

OPTIMIZATION OF STOCHASTIC SYSTEMS
VIA SIMULATION

Peter W. Glynn
Department of Operations Research
Stanford University
Stanford, CA 94305

ABSTRACT

In this paper, we discuss some research issues related to the general topic of optimizing a stochastic system via simulation. In particular, we devote extensive attention to finite-difference estimators of objective function gradients and present a number of new limit theorems. We also discuss a new family of orthogonal function approximations to the global behavior of the objective function. We show that if the objective function is sufficiently smooth, the convergence rate can be made arbitrarily close to $n^{-1/2}$ in the number of observations required. The paper concludes with a brief discussion of how these ideas can be integrated into an optimization algorithm.

1. INTRODUCTION

In recent years, considerable attention has been devoted, in the simulation literature, to the development of algorithms for optimizing complex stochastic systems. In this paper, we shall focus on describing some of the basic issues that arise in the study of numerical optimization routines for finite-dimensional continuous parameter optimization problems.

To precisely describe the class of problems that we shall consider, let $\theta \in \Lambda$ be the decision parameter over which the optimization is to occur; the set $\Lambda \subseteq \mathbb{R}^d$ is the admissible set of decision parameters. For each $\theta \in \Lambda$, let $(\Omega, \mathcal{F}, P_\theta)$ be the associated probability space. The probability measure P_θ describes how the random environment is affected by the choice of θ . For each $\theta \in \Lambda$, let $X(\theta)$ be a real-valued random variable corresponding to the "cost" of running the system under θ . Then

$$(1.1) \quad \alpha(\theta) \triangleq \int_{\Omega} X(\theta, \omega) P(d\omega)$$

is the expected cost of running the system under parameter θ . Assuming Λ is some open subset of \mathbb{R}^d , the general finite-dimensional continuous parameter stochastic optimization problem involves finding $\theta^* \in \Lambda$ to minimize $\alpha(\theta)$, subject (possibly) to constraints of the form

$$\beta_i(\theta) \triangleq \int_{\Omega} Y_i(\theta, \omega) P(d\omega) \geq 0,$$

$1 \leq i \leq m$, where $\{Y_i(\theta) : 1 \leq i \leq m\}$ is a collection of "random constraints."

In most practical applications, the objective function $\alpha(\theta)$ and the constraints $\beta_i(\theta)$ are "smooth" functions of the decision parameter θ (even though, typically, $X(\theta, \omega)$ and the $Y_i(\theta, \omega)$'s are not globally smooth in θ , for fixed ω). Given that the functions $\alpha(\theta)$ and $\beta_i(\theta)$ ($1 \leq i \leq m$) can be cheaply evaluated without error, deterministic mathematical programming techniques may be applied to the above optimization problem. Such methods typically take advantage of derivative information of some kind (often evaluated through numerically stable finite-difference approximations). Of course, in the context of a complex stochastic system, the objective function $\alpha(\theta)$ and the constraints $\beta_i(\theta)$ will typically be evaluated via Monte Carlo simulation. As a consequence, there will be random error associated with the corresponding function evaluations. In spite of the presence of such error, it is to be expected that derivative information will continue to play an important role in the development of successful optimization algorithms based on simulation. A significant portion of this paper is therefore devoted to a discussion of the various approaches that may be used to calculate derivatives (or, more generally, gradients) via simulation.

Section 2 is devoted to a discussion of the convergence characteristics of finite-difference estimators; much of this material appears here for the first time. In Section 3, we describe a class of unbiased gradient estimators that are based on likelihood ratio ideas. Section 4 focuses on a class of gradient estimation techniques for discrete-event systems known as perturbation analysis methods. The estimators of Sections 3 and 4 both typically attain a somewhat faster convergence rate than that available through the finite-difference methods of Section 2. The discussion of Sections 2 through 4 emphasizes the scalar setting in which $d = 1$; Section 5 is therefore devoted to describing the extension of these ideas to the case in which the decision parameter θ is vector-valued.

In Section 6, we discuss some new results related to global approximation of the objective function (and/or constraints) by orthogonal functions (specifically, trigonometric polynomials). One way to apply such “surface fitting” techniques is to optimize the fitted surface (using deterministic methods) and to use the resulting optimizer as an approximation to the optimizer of the true surface.

Section 7 is devoted to a discussion of how the results of Section 2 through 6 can be used in an optimization setting. Specifically, we discuss some of the convergence theory for the Robbins-Monro and Kiefer-Wolfowitz algorithms.

2. FINITE-DIFFERENCE ESTIMATORS

In this section, we describe some of the finite-difference approximations that can be used to numerically calculate the derivative of a function $\alpha(\theta)$ of the form (1.1), when θ is scalar (i.e., $\theta \in \mathbb{R}$). In Section 5, we discuss the special considerations that arise in dealing with gradients of functions in which the parameter θ is vector-valued (i.e., $\theta \in \mathbb{R}^d$).

2.1 Forward-Difference Estimators

Suppose that we wish to estimate $\alpha'(\theta_0)$. The idea here is to estimate, via simulation, the function values $\alpha(\theta_0 + h)$ and $\alpha(\theta_0)$ and to form a corresponding finite-difference approximation to $\alpha'(\theta_0)$. More specifically, let $X_1(\theta_0 + h), X_2(\theta_0 + h), \dots$ be i.i.d. replicates of the r.v. $X(\theta_0 + h)$, simulated under common distribution $P_{\theta_0 + h}$. Similarly, we let $X_1(\theta_0), X_2(\theta_0), \dots$ be an independent stream of i.i.d. replicates of the r.v. $X(\theta_0)$, generated under common distribution P_{θ_0} . Consider the forward-difference estimator

$$\Delta\alpha_1(n, h) = \frac{1}{n} \sum_{i=1}^n \left(\frac{X_i(\theta_0 + h) - X_i(\theta_0)}{h} \right).$$

The determination of the best possible difference increment h introduces a trade-off between the variance of the estimator and its bias. If h is chosen too small (relative to n), the variance contribution to the mean square error will dominate, whereas if h is chosen too large (relative to n), the bias will govern the convergence rate. It turns out that the optimal difference increment $h = h_n$, in this setting, is typically of order $n^{-1/4}$. To rigorously state the result, we assume that:

- (2.1)
- i) $P_\theta\{X(\theta)\varepsilon\cdot\} \Rightarrow P_{\theta_0}\{X(\theta_0)\varepsilon\cdot\}$ as $\theta \rightarrow \theta_0$ (\Rightarrow denotes convergence in distribution),

- ii) $0 < \text{var}_{\theta_0} X(\theta_0) \triangleq E_{\theta_0}(X(\theta_0) - \alpha(\theta_0))^2 < \infty$ ($E_\theta(\cdot)$ denotes the expectation operator corresponding to P_θ),
- iii) $\sigma^2(\theta) \triangleq \text{var}_\theta X(\theta)$ is continuous in an open neighborhood of θ_0 ,
- iv) $\alpha(\theta) = E_\theta X(\theta)$ is infinitely differentiable in an open neighborhood of θ_0 .

The following theorem states that if the difference increment is chosen optimally, then the convergence rate of $\Delta\alpha_1(n, h_n)$ to $\alpha'(\theta_0)$ is $n^{-1/4}$.

- (2.2) **THEOREM.** Assume (2.1). If $\alpha''(\theta_0) \neq 0$, then:
- a) if $n^{1/4}h_n \rightarrow \infty$ with $h_n \rightarrow 0$, $n^{1/4}|\Delta\alpha_1(n, h_n) - \alpha'(\theta_0)| \Rightarrow \infty$ as $n \rightarrow \infty$ (we say that $Z_n \Rightarrow \infty$ as $n \rightarrow \infty$ if, for every $K > 0$, $P\{Z_n > K\} \rightarrow 1$ as $n \rightarrow \infty$),
 - b) if $n^{1/4}h_n \rightarrow 0$, $n^{1/4}|\Delta\alpha_1(n, h_n) - \alpha'(\theta_0)| \Rightarrow \infty$ as $n \rightarrow \infty$,
 - c) if $n^{1/4}h_n \rightarrow h > 0$, then $n^{1/4}(\Delta\alpha_1(n, h_n) - \alpha'(\theta_0)) \Rightarrow \frac{\sqrt{2}\sigma(\theta_0)}{h}N(0, 1) - h\alpha''(\theta_0)/2$ as $n \rightarrow \infty$.

The proof of this result appears in the Appendix. (A similar theorem, under different hypotheses, appears in FOX and GLYNN (1989).) We note that the value of h which minimizes the second moment of the limiting r.v. appearing in c) is

$$h^* = \left(\frac{8\sigma^2(\theta_0)}{\alpha''(\theta_0)^2} \right)^{1/4}.$$

Thus, the difference increment that minimizes asymptotic mean square error is $h_n = h^*n^{-1/4}$. This result was obtained previously by ZAZANIS and SURI (1986). It is worth observing that if one wishes to minimize the mean absolute error of the estimator, then the optimal difference increment takes the form $h_n = h_*n^{-1/4}$, where (typically) $h_* \neq h^*$. (To see this, observe that h_* would be obtained by minimizing the first absolute moment of the limiting normal r.v. appearing in c.) Stated more abstractly, the L^2 and L^1 error criteria do not yield precisely the same sequence of optimal difference increments.

2.2 Central-Difference Estimators

Theorem 2.2 states that the forward-difference estimator converges (at best) at rate $n^{-1/4}$ to the derivative $\alpha'(\theta_0)$. One way to improve upon this poor convergence rate is to instead use a central-difference approximation to the derivative. When function evaluations are made without error, this is known to be a numerically more accurate approximation to the derivative.

