MANAGEMENT SCIENCE
Vol. 36, No. 11, November 1990
Printed in U.S.A.

A UNIFIED VIEW OF THE IPA, SF, AND LR GRADIENT ESTIMATION TECHNIQUES*

PIERRE L'ECUYER
Département d'informatique, Université Laval, Ste-Foy, Québec, Canada G1K 7P4

We study the links between the likelihood-ratio (LR) gradient-estimation technique (sometimes called the score-function (SF) method), and infinitesimal perturbation analysis (IPA). We show how IPA can be viewed as a (degenerate) special case of the LR and SF techniques by selecting an appropriate representation of the underlying sample space for a given simulation experiment. We also show how different definitions of the sample space yield different variants of the LR method, some of them mixing IPA with more straightforward LR. We illustrate this by many examples. We also give sufficient conditions under which the gradient estimators are unbiased.
(DISCRETE EVENT SYSTEMS; SIMULATION; GRADIENT ESTIMATION; PERTURBATION ANALYSIS; LIKELIHOOD RATIO; SCORE FUNCTION)

1. The LR Gradient Estimation Technique

Consider a stochastic simulation model parameterized by a real vector $\theta \in \Theta$ of continuous parameters, where Θ is some open subset of $\mathbb{R}^d$. Suppose we want to estimate the gradient $\nabla\alpha(\theta)$ of the (differentiable) expected value $\alpha(\theta)$ of some real-valued objective function. Infinitesimal perturbation analysis (IPA) and a likelihood ratios (LR) approach, sometimes called the score function (SF) method, are among the proposed techniques to estimate such a gradient by simulation. For IPA, see Glasserman (1988a, b, c), Gong and Ho (1987), Heidelberger et al. (1989), Ho (1987), and Suri (1987, 1989). For LR (or SF), see Aleksandrov, Sysoyev, and Shemeneva (1968), Arsham et al. (1989), Bratley, Fox, and Schrage (1987), Glynn (1987), Reiman and Weiss (1989), and Rubinstein (1986, 1989, 1990). Other techniques are based on finite differences (see, for example, L'Ecuyer, Giroux, and Glynn 1990; Meketon 1987; Suri 1989; and Arsham et al. 1989). IPA and LR (or SF) have often been viewed as distinct (competing) techniques. The aim of this paper is to present them in a unified framework and show the strong links that exist between them.

The basic idea of LR is that $\alpha(\theta)$ can usually be viewed as the expectation of some function of θ and of the "sample path" ω, say $h(\theta, \omega)$, with respect to a probability measure P_θ over some measurable space (Ω, Σ). Here, Ω is the sample space, and $\omega \in \Omega$ represents all the "random elements" in the simulation, so that when ω is fixed, the evolution of the system becomes deterministic. More specifically, we assume that $h(\theta, \cdot)$ is Σ-measurable. Usually, one cannot differentiate this expectation directly by differentiating inside the integral: one reason is that P_θ typically depends on θ. That dependence can be eliminated if one can take, on the same measurable space, a probability measure G independent of θ that dominates the P_θ's for $\theta \in \Theta$; this means that P_θ is absolutely continuous with respect to G—that is for every measurable set B, $G(B) = 0$ implies $P_\theta(B) = 0$. In that case, one can rewrite

$$\alpha(\theta) = \int_\Omega h(\theta, \omega) dP_\theta(\omega) = \int_\Omega \left[h(\theta, \omega) \frac{dP_\theta}{dG}(\omega) \right] dG(\omega)$$

$$= \int_\Omega h(\theta, \omega) L(G, \theta, \omega) dG(\omega), \qquad (1)$$

* Accepted by James R. Wilson; received March 16, 1989. This paper has been with the author 3 months for 2 revisions.

0025-1909/90/3611/1364$01.25

where $L(G, \theta, \omega) = (dP_\theta/dG)(\omega)$ is the Radon-Nikodym derivative of P_θ with respect to G. By the Radon-Nikodym theorem (Billingsley 1986, Theorem 32.2), $L(G, \theta, \omega)$ exists and (1) is valid if and only if G dominates P_θ. If sampling is done using G, then $h(\theta, \omega)L(G, \theta, \omega)$ can be used as an estimator of $\alpha(\theta)$. This "change of measure" approach, called importance sampling, is also a variance reduction technique (Bratley, Fox, and Schrage 1987; Glynn and Iglehart 1988; Rubinstein 1989; and Goyal et al. 1989).

Under appropriate regularity conditions (that permit interchanging the derivative and expectation), one can differentiate α by differentiating the bracketed term with respect to θ inside the integral:

$$\nabla_\theta \alpha(\theta) = \int_\Omega \psi(\theta, \omega) dG(\omega), \tag{2}$$

where

$$\begin{aligned}\psi(\theta, \omega) &= \nabla_\theta (h(\theta, \omega)L(G, \theta, \omega)) \\ &= L(G, \theta, \omega)\nabla_\theta h(\theta, \omega) + h(\theta, \omega)\nabla_\theta L(G, \theta, \omega)\end{aligned} \tag{3}$$

when this gradient exists. Sufficient regularity conditions are given in Glasserman (1988a, b, c), Reiman and Weiss (1989), and Rubinstein (1989), either in very general form, or for special cases. In §3 of this paper, we give conditions that are not the most general, but are practical and general enough for most applications. The need for these regularity conditions and for the existence of the Radon-Nikodym derivative certainly impose some limitations on gradient estimation methods that are based on (2) and (3) (see the examples in §4), but there are also many practical cases where these methods work well.

When (2) holds, $\psi(\theta, \omega)$ can be used to estimate $\nabla\alpha(\theta)$. Only one simulation experiment (using G) is required to estimate the gradient. However, the work to compute $\psi(\theta, \omega)$ could be much higher than the work to generate ω and compute $h(\theta, \omega)$, especially if d is large. In principle, $\psi(\cdot, \omega)$ can be evaluated at any value of θ for which (2) holds, permitting estimation of the gradient at different points with the same G (Arsham et al. 1989; Rubinstein 1986 and 1990). But the variance of the gradient estimator is sometimes dramatically high for some values of θ. In some cases, for any fixed G and ω, the estimator $\psi(\theta, \omega)$ can be written down explicitly, after observing the simulation, as a function of θ. This permits one to estimate the gradient everywhere in Θ by a single simulation. The (sample) function $\psi(\cdot, \omega)$ can then be used directly in a standard (nonlinear) optimization algorithm, yielding an estimator of the optimum by a single simulation. Rubinstein (1990) discusses this approach and gives numerical results.

How do we choose G? Among those G for which (2) and (1) hold, one would like to choose one for which the variance is low, $\psi(\theta, \omega)$ is easy to compute, and ω is easy to generate. But this is not always easy to do.

To estimate $\nabla\alpha(\theta_0)$ at a single point θ_0, an easy choice for G, when (1) and (2) hold for it, is P_{θ_0}. In that case, $L(P_{\theta_0}, \theta_0, \omega) = 1$ and at $\theta = \theta_0$, we obtain

$$\psi(\theta, \omega) = \nabla_\theta h(\theta, \omega) + h(\theta, \omega)S(\theta, \omega), \tag{4}$$

where

$$S(\theta, \omega) = \nabla_\nu L(P_\theta, \nu, \omega)|_{\nu=\theta}. \tag{5}$$

Since θ_0 can vary in Θ, (4) can be viewed in that case as defining $\psi(\theta, \omega)$ for all $\theta \in \Theta$, with the interpretation that ω is generated according to P_θ.

Typically, ω can be viewed as the set of values taken by a finite sequence of independent (possibly multivariate) random variables. For example, let $\omega = (x_1, \ldots, x_p, y_1, \ldots, y_q)$, where for $1 \le i \le p$ and $1 \le j \le q$, x_i is the value taken by a continuous random variable with density $f_{i,\theta}$ and y_j is the value taken by a discrete random variable with

probability mass $p_{j,\theta}$. In that case, with $G = P_{\theta_0}$, the Radon-Nikodym derivative becomes the likelihood ratio

$$L(P_{\theta_0}, \theta, \omega) = \left(\prod_{i=1}^{p} \frac{f_{i,\theta}(x_i)}{f_{i,\theta_0}(x_i)}\right)\left(\prod_{j=1}^{q} \frac{p_{j,\theta}(y_j)}{p_{j,\theta_0}(y_j)}\right) \tag{6}$$

and its gradient is

$$S(\theta, \omega) = \nabla_\theta \ln\left(\prod_{i=1}^{p} f_{i,\theta}(x_i) \prod_{j=1}^{q} p_{j,\theta}(y_j)\right) = \sum_{i=1}^{p} \nabla_\theta \ln f_{i,\theta}(x_i) + \sum_{j=1}^{q} \nabla_\theta \ln p_{j,\theta}(y_j). \tag{7}$$

Note that $S(\theta, \omega)$ is a sum whose number of terms $(p + q)$ is the number of basic density and probability mass functions that depend on θ. Of course, this number can be random (that is, p and q can be different for different ω's). Assuming that the variances of the sample performance measure $h(\theta, \omega)$ and of the gradient $\nabla_\theta h(\theta, \omega)$ are bounded, we see that the variance of the gradient estimator (4) increases linearly (in general) with $(p + q)$ (which is typically a linear function of the simulation length). From this reasoning, we should expect this technique to work much better for finite-horizon simulations for which only a small number of random variates are generated with probability laws that depend on θ. L'Ecuyer and Glynn (1990) have also developed a control variate approach that reduces the variance from $O(p + q)$ to $O(1)$.

