Sensitivity analysis via likelihood ratios. In: *Proceedings of the 1986 Winter Simulation Conference*, p. 285–289.

Rubinstein, R. Y. (1981). Sensitivity analysis and performance extrapolation for computer simulation models, Technical Report, Harvard University, Cambridge, MA.

AUTHOR'S BIOGRAPHY

PETER W. GLYNN is an associate professor in the Department of Industrial Engineering at the University of Wisconsin-Madison. He received a B.Sc. in mathematics from Carleton University in 1978 and a Ph.D. in operations research from Stanford University in 1982. His current research interests include Monte Carlo simulation, computational probability, and queueing theory. He is a member of ORSA, TIMS, the IMS, and the Statistical Society of Canada.

Peter W. Glynn
Department of Industrial Engineering
University of Wisconsin-Madison
1513 University Avenue
Madison, WI 53706
608/263-6790