To precisely define the estimator, we let $X_1(\theta_0 + h)$, $X_2(\theta_0 + h)$, \dots be i.i.d. replicates of the r.v. $X(\theta_0 + h)$, simulated under common distribution P_{θ_0+h} . Similarly, we let $X_1(\theta_0 - h)$, $X_2(\theta_0 - h)$, \dots be an independent stream of i.i.d. replicates of the r.v. $X(\theta_0 - h)$, generated under P_{θ_0-h} . The central-difference estimator is defined as

$$\Delta\alpha_2(n, h) = \frac{1}{n} \sum_{i=1}^n \left(\frac{X_i(\theta_0 + h) - X_i(\theta_0 - h)}{2h} \right).$$

The following theorem summarizes the behavior of $\Delta\alpha_2(n, h)$; the proof is similar to that of Theorem 2.2 and is omitted.

- (2.3) **THEOREM.** Assume (2.1). If $\alpha^{(3)}(\theta_0) \neq 0$, then:
- if $n^{1/6}h_n \rightarrow \infty$ with $h_n \rightarrow 0$, $n^{1/3}|\Delta\alpha_2(n, h_n) - \alpha'(\theta_0)| \Rightarrow \infty$ as $n \rightarrow \infty$,
 - if $n^{1/6}h_n \rightarrow 0$, $n^{1/3}|\alpha_2(n, h_n) - \alpha'(\theta_0)| \Rightarrow \infty$ as $n \rightarrow \infty$,
 - if $n^{1/6}h_n \rightarrow h > 0$, then $n^{1/3}(\Delta\alpha_2(n, h_n) - \alpha'(\theta_0)) \Rightarrow \sqrt{2} \frac{\alpha(\theta_0)}{h} N(0, 1) - \frac{h^2}{3} \alpha^{(3)}(\theta_0)$ as $n \rightarrow \infty$.

The improved convergence rate (of order $n^{-1/3}$, as opposed to $n^{-1/4}$ for forward differences) is obtained here because of the fact that central differences are less biased than forward differences. This permits the difference increment to be chosen larger (of order $n^{-1/6}$, as opposed to $n^{-1/4}$ for forward differences) which, in turn, reduces the variability of the estimator.

The choice of h in c) that minimizes the asymptotic mean square error of the central difference estimator is (see also ZAZANIS and SURI (1986))

$$h^* = \left(\frac{9\sigma^2(\theta_0)}{\alpha^{(3)}(\theta_0)^2} \right)^{1/6}.$$

2.3 Finite-Difference Estimators Using Common Random Numbers

The central-difference estimator improves upon the convergence rate of the forward-difference estimator by reducing its bias (for fixed h). The method that we shall describe here improves upon the convergence rate of the forward-difference estimator by reducing its variability (for fixed h). The idea is to generate the replicates of $X(\theta_0 + h)$ using the same stream of random numbers that were used to obtain the replicates of $X(\theta_0)$. This, of course, is nothing more than the

method of common random numbers, as applied to derivative estimation.

Suppose that the r.v. $Y(\theta_0 + h)$ is produced from the same stream of random numbers as is $X(\theta_0)$ and shares the same distribution as does $X(\theta_0 + h)$ under P_{θ_0+h} . By convention, we set $Y(\theta_0) = X(\theta_0)$. Let $\Delta Y(h) = Y(\theta_0 + h) - Y(\theta_0)$. We make the following assumptions about our common random number scheme:

- (2.4)
- $P\{Y(h)\varepsilon\cdot\} = P_{\theta_0+h}\{X(\theta_0 + h)\varepsilon\cdot\}$,
 - $E\Delta Y(h)^2 = h\sigma_1^2 + o(h)$ as $h \downarrow 0$, where $\sigma_1^2 > 0$,
 - there exists $\varepsilon > 0$ such that $E|\Delta Y(h)|^{2+\varepsilon} = h\beta_1 + o(h)$ as $h \downarrow 0$,
 - $\alpha(\theta) = E_\theta X(\theta)$ is two times continuously differentiable in an open neighborhood of θ_0 .

Let $\Delta Y_1(h), \Delta Y_2(h), \dots$ be an i.i.d. sequence of replicates of the r.v. $\Delta Y(h)$. The forward-difference common random numbers estimator for $\alpha'(\theta_0)$ is then given by

$$\Delta\alpha_3(n, h) = \frac{1}{n} \sum_{i=1}^n \frac{\Delta Y_i(h)}{h}.$$

Before proceeding to a statement of the convergence rate theorem for $\Delta\alpha_3(n, h)$, we pause to discuss our assumptions further. Consider the typical discrete-event system. Let $A(h)$ be the event that $Y(\theta_0 + h)$ experiences a change in the order of events from that experienced by $Y(\theta_0)$. On the event $A(h)$, $Y(\theta_0 + h) - Y(\theta_0)$ is typically of unit magnitude. On the other hand, on the complement of $A(h)$, $Y(\theta_0 + h) - Y(\theta_0)$ is typically of order h in magnitude. Also, for most discrete-event systems $P(A(h)) = \lambda h + o(h)$ for some $\lambda > 0$. For $p > 0$, write

$$E\Delta Y(h)^p = E\Delta Y(h)^p I(A(h)) + E\Delta Y(h)^p I(A(h)^c).$$

This decomposition suggests that for most discrete-event systems, $E\Delta Y(h)^p = \beta_p h + o(h)$ for $p \geq 1$. This explains the form of (2.4) ii) and iii).

- (2.5) **THEOREM.** Assume (2.4). If $\alpha''(\theta_0) \neq 0$ and $nh_n \rightarrow \infty$, then:

- if $n^{1/3}h_n \rightarrow \infty$ with $h_n \rightarrow 0$, $n^{1/3}|\Delta\alpha_3(n, h_n) - \alpha'(\theta_0)| \Rightarrow \infty$ as $n \rightarrow \infty$,
- if $n^{1/3}h_n \rightarrow 0$, $n^{1/3}|\alpha_3(n, h_n) - \alpha'(\theta_0)| \Rightarrow \infty$ as $n \rightarrow \infty$,
- if $n^{1/3}h_n \rightarrow h > 0$, $n^{1/3}(\Delta\alpha_3(n, h_n) - \alpha'(\theta_0)) \Rightarrow \frac{\sigma_1}{\sqrt{h}} N(0, 1) - \frac{h}{2} \alpha''(\theta_0)$ as $n \rightarrow \infty$.

We note that the convergence rate here is of order $n^{-1/3}$, the same as that obtained earlier for the central-difference estimator. Observe that the optimal difference increment now is of order $n^{-1/3}$, which is much smaller than the difference increment of order $n^{-1/6}$ derived for central differences. The lower variability of the common random numbers estimator is what permits us to choose the smaller increment. We further note that the value of h appearing in c) which minimizes the asymptotic mean square error of the estimator $\Delta\alpha_3(n, h_n)$ is

$$h^* = \left(\frac{2\sigma^2}{\alpha''(\theta_0)^2} \right)^{1/3}.$$

Of course, we can also combine central differences and common random numbers. Let $\Delta_c Y(h) = Y(\theta_0 + h) - Y(\theta_0 - h)$. The following assumption is the analogue of (2.4);

(2.6)

- i) $P\{Y(h)\varepsilon\cdot\} = P_{\theta_0+h}\{X(\theta_0 + h)\varepsilon\cdot\}$,
- ii) $E\Delta_c Y(h)^2 = h\sigma_2^2 + o(h)$ as $h \downarrow 0$, where $\sigma_2^2 > 0$,
- iii) there exists $\varepsilon > 0$ such that $E|\Delta Y_c(h)|^{2+\varepsilon} = h\beta_2 + o(h)$ as $h \downarrow 0$,
- iv) $\alpha(\theta) = E_\theta X(\theta)$ is three times continuously differentiable in an open neighborhood of θ_0 .

Let $\Delta_c Y_1(h), \Delta_c Y_2(h), \dots$ be an i.i.d. sequence of replicates of the r.v. $\Delta_c Y(h)$. The central-difference common random numbers estimator of $\alpha'(\theta_0)$ is given by

$$\Delta\alpha_4(n, h) = \frac{1}{n} \sum_{i=1}^n \frac{\Delta_c Y_i(h)}{2h}.$$

The proof of the following convergence rate theorem for $\Delta\alpha_4(n, h)$ follows the same lines as that for $\Delta\alpha_3(n, h)$; the proof is therefore omitted.

(2.7) **THEOREM.** Assume (2.6). If $\alpha^{(3)}(\theta_0) \neq 0$ and $nh_n \rightarrow \infty$, then:

- a) if $n^{1/5}h_n \rightarrow \infty$ with $h_n \rightarrow 0$, $n^{2/5}|\Delta\alpha_4(n, h_n) - \alpha'(\theta_0)| \Rightarrow \infty$ as $n \rightarrow \infty$,
- b) if $n^{1/5}h_n \rightarrow 0$, $n^{2/5}|\Delta\alpha_4(n, h_n) - \alpha'(\theta_0)| \Rightarrow \infty$ as $n \rightarrow \infty$,
- c) if $n^{1/5}h_n \rightarrow h > 0$, $n^{2/5}(\Delta\alpha_4(n, h_n) - \alpha'(\theta_0)) \Rightarrow \frac{\sigma_2}{\sqrt{h}}N(0, 1) - \frac{h^2}{3}\alpha^{(3)}(\theta_0)$ as $n \rightarrow \infty$.