Since regenerative simulations can be analyzed in a way very similar to finite-horizon simulations (Glynn 1987; Glynn and Iglehart 1988), the above remark also applies to steady-state regenerative simulations for which a small number of θ-dependent variates are generated per regenerative cycle. A version of the method specially adapted for regenerative systems is presented in Glynn (1987). Arsham et al. (1989) present another version based on the "batch means" technique, for dealing with steady-state systems.

The method can be generalized to higher order derivatives (see Reiman and Weiss 1989; Rubinstein 1989 and 1990 for special cases). Let $\nabla_\theta^2 \alpha(\theta)$ denote the $d \times d$ Hessian matrix, whose element (i, j) is the (second order) partial derivative of $\alpha(\theta)$, with respect to the components i and j of θ. More generally, let $\nabla_\theta^r \alpha(\theta)$ denote the $d \times d \times \cdots \times d$ (r times) array (that is, r-dimensional tensor) that contains all rth order partial derivatives. To obtain an estimator of $\nabla_\theta^r \alpha(\theta)$, take the $(r - 1)$th derivative of (3). This yields

$$\widehat{\nabla_\theta^r \alpha}(\theta) = \nabla_\theta^{r-1} \psi(\theta, \omega) = \sum_{i=0}^{r} \binom{r}{i} (\nabla_\theta^i h(\theta, \omega)) \otimes (\nabla_\theta^{r-i} L(G, \theta, \omega)), \tag{8}$$

where $\otimes$ denotes the external product (the product of an i-dimensional tensor by a $(r - i)$-dimensional one gives an r-dimensional tensor). Sufficient conditions for unbiasedness can be obtained by adapting Assumption A1 of §3. As a special case, if $G = P_{\theta_0}$ and $h(\theta, \omega)$ is independent of θ, then the estimator of $\nabla_\theta^r \alpha(\theta)$ becomes

$$h(\theta, \omega) \nabla_\theta^r L(P_{\theta_0}, \theta, \omega) = h(\theta, \omega) \nabla_\theta^{r-1} S(\theta, \omega).$$

For example, an unbiased estimator of the Hessian is obtained, under appropriate regularity conditions, by differentiating (3) with respect to θ, which yields

$$\widehat{\nabla_\theta^2 \alpha}(\theta) = h(\theta, \omega) \nabla_\theta^2 L(G, \theta, \omega) + 2 \nabla_\theta L(G, \theta, \omega) (\nabla_\theta h(\theta, \omega))' + L(G, \theta, \omega) \nabla_\theta^2 h(\theta, \omega), \tag{9}$$

where $'$ denotes the transposition into a line vector. If $G = P_{\theta_0}$, that Hessian estimator becomes

$$\widehat{\nabla_\theta^2 \alpha}(\theta) = h(\theta, \omega) \nabla_\theta S(\theta, \omega) + 2 S(\theta, \omega) (\nabla_\theta h(\theta, \omega))' + \nabla_\theta^2 h(\theta, \omega) + h(\theta, \omega) [S(\theta, \omega)]^2. \tag{10}$$

Using the right-hand-side expression in (4) as a gradient estimator has been proposed (sometimes in a less general setting) by, e.g., Aleksandrov, Sysoyev, and Shemeneva (1968), Arsham et al. (1989), Glynn (1987), Reiman and Weiss (1989), and Rubinstein (1986, 1989, 1990). Some authors call this technique the score function (SF) method, because in statistics, the expression (7) is called the score function. Others call it the likelihood ratio (LR) method, because $L(G, \theta, \omega)$ is typically a likelihood ratio. Arsham et al. (1989) reserve the term LR for the more general estimators, like (3), that incorporate importance sampling. Another name for that (general) technique could be "common probability space approach." In this paper, we adopt the following definition.

DEFINITION 1. An LR gradient estimator for $\nabla_\theta \alpha(\theta)$ is one that is defined by the right-hand-side expression in (3), where ω obeys G, provided that this expression exists for almost all ω and $h(\theta, \cdot)$ is a measurable function of ω whose expectation is $\alpha(\theta)$. We call such an estimator an SF gradient estimator if it is also defined by (4), where ω obeys P_θ. Under the same conditions, the estimators (9) and (10) are called LR and SF Hessian estimators respectively. □

2. Choosing What ω Should Represent

2.1. *The Underlying Sample Space*

A key point of this paper concerns the choice of the sample space Ω. In fact, for a given simulation model, there can be different ways of defining the sample space and the meaning of ω. Generally, all random variables are generated by generating $U(0, 1)$ variates as needed and transforming them in the appropriate way. Hence, ω can be viewed as a sequence of independent $U(0, 1)$ variates, and the value of the objective function $h(\theta, \omega)$ is a (measurable) function of this sequence. But this is only one way of viewing it. In fact, there is no need to assume that $U(0, 1)$ variates are used to drive the simulation in the first place. For instance, the SF technique can be used to estimate the gradient not only for a simulation model, but also for a real system (provided P_θ is known and ω can be observed). In that case, ω will usually not be a sequence of $U(0, 1)$ variates. Consider for instance an M/G/1 queue. One possibility is to view ω as the sequence of interarrival and service times, and define $\Omega = [0, \infty)^\infty$ as the sample space. In that case, the distribution of ω depends on θ, whereas it does not when ω is defined as a sequence of $U(0, 1)$ variates.

So, there might be different ways of defining the sample space (and the associated probability space $(\Omega, \Sigma, P_\theta)$). Different choices may lead to different gradient estimators, some being more efficient than others. This will be discussed further in the next subsection.

Some might feel more comfortable with an underlying probability space in which the "basic" random element is a sequence of independent $U(0, 1)$ variates. Let $(\tilde{\Omega}, \tilde{\Sigma}, \tilde{P})$ be such a space. If $(\Omega, \Sigma, P_\theta) \neq (\tilde{\Omega}, \tilde{\Sigma}, \tilde{P})$, we assume that there is a measurable transformation $\phi_\theta: \tilde{\Omega} \to \Omega$ such that $\omega = \phi_\theta(\tilde{\omega})$ and such that $P_\theta(B) = \tilde{P}(\phi_\theta^{-1}(B))$ for every measurable set B. Note that ω may contain less information than $\tilde{\omega}$, and that its probability law may depend on θ. We can also define the $\tilde{\Sigma}$-measurable function $\tilde{h}(\theta, \cdot)$ by $\tilde{h}(\theta, \tilde{\omega}) = h(\theta, \phi_\theta(\tilde{\omega}))$ (note that $h(\theta, \cdot)$ is Σ-measurable). We have

$$\alpha(\theta) = \int_\Omega h(\theta, \omega) dP_\theta(\omega) = \int_{\tilde{\Omega}} \tilde{h}(\theta, \tilde{\omega}) d\tilde{P}(\tilde{\omega}).$$

In fact, this notion of an underlying sample space $\tilde{\Omega}$ and transformation ϕ_θ is not really necessary. We introduced it here just to clarify some links with the common practice (in simulation) of viewing the sample space that way. What we really have in mind is to just define the sample space as Ω and forget about $\tilde{\Omega}$.

2.2. *Infinitesimal Perturbation Analysis*

Suppose that P_θ is defined in such a way that ω can be viewed as a sequence of independent $U(0, 1)$ variates. Then, P_θ is independent of θ, $L(P_{\theta_0}, \theta, \omega) = 1$, and the last term in (4) vanishes. In that case, under the appropriate regularity conditions, we obtain

$$\nabla\alpha(\theta) = \int_\Omega \nabla_\theta h(\theta, \omega) dP_\theta(\omega) \tag{11}$$

and $\nabla_\theta h(\theta, \omega)$ is the usual IPA gradient estimator (Cao 1987; Glasserman 1988a, b; Heidelberger et al. 1989; Ho 1987; Rubinstein 1989; Suri 1987, 1989).

DEFINITION 2. An IPA estimator for $\nabla\alpha(\theta)$ is defined as $\nabla_\theta h(\theta, \omega)$ (provided that this quantity exists for almost all ω), where ω is a sequence of values taken by independent $U(0, 1)$ random variables, and $h(\theta, \cdot)$ is a measurable function of ω whose expectation is $\alpha(\theta)$. Likewise, $\nabla^2_\theta h(\theta, \omega)$ is an IPA Hessian estimator. □

When (11) is satisfied, we have an unbiased IPA gradient estimator (and similarly for the Hessian). There often exist different functions $h(\theta, \cdot)$ that satisfy the above definition, and thus different IPA gradient (or Hessian) estimators for the same α. Often, some are unbiased while others are not. In practice, the function $h(\theta, \cdot)$ is usually defined or implied by the simulation model.

The basic idea of IPA is to generate a sample path ω, viewed as a sequence of $U(0, 1)$ variates; and then for ω fixed, to observe the effect of an infinitesimal perturbation on θ (around θ_0) by propagating it over the sample path, assuming that the sequence of events does not change and that the events can only "slide" in time. The gradient estimator is taken as the gradient of the objective function for that fixed value of ω. The propagation rules permit evaluating $\nabla_\theta h(\theta, \omega)$ only at $\theta = \theta_0$, and thus estimating the gradient only at θ_0. If that definition of ω is used in (2) with $G \neq P_{\theta_0}$, one gets a combination of IPA with importance sampling.