Thus, combining common random numbers and central differences improves the convergence rate of the derivative estimator to order $n^{-2/5}$. Furthermore, the difference increment that minimizes the asymptotic mean square error of the estimator $\alpha_4(n, h_n)$ is $h_n = h^*n^{-1/5}$, where

$$h^* = \left(\frac{9\sigma_2^2}{4\alpha^{(3)}(\theta_0)^2} \right)^{1/5}.$$

2.4 Finite-Difference Estimators With A Near-Optimal Convergence Rate

In the preceding three sections, we have discussed four different finite-difference estimators. The convergence rate was improved from order $n^{-1/4}$, in the case of forward differences with independent streams of random numbers, to order $n^{-2/5}$, which was achieved by a central difference estimator that used a common stream of random numbers. A natural question that arises here is whether any further improvement is possible. In particular, can one obtain finite-difference estimators for the derivative that achieve a convergence rate that is arbitrarily close to the best possible rate for a Monte Carlo procedure, namely $n^{-1/2}$? We will now answer this question in the affirmative by developing such a class of estimators.

To produce the type of estimator that we have in mind, we need to obtain a finite-difference approximation to $\alpha'(\theta_0)$ that is as unbiased as possible. Suppose, for the moment, that α is an analytic function in θ . Then

$$(2.8) \quad \alpha(\theta + h) = \sum_{n=0}^{\infty} \alpha^{(n)}(\theta) \frac{h^n}{n!}.$$

Let $T_h\alpha$ be the "shifted" function defined by $(T_h\alpha)(\theta) = \alpha(\theta + h)$. We further let $D\alpha$ be the derivative function specified by $(D\alpha)(\theta) = \alpha^{(1)}(\theta)$. The expansion (2.8) may then be written as

$$(2.9) \quad T_h\alpha = \sum_{n=0}^{\infty} \frac{h^n}{n!} D^n \alpha.$$

Proceeding formally, we may rewrite (2.9) in terms of the operators T_h and D as

$$(2.10) \quad T_h = \sum_{n=0}^{\infty} \frac{h^n}{n!} D^n = \exp(hD).$$

We now wish to express the operator D in terms of the shift operator T_h :

$$D = \frac{1}{h} \log(T_h).$$

Expanding the logarithm in a formal power series, we obtain

$$(2.11) \quad D = \frac{1}{h} \sum_{k=1}^{\infty} (T_h - I)^k \frac{(-1)^{k+1}}{k}.$$

To obtain a finite-difference approximation to $\alpha'(\theta_0) = (D\alpha)(\theta_0)$, we truncate the series (2.11) at the m 'th term:

$$\begin{aligned} D &\approx \frac{1}{h} \sum_{k=1}^m (T_h - I)^k k \frac{(-1)^{k+1}}{k} \\ &= \frac{1}{h} \sum_{k=1}^m \sum_{\ell=0}^k \binom{k}{\ell} (T_h)^\ell (-1)^{k-\ell} \frac{(-1)^{k+1}}{k}. \end{aligned}$$

Noting that $(T_h)^\ell = T_{h\ell}$, we obtain the following approximation to $\alpha'(\theta_0)$:

$$(2.12) \quad \alpha'(\theta_0) \approx \frac{1}{h} \sum_{k=1}^m \sum_{\ell=0}^k \binom{k}{\ell} \frac{(-1)^{1-\ell}}{k} \alpha(\theta_0 + h\ell).$$

To obtain a finite-difference estimator for $\alpha'(\theta_0)$, we let $X_1(\theta_0 + h\ell), X_2(\theta_0 + h\ell), \dots$ be i.i.d. replicates of the r.v. $X(\theta_0 + h\ell)$, simulated under common distribution $P_{\theta_0+h\ell}$ ($0 \leq \ell \leq m$). We further generate each of the $m+1$ sequences independently of one another (i.e., the sequences $(X_i(\theta_0 + h\ell) : i \geq 1)$ are mutually independent for $0 \leq \ell \leq m$). Set

$$Z_i(h) = \frac{1}{h} \sum_{k=1}^m \sum_{\ell=0}^k \binom{k}{\ell} \frac{(-1)^{1-\ell}}{k} X_i(\theta_0 + h\ell).$$

The expectation of $Z_i(h)$ then matches the right-hand side of (2.12). We then obtain a finite-difference derivative estimator by setting

$$\alpha_5(h, n) = \frac{1}{n} \sum_{i=1}^n Z_i(h).$$

Our next theorem describes the convergence rate of $\alpha_5(h_n, n)$, when the difference increment is chosen appropriately.

(2.13) **THEOREM.** Assume (2.1). If $m \geq 1$ and $h_n n^{1/2m} \rightarrow h > 0$ as $n \rightarrow \infty$, then

$$n^{\frac{1}{2} - \frac{1}{2m}} (\alpha_5(h_n, n) - \alpha'(\theta_0)) \Rightarrow \sigma_3 N(0, 1)$$

as $n \rightarrow \infty$, where $\sigma_3^2 = \sigma^2(\theta_0) \chi_m / h^2$ and

$$\chi_m = \left(\sum_{k=1}^m \frac{1}{k} \right)^2 + \sum_{\ell=1}^m \left(\sum_{k=\ell}^m \binom{k}{\ell} \frac{1}{k} \right)^2.$$

According to Theorem 2.13, the convergence rate of $\alpha_5(h_n, n)$ may be made as close as we wish to $n^{-1/2}$, by choosing m sufficiently large. To some extent, this convergence rate is deceptive. Note, in particular, that the constant χ_m is increasing in m . Furthermore, the construction of each observation $Z_i(h)$ that enters $\alpha_5(h, n)$ requires m independent simulations. As a consequence, the computational effort required to generate $\alpha_5(h_n, n)$ is sensitive to the choice of m . Thus, although the convergence rate promised by Theorem 2.13 is significantly better than those described earlier in this section, the run-lengths required to see such an improvement may be quite large.

3. LIKELIHOOD RATIO DERIVATIVE ESTIMATORS

In certain settings, it is possible to construct derivative estimators that achieve the best possible rate of convergence for a Monte Carlo estimator, namely $n^{-1/2}$ in the number of observations n that are generated.

Suppose, for the moment, that the distribution defining $\alpha(\theta)$ is independent of θ . Then, all the θ -dependence of α sits in the r.v. $X(\theta)$, so that $\alpha(\theta) = EX(\theta)$. Assuming that we can interchange the derivative operator and the expectation, we get

$$\alpha'(\theta_0) = EW(\theta_0),$$

where $W(\theta_0) = X'(\theta_0)$. Then, by generating i.i.d. replicates of $W(\theta_0)$, we obtain an estimator which (use the standard central limit theorem) possesses the canonical convergence rate $n^{-1/2}$. The idea behind the likelihood ratio method (and the perturbation analysis approach of the next section) is to structure the representation of α so that the driving distribution is rendered independent of θ .

Suppose that the distribution P_θ defining α has density $L(\theta)$ with respect to some common distribution P , so that

$$(3.1) \quad P_\theta(d\omega) = L(\theta, \omega)P(d\omega).$$

The r.v. $L(\theta)$ is called the likelihood ratio of P_θ (with respect to P). Under this assumption,

$$\begin{aligned} \alpha(\theta) &= \int_{\Omega} X(\theta, \omega)P_\theta(d\omega) \\ &= \int_{\Omega} X(\theta, \omega)L(\theta, \omega)P(d\omega) \\ &= \int_{\Omega} Y(\theta, \omega)P(d\omega) \\ &= EY(\theta), \end{aligned}$$

where $Y(\theta) = X(\theta)L(\theta)$. This is the desired representation of α . Assuming that the derivative-expectation interchange is valid (and it typically is), we obtain $\alpha'(\theta_0) = EW(\theta_0)$, where $W(\theta_0) = X'(\theta_0)L(\theta_0) + X(\theta_0)L'(\theta_0)$. Hence, the key to obtaining likelihood ratio derivative estimators is finding a distribution P and a r.v. $L(\theta)$ such that (3.1) holds (at least for θ in an open neighborhood of θ_0).

This idea is easily illustrated when the basic sample space Ω is the real line. Suppose that the distribution P_θ takes the form $P_\theta(dx) = f(\theta, x)\mu(dx)$. For example, if $\mu(dx) = dx$, we are saying that P_θ has a (Lebesgue) density for each θ . Choose $g(x) > 0$ so that

$$\int_{\mathbf{R}} g(x)\mu(dx) = 1.$$

(This can always be done if μ is σ -finite.) Set $P(dx) = g(x)\mu(dx)$ and observe that (3.1) holds with $L(\theta, x) = f(\theta, x)/g(x)$.

Suppose that we are interested in estimating both $\alpha(\theta_0)$ and $\alpha'(\theta_0)$. Assume that the set $\Lambda(\theta) = \{x : f(\theta, x) > 0\}$ is independent of θ in a neighborhood of θ_0 . A particularly convenient choice of $g(x)$, in this case, is $g(x) = f(\theta_0, x)$. Because $\Lambda(\theta)$ is independent of θ , the likelihood ratio $L(\theta, x) = f(\theta, x)/f(\theta_0, x)$ is well-defined in a neighborhood of θ . This choice of g has several advantages. Note that to generate $W(\theta_0)$, the simulation involves generating outcomes under the distribution P_{θ_0} . The distribution P_{θ_0} is typically the "natural" distribution for estimating $\alpha(\theta_0)$. Hence, $\alpha(\theta_0)$ and $\alpha'(\theta_0)$ can easily be estimated from the same set of sampling experiments. A second advantage is that by choosing g in this way, $L(\theta_0) = 1$ and the formula for $W(\theta_0)$ simplifies. In fact, in most applications, the calculation of $L'(\theta_0)$ also simplifies considerably, when g is chosen so that $g(x) = f(\theta_0, x)$. Furthermore, in

many sampling settings, this choice of g leads to an estimator $W(\theta_0)$ that has desirable variance properties.

As the above discussion suggests, an important issue in the development of likelihood ratio gradient estimators is the construction of a likelihood ratio (for a given class of discrete-event simulations) that has desirable computational and variability characteristics. For example, it turns out that in a discrete-event simulation context, the likelihood ratio typically exists only for terminating simulation problems. Of course, steady-state characteristics can be analyzed as a limit of finite-horizon estimation problems. Unfortunately, the associated likelihood ratios become successively more unstable as the time horizon gets large. However, this problem can be avoided if the discrete-event system has the right kind of structure (typically, regenerative structure). REIMAN and WEISS (1986) discuss some of the relevant ideas.