According to the above definitions, IPA can be viewed as a special case of SF (and LR). One big advantage of IPA is that since no component of ω depends on θ, the variance no longer increases with the simulation length. But the function $h(\theta, \omega)$ must absorb all the transformations and may become overly complex, sometimes making the actual computation of $\nabla_\theta h(\theta, \omega)$ intractable, or invalidating (2). Typically, in that case, one of the requirements of assumption $A1$ (§3) is no longer satisfied. In §4, we give examples of this. In many cases, $h(\cdot, \omega)$ remains differentiable almost everywhere, but becomes discontinuous. Sometimes, a clever definition of $h(\theta, \omega)$ might yield an unbiased IPA estimator, while more straightforward definitions lead to biased estimators. See, for example, Glasserman (1989) and Gong and Ho (1987). In fact, a large part of the IPA literature deals with the development of effective techniques to compute $\nabla_\theta h(\theta, \omega)$ during the simulation (see, e.g., Glasserman 1989; Gong and Ho 1987; Ho 1987; Suri 1987 and 1989). One might even associate the term IPA more with these techniques than with equation (11), and these "IPA" techniques can be used to implement LR as well. Some people might argue that since the likelihood ratio has disappeared in (11), this is no longer LR. In fact, we just view it as a special case (or degenerate case, in some sense).

As we said before, ω can also represent something other than a sequence of $U(0, 1)$ variates. For example, ω can be viewed as representing the whole history of the system, including all events with their types and occurrence times, etc. In many cases, ω can carry enough θ-dependent high level information so that for a given value of ω, $h(\theta, \omega)$ does not depend on θ any more, and the first term of the right-hand-side expression in (3) vanishes. In that case, we will say that we have pure LR (or SF). But if ω carries too much information, one might be unable to write down P_θ explicitly, preventing the actual computation of the score function and of the gradient estimator. For this reason, ω is

usually taken as a sequence of independent random variates. In Glynn (1987), for example, when simulating a Markov chain, ω is taken as the sequence of transitions of the chain.

For the extreme case where ω is defined as the value of the objective function itself, i.e., $h(\theta, \omega) = \omega$ (and $\Omega = \mathbb{R}$) by definition, P_θ is actually the distribution function of the cost. When we can write down P_θ and the associated likelihood ratio, there is usually no need to simulate, since $\nabla\alpha(\theta)$ can be computed directly. Monte-Carlo methods are precisely useful for the cases where we cannot efficiently compute the expression directly.

Between these extremes, there are often other possibilities. For instance, if a set of $U(0, 1)$ values must go through many levels of transformation, one may choose any one of the levels to define ω. See §4.5 for an example. Also, ω might contain the original $U(0, 1)$ values for some of the generated random variables, and the transformed values for others. This gives rise to hybrid methods "mixing" in some way IPA with SF. According to our definitions, this is still SF (and LR). In §4, we give examples for which one might think that SF does not apply, but for which SF effectively applies if the sample space and ω are defined appropriately.

But what is the best way, then, to define ω? There is no easy answer to this question. There is no straightforward recipe. Of course, one would like (2) to be valid. There are examples for which (2) is valid if ω represents higher-level information, e.g., the set of actual interarrival times, service times, and transitions between nodes in a queueing network, while (2) is not valid if ω represents the sequence of $U(0, 1)$ variates. But for other examples, the opposite is true (see Reiman and Weiss 1989; Glasserman 1988a, b; and the examples in §4). In certain situations, (2) might be valid for none of the extreme cases, but for some intermediate definition of ω (see §4.5). If (2) is valid for many possible definitions of ω, one will then try to minimize the variance multiplied by the expected work per run (see Fox and Glynn 1990). This is certainly problem-dependent, but according to the remarks following equation (7), trying to put the least number of θ-dependent components in ω appears to be a good strategy.

One consequence of the above discussion is that many properties of the IPA method also apply to LR, and vice-versa. For example, the validity of interchanging the derivative and expectation is a problem for the LR method in general. Various (problem dependent) devices have been suggested to "smooth out" or transform some problems for which IPA does not apply directly, into problems for which IPA will work correctly (see, e.g., Glasserman 1989; Gong and Ho 1987; the references in Glasserman 1988a, b; Heidelberger et al. 1989; and Suri 1989). In principle, one could think of developing such devices for LR in general.

When $\alpha(\theta)$ is a steady-state performance measure and ω contains an infinite sequence of θ-dependent random variates, the Radon-Nikodym derivative in (1) typically does not exist. However, (11) might be valid in that case, and then, one typically has $\nabla_\theta h(\theta, \omega) = \nabla\alpha(\theta)$ with probability one (this is when IPA is strongly consistent).

3. Interchanging the Derivative and Expectation

In this section, we give sufficient validity conditions for the interchange of derivative and expectation that is required to justify (2). Conditions for specific cases are also given in Glasserman (1988a, b, c), Reiman and Weiss (1989), and Rubinstein (1989). Verifying directly the most general conditions (uniform integrability) is usually not straightforward. But often, the conditions below can be verified directly, as we will see in the examples of the next section. See also the discussion in Glasserman (1988a).

Each component of the gradient can be dealt with separately. For $i = 1, \ldots, d$, to study the ith component of the gradient, we look at what happens when only the component i of θ is allowed to change and all other components of θ are fixed. To simplify

things, in this section, we assume that $d = 1$. For the more general case, just apply the results below to each component of θ (while the other components are fixed). All probabilistic statements in this section are made with respect to the probability measure G. The theorem below is an adaptation of Lemma 1 in Glasserman (1988a). It uses the following assumption:

ASSUMPTION A1. Let $d = 1$ and $H(\theta, \omega) = h(\theta, \omega)L(G, \theta, \omega)$. There is an open neighborhood Υ of θ_0, $\Upsilon \subseteq \Theta$, and a measurable set $\Xi \subseteq \Omega$, such that $G(\Xi) = 1$, and for all $\omega \in \Xi$, $H(\cdot, \omega)$ exists and is continuous everywhere in Υ, and is also differentiable in $D(\omega) \subseteq \Upsilon$, where $\Upsilon \backslash D(\omega)$ is at most a denumerable set. Assume that (1) holds for all $\theta \in \Upsilon$. Define

$$\Psi(\omega) = \begin{cases} \sup_{\theta \in D(\omega)} |\psi(\theta, \omega)| & \text{if} \quad \omega \in \Xi; \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

Assume that there exists a function $\Gamma\colon \Omega \to \mathbb{R}$, integrable with respect to G, such that $\Psi(\omega) \leq \Gamma(\omega)$ for all $\omega \in \Omega$. □

THEOREM 1. *Under* A1, *equation* (2) *is valid for* $\theta = \theta_0$.

PROOF. The proof is largely inspired by the proof of Lemma 1 in Glasserman (1988a). From a generalized version of the mean value theorem (see Theorem 8.5.3 in Dieudonné 1969), if θ_0 and $\theta_0 + h$ are in Υ and $\omega \in \Xi$,

$$\left|\frac{H(\theta_0 + h, \omega) - H(\theta_0, \omega)}{h}\right| \leq \sup_{\theta \in [\theta_0, \theta_0 + h] \cap D(\omega)} |\psi(\theta, \omega)| \leq \Gamma(\omega).$$

Hence, from the dominated convergence theorem,

$$\begin{aligned} \int_\Omega \psi(\theta_0, \omega) dG(\omega) &= \int_\Omega \left(\lim_{h \to 0} \frac{H(\theta_0 + h, \omega) - H(\theta_0, \omega)}{h}\right) dG(\omega) \\ &= \lim_{h \to 0} \left(\frac{\alpha(\theta_0 + h) - \alpha(\theta_0)}{h}\right) \\ &= \nabla\alpha(\theta_0). \qquad \square \end{aligned}$$

In Assumption A1, we do not need continuous differentiability, nor differentiability everywhere in Υ. But continuity is usually essential. Saying that (1) holds for all $\theta \in \Upsilon$ means that the first two equalities in (1) hold, and implies in particular that G dominates P_θ for all $\theta \in \Upsilon$. We only require Ψ to be bounded by an integrable function instead of requiring its own integrability. This way, we do not need to check if Ψ is measurable. Checking the last condition of A1 using (12) is typically easier than directly bounding $(H(\theta_0 + h, \omega) - H(\theta_0, \omega))/h$ uniformly in h for h in some neighborhood of 0. If for all $\omega \in \Xi$, $H(\cdot, \omega)$ is continuously differentiable in Υ and Υ is bounded, then one can trivially take $\Gamma(\omega) = \Psi(\omega) = \sup_{\theta \in \Upsilon} |\nabla_\theta H(\theta, \omega)|$ as the integrable bounding function. (See also Fox. 1989).

4. Examples

In this section, we give simple examples, some with numerical results, to illustrate the main ideas of the paper. The first five examples deal with a simple $M/G/1$ queue that evolves until a certain (fixed) number of departures have occurred (this is slightly simpler than the case of a fixed-time horizon, where the SF has a random number of terms and the IPA estimator is a bit more complicated to compute). The next two examples consider the lifetime of a k-out-of-N reliability system, without repairs. The following one discusses

a general continuous-time Markov chain, while the last one looks at sensitivity with respect to thresholds. In particular, we look at replacement policies defined by thresholds in a multicomponent system. In the latter case, we actually do not know how LR can be used efficiently to estimate the gradient.

Consider a $M/G/1$ queue, initially empty, and let $\alpha(\theta)$ be the expected mean system time (waiting + service times) for the first T customers in the system, where θ is a parameter of the service time distribution. The arrival rate is $\lambda = 1$. We want to estimate the derivative $\alpha'(\theta)$ at a given point $\theta = \theta_0$ by simulating at that point. For a given realization ω, $h(\theta, \omega)$ represents the observed average system time for the T customers. We have

$$h(\theta, \omega) = \frac{1}{T} \sum_{i=1}^{T} (W_i + S_i) \tag{13}$$

where W_i and S_i are respectively the actual (observed) waiting time and service time for customer i (these are [deterministic] functions of θ and ω). Let A_i denote the interarrival time between customers $i-1$ and i (A_1 is the arrival time of customer 1 and the system starts at time 0). We have $W_1 = 0$, and $W_{i+1} = \max(0, W_i - A_{i+1} + S_i)$ for $i > 0$. The first five examples below are variants of this one; only the service time distributions differ. Application of IPA to this system has been analyzed in Suri and Zazanis (1988) when the objective function is the steady-state average system time per customer. It has been shown that under some conditions on the service time distribution, IPA gives an asymptotically unbiased and strongly consistent gradient estimate. For the case of a finite number of customers, the validity of IPA has been analyzed, e.g., in Glasserman (1988a, Example 4).