4. PERTURBATION ANALYSIS DERIVATIVE ESTIMATORS

In Section 3, we described the likelihood ratio approach to derivative estimation. The basic idea was to use the method of likelihood ratios so as to obtain a representation of α in which the driving distribution is independent of θ . In this section, we describe an alternative technique for obtaining such a representation. The idea is to return to the common random numbers technique described in Section 2.3. Suppose that we can find a probability space $(\Omega, \mathcal{F}, \mathcal{P})$ and a family of r.v.'s $\{Y(h) : |h| < \varepsilon\}$ such that:

$$(4.1) \quad P\{Y(h)\varepsilon\} = P_{\theta_0+h}\{X(\theta_0+h)\varepsilon\}.$$

Under this assumption, it follows that

$$\alpha(\theta_0+h) = EY(h)$$

for $|h| < \varepsilon$. Assuming that we can interchange the derivative and expectation operators, we find that $\alpha'(\theta_0) = EW(\theta_0)$, where $W(\theta_0) = Y'(0)$. Hence, by generating i.i.d. replicates of the r.v. $W(\theta_0)$, we obtain an estimator that achieves the $n^{-1/2}$ convergence rate that is best possible for a Monte Carlo procedure.

As in the case of the likelihood ratio method, this technique is best illustrated when the basic sample space is the unit interval. Let $\Omega = [0, 1]$, \mathcal{F} = Borel sets of $[0, 1]$, and let P be uniform distribution on Ω . Set $U(\omega) = \omega$ and observe that

$$Y(h) = F_h^{-1}(U)$$

satisfies (4.1) (where $F_h(x) = P_{\theta_0+h}\{X(\theta_0+h) \leq x\}$). To calculate $W(\theta_0)$, we need to determine $\frac{d}{dh}F_h^{-1}(x)|_{h=0}$. Assume, for the moment, that F_h has a (Lebesgue) density f_h for each h . By definition of the inverse distribution function F_h , we have

$$F_h(F_h^{-1}(x)) = x.$$

Differentiating both sides of the above expression with respect to h , we get

$$\frac{\partial}{\partial h}F_h(F_h^{-1}(x)) + f_h(F_h^{-1}(x))\frac{d}{dh}F_h^{-1}(x) = 0,$$

from which we obtain

$$\frac{d}{dh}F_h^{-1}(x) = -\frac{\partial}{\partial h}F_h(F_h^{-1}(x))/f_h(F_h^{-1}(F_h^{-1}(x))).$$

Thus, in this setting, we find that

$$W(\theta_0) = -\frac{\partial}{\partial h}F_h(F_h^{-1}(U))/f_h(F_h^{-1}(U))|_{h=0}.$$

One undesirable feature of the approach that we have just outlined is that since we have taken our basic probability space as uniform distribution on $[0, 1]$, the generation of $W(\theta_0)$ appears to require inversion (i.e., calculation of $F_h^{-1}(\cdot)$). Recall, however, that $F_h^{-1}(U)$ has the same distribution as $Y(h)$ (or, alternatively, $X(\theta_0+h)$ under P_{θ_0+h}). Hence, $W(\theta_0)$ (under P) shares the same distribution as

$$(4.2) \quad -\frac{\partial}{\partial h}F_h(X(\theta_0))/f_h(X(\theta_0))|_{h=0}$$

(under P_{θ_0}). The advantage of the representation (4.2) is that we can generate the derivative observations using precisely the same algorithm as that used to estimate $\alpha(\theta_0)$ itself (since $\alpha(\theta_0)$ is typically estimated by generating i.i.d. replicates of $X(\theta_0)$ under P_{θ_0}). See GLYNN (1987) for additional details.

The argument that led to (4.2) appears to require existence of a density f_h . It turns out that in many settings, the

common random numbers technique can still be applied, in spite of the non-existence of a density. For example, suppose that $X(\theta) = g(\theta, Y)$ and Y has distribution $F(y/\theta)$ under P_θ . If we set $Y(h) = (\theta_0+h)Y/\theta_0$ ($\theta_0 > 0$), then (4.1) is satisfied with $P = P_{\theta_0}$, and we find that

$$\alpha(\theta_0+h) = E_{\theta_0}g(\theta_0+h, (\theta_0+h)Y/\theta_0).$$

If g is smooth, it is clear that the derivative of $g(\theta_0+h, (\theta_0+h)Y/\theta_0)$ exists, regardless of the nature of the distribution F .

The representation (4.2) (for $W(\theta_0)$) can be derived via an alternative argument. Recall that $\alpha(\theta) = E_\theta X(\theta)$. As a consequence, if $X(\theta)$ is non-negative, we find that

$$(4.3) \quad \begin{aligned} \alpha(\theta_0+h) &= E_{\theta_0+h} \int_0^\infty I(X(\theta_0+h) > x) dx \\ &= \int_0^\infty (1 - F_h(x)) dx. \end{aligned}$$

Hence, assuming that the derivative and integral operators can be interchanged, we obtain

$$(4.4) \quad \alpha'(\theta_0) = -\int_0^\infty \frac{\partial}{\partial h}F_h(x)dx|_{h=0}.$$

In order to apply the Monte Carlo method to the numerical evaluation of the integral appearing in (4.4), we need to represent it as an expectation. One way to do this is as follows (assuming F_h has a density f_h):

$$(4.5) \quad \begin{aligned} \alpha'(\theta_0) &= -\int_0^\infty \frac{\frac{\partial}{\partial h}f_h(x)}{f_h(x)} \cdot f_h(x)dx|_{h=0} \\ &= -E_{\theta_0} \frac{\partial}{\partial h}F_h(X(\theta_0))/f_h(X(\theta_0))|_{h=0}, \end{aligned}$$

which is just (4.2). It is interesting to note that an alternative representation of the expectation $\alpha(\theta_0+h)$ exists:

$$(4.6) \quad \alpha(\theta_0+h) = \int_0^\infty xF_h(dx).$$

If F_h has a density, this becomes

$$\alpha(\theta_0 + h) = \int_0^\infty x f_h(x) dx.$$

Assuming that the derivative-integral interchange is valid, we get

$$(4.7) \quad \alpha'(\theta_0) = \int_0^\infty x \frac{\partial}{\partial h} f_h(x) dx \Big|_{h=0}.$$

To represent the integral in (4.7) as an expectation, we use the same idea as in (4.5):

$$\alpha'(\theta_0) = E_{\theta_0} X(\theta_0) \frac{\frac{\partial}{\partial h} f_h(X(\theta_0))}{f_h(X(\theta_0))} \Big|_{h=0}.$$

In particular, if $X(\theta) \equiv X$ (in which case $\alpha(\theta) = E_\theta X$), we obtain

$$(4.8) \quad \alpha'(\theta_0) = E_{\theta_0} X \frac{\frac{\partial}{\partial h} f_h(X)}{f_h(X)} \Big|_{h=0}.$$

It turns out that (4.8) is precisely the likelihood ratio derivative estimator of Section 3. Hence, in this simple setting, the common random numbers approach and the likelihood ratio method derive from the two analytical representations (4.3) and (4.6) for the mean of a non-negative r.v. Since (4.3) and (4.6) are usually obtained from one another by an integration-by-parts, it follows that the likelihood ratio and common random numbers methods are related through an integration-by-parts in this simple context. We note, parenthetically, that if $X(\theta)$ is non-negative, then

$$E_\theta X(\theta) = p \int_0^\infty t^{p-1} P_\theta \{X(\theta) > t^p\} dt$$

for $p > 0$, from which it follows that

$$(4.9) \quad \alpha'(\theta_0) = -p E_{\theta_0} X(\theta_0)^{p-1} \frac{\frac{\partial}{\partial h} F_h(X(\theta_0)p)}{f_h(X(\theta_0))} \Big|_{h=0}.$$

Formula (4.9) generalizes (4.5). In principle, one could optimize over p in order to determine that p -value which yields a derivative estimator with the minimum variance.

It turns out that the common random numbers derivative estimation method described above can be applied to calculate derivatives of performance measures for discrete-event dynamical systems. The subject of **perturbation analysis** is concerned with the study and development of the resulting estimators. For example, consider a discrete-event system in which the measure P_θ characterizes the distribution (over sample trajectories) when the event-scheduling distributions are indexed by a scale parameter θ . Now, because θ appears as a scale parameter in the event-scheduling distribution, we can view the event-scheduling r.v.'s that are generated as taking the form $\theta X_1, \theta X_2, \dots$ for r.v.'s X_1, X_2, \dots having distribution independent of θ . For discrete-event systems in which the probability of two events occurring simultaneously is zero, a small perturbation of the event times will have no effect on the order of the state transitions experienced by the discrete-event system. The effect of the parameter θ will reflect itself only in the timing of the sequence of state transitions. Furthermore, as HO and CAO (1983) point out, the manner in which the perturbation propagates itself through the sequence of event timings is suitable to a highly efficient recursive computation (i.e., the perturbation of the n 'th state transition epoch is easily calculated from that of the $(n-1)$ 'st). These ideas lead to an easily calculated sample path derivative for discrete-event systems in which the event-scheduling distributions are parameterized; see SURI (1987) for additional details on the nature of the infinitesimal perturbation analysis (IPA) derivative computation.

As described above, the IPA approach to derivative estimation focuses on derivative estimation relative to perturbations in the event-scheduling distributions. In many queueing settings, one wishes to optimize over routing probabilities, however. Likelihood ratio methods are highly flexible and can be applied in a straightforward manner to such problems. Recent extensions of IPA to such routing probability derivative estimation problems hold significant promise, however (see HO and CAO (1985)).

Empirical evidence, gathered to date, appears to suggest that when both IPA and likelihood ratio methods apply to a given problem, the IPA estimator will typically be more efficient (in the sense of having lower variability). This conclusion stems, in part, from the fact that likelihood ratio estimators are known to have a variability that increases in a roughly linear fashion with the time horizon of the simulation; see REIMAN and WEISS (1986).

Some care must be taken in applying IPA techniques to a given problem, however. The difficulty is that the interchange of derivative and expectation operators that is needed to rigorously justify the IPA estimator (see (4.1)) may sometimes be invalid. In such settings, the IPA estimator can

converge to the wrong quantity. To get some sense of the problem, we note that if $Y(h)$ has a well-behaved derivative $Y'(0)$ at $h = 0$, then we would expect that

$$(4.10) \quad h^{-2} \text{var}[Y(h) - Y(0)] \rightarrow \text{var}Y'(0)$$

as $h \downarrow 0$. (In fact, (4.10) is a sufficient condition for permitting the interchange of derivative and expectation.) Hence,

$$(4.11) \quad E\Delta Y(h)^2 = h^2 EY'(0)^2 + o(h^2)$$

as $h \downarrow 0$. Recall, however, that in Section 2.3, we argued that the typical behavior of a discrete-event system was governed by (2.4) ii), which contradicts (4.11). The difficulty is that while the effect of the perturbations on the state transition sequence may be ignored in calculating $W'(\theta_0)$, it cannot be typically ignored in calculating $\alpha'(\theta_0)$. In HEIDELBERGER et al. (1988), this point is analyzed further. It is shown that conventional IPA can be inconsistent (in the sense of convergence to an incorrect answer) for multiple customer-type queueing networks. However, conventional IPA turns out to be consistent for a large number of performance measures associated with single customer type networks.

Furthermore, a number of extensions in the basic IPA algorithm hold significant promise for overcoming the difficulties that arise in the multiple customer context. In particular, a new version of IPA, known as smoothed perturbation analysis (SPA), is now under development. The idea is that, rather than work with the “raw” sample path $Y(h)$ itself, one considers instead the conditional expectation of $Y(h)$ with respect to some appropriately chosen conditioning variable Z (appropriate in the sense that $E(Y(h)|Z)$ is easily calculated). Since a conditional expectation involves an integration operation, the conditioning ought to yield a process $E(Y(h)|Z)$ which is smoother in h than is $Y(h)$ itself. As a consequence, SPA has the potential to deal with estimation problems for which classical IPA does not work; see GLASSERMAN and GONG (1989) for further details.

5. GRADIENT ESTIMATION

In the previous three sections, we have described derivative estimation techniques that are applicable to problems in which the decision parameter θ is scalar-valued. The methods of Section 2 gave rise to estimators for which their respective convergence rates were slower than $n^{-1/2}$ in the number n of observations n that were generated. On the

other hand, the likelihood ratio and perturbation analysis techniques that were described in Sections 3 and 4 attained the canonical convergence rate of $n^{-1/2}$.

The generalization of these ideas to the setting in which θ is vector-valued is straightforward. The partial derivatives with respect to each of the co-ordinates θ_i is easily calculated in the same way as the scalar derivatives were estimated earlier. However, the computational complexity of calculating a d -dimensional gradient is highly sensitive to d and is an issue which is specific to the setting in which θ is vector-valued (as opposed to scalar valued). For example, note that a forward difference estimator for the d -dimensional gradient $\nabla\alpha(\theta_0)$ involves performing simulations at the $d+1$ parameter points $\theta_0, \theta_0 + h_1 e_1, \dots, \theta_0 + h_d e_d$, where e_i is the i 'th unit vector. On the other hand, a central difference approximation requires simulating at the $2d$ points $\theta_0 \pm h_d e_d$. Thus, a central difference estimator for a d -dimensional gradient requires roughly twice as much computational effort as a forward difference estimator to obtain the same number of observations. This, however, is balanced by the fact that the convergence rate of a central difference estimator is more rapid than that of a forward difference estimator. As a consequence, we see that if d is large, a forward difference estimator may be more efficient for small n . If n is large enough, however, the central difference estimator always wins. This dimensionality effect becomes even more pronounced for the “near optimal” difference estimators of Section 2.4. Note that to estimate a d -dimensional gradient, simulations at the $md + 1$ points $\theta_0, \theta_0 + h_1 \ell e_1, \dots, \theta_0 + h_d \ell e_d$ ($1 \leq \ell \leq m$) are needed. Hence, the dimensionality degradation that occurs with this estimator is even more serious than that experienced by the central difference estimators discussed earlier. An additional disadvantage of this class of estimators is that they can be quite sensitive to numerical round-off error when m is large. (The presence of the alternating sign $(-1)^{1-\ell}$ can lead to numerical instability.)

Turning now to the likelihood ratio and perturbation analysis estimators, we note that both of these estimators, when applied to estimation of the gradient, require only a single simulation at the parameter point θ_0 . Of course, the additional computer time required to calculate the d partial derivatives from the single simulation imply that the computational effort to compute a d -dimensional gradient is still increasing in d . However, one would expect that these estimation algorithms would be less sensitive to d than are the finite-difference estimators of Section 2. Thus, the likelihood ratio and perturbation analysis estimators improve upon finite difference estimators in two ways: computation time is less sensitive to the dimension d , and the convergence rate is $n^{-1/2}$.

6. ORTHOGONAL FUNCTION APPROXIMATIONS

One of the reasons that gradient estimation plays a key role in optimization is that the gradient gives information about the shape of the objective function. When such shape information is added to that supplied by a function evaluation, we are essentially being given an affine approximation to the function in a neighborhood of the point at which the evaluations occurred. More generally, if all the partial derivatives of an analytic function are given at a single (fixed) point, the entire global behavior of the function is then determined. The ability to obtain global information about the behavior of the objective function is clearly useful in an optimization context.

As indicated above, one way to cheaply infer global behavior is via a Taylor series expansion that is determined by the partial derivatives of the function. Another approach involves attempting to expand the function in an orthogonal expansion of some kind. We shall now illustrate this idea in the case that the decision parameter θ is scalar valued and the orthogonal functions are the trigonometric functions. In this case, we will then obtain a Fourier-like expansion of the objective function.

Suppose that we are interested in studying the behavior of the objective function over the interval $[0, \pi]$. (By transforming the interval if necessary, this is equivalent to studying α over an arbitrary compact interval of the form $[a, b]$.) We can then make α 2π -periodic by extending α to $[-\pi, 0)$ via the even extension $\alpha(\theta) = \alpha(-\theta)$ and then letting $\alpha(\theta + 2\pi) = \alpha(\theta)$. Assuming α is continuously differentiable on $[0, \pi]$, it is well known (see FULKS (1969), p. 547) that for each $\theta \in [-\pi, \pi]$,

$$(6.1) \quad \alpha(\theta) = \sum_{k=0}^{\infty} \langle \alpha, e_k \rangle e_k(\theta),$$

where the e_k 's are the normalized cosine functions defined by

$$e_k(\theta) \triangleq \begin{cases} \frac{1}{\sqrt{2\pi}}, & k = 0 \\ \frac{1}{\sqrt{\pi}} \cos k\theta, & k \geq 1 \end{cases}$$

and

$$(6.2) \quad \langle x, y \rangle \triangleq \int_{-\pi}^{\pi} x(\theta)y(\theta)d\theta.$$

(The Fourier series for α is a cosine series because of the fact that α is an even function.) The functions e_0, e_1, \dots are orthogonal with respect to the inner product (6.2) that we have defined, in the sense that $\langle e_k, e_\ell \rangle = 0$ for $k \neq \ell$. (In fact, they are orthonormal since $\langle e_k, e_k \rangle = 1$ for $k \geq 0$.) Hence, (6.1) expresses α as a linear combination of the orthonormal "vectors" e_0, e_1, e_2, \dots . Thus, we can estimate α by estimating each of the inner products $a_k \triangleq \langle \alpha, e_k \rangle$ for $k \geq 0$. In contrast to expanding a function in a Taylor series, each of the coefficients a_k is defined by an integral, namely

$$a_k = \begin{cases} \sqrt{\frac{2}{\pi}} \int_0^{\pi} \alpha(\theta) d\theta, & k = 0 \\ \frac{2}{\sqrt{\pi}} \int_0^{\pi} \alpha(\theta) \cos k\theta d\theta, & k \geq 1. \end{cases}$$

Monte Carlo methods are well suited to estimating integrals. In particular, suppose we generate U as a uniformly distributed r.v. on the interval $[0, \pi]$ and then simulate $X(U)$ under the distribution P_u . Then, a_k can be represented as

$$a_k = \begin{cases} \sqrt{2\pi} EX(U), & k = 0 \\ 2\sqrt{\pi} EX(u) \cos kU, & k \geq 1. \end{cases}$$

Hence, the r.v.

$$\Lambda(m, \theta) = X(U) \left(1 + 2 \sum_{k=1}^m \cos(kU) \cdot \cos k\theta \right)$$

has expectation

$$E \Lambda(m, \theta) = \sum_{k=0}^m \langle \alpha, e_k \rangle e_k(\theta).$$

Thus, $\Lambda(m, \theta)$ is an unbiased estimator for the first $m+1$ terms in the Fourier series of α . Note that only one simulation is required to estimate the first $m+1$ Fourier coefficients of α .

Suppose that we generate n i.i.d. copies $\Lambda_1(m, \theta), \dots, \Lambda_n(m, \theta)$ of the r.v. $\Lambda(m, \theta)$. We can then form the estimator

$$\alpha_n(\theta) = \frac{1}{n} \sum_{i=1}^n \Lambda_i(m_n, \theta);$$

we permit $m = m_n$ to be a function of the sample size n (since m will have to grow with n in order to asymptotically remove the bias of the estimator).