4.1. *Applying IPA or SF to an $M/M/1$ Queue*

Let the service time distribution be exponential with mean θ, $0 < a < \theta < b$. IPA is known to work for that case: assuming that the interarrival and service times are generated by inversion, one can take ω as the sequence of $U(0, 1)$ values used to generate them, i.e., $\omega = (U_1, \ldots, U_{2T})$, $A_i = -\ln(1 - U_{2i-1})$ and $S_i = -\theta \ln(1 - U_{2i})$. An infinitesimal perturbation on S_i affects the system time of customer i and of all the customers (if any) that follow him in the same busy period. Therefore,

$$\psi(\theta, \omega) = \nabla_\theta h(\theta, \omega) = \frac{1}{T} \sum_{i=1}^{T} \sum_{j \in B_i} \frac{\partial S_j}{\partial \theta} \tag{14}$$

where B_i is the set containing customer i and all the customers that precede him in the same busy period (if any), and $\partial S_i/\partial\theta = S_i/\theta$. This can be computed during the simulation as described in Suri and Zazanis (1988). We can easily verify assumption A1. In fact, for any ω, S_i is continuous and differentiable in θ, and $h(\cdot, \omega)$ is continuous in the S_i's (and in the W_i's, which are continuous in the S_i's). Also, $h(\cdot, \omega)$ is not differentiable at θ only when two events (arrival or departure) occur simultaneously, and this happens at most for a finite number of values of θ. Since

$$\sup_{a<\theta<b} |\psi(\theta, \omega)| \le (1/T) \sum_{i=1}^{T} \sum_{j=1}^{i} (-\ln(1 - U_{2j})) \le -\sum_{j=1}^{T} \ln(1 - U_{2j})$$

and the latter is integrable, Lemma 1 applies and IPA provides an unbiased gradient estimator for that case. However, A1 is not satisfied for the IPA Hessian estimator, since $\nabla_\theta h(\cdot, \omega)$ is usually discontinuous. In fact, $\nabla_\theta^2 h(\theta, \omega)$ is zero whenever it exists.

Another choice takes ω as the set of actual interarrival and service times: $\omega = (A_1, S_1, \ldots, A_T, S_T)$ and $G = P_{\theta_0}$. In this case, $dP_\theta(\omega)/d\omega$ is the product of their densities:

$$dP_\theta(\omega) = \prod_{i=1}^{T} \left(\frac{1}{\theta} e^{-S_i/\theta} e^{-A_i} dS_i dA_i \right),$$

and $\nabla_\theta h(\theta, \omega) = 0$. This is pure SF. Only the service time densities appear in the likelihood ratio, since the interarrival times are independent of θ (in fact, taking either the actual interarrival times or the corresponding $U(0, 1)$ values in ω makes no difference here). One has

$$L(P_{\theta_0}, \theta, \omega) = (\theta_0/\theta)^T \prod_{i=1}^{T} e^{-S_i/\theta + S_i/\theta_0}$$

and

$$S(\theta, \omega) = \sum_{i=1}^{T} \frac{\partial}{\partial \theta} \ln \left(\frac{1}{\theta} e^{-S_i/\theta} \right) = \frac{1}{\theta^2} \sum_{i=1}^{T} (S_i - \theta).$$

$S(\theta, \omega)$ can be computed easily together with $h(\theta, \omega)$ during the simulation. For any ω, $H(\cdot, \omega)$ is continuous and differentiable in (a, b) (note that $H(\theta, \omega)$ depends on θ only through $L(P_{\theta_0}, \theta, \omega)$). The gradient estimator (4), for all $\theta \in (a, b)$, is

$$\psi(\theta, \omega) = \frac{1}{T\theta^2} \left(\sum_{i=1}^{T} (W_i + S) \right) \left(\sum_{i=1}^{T} (S_i - \theta) \right).$$

Here, for any ω,

$$\Psi(\omega) \le \left(\sum_{i=1}^{T} (W_i + S_i) \right) \left(\sum_{i=1}^{T} (S_i + b) \right) / (Ta^2).$$

For each $\theta_0 \in (a, b)$, the latter is clearly P_{θ_0}-integrable, so that Lemma 1 applies. The corresponding Hessian estimator

$$h(\theta, \omega) \nabla_\theta S(\theta, \omega) = \frac{1}{T\theta^3} \left(\sum_{i=1}^{T} (W_i + S_i) \right) \left(T\theta - 2 \sum_{i=1}^{T} S_i \right)$$

is also unbiased.

When the cost function is a sum of terms like in (13), one can estimate the gradient of each individual term separately. For a given term, it is not always necessary to use the "global" score function; one can use a score function based only on the random variables that influence that term. This usually reduces variance. For example, to estimate the gradient of

$$\alpha_i(\theta) = E[h_i(\theta, \omega)] = E[(W_i + S_i)/T],$$

using "pure" SF as above, one can use the score function

$$S_i(\theta, \omega) = \sum_{j=1}^{i} (S_j - \theta)/\theta^2,$$

in which case the estimator becomes

$$\psi_i(\theta, \omega) = (W_i + S_i) \sum_{j=1}^{i} (S_j - \theta)/(T\theta^2).$$

An estimator of $\nabla\alpha(\theta)$ is then

$$\sum_{i=1}^{T} \psi_i(\theta, \omega) = \frac{1}{T\theta^2} \sum_{i=1}^{T} \left((W_i + S_i) \sum_{j=1}^{i} (S_j - \theta) \right).$$

This SF gradient estimator has approximately half the number of terms and half the variance of the previous one.

In principle, one can also combine the IPA and pure SF approaches. For instance, one can take S_i for some of the customers and the corresponding $U(0, 1)$ values for others. This would give rise to more complex expressions but could be implemented in practice without too much difficulty, and Lemma 1 would still apply. There might be no practical advantage of doing such a combination in that case, but there are other examples where it can be helpful. See, e.g., §4.4. Suppose now that ω contains only the set of waiting times. Then, one faces the problem of expressing P_θ. In fact, the waiting times are dependent random variables whose distributions are quite complex in general. Except for small values of T, this is not practical. Finally, for the extreme case where $\omega = h(\theta, \omega)$, one has to compute the distribution of the average cost (and generate ω) using numerical methods!

4.2. *An $M/G/1$ Queue whose Service Time is Discrete with Parameter-Dependent Probability Mass*

Let $0 < a < b, 0 < c < \theta < d < 1$, and suppose that the service time is b with probability θ, and a with probability $1 - \theta$. In this case, IPA does not apply (see also Suri and Zazanis 1988), but if ω contains the set of actual service times, then SF does apply.

Suppose $U_1, \ldots, U_T$ are the $U(0, 1)$ variates used to generate the service times. Let $C_i = 1$ and $S_i = b$ if $U_i \le \theta$, $C_i = 0$ and $S_i = a$ otherwise (C_i is Bernoulli (θ)). For IPA, ω contains $(U_1, \ldots, U_T)$, and $H(\cdot, \omega) = h(\cdot, \omega)$ is discontinuous since S_i jumps from b to a at $\theta = U_i$. This is why IPA does not work. But suppose $\omega = (A_1, S_1, \ldots, A_T, S_T)$. In this case, the likelihood ratio can be expressed in terms of the variables $C_1, \ldots, C_T$. Their joint probability mass is

$$p_\theta(C_1, \ldots, C_T) = \prod_{i=1}^{T} \theta^{C_i}(1 - \theta)^{1-C_i},$$

and $H(\theta, \omega) = K(\omega)p_\theta(C_1, \ldots, C_T)$, where $K(\omega) = h(\theta, \omega)/p_{\theta_0}(C_1, \ldots, C_T)$ does not depend on θ (because ω contains all the information to compute $h(\theta, \omega)$ independently of θ). $H(\cdot, \omega)$ is continuous and differentiable on $(0, 1)$. Since

$$|\psi(\theta, \omega)| \le (1/T) \sum_{i=1}^{T} (W_i + S_i) \sum_{i=1}^{T} |C_i - \theta|/(\theta(1 - \theta)) \le \sum_{i=1}^{T} (W_i + S_i)/(c(1 - d)),$$

Ψ is bounded by an integrable function, so that A1 is satisfied and we get an unbiased estimator.

4.3. *An $M/G/1$ Queue whose Service Time is Discrete with Parameter-Dependent Support*

Suppose that the service time is θ with probability p, and 2θ with probability $1 - p$, where $\theta > 0$ is the parameter and p is a constant, $0 < p < 1$. Here, the "naive" application of LR, where ω contains the set of actual service times and $G = P_{\theta_0}$, does not apply because there is no neighborhood of θ_0 in which the Radon-Nikodym derivative exists (the event $\{S_i \ne \theta_0, S_i \ne 2\theta_0\}$ has P_{θ_0}-measure 0, but not P_θ-measure 0 for $\theta \ne \theta_0$). Even if one defines L directly as in (6), one gets $H(\theta, \omega) = 0$ everywhere except at $\theta = \theta_0$, which is discontinuous. For IPA, one can use $U_1, \ldots, U_T$ to generate the service times: $S_i = \theta$ if $U_i < p$, $S_i = 2\theta$ otherwise. $\psi(\theta, \omega)$ can be computed as in (14), again with $\partial S_i/\partial\theta = S_i/\theta$. The arguments to verify A1 are the same as in §4.1, and so IPA applies.