In order to measure the distance of the estimator $\alpha_n(\cdot)$ from the function $\alpha(\cdot)$, we use the norm

$$\|x\| \triangleq ((x, x))^{1/2} = \left(\int_{-\pi}^{\pi} x^2(\theta) d\theta \right)^{1/2}.$$

Our goal is to describe the magnitude of the distance $\|\alpha_n - \alpha\|$. Let

$$\hat{a}_k(n) = \begin{cases} \sqrt{2\pi} n^{-1} \sum_{i=1}^n X_i(U_i), & k = 0 \\ 2\sqrt{\pi} n^{-1} \sum_{i=1}^n X_i(U_i) \cos(kU_i), & k \geq 1 \end{cases}$$

and note that

$$\alpha_n(\theta) = \sum_{k=0}^{m_n} \hat{a}_k(n) e_k(\theta).$$

If α is continuously differentiable on $[0, \pi]$, then (6.1) is valid, so that

$$\alpha_n(\theta) - \alpha(\theta) = \sum_{k=0}^{m_n} (\hat{a}_k(n) - a_k) e_k(\theta) - \sum_{k > m_n} a_k e_k(\theta)$$

where $a_k = \langle \alpha, e_k \rangle$. Then, the orthogonality of the e_k 's guarantees that

$$\|\alpha_n - \alpha\|^2 = \sum_{k=0}^{m_n} (\hat{a}_k(n) - a_k)^2 + \sum_{k > m_n} a_k^2.$$

Hence,

$$\begin{aligned} (6.3) \quad E\|\alpha_n - \alpha\|^2 &= \sum_{k=0}^{m_n} \text{var} \hat{a}_k(n) + \sum_{k > m_n} a_k^2 \\ &= 2\pi n^{-1} \left(EX^2(U) + 2 \sum_{k=1}^{m_n} EX^2(U) \cos^2 kU \right) \\ &\quad - n^{-1} \sum_{k=0}^{m_n} a_k^2 + \sum_{k > m_n} a_k^2. \end{aligned}$$

It is worth observing that if each of the Fourier coefficients were to be estimated independently (so that $\wedge(m, \theta)$ takes

the form $X_0(U_0) + 2 \sum_{k=1}^m X_k(U_k) \cos(kU_k) \cos(k\theta)$), this would have no effect on (6.3). In other words, $E\|\alpha_n - \alpha\|^2$ is unaffected by whether the coefficients are estimated independently or not.

We are now ready to state a limit theorem for $E\|\alpha_n - \alpha\|^2$.

(6.4) **THEOREM.** Let $b(\theta) = E_{\theta} X^2(\theta)$. Suppose that $b(\theta)$ is continuous on $[0, \pi]$ and $\alpha(\theta)$ has a continuous p 'th derivative on $[0, \pi]$ ($p \geq 1$). If $m_n = n^{1/2p}$, then:

- $EX^2(U) + 2 \sum_{k=1}^{m_n} EX^2(U) \cos^2(kU) \sim m \int_0^{\pi} b(\theta) d\theta / \pi$ as $m \rightarrow \infty$,
- $n^{1-\gamma} E\|\alpha_n - \alpha\|^2 \rightarrow 0$ as $n \rightarrow \infty$ for any $\gamma > 1/2p$,
- $n^{1/2-\gamma/2} \|\alpha_n - \alpha\| \Rightarrow 0$ as $n \rightarrow \infty$ for any $\gamma > 1/2p$,
- for $\varepsilon > 0$, $m\{\theta \in [0, \pi] : |\alpha_n(\theta) - \alpha(\theta)| \geq \varepsilon n^{\gamma/2-1/2}\} \Rightarrow 0$ as $n \rightarrow \infty$, for $\gamma > 1/2p$. (m is Lebesgue measure.)

Thus, if α is sufficiently smooth, we can obtain a global convergence rate arbitrarily close to $n^{-1/2}$ in the number of observations n that are simulated. (We note that because α is 2π -periodic and is an even function, α 's first p derivatives must vanish at 0 and π in order to satisfy the smoothness hypothesis of Theorem 6.4. If α does not satisfy the condition, one can shrink α 's domain of definition to $[\varepsilon, \pi - \varepsilon]$ and then smoothly extend α to $[0, \pi]$ in order to satisfy the smoothness hypothesis.)

7. STOCHASTIC APPROXIMATION ALGORITHMS

In this section, we briefly describe how the results of the previous sections can be integrated into an optimization algorithm.

Consider the unconstrained problem in which the goal is to minimize the objective function $\alpha(\theta)$ over $\theta \in \mathbb{R}^d$. The idea is to develop a recursive procedure in which the $(n+1)$ 'st iterate is likely to be closer to the minimizer θ^* than is the n 'th iterate. Specifically, suppose that it is possible to generate r.v.'s $W(\theta)$ such that $EW(\theta) \approx \nabla \alpha(\theta)$; as discussed in Sections 2-5, this can be done either through a finite-difference approximation or through the likelihood ratio and perturbation analysis gradient estimators (in which case $EW(\theta)$ is typically equal to $\nabla \alpha(\theta)$). Assuming existence of such r.v.'s $W(\theta)$, consider the recursion

$$(7.1) \quad \theta_{n+1} = \theta_n - n^{-1} \Gamma V_{n+1},$$

where Γ is a given $d \times d$ matrix and $P\{V_{n+1}\varepsilon A|\theta_0, V_0, \dots, \theta_n, V_n\} = P\{W(\theta_n)\varepsilon A\}$. In other words, the r.v. V_{n+1} is generated by simulating a copy of $W(\theta_n)$.

In the case that $E\{V_{n+1}|\theta_0, V_0, \dots, \theta_n, V_n\} = \nabla\alpha(\theta_n)$ (as would occur if the methods of Sections 3 and 4 were used), algorithm (7.1) is known as the **Robbins-Monro algorithm**. Assuming that α is twice continuously differentiable, the optimal choice of the matrix Γ then turns out to be

$$\Gamma = H(\theta^*)^{-1},$$

where H is the Hessian of second derivatives; see POLYAK and TSYPKIN (1980) for details.

It is particularly illuminating to consider (7.1) in the case that θ is scalar-valued. In this case, (7.1) takes the form

$$(7.2) \quad \theta_{n+1} = \theta_n - n^{-1}cV_{n+1}.$$

Note that if $c > 0$, θ_{n+1} has a tendency to be smaller than θ_n when $\alpha'(\theta_n) > 0$, and has a tendency to increase when $\alpha'(\theta_n) < 0$. As a consequence, the sequence $(\theta_n : n \geq 1)$ has a tendency to move towards a point θ^* for which $\alpha'(\theta^*) = 0$ and $\alpha'(\theta) > 0$ ($\alpha'(\theta) < 0$) for θ in a neighborhood to the right (left) of θ^* . Any such θ^* must necessarily be a local minimizer of α . Thus, the algorithm (7.1) appears intuitively reasonable.

In fact, (7.1) has good convergence characteristics. If $E\{V_{n+1}|\theta_0, V_0, \dots, \theta_n, V_n\} = \alpha'(\theta_n)$ and $d = 1$, RUPPERT (1982) has shown that under suitable regularity conditions,

$$(7.3) \quad n^{1/2}t^{1/2}(\theta_{\lfloor nt \rfloor} - \theta^*) \Rightarrow at^{-D-\frac{1}{2}}B(t^{2D+1})$$

as $n \rightarrow \infty$ (in the Skorohod space $D[\varepsilon, \infty)$, $\varepsilon > 0$), where $B(\cdot)$ is standard Brownian motion and a and d are certain problem-dependent constants. Setting $t = 1$ in (7.3), we conclude that θ_n converges to θ^* at rate $n^{-1/2}$ when unbiased estimators of the gradient are available.

On the other hand, in certain applications, only finite-difference approximations to the gradient may be present. For example, suppose $d = 1$ and that $P\{V_{n+1}\varepsilon A|\theta_0, V_0, \dots, \theta_n, V_n\} = P\{[X(\theta_n + cn^{1/6}) - X(\theta_n - cn^{-1/6})]/2cn^{-1/6}\varepsilon A\}$, where $X(\theta_n + cn^{-1/6})$ is simulated (under $P_{\theta_n+cn^{-1/6}}$) independently of $X(\theta_n - cn^{-1/6})$ (under

$P_{\theta_n-cn^{-1/6}}$). Here a central difference approximation to the derivative is being used (recall that $n^{-1/6}$ is the optimal difference increment as specified by Theorem 2.3). The resulting minimization algorithm is known as the **Kiefer-Wolfowitz procedure**. As one might expect, some degradation in the convergence rate occurs as a consequence of the finite-difference approximation. Specifically, RUPPERT (1982) shows that under suitable regularity hypotheses,

$$(7.4) \quad n^{1/3}t^{1/3}(\theta_{\lfloor nt \rfloor} - \theta^*) \Rightarrow b_1 + b_2t^{-A-\frac{1}{2}}B(t^{2A+1})$$

as $n \rightarrow \infty$ (in $D[\varepsilon, \infty)$), where b_1, b_2 , and A are problem-dependent constants. Thus, the convergence rate of the Kiefer-Wolfowitz procedures in which a central difference approximation is used to estimate the gradient, is $n^{-1/3}$ in the number of observations generated. Note that this convergence rate is identical to that discovered in Theorem 2.3. The fact that the convergence rates for the optimization algorithms (7.3) and (7.4) match the convergence rates of the corresponding gradient estimators indicates the pivotal role that gradient estimation plays in the optimization setting.

While the above discussion has focused on unconstrained optimization, constrained variants of (7.1)-(7.2) are also available. Among the approaches that have been studied are penalty function methods and Lagrange multiplier techniques; see RUBINSTEIN (1986) for a more extensive description.

A somewhat different philosophy for optimizing stochastic systems via simulation involves the idea of using simulation to develop a description of the global behavior of the objective function and constraints. The orthogonal function approximations of Section 6 would represent one way to obtain such global descriptions. Having fitted functional approximations to the objective function and constraints, one can then use deterministic techniques to optimize the fitted surface. One then uses the optimizer of the fitted surface as an approximation to the optimizer of the original stochastic system.

ACKNOWLEDGMENTS

This research was supported by the IBM Corporation under the SUR-SST Contract 2480042 and by the U.S. Army Research Office under Contract DAAL-03-88-K-0063.

APPENDIX

Proof of Theorem 2.2. Let $\hat{X}_i(\theta) = X_i(\theta) - \alpha(\theta)$ be the centered version of $X_i(\theta)$. We first wish to show that when $h_n \rightarrow 0$,

$$(A.1) \quad n^{-1/2} \left(\sum_{i=1}^n \hat{X}_i(\theta_0 + h_n) - \hat{X}_i(\theta_0) \right) \\ \Rightarrow \sqrt{2}\sigma(\theta_0)N(0, 1)$$

as $n \rightarrow \infty$. This follows from the central limit theorem for triangular arrays (see, for example, CHUNG (1974), pp. 205–209). In particular, to verify Lindeberg's condition, we observe that (2.1) ii) and iii) together imply that $\hat{X}_i(\theta_0 + h_n) \Rightarrow \hat{X}_i(\theta_0)$ and $E\hat{X}_i(\theta_0 + h_n)^2 \rightarrow E\hat{X}_i(\theta_0)^2$, from which it follows that $\{\hat{X}_i(\theta_0 + h_n)^2 : n \geq n_0\}$ is uniformly integrable (see Theorem 4.5.4, p. 97, CHUNG (1974)). The uniform integrability of $\{\hat{X}_i(\theta_0 + h_n)^2 : n \geq n_0\}$ implies that of $\{(\hat{X}_i(\theta_0 + h_n) - \hat{X}_i(\theta_0))^2 : n \geq n_0\}$, from which Lindeberg's condition is an easy consequence. This establishes (A.1).

We now note that

$$(A.2) \quad n^{1/4}(\Delta\alpha_1(n, h_n) - \alpha'(\theta_0)) \\ = \frac{n^{-1/2}}{n^{1/4}h} \left(\sum_{i=1}^n \hat{X}_i(\theta_0 + h_n) - \hat{X}_i(\theta_0) \right) \\ + n^{1/4}b(h_n)$$

where $b(h) = (\alpha(\theta_0 + h) - \alpha(\theta_0))/h - \alpha'(\theta_0)$. It is evident that by (2.1) iv), $\alpha(\theta_0 + h) = \alpha(\theta_0) + h\alpha'(\theta_0) + h^2\alpha''(\theta_0)/2 + o(h^2)$, so that $b(h) = h\alpha''(\theta_0)/2 + o(h)$. The theorem then follows immediately from (A.1) and (A.2).

Proof of Theorem 2.5. The proof proceeds along the same basic lines as in Theorem 2.2. We first note that (2.4) ii) and iv) together imply that $\sigma_n^2 \triangleq \text{var}\Delta Y(h_n) = h_n\sigma^2 + o(h_n) - (\alpha'(\theta_0)h_n + o(h_n))^2 = h_n\sigma^2 + o(h_n)$. Let $\Delta\hat{Y}_i(h) = \Delta Y_i(h) - E\Delta Y_i(h)$. We wish to show that if $h_n \rightarrow 0$ with $nh_n \rightarrow \infty$, then

$$(A.3) \quad \frac{1}{\sqrt{nh_n}} \sum_{i=1}^n \Delta\hat{Y}_i(h_n) \Rightarrow \sigma N(0, 1)$$

as $n \rightarrow \infty$. To obtain (A.3), we need to verify Lindeberg's condition for the triangular array $\{\Delta\hat{Y}_i(h_n)/\sqrt{nh_n} : 1 \leq i \leq n, n \geq 1\}$. But the Lindeberg condition reduces here to verifying that for $K > 0$,

$$(A.4) \quad E\{\Delta\hat{Y}_1(h_n)^2/\sigma_n^2; \Delta\hat{Y}_1(h_n)^2 > K\sigma_n^2 \cdot n\} \rightarrow 0$$

as $n \rightarrow \infty$. The left-hand side of (A.4) can be bounded by

$$E \left\{ \frac{|\Delta\hat{Y}_1(h_n)|^{2+\varepsilon}}{\sigma_n^{2+\varepsilon} K^{\varepsilon/2} n^{\varepsilon/2}}; \Delta\hat{Y}_1(h_n)^2 > K\sigma_n^2 n \right\} \\ \leq E|\Delta\hat{Y}_1(h_n)|^{2+\varepsilon}/(\sigma_n^{2+\varepsilon} K^{\varepsilon/2} n^{\varepsilon/2}) \\ \sim \frac{\beta}{\sigma^{2+\varepsilon} K^{\varepsilon/2}} (nh_n)^{-\varepsilon/2}$$

as $n \rightarrow \infty$; this yields Lindeberg's condition (since $nh_n \rightarrow \infty$). The proof is completed by writing

$$n^{1/3}(\Delta\alpha_3(n, h_n) - \alpha'(\theta_0)) \\ = \frac{1}{\sqrt{h_n} n^{1/3}} \sum_{i=1}^n \Delta\hat{Y}_i(h_n)/\sqrt{nh_n} + n^{1/3}b(h_n),$$

where $b(h) = E\Delta Y_i(h)/h - \alpha'(\theta_0)$. (To obtain the desired limit theorem, use (A.3) and $b(h_n) = \alpha''(\theta_0)h_n/2 + o(h_n)$.)

Proof of Theorem 2.13. We first observe that there exists $\varepsilon > 0$ such that $\{hZ_i(h) : 0 < h < \varepsilon\}$ is uniformly integrable. (Use the same argument as that employed in the proof of Theorem 2.2). Hence, Lindeberg's condition may be verified, to obtain the central limit theorem

$$(A.5) \quad n^{-1/2} \sum_{i=1}^n h_n Z_i(h_n) \Rightarrow \sigma_3 N(0, 1)$$

as $n \rightarrow \infty$. The next step is to study the bias term $\beta(h) = EZ_i(h) - \alpha'(\theta_0)$. We start by observing that since α is m times continuously differentiable, it follows that

$$(A.6) \quad EZ_i(h) = \frac{1}{h} \sum_{k=1}^m \sum_{\ell=0}^k \binom{k}{\ell} \frac{(-1)^{1-\ell}}{k} \alpha(\theta_0 + h\ell) \\ = \frac{1}{h} \sum_{k=1}^m \sum_{\ell=0}^k \binom{k}{\ell} \frac{(-1)^{1-\ell}}{k} \\ \left(\sum_{r=0}^m \alpha^{(r)}(\theta_0) \frac{(h\ell)^r}{r!} + o(h^m) \right) \\ = \frac{1}{h} \sum_{r=0}^m \alpha^{(r)}(\theta_0) \frac{h^r}{r!} \sum_{k=1}^m \sum_{\ell=0}^k \binom{k}{\ell} \\ \frac{(-1)^{1-\ell}}{k} \ell^r + o(h^{m-1}) \\ = \frac{1}{h} \sum_{r=0}^m \alpha^{(r)}(\theta_0) h^r \gamma_r + o(h^{m-1}),$$

where

$$\gamma_r = \sum_{k=1}^m \sum_{\ell=0}^k \binom{k}{\ell} \frac{(-1)^{1-\ell}}{kr!} \ell^r.$$

To get a handle on the γ_r 's, we note that

$$\begin{aligned} x &= \log(1 + e^x - 1) \\ &= \sum_{k=1}^m \frac{(-1)^{k+1}}{k} (e^x - 1)^k + O(x^{m+1}) \\ &= \sum_{k=1}^m (-1)^{k+1} \sum_{\ell=0}^k \binom{k}{\ell} (-1)^{k-\ell} \frac{e^{\ell x}}{k} + O(x^{m+1}) \\ &= \sum_{k=1}^m \sum_{\ell=0}^k \binom{k}{\ell} \frac{(-1)^{1-\ell}}{k} \sum_{r=0}^m \frac{\ell^r x^r}{r!} + O(x^{m+1}) \\ &= \sum_{r=0}^m \frac{x^r}{r!} \sum_{k=1}^m \sum_{\ell=0}^k \binom{k}{\ell} \frac{(-1)^{1-\ell}}{k} \ell^r + O(x^{m+1}) \\ &= \sum_{r=0}^m x^r \gamma_r + O(x^{m+1}). \end{aligned}$$

Comparing coefficients in x , we conclude that $\gamma_r = 0$ for $r \neq 1$ and $\gamma_1 = 1$. Substituting in (A.6), we conclude that $EZ_i(h) = \alpha^{(1)}(\theta_0) + o(h^{m-1})$. Hence, $\beta(h) = o(h^{m-1})$. To obtain the desired central limit theorem, we write

$$\begin{aligned} &n^{\frac{m-1}{2m}} (\alpha_5(h_n, n) - \alpha'(\theta_0)) \\ &= \frac{1}{n^{1/2m} h_n} \cdot n^{-1/2} \sum_{i=1}^n h_n Z_i(h_n) \\ &\quad + n^{\frac{m-1}{2m}} \cdot \beta(h_n), \end{aligned}$$

and use (A.5) and the estimate $\beta(h_n) = o(h_n^{m-1})$.