4.4. *Applying a Mixture of IPA and Pure SF to an $M/G/1$ Queue*

Let q be a constant, $0 < q < 1$, and suppose that the service time is generated as in §4.2 with probability q, and as in §4.3 with probability $1 - q$. Let $D_1, \ldots, D_T$ be the

corresponding Bernoulli (q) random variables, i.e., $D_i = 1$ if the service time of the ith customer is generated from the first distribution (Example 4.2), $D_i = 0$ otherwise. When $D_i = 1$, let C_i be a Bernoulli (θ) random variable and $S_i = a + C_i(b - a)$, as in §4.2; otherwise let C_i be Bernoulli (p) and $S_i = 2\theta - \theta C_i$, as in §4.3. In that case, we can define ω as the sequence of values of D_i and C_i, plus the $U(0, 1)$ values used to generate the interarrival times. The likelihood ratio depends on ω only through the D_i's and C_i's, whose joint probability mass is

$$p_\theta(D_1, C_1, \ldots, D_T, C_T) = \prod_{i=1}^{T} (q\theta^{C_i}(1-\theta)^{1-C_i})^{D_i}((1-q)p^{C_i}(1-p)^{1-C_i})^{1-D_i},$$

and one obtains

$$S(\theta, \omega) = \sum_{i=1}^{T} D_i\left(\frac{C_i}{\theta} - \frac{1-C_i}{1-\theta}\right) = \sum_{i=1}^{T} \frac{D_i(C_i - \theta)}{\theta(1-\theta)}.$$

This can be computed along with $h(\theta, \omega)$ during the simulation. To compute $\nabla_\theta h(\theta, \omega)$ for a fixed ω, one applies a mixture of the more traditional IPA and SF techniques: the service times of the customers for which $D_i = 0$ are "perturbed" using the usual IPA technique, while the perturbations for the other service times are considered to be zero. More specifically, for a fixed ω, $\nabla_\theta h(\theta, \omega)$ is computed using the right-hand side of (14), but with

$$\frac{\partial S_i}{\partial \theta} = \begin{cases} 0 & \text{if} \quad D_i = 1; \\ 1 & \text{if} \quad D_i = 0 \quad \text{and} \quad S_i = \theta; \\ 2 & \text{if} \quad D_i = 0 \quad \text{and} \quad S_i = 2\theta. \end{cases}$$

Again, assumption A1 is easily verified, by combining the arguments of the two previous examples. Therefore, by mixing IPA with SF, we obtain an unbiased gradient estimate, despite the fact that neither IPA alone nor "pure" SF (putting all the S_i's in ω) works.

Table 1 gives the results of a numerical experiment for this example. We used $T = 10$, $q = p = a = \frac{1}{2}$, $b = \frac{3}{2}$, and estimated the derivative at $\theta = 0.2$, 0.5 and 0.9. We used two gradient estimation techniques: symmetric finite differences with common random numbers (FDC) (see L'Ecuyer, Giroux, and Glynn 1990, Meketon 1987, Suri 1989), and the "hybrid" SF method described above (SF-IPA). For FDC, simulations were made at $\theta \pm 0.01$, starting from the same (empty) state, and the same $U(0, 1)$ values were used on both sides, with proper synchronization. In each case, we made 100000 replications and computed a 95% confidence interval. The same streams of random numbers were used for the six different entries of Table 1. As expected, the results from the two techniques agree. They also agree with the exact values. For FDC, the gradient estimator has some bias, due to the finite differences, but here, that bias is "lost in the noise", since the confidence intervals cover the exact values. The FDC estimates took approximately twice the time of the SF estimates to compute. The Appendix explains how to compute the exact gradient values.

TABLE 1

Numerical Results for Example 4.4: 95% *Confidence Intervals for the Gradient, Based on* 10^5 *Replications*

θ	True Grad.	SF-IPA	FDC
0.2	2.551	2.53 ± .04	2.55 ± .03
0.5	3.980	3.96 ± .06	4.01 ± .04
0.9	5.525	5.53 ± .16	5.53 ± .05

4.5. *An $M/G/1$ Queue whose Service Time is $k\theta$ Times a Bernoulli (θ)*

Suppose that the service time is $k\theta$ with probability θ, and 0 with probability $1 - \theta$, for some constant $k > 0$. For each service time, one can generate a $U(0, 1)$ variate U_i, then define $C_i = 1$ and $S_i = k\theta$ if $U_i < \theta$, $C_i = S_i = 0$ otherwise. In that case, it is easy to see that neither $\omega = (A_1, U_1, \ldots, A_T, U_T)$ nor $\omega = (A_1, S_1, \ldots, A_T, S_T)$ works. However, if one takes $\omega = (A_1, C_1, \ldots, A_T, C_T)$, then SF works. In fact, C_i is Bernoulli (θ). The score function is

$$S(\theta, \omega) = \sum_{i=1}^{T} \frac{(C_i - \theta)}{\theta(1 - \theta)},$$

and $\nabla_\theta h(\theta, \omega)$ can be computed using the right-hand side of (14), with $\partial S_i/\partial\theta = kC_i = S_i/\theta$.

For this example, we performed the same numerical experiment as for the previous example, with $k = 2$, and the results appear in Table 2. Again, they agree very well with the exact values (computed as explained in the Appendix). The FDC estimates took approximately twice the time to compute.

4.6. *A k-out-of-N Reliability System with Parameter-Dependent Component-Lifetime Densities*

Consider a k-out-of-N reliability system with identical components. The N components have independent (random) lifetimes $X_1, \ldots, X_N$, each with distribution F_θ. For simplicity, assume that this distribution is continuous, with density f_θ, that the support of f_θ is independent of θ for $\theta \in \Theta$, and that for each $x > 0$, both $F_\theta(x)$ and $f_\theta(x)$ are differentiable w.r.t. θ. The system is down (failed) when less than k components are still alive. For a given realization ω, $h(\theta, \omega)$ represents the system's lifetime and $\alpha(\theta)$ its expectation. We examine four of the (many) possible choices for ω. We want to estimate $\nabla\alpha(\theta_0)$, using (4). Let s and X_s denote the number and lifetime of the last component that fails (the $(N - k + 1)$th failure). One has $h(\theta, \omega) = X_s$.

One can take $\omega = (X_1, \ldots, X_N)$, in which case $\nabla_\theta h(\theta, \omega) = 0$, and

$$S(\theta, \omega) = \sum_{i=1}^{N} \frac{\partial}{\partial\theta} \ln f_\theta(X_i).$$

This is pure SF in its most straightforward form.

As a second choice, one can define $\omega = (s, X_s, A)$, where $A = \{j \mid X_j > X_s\}$ is the set of $k - 1$ components that are still alive when the system fails. For every $i \in \{1, \ldots, N\}$, $A \subseteq \{1, \ldots, N\} \setminus \{i\}$ such that $|A| = k - 1$, $x \geq 0$ and $\delta > 0$, the event $\{s = i, x < X_s \leq x + \delta, X_j > x + \delta \text{ for } j \in A, X_j \leq x \text{ for } j \notin A \cup \{i\}\}$ has probability

$$P_\theta(x, \delta) = [F_\theta(x + \delta) - F_\theta(x)][F_\theta(x)]^{N-k}[1 - F_\theta(x + \delta)]^{k-1}.$$

The density of ω at $X_s = x$ is therefore

$$\lim_{\delta \to 0^+} P_\theta(x, \delta)/\delta = f_\theta(x)[F_\theta(x)]^{N-k}[1 - F_\theta(x)]^{k-1}$$

TABLE 2

Numerical Results for Example 4.5: 95% *Confidence Intervals for the Gradient, Based on* 10^5 *Replications*

θ	True Grad.	SF	FDC
0.2	1.038	1.03 ± .01	1.03 ± .01
0.5	4.529	4.49 ± .05	4.51 ± .04
0.9	16.671	16.61 ± .30	16.75 ± .12

if $|A| = k - 1$, 0 otherwise. This family of densities is dominated by any one of them (i.e., any $\theta_0 \in \Theta$). We still have $\nabla_\theta h(\theta, \omega) = 0$ and the score function becomes

$$S(\theta, \omega) = \frac{\partial}{\partial\theta}[\ln f_\theta(X_s) + (k-1)\ln(1 - F_\theta(X_s)) + (N-k)\ln F_\theta(X_s)].$$

A third choice uses the IPA technique: suppose $X_i = F_\theta^{-1}(U_i)$ where the U_i's are independent $U(0, 1)$ variates, for $i = 1, \ldots, N$, and take $\omega = (U_1, \ldots, U_N)$. Here, $\nabla_\theta h(\theta, \omega) = \partial F_\theta^{-1}(U_s)/\partial\theta$, where s is defined as above, and $S(\theta, \omega) = 0$. There are many cases where we can obtain equivalent results by taking ω as something else than the U_i's. Consider for example the Weibull distribution: $F(x) = 1 - \exp(-\lambda x^\alpha)$ and $F^{-1}(U) = (-\ln(1-U)/\lambda)^{1/\alpha}$. In this case, one can take $\omega = (V_1, \ldots, V_N)$ where $V_i = -\ln(1 - U_i)$ and this is equivalent to IPA (for $\theta = \lambda$ or $\theta = \alpha$). The expression for $\nabla_\theta h(\theta, \omega)$ can be obtained with slightly less manipulations in the latter case.