Proof of Theorem 6.4. To prove a), we will show that $EX^2(U) \cos^2(kU) \rightarrow 2^{-1} \int_0^\pi b(\theta) d\theta / \pi$ as $k \rightarrow \infty$. We start by observing that

$$(A.7) \quad EX^2(U) \cos^2(kU) = \int_0^\pi b(\theta) \cos^2(k\theta) d\theta / \pi.$$

Since $b(\cdot)$ is continuous, it is evident that $b(\cdot)$ is uniformly continuous on $[0, \pi]$. Hence, for every $\varepsilon > 0$, there exists

$N = N(\varepsilon)$ such that $|b(\theta) - b_N(\theta)| < \varepsilon$, where $b_N(\cdot)$ is the piecewise constant function defined by

$$b_N(\theta) = b\left(\frac{\pi}{N} \left\lfloor \frac{N\theta}{\pi} \right\rfloor\right).$$

Since $\cos^2(k\theta) \leq 1$, it follows that

$$(A.8) \quad \left| \int_0^\pi b_N(\theta) \cos^2(k\theta) d\theta - \int_0^\pi b(\theta) \cos^2(k\theta) d\theta \right| < \varepsilon,$$

$$(A.9) \quad \left| \int_0^\pi b_N(\theta) d\theta - \int_0^\pi b(\theta) d\theta \right| < \varepsilon.$$

Choose k so large that $2\pi N/k < \varepsilon$. Then, setting $\ell = \lfloor k/2N \rfloor$, we have

$$\begin{aligned} &\int_0^\pi b_N(\theta) \cos^2(k\theta) d\theta \\ &= \sum_{j=0}^{N-1} \int_{\pi j/N}^{\pi(j+1)/N} b_N(\theta) \cos^2(k\theta) d\theta \\ &= \sum_{j=0}^{N-1} b_N(\pi j/N) \int_{\pi j/N}^{\pi(j+1)/N} \cos^2(k\theta) d\theta \\ &= \sum_{j=0}^{N-1} b_N(\pi j/N) \frac{1}{k} \int_{\pi k j/N}^{\pi k(j+1)/N} \cos^2(u) du \\ &= \sum_{j=0}^{N-1} b_N(\pi j/N) \frac{1}{k} \left[\int_{\pi k j/N}^{\pi k(j+1)/N} \cos^2(u) du \right. \\ &\quad \left. + \int_0^{\pi(k/N-2\ell)} \cos^2\left(v + \frac{2\pi j}{N} + 2\pi\ell\right) dv \right] \\ &= \sum_{j=0}^{N-1} b_N(\pi j/N) \left[\frac{\pi\ell}{k} + \frac{1}{k} \int_0^{\pi(k/N-2\ell)} \cos^2\left(v + \frac{2\pi j}{N} + 2\pi\ell\right) dv \right] \\ &= \sum_{j=0}^{N-1} b_N(\pi j/N) \frac{1}{k} \int_{\pi k j/N}^{\pi k(j+1)/N} dv / 2 \\ &\quad + \sum_{j=0}^{N-1} b_N(\pi j/N) \frac{1}{k} \int_0^{\pi(k/N-2\ell)} \left[\cos\left(v + \frac{2\pi j}{N} + 2\pi\ell\right) - \frac{1}{2} \right] dv \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{2} \int_0^\pi b_N(\theta) d\theta \\
(A.10) \quad &+ \sum_{j=0}^{N-1} b_N(\pi j/N) \frac{1}{k} \int_0^{\pi(k/N-2\ell)} \\
&\left[\cos^2 \left(v + \frac{2\pi j}{N} + 2\pi\ell \right) - \frac{1}{2} \right] dv.
\end{aligned}$$

(We've used the fact that $\int_\theta^{\theta+2\pi\ell} \cos^2(u) du = \int_\theta^{\theta+2\pi\ell} \sin^2(u) du$, so that $\int_\theta^{\theta+2\pi\ell} \cos^2(u) du = 2^{-1} \int_\theta^{\theta+2\pi\ell} [\cos^2(u) + \sin^2(u)] du = \pi\ell$.) But $|k/N - 2\ell| \leq 2$, so

$$\begin{aligned}
&\left| \frac{1}{k} \int_0^{\pi(k/N-2\ell)} \left[\cos^2 \left(v + \frac{2\pi j}{N} + 2\pi\ell \right) - \frac{1}{2} \right] dv \right| \\
&\leq \frac{2\pi}{k}
\end{aligned}$$

and thus

$$\begin{aligned}
(A.11) \quad &\left| \sum_{j=0}^{N-1} b_N(\pi j/N) \frac{1}{k} \int_0^{\pi(k/N-2\ell)} \right. \\
&\left. \left[\cos^2 \left(v + \frac{2\pi j}{N} + 2\pi\ell \right) - \frac{1}{2} \right] dv \right| \leq M\varepsilon
\end{aligned}$$

where $M = \max\{|B(\theta)| : \theta \in [0, \pi]\}$. Combining (A.8)–(A.11), we obtain

$$\left| \int_0^\pi b(\theta) \cos^2(k\theta) d\theta - 2^{-1} \int_0^\pi b(\theta) d\theta \right| \leq (M+2)\varepsilon.$$

Since ε is arbitrary, we obtain our desired conclusion. Turning now to b), we note that the smoothness of α implies that (see FULKS (1969), p. 551)

$$\sup_{0 \leq \theta \leq \pi} \left| \alpha(\theta) - \sum_{k=0}^m a_k e_k(\theta) \right| = O(m^{\frac{1}{2}-p})$$

so that

$$\left\| \alpha - \sum_{k=0}^m a_k e_k \right\|^2 = O(m^{1-2p})$$

i.e.,

$$\sum_{k=m+1}^{\infty} a_k^2 = O(m^{1-2p}).$$

Hence, if $m_n \doteq n^{\frac{1}{2}p}$, it follows that $\sum_{k>m_n} a_k^2 = O(n^{\frac{1}{2}p-1})$. Furthermore, part a) implies that

$$\begin{aligned}
n^{-1} &\left(EX^2(u) + 2 \sum_{k=1}^{m_n} EX^2(u) \cos^2(ku) - \sum_{k=0}^{m_n} a_k^2 \right) \\
&= O(n^{\frac{1}{2}p-1});
\end{aligned}$$

b) then follows immediately from these estimates. Result c) is a well-known consequence of b). For d), we note that

$$\begin{aligned}
&n^{1-\gamma} \|\alpha_n - \alpha\|^2 \\
&= n^{1-\gamma} \int_{-\pi}^{\pi} (\alpha_n(\theta) - \alpha(\theta))^2 d\theta \\
&\geq \varepsilon m \left\{ \theta \in [-\pi, \pi] : |\alpha_n(\theta) - \alpha(\theta)| \geq \varepsilon n^{\frac{\gamma}{2}-\frac{1}{2}} \right\}.
\end{aligned}$$

Part b) implies that $Em\{\theta \in [-\pi, \pi] : |\alpha_n(\theta) - \alpha(\theta)| \geq \varepsilon n^{\frac{\gamma}{2}-\frac{1}{2}}\} \rightarrow 0$ as $n \rightarrow \infty$, yielding d).

REFERENCES

- CHUNG, K.L. (1974). *A First Course in Probability Theory*, Academic Press, New York.
- FOX, B.L. and GLYNN, P.W. (1989). Replication schemes for limiting replications. To appear in *Probability Theory of the Engineering and Informational Sciences*.
- FULKS, W. (1969). *Advanced Calculus*. John Wiley, New York.
- GLASSERMAN, P. and GONG, W. (1989). Derivative estimates from discontinuous realizations: smoothing techniques. To appear in *Proceedings of the 1989 Winter Simulation Conference*.
- GLYNN, P.W. (1987). Construction of process-differentiable representations for parametric families of distributions. Technical Report, Mathematics Research Center, University of Wisconsin, Madison.
- HEIDELBERGER, P., CAO, X.R., ZAZANIS, M.A., and SURJ, R. (1988). Convergence properties of infinitesimal

- perturbation analysis estimates. *Management Science*, **34**, 1281–1302.
- HO, Y.C. and CAO, X.R. (1983). Perturbation analysis of queueing networks. *Journal of Optimization Theory and Applications* **40**, 559–582.
- HO, Y.C. and CAO, X.R. (1985). Performance sensitivity to routing changes in queueing networks and flexible manufacturing systems using perturbation analysis. *IEEE Journal of Robotics and Automation*, **RA-1**, 615–672.
- POLYAK, B.T. and TSYPKIN, Y.Z. (1980). Optimal pseudogradient adaptive algorithms. *Automat. Remote Control* **8**, 74–84.
- REIMAN, M.I. and WEISS, A. (1986). Sensitivity analysis via likelihood ratios. *Proceedings of the 1986 Winter Simulation Conference*, 285–289.
- RUBINSTEIN, R.Y. (1986). *Monte Carlo Optimization, Simulation and Sensitivity of Queueing Networks*. John Wiley, New York.
- RUPPERT, D. (1982). Almost sure approximations to the Robbins-Monro and Kiefer-Wolfowitz processes with dependent noise. *Annals of Probability* **10**, 178–187.
- SURI, R. Infinitesimal perturbation analysis for general discrete event systems. *J. Assoc. Comput. Mach.* **34**, 686–717.
- ZAZANIS, M.A. and SURI, R. (1986). Comparison of perturbation analysis with conventional sensitivity estimates for regenerative stochastic systems. Technical Report, Harvard University.

AUTHOR'S BIOGRAPHY

PETER W. GLYNN received the Ph.D. degree from Stanford University, Stanford, CA in 1982. He spent five years as an Assistant Professor in the Department of Industrial Engineering at the University of Wisconsin-Madison. He is currently an Associate Professor in the Department of Operations Research at Stanford University. His research interests include stochastic systems, computational probability, and simulation.