For exponential lifetimes, there is the following fourth choice. Let $1/\theta$ be the failure rate for each component, and let the state of the system be defined as the number of components that are still alive. That system evolves as a continuous-time Markov chain. It goes from N to $N-1, \ldots,$ to k, and finally to $k-1$ where it dies. The jump rate from $N - i$ to $N - i - 1$ is $(N-i)/\theta$. Hence, ω can be defined as $\omega = (Y_N, \ldots, Y_k)$, where Y_j is exponential with mean θ/j, and $h(\theta, \omega) = Y_N + \cdots + Y_k$. One has in this case

$$S(\theta, \omega) = \sum_{j=k}^{N}\left(\frac{jY_j - \theta}{\theta^2}\right)$$

and $\nabla_\theta h(\theta, \omega) = 0$.

For the exponential case, analytical formulas are obtained readily:

$$\alpha(\theta) = E[h(\theta, \omega)] = E[Y_N + \cdots + Y_k] = \theta(1/N + \cdots + 1/k),$$

and $\nabla\alpha(\theta) = (1/N + \cdots + 1/k)$. This is in fact discrete-time conversion (Fox and Glynn 1990) applied to the above continuous-time Markov chain. In this case, it yields zero variance. For many other distributions, analytical formulas can also be obtained by exploiting the fact that X_s is an order statistic. Often, one can write down the density (or probability mass) and the expectation of X_s explicitly, and differentiate. This is what we did to compute the exact gradient for the Weibull case in Table 3.

TABLE 3

95% Confidence Intervals for the Gradient, Based on 10^5 Replications (Example 4.6)

N	k	Distrib.	θ	Exact	SF1	SF2	IPA	MC
8	2	Expon.	1.0	1.7179	1.694 ± .042	1.705 ± .036	1.718 ± .004	1.732 ± .040
8	6	Expon.	1.0	0.4346	0.430 ± .011	0.436 ± .009	0.436 ± .001	0.440 ± .009
8	2	Expon.	5.0	1.7179	1.694 ± .042	1.705 ± .036	1.718 ± .004	1.732 ± .040
20	5	Expon.	1.0	1.5145	1.485 ± .048	1.483 ± .041	1.513 ± .002	1.523 ± .045
20	10	Expon.	1.0	0.7688	0.755 ± .025	0.764 ± .020	0.769 ± .001	0.771 ± .020
20	15	Expon.	1.0	0.3462	0.340 ± .012	0.347 ± .008	0.347 ± .001	0.350 ± .008
50	10	Expon.	1.0	1.6703	1.669 ± .024	1.674 ± .020	1.671 ± .001	1.650 ± .022
50	40	Expon.	1.0	0.2457	0.246 ± .004	0.246 ± .002	0.246 ± .001	0.246 ± .002
1	1	Weibull	1.0	−0.4228	−0.404 ± .033	−0.404 ± .033	−0.420 ± .010	
8	2	Weibull	1.0	−1.0756	−1.027 ± .057	−1.052 ± .041	−1.074 ± .008	
8	6	Weibull	1.0	0.2930	0.308 ± .012	0.296 ± .004	0.294 ± .001	
20	5	Weibull	1.0	−0.6843	−0.618 ± .063	−0.669 ± .031	−0.682 ± .004	
20	10	Weibull	1.0	0.1662	0.197 ± .030	0.169 ± .007	0.167 ± .001	
20	15	Weibull	1.0	0.3389	0.350 ± .014	0.342 ± .006	0.340 ± .001	

All this can be adapted to more general reliability networks, with more complex structures, components that have different lifetime distributions (with possibly different parameters), repair possibilities, etc. For the last case (exponential lifetimes and fourth choice of ω), the state description of the Markov chain would get more complicated in general. A "practical" analytical formula is not always available for the exponential case, but replacing transition times by their (conditional) expectations, called discrete-time conversion (Fox and Glynn 1990), always reduces variance without increasing work. See also §4.8. In practice, importance sampling (when it works) can also be very effective when simulating such networks (see Goyal et al. 1989).

We performed some numerical experiments, first with the exponential distribution with mean θ: $F_\theta(x) = 1 - \exp(-x/\theta)$, then with a Weibull distribution with θ as the shape parameter: $F_\theta(x) = 1 - \exp(-x^\theta)$, $x \geq 0$. We tried different values of N and k. The results appear in Table 3, in the form of 95% confidence intervals on $\nabla\alpha(\theta)$. The "methods" SF1, SF2, IPA, and MC correspond to the four choices of ω described above, in that order. "Exact" refers to the exact value of the gradient. Note that MC does not apply for the Weibull distribution, but for the other cases, A1 can be verified easily for all $\theta_0 > 0$. We leave that as an (easy) exercise for the reader. Expressions used in the estimators are given in Table 4 for these two distributions. For the Weibull case, at $\theta = \theta_0 = 1$, the exact value is given by

$$\frac{\partial}{\partial\theta} E[X_s] = \frac{N!}{(k-1)!} \sum_{i=0}^{N-k} \frac{(-1)^i(\ln(k+i) + \gamma - 1)}{i!(N-k-i)!(k+i)^2}$$

where $\gamma = 0.5772157$ is the Euler's constant. In our implementations, the four estimators (four methods) were computed simultaneously within the same simulations.

From these numerical results, SF2 appears to be generally better than SF1, and MC is not much better than even SF1. The most efficient by far is certainly IPA. This is not surprising, since the score function has zero variance in that case. In fact, in most cases where IPA works in practice, it is typically the most efficient gradient estimation technique.

4.7. *A Reliability System with Parameter-Dependent Support for the Component-Lifetime Densities*

In the previous example, suppose that the component lifetimes are uniformly distributed, between 0 and θ. The density of X_i is $f_\theta(x) = 1/\theta$ for $0 \leq x \leq \theta$. Since the support of f_θ depends on θ, if we take $\omega = (X_1, \ldots, X_N)$, the Radon-Nikodym derivative dP_θ/dP_{θ_*} does not exist for $\theta > \theta_*$, but exists for $\theta \leq \theta_*$. For (1) to be valid, one can take $G = P_{\theta_*}$ for some (fixed) θ_* that is larger than all values of θ for which we might

TABLE 4

Expressions Used in the Gradient Estimators, for Two Distributions

	Exponential	Weibull
$f_\theta(x)$	$(1/\theta)e^{-x/\theta}, \quad x > 0$	$\lambda\theta x^{\theta-1}e^{-\lambda x^\theta}, \quad x > 0$
$F_\theta(x)$	$1 - e^{-x/\theta}$	$1 - e^{-\lambda x^\theta}$
$F_\theta^{-1}(U)$	$-\theta \ln(1-U)$	$(-\ln(1-U)/\lambda)^{1/\theta}$
$\frac{\partial}{\partial\theta} X$ for $X = F_\theta^{-1}(U)$	X/θ	$-X(\ln X)/\theta$
$\frac{\partial}{\partial\theta} \ln f_\theta(x)$	$(x-\theta)/\theta^2$	$1/\theta + (1 - \lambda x^\theta)\ln x$
$\frac{\partial}{\partial\theta} \ln F_\theta(x)$	$\frac{-xe^{-x/\theta}}{\theta^2(1-e^{-x/\theta})}$	$\frac{\lambda x^\theta (\ln x) e^{-\lambda x^\theta}}{1 - e^{-\lambda x^\theta}}$
$\frac{\partial}{\partial\theta} \ln(1 - F_\theta(x))$	x/θ^2	$-\lambda x^\theta \ln x$

want to estimate $\nabla\alpha(\theta)$. But then, a different problem emerges. Since $f_\theta(x) = 1/\theta$ for $x < \theta$ and 0 for $x > \theta$, for any fixed $\omega = (X_1, \ldots, X_N)$,

$$H(\theta, \omega) = X_s \prod_{i=1}^{N} (f_\theta(X_i)/f_{\theta_*}(X_i))$$

is discontinuous in θ at $\theta = X_i$, for each i. In fact, if $G = P_{\theta_*}$ and Υ is a neighborhood of θ_*, $\max_i X_i$ falls into Υ with positive probability, and whenever this happens, $H(\cdot, \omega)$ is not continuous in Υ. Therefore, A1 is not satisfied. This illustrates the fact that even when $\nabla_\theta h(\theta, \omega) = 0$, the existence of the Radon-Nikodym derivative is not a sufficient condition for LR to apply.

Note that IPA applies for this case: the gradient estimator is X_s/θ, the same as for the exponential case.

4.8. *Continuous-Time Markov Chains*

Consider a continuous-time Markov chain with finite state space S. Let λ_i denote the jump rate out of state i, and p_{ij} be the transition probability from i to j. There is also a cost incurred continuously at rate c_i when in state i. Suppose that these quantities depend on some parameter vector $\theta \in \Theta \subseteq \mathbb{R}^d$. Let $\alpha(\theta)$ be the total expected cost for the first T transitions, where T is fixed.

Simulation is often the most convenient tool to analyze such chains, particularly for very large state spaces (see, e.g., Glynn and Iglehart 1988; Fox 1989). Here, a future-event list is not necessary to run the simulation; one can just use the transition probabilities to jump from state to state. Typically, the transition matrix is very sparse, and from any given state i, the number of reachable states is small. In fact, there is usually no need to write down that matrix or to enumerate S. Take for instance a closed Jackson network with say m nodes and n customers (one server per node, one class of customers): if $m = 20$ (say) and $n = 100$ (say), the state space is huge, but from any given state with say B busy nodes, there are only B possibilities for the node where the departure occurs and at most 20 possibilities for the destination node of the departing customer. It is quite easy to generate the two corresponding discrete variables and there is no need to generate explicitly even a row of the transition matrix. As discussed in Fox (1990), the second variable (destination node) can be generated in $O(1)$ time via the alias method (see Bratley, Fox, and Schrage 1987), after some initial setup time to compute an alias table for each row of the routing matrix. Using the alias method for the first variable (departing node) is much more complicated, since its probability law varies with the set of nodes that are busy. Of course, this variable can be generated directly, via "inversion", or indirectly, via a future-event schedule. If B is typically modest (say at most 20), then either method is adequate. However, Fox (1990) proposes a clever method that beats both of these when m is large and the traffic is heavy. His method exploits efficiently the similarity in successive, unnormalized departure-rate vectors, and uses the alias method as a subroutine. It also applies to simulations of highly-dependably reliability systems (§4.6).

Let $\{X_n, n \geq 0\}$ be the embedded Markov chain (the sequence of visited states), $0 = \tau_0 \leq \tau_1 \leq \tau_2 \leq \cdots$ the transition times (the system jumps into X_n at time τ_n), and for $n \geq 0$, $\zeta_n = \tau_{n+1} - \tau_n$. Note that ζ_n is exponential with mean $1/\lambda_{X_n}$.

A simple choice for ω, when A1 is satisfied for it, is

$$\omega = (X_0, \zeta_0, X_1, \zeta_1, \ldots, X_{T-1}, \zeta_{T-1}).$$

Except for very small T, the variance of the gradient estimator is then usually quite high and, as the previous examples suggest, one would usually prefer to use IPA if it is applicable. Unfortunately, IPA does not always work for the transition probabilities (see Glasserman

1989), but it can be used here for the times between jumps: take

$$\omega = (X_0, U_0, \ldots, X_{T-1}, U_{T-1}),$$

where U_i is the $U(0, 1)$ variate used to generate ζ_i (by inversion). This could certainly help, but there is a better choice: just take $\omega = (X_0, \ldots, X_{T-1})$. The ζ_i's can simply be replaced by their conditional expectations. This is discrete-time conversion (Fox and Glynn 1990). It reduces simultaneously the variance of the cost estimate and the variance of the likelihood ratio. In fact, there is no need to generate any ζ_i. The cost estimate is simply.

$$\sum_{n=0}^{T-1} c_{X_n}/\lambda_{X_n}. \tag{15}$$

If all the p_{ij}'s depend on θ, there are still $T - 1$ terms in the score function (assuming X_0 fixed) and the variance could kill us for large T. There might be cases, however, where only *some* of the p_{ij}'s depend on θ and this can make a big difference in the variance of the gradient estimator. If only the transition times (and not the transition probabilities) are influenced by θ, then the gradient estimator is readily obtained by differentiating (15).

On the other hand, terms in the likelihood ratios cannot be replaced (in general) by their expectations. For example, the score function in (4) has zero expectation, but is correlated with $h(\theta, \omega)$ and the expectation of the product is not zero. There are special cases, however, where it can be done: see, e.g., Glasserman (1989) and Algorithm B in Glynn (1987).

With adaptations, the above approach also applies to semi-Markov processes and (though this is not necessarily trivial) to the case where T is a random stopping time. Glasserman (1989) has developed a technique to apply IPA to a continuous-time Markov chains when the transition probabilities depend on θ. It is based on a clever definition of $h(\theta, \omega)$, i.e., a clever way of transforming the $U(0, 1)$ variates into the performance measure. Straightforward (naïve) definitions typically do not work. Glasserman's method works under certain conditions, the most important being that any two states directly accessible from a given common state also have a common state directly accessible from them, and that the objective function is made up of continuously cumulated costs.

4.9. *A Multi-Component Replacement System*

(Example taken from Haurie and L'Ecuyer 1986.) Consider a system comprised of N identical components, that evolve independently. Each component has a random lifetime distribution with increasing failure rate. Whenever a component fails, it must be replaced instantly by a new one. Other components may be replaced (preventively) on the same occasion. The repairman can also halt the system at any moment and replace preventively any number of working components. All replacements are assumed instantaneous. A failure cost c_f is incurred every time a component fails. At each intervention, there is also a fixed cost c_i, and a replacement cost which is c_r times the number of components replaced. Preventive replacements are made to avoid some of the failures, and replacements are sometimes lumped together to pay the fixed cost less often.

Here, we restrict out attention to the (generally suboptimal) class of policies defined by two thresholds: $\theta_1 > \theta_2 > 0$. Whenever a component fails or reaches age θ_1, the repairman intervenes and replaces all components older than θ_2. We are interested in the total cost for a fixed duration T, assuming that all components are new at the beginning. The parameter here is $\theta = (\theta_1, \theta_2)$.

Suppose ω is the set of generated component lifetimes (for fixed θ, this is enough for computing the cost). Then, the likelihood ratio is always one, since the component

lifetimes do not depend on θ, but $h(\theta, \omega)$ is discontinuous in θ. More specifically, for any neighborhood Υ of $\theta = (\theta_1, \theta_2)$, $h(\cdot, \omega)$ will be discontinuous in Υ if some component lifetimes are near enough to θ_1 or θ_2 to change the sequence of failures when θ changes inside Υ (for fixed ω), and the set of values of ω for which this happens has positive probability. Exactly the same problem occurs with IPA (ω is the set of $U(0, 1)$ values used to generate the component lifetimes). Suppose now that ω includes the sequence of all failures and replacements, with their times. Now, $h(\theta, \omega)$ becomes independent of θ, but in general, for any fixed value θ_0 of θ and any G, the Radon-Nikodym derivative dP_θ/dG does not exist in a neighborhood of θ_0, since a typical ω will have nonzero P_θ-probability (or density) for only one value of θ. More precisely, let Υ be a neighborhood of θ_0. For any $\theta = (\theta_1, \theta_2) \in \Upsilon$, let B_θ be the event $\{\omega \mid$ at every preventive replacement, the age of the oldest component is $\theta_1\}$. Note that $P_\theta(B_\theta) = 1$, while for any $\theta' = (\theta'_1, \theta'_2) \in \Upsilon$ such that $\theta'_1 \neq \theta_1$, one has $B_\theta \cap B_{\theta'} = \varnothing$ and $P_{\theta'}(B_\theta) = 0$. For dP_θ/dG to exist in Υ, it is necessary that $G(B_\theta) > 0$ for all $\theta \in \Upsilon$. But since Υ is an interval, this is impossible.

At this point, we do not know how to apply LR (and in particular IPA) to estimate the gradient for this example. Many other examples, most of them involving "threshold" parameters, fall into this category. For instance, think of a (s, S) inventory systems, where $\theta = (s, S)$, or a time-sharing computer system where the parameter is the quantum size (see Kleinrock 1976, p. 160), or a checkpoint-rollback-recovery system (for databases; see L'Ecuyer and Malenfant 1988) where θ is the time between checkpoints (or is used in a rule to decide the next checkpoint time, based on the state of the system), etc. At present, for all these examples, to the best of our knowledge, a "finite-differences" approach (preferably with common random numbers) must be used.

5. Conclusion

We pointed out the strong relationship that exists between IPA, SF, and LR. In §3, we gave easily verifiable conditions under which LR (and/or its special case IPA) does apply. When LR does not apply, these conditions sometimes permit us to understand why. We have illustrated with examples some ideas related to this approach. In particular, there are often many different ways to implement LR, some being much more efficient than others. IPA and "pure" SF (i.e., as used for instance in Glynn 1987; Reiman and Weiss 1989; Rubinstein 1986 and 1989) can sometimes be combined on the same problem, and for the same parameter. However, when IPA applies, it is typically by far the most efficient method.

In practice, the change of measure used to define (1) through (3) can sometimes be used to reduce the variance (this is importance sampling; see Goyal et al. 1989). We have not explored that issue in this paper. For all the examples in §4, we have used $G = P_{\theta_0}$ to estimate $\nabla\alpha(\theta_0)$, but substantial variance reductions can sometimes be obtained by using a different (and appropriate) G.

For some examples, it appears that the method of finite differences (FD) remains the only applicable approach at this time. Some experimental evidence (L'Ecuyer, Giroux, and Glynn 1990, and §4 in this paper) suggests that in many cases where IPA applies, FD with common random numbers might be practically as good as IPA when θ has only one component ($d = 1$). But for a large number of parameters (large value of d), performing all the simulations required for FD becomes rather time consuming, and IPA or a good LR implementation might beat FD significantly.[1]

[1] Most of this work was done while the author was enjoying the hospitality of the Operations Research Department at Stanford University. It has been supported by NSERC-Canada grant #A5463 and FCAR-Québec grant #EQ2831. I wish to thank B. L. Fox, P. W. Glynn, M. Nakayama, A. Haurie, R. Rubinstein, and R. Suri for helpful discussions and comments. Many suggestions and corrections by the Department Editor James R. Wilson have led to important improvements of the paper.

Appendix: Computing Exact Gradient Values for §4.1 to §4.5

To compute the exact gradient values for §4.1–§4.5, we can write recursive equations to compute $V_n(s)$ and $\nabla_\theta V_n(s)$ in terms of V_{n+1} and $\nabla_\theta V_{n+1}$, where $V_n(s)$ represents the expected total system time spent from now on by the next $T - n$ customers to depart, given that there are s customers in the system, $1 \le s \le T - n$, one of which is beginning its service. $V_{T-1}(s)$ is the expected service time (independently of s), and $V_0(1)/T$ is the expected average system time for the first T customers. $V_T(s) = \nabla_\theta V_T(s) = 0$ for all s, and $V_n(0) = V_n(1)$ for $0 \le n < T$.

We now derive these recursive equations. Assume that the arrival rate is λ and that the service time is a discrete random variable, with probability mass p_θ, over a denumerable sample space $\Omega' \subset [0, \infty)$. Let ζ denote the next service time duration, and i the number of arrivals during that service time (which is the value of a Poisson $(\lambda\zeta)$ random variable). For $1 \le s \le T - n$, one has

$$V_n(s) = \sum_{\zeta \in \Omega'} \varphi_0(\zeta) p_\theta(\zeta),$$

where

$$\varphi_0(\zeta) = \zeta s + \sum_{i=0}^{\infty} \frac{e^{-\lambda\zeta}(\lambda\zeta)^i}{i!} (\varphi_1(n, s, i) + \zeta\varphi_2(n, s, i)).$$

Here, ζs represents the waiting time during the next transition (of duration ζ) for the customers already in the system, $\zeta\varphi_2(n, s, i)$ is the expected waiting time during that same period for the min $(i, T - n - s)$ customers who arrive during it and that are taken (we stop taking customers when a total of T have arrived), and $\varphi_1(n, s, i)$ is the total expected waiting time from the next transition on.

If $i > T - n - s$, only the first $T - n - s$ customers are taken. In that case,

$$\zeta\varphi_2(n, s, i) = E[r\zeta - Y_1 - \cdots - Y_r]$$

where $r = T - n - s$ and $Y_1, \ldots, Y_i$ are the order statistics of i i.i.d. $U(0, \zeta)$ variates. Observe that Y_j/ζ has the beta distribution with shape parameters $\alpha = j$ and $\beta = i - j + 1$. Therefore, $E[Y_j/\zeta] = j/(i + 1)$, so that

$$\zeta\varphi_2(n, s, i) = r\zeta - \sum_{j=1}^{r} \zeta j/(i + 1) = r\zeta - \zeta r(r + 1)/(2(i + 1)).$$

One then has

$$\varphi_1(n, s, i) = \begin{cases} V_{n+1}(s + i - 1) & \text{if} \quad i \le T - n - s; \\ V_{n+1}(T - n - 1) & \text{otherwise;} \end{cases}$$

$$\varphi_2(n, s, i) = \begin{cases} i/2 & \text{if} \quad i \le T - n - s; \\ (T - n - s)(1 - (T - n - s + 1)/(2(i + 1))) & \text{otherwise.} \end{cases}$$

Also,

$$\nabla_\theta V_n(s) = \sum_{\zeta \in \Omega'} [\nabla_\theta p_\theta(\zeta)\varphi_0(\zeta) + p_\theta(\zeta)\varphi_3(\zeta)],$$

where

$$\varphi_3(\zeta) = s\frac{\partial\zeta}{\partial\theta} + \sum_{i=0}^{\infty} \frac{e^{-\lambda\zeta}(\lambda\zeta)^i}{i!} [\nabla_\theta\varphi_1(n, s, i) + \frac{\partial\zeta}{\partial\theta}(\varphi_2(n, s, i) + (i/\zeta - \lambda)[\varphi_1(n, s, i) + \zeta\varphi_2(n, s, i)])].$$

Typically, the (inner) infinite sum converges quickly and can be approximated by taking a finite number of terms. For our examples, a better precision than necessary was obtained by taking the first 30 terms.

TABLE 5

Expressions to Compute $\nabla_\theta V_n(s)$ in Example 4.4

ζ	$\partial\zeta/\partial\theta$	$p_\theta(\zeta)$	$\partial p_\theta(\zeta)/\partial\theta$
b	0	$q\theta$	q
a	0	$q(1 - \theta)$	$-q$
θ	1	$(1 - q)p$	0
2θ	2	$(1 - q)(1 - p)$	0

TABLE 6
Expressions to Compute $\nabla_\theta V_n(s)$ in Example 4.5

ζ	$\partial\zeta/\partial\theta$	$p_\theta(\zeta)$	$\partial p_\theta(\zeta)/\partial\theta$
$k\theta$	k	θ	1
0	0	$1-\theta$	-1

For §4.4, $\lambda = 1$ and the set of possible values for ζ is $\Omega' = \{b, a, \theta, 2\theta\}$. One gets the expressions given in Table 5. The corresponding expressions for §4.5 are given in Table 6.

References

ALEKSANDROV, V. M., V. I. SYSOYEV AND V. V. SHEMENEVA, "Stochastic Optimization," *Engineering Cybernetics*, 5 (1968), 11–16.

ARSHAM, H., A. FEUERVERGER, D. L. MCLEISH, J. KREIMER AND R. Y. RUBINSTEIN, "Sensitivity Analysis and the "What If" Problem in Simulation Analysis," *Math. Comput. Modelling*, 12, 2 (1989), 193–219.

BILLINGSLEY, P., *Probability and Measure* (2nd ed.), Wiley, New York, 1986.

BRATLEY, P., B. L. FOX AND L. E. SCHRAGE, *A Guide to Simulation* (2nd ed.), Springer-Verlag, New York, 1987.

CAO, X. R., "Sensitivity Estimates Based on One Realization of a Stochastic System," *J. Stat. Comput. Simul.*, 27 (1987), 211–232.

DIEUDONNÉ, J., *Foundations of Modern Analysis* (2nd ed.), Academic Press, New York, 1969.

FOX, B. L., "Computing Expected Cumulative Reward up to Absorption and its Gradient," Technical Report, Math. Dept., Univ. of Colorado, 1989.

———, "Generating Markov-Chain Transitions Quickly: I," to appear in *ORSA J. Computing* (1990).

——— AND P. W. GLYNN, "Discrete-Time Conversion for Simulating Finite-Horizon Markov Processes," to appear in *SIAM J. Appl. Math.* (1990).

GLASSERMAN, P., "Performance Continuity and Differentiability in Monte Carlo Optimization," Proc. Winter Simulation Conf. 1988, Society for Computer Simulation, 1988a, 518–524.

———, "Structural Conditions for Perturbation Analysis Derivative Estimation: Finite Time Performance Indices," Manuscript, Harvard Univ., 1988b.

———, "Structural Conditions for Perturbation Analysis of Queueing Systems," Manuscript, Harvard Univ., 1988c.

———, "Derivative Estimates from Simulation of Continuous Time Markov Chains," Technical report, AT&T Bell Laboratories, Holmdel, New Jersey, 1989.

GLYNN, P. W., "Likelihood Ratio Gradient Estimation: an Overview," Proc. Winter Simulation Conf. 1987, Society for Computer Simulation, 1987, 366–375.

——— AND D. L. IGLEHART, "Simulation Methods for Queues: an Overview," *Queueing Systems*, 3 (1988), 221–255.

GONG, W. B. AND Y.-C. HO, "Smoothed (Conditional) Perturbation Analysis of Discrete Event Dynamical Systems," *IEEE Trans. on Automatic Control*, AC-32, 10 (1987), 858–866.

GOYAL, A., P. SHAHABUDDIN, P. HEIDELBERGER, V. F. NICOLA AND P. W. GLYNN, "A Unified Framework for Simulating Markovian Models of Highly Dependable Systems," Technical report RC 14772, IBM Research Division, Yorktown Heights, 1989.

HAURIE, A. AND P. L'ECUYER, "Approximation and Bounds in Discrete Event Dynamic Programming," *IEEE Trans. on Automatic Control*, AC-31, 3 (1986), 227–235.

HEIDELBERGER, P., X.-R. CAO, M. A. ZAZANIS AND R. SURI, "Convergence Properties of Infinitesimal Perturbation Analysis Estimates," *Management Sci.*, 34, 11 (1989), 1281–1302.

HO, Y.-C., "Performance Evaluation and Perturbation Analysis of Discrete Event Dynamic Systems," *IEEE Trans. on Automatic Control*, AC-32, 7 (1987), 563–572.

KLEINROCK, L., *Queueing Systems*, Computer Applications, Vol. 2, Wiley, New York, 1976.

L'ECUYER, P., N. GIROUX AND P. W. GLYNN, "Stochastic Optimization by Simulation: Convergence Proofs and Experimental Results for the GI/G/1 Queue in Steady-State," submitted for publication (1990).

——— AND P. W. GLYNN, "A Control Variate Scheme for Likelihood Ratio Gradient Estimation," in preparation (1990).

——— AND J. MALENFANT, "Computing Optimal Checkpointing Strategies for Rollback and Recovery Systems," *IEEE Trans. on Computers*, 37, 4 (1988), 491–496.

MEKETON, M. S., "Optimization in Simulation: A Survey of Recents Results," Proc. Winter Simulation Conf. 1987, Society for Computer Simulation, 1987, 58–67.

REIMAN, M. I. AND A. WEISS, "Sensitivity Analysis for Simulation via Likelihood Ratios," *Oper. Res.*, 37, 5 (1989), 830–844.

RUBINSTEIN, R. Y.,"The Score Function Approach for Sensitivity Analysis of Computer Simulation Models," *Math. and Computers in Simulation*, 28 (1986), 351–379.

———, "Sensitivity Analysis and Performance Extrapolation for Computer Simulation Models," *Oper. Res.*, 37, 1 (1989), 72–81.

———, "Monte Carlo Methods for Performance Evaluation, Sensitivity Analysis and Optimization of Stochastic Systems," to appear in *The Encyclopedia of Computer Science and Technology* (Kent and Williams, Eds.), 1990.

SURI, R., "Infinitesimal Perturbation Analysis for General Discrete Event Systems," *J. ACM*, 34, 3 (1987), 686–717.

———, "Perturbation Analysis: The State of the Art and Research Issues Explained via the $GI/G/1$ Queue," *Proc. IEEE*, 77, 1 (1989), 114–137.

——— AND M. ZAZANIS, "Perturbation Analysis Gives Strongly Consistent Estimates for the $M/G/1$ Queue," *Management Sci.*, 34, 1 (1988), 39–64.