Discrete Event Dynamic Systems: Theory and Applications, *, 1–24 (Draft: April 18, 1996)

Combining the Stochastic Counterpart and Stochastic Approximation Methods

JEAN-PIERRE DUSSAULT dussault@dmi.usherb.ca
Département de mathématiques et informatique, Université de Sherbrooke, Sherbrooke J1K 2R1, Canada

DONALD LABRECQUE AND PIERRE L'ECUYER lecuyer@iro.umontreal.ca
Département d'IRO, Université de Montréal, C.P. 6128, Succ. Centre-Ville, Montréal H3C 3J7, Canada

REUVEN Y. RUBINSTEIN ierrr01@ie.technion.ac.il
Faculty of Industrial Engineering and Management Technion—Israel Institute of Technology, Haifa 32000, Israel and Department of Mathematics, EPFL-Ecublens CH-1015 Lausanne, Switzerland

Abstract. In this work, we examine how to combine the score function method with the standard crude Monte Carlo and experimental design approaches, in order to evaluate the expected performance of a discrete event system and its associated gradient *simultaneously* for different scenarios (combinations of parameter values), as well as to optimize the expected performance with respect to two parameter sets, which represent parameters of the underlying probability law (for the system's evolution) and parameters of the sample performance measure, respectively. We explore how the stochastic approximation and stochastic counterpart methods can be combined to perform optimization with respect to both sets of parameters at the same time. We outline three combined algorithms of that form, one sequential and two parallel, and give a convergence proof for one of them. We discuss a number of issues related to the implementation and convergence of those algorithms, introduce averaging variants, and give numerical illustrations.

Keywords: Score function, sensitivity analysis, optimization, stochastic counterpart, stochastic approximation.

1. Introduction

Let

$$\ell(v,\theta) = E_v\{L(Y,\theta)\} \tag{1}$$

be the expected performance of a discrete event system (DES), where L is the sample performance driven by an input vector Y, which has a probability density function (pdf) $f(y,v)$. In (1), f and L depend on the vectors of parameters v and θ, respectively, and the subscript v in E_v means that the expectation is taken with respect to $f(\cdot,v)$. In other words, v is a parameter of the probability law, while θ is a parameter of the sample performance. We assume that $\ell(v,\theta)$ is not available analytically and that we need to resort to Monte Carlo simulation methods for its estimation. We are concerned with the following questions:

(i) Solve the so-called "what-if" problem; that is, to estimate $\ell(v,\theta)$ and its gradient w.r.t. v and θ, $\nabla_v \ell(v,\theta)$ and $\nabla_\theta \ell(v,\theta)$, in functional form, or simultaneously for different values of v and θ;

(ii) Combine the Crude Monte Carlo (CMC) and the score function (SF) methods, to deal with parameters θ and v, respectively;

(iii) Solve an optimization problem associated with $\ell(v,\theta)$, where both v and θ are decision parameters.

As a motivating example, consider a queueing network containing $GI/D/1$ or $GI/G/c/m$ queues, where c and m denote the number of parallel servers and the buffer size, respectively. In the first case, v might be a parameter (vector) of the interarrival pdf $f(y,v)$ and θ might be the vector of (fixed) service times, while in the second case, v might be the vector of the interarrival and service rates in the joint pdf $f(y,v)$ and θ might be the vector of the buffer size and the number of parallel servers, respectively.

In its *original form*, the likelihood ratio (LR) or score function (SF) method [10], [28], [29], [30], [31], [32] permits the solution of the "what-if" problem from a *single simulation run* (single sample path) with respect to v *alone*, that is, when θ is *fixed*. Roughly, the use of the score function transforms an estimator of $\ell(v,\theta)$ into an estimator of $\nabla_v \ell(v,\theta)$, whereas the likelihood ratio transforms point estimators into functional estimators, thereby allowing the estimation of the entire functions $\ell(\cdot,\theta)$ and $\nabla_v \ell(\cdot,\theta)$ from a single simulation run, for any given value of θ. The latter "likelihood ratio" or "change of measure" technique is in fact exactly the same as that used in *importance sampling* for variance reduction [9]. Unlike SF, the CMC method permits the solution of the "what-if" problem with respect to both v and θ, simply by performing separate simulations at each parameter values of interest. Since it requires *multiple runs* (at least a separate run for each point (v,θ)), it is typically time-consuming. Note that a modification of the SF method, the so-called "*push out*" method [32], as well as the perturbation analysis (PA) method [8], also called "*push in*" method in [32], combined with the use of a likelihood ratio, permit (in many cases) the solution of the "what-if" problem *simultaneously from a single simulation* with respect to both v and θ; see [19] for examples. Here, we shall not deal with the latter approaches. We should mention that the SF method sometimes suffers from a variance explosion problem (the variance of the estimator may become huge at some values), especially when the values of v of interest span a large area (see [19], [32] for details). But there are ways of dealing with such problems (e.g., break the area of interest into smaller subareas), at least for certain classes of applications [32].

Suppose now that we want to minimize $\ell(v,\theta)$ with respect to v and θ. One approach for minimizing $\ell(v,\theta)$ w.r.t. v for fixed θ is to compute an estimator of $\ell(\cdot,\theta)$ in functional form, using a likelihood ratio, based on N replicates of the simulation, and then minimize the (sample) value of that estimator w.r.t. v using conventional mathematical programming tools. The latter minimization problem is called the *stochastic counterpart* (SC). That SC optimization approach is studied in much detail in Rubinstein and Shapiro [32], where it is shown that the sample optimizer converges to the true optimizer with probability one (w.p. 1), and obeys a central-limit theorem, as the sample size N goes to ∞. If the number of values of θ of interest is finite and not too large, then the SC approach can be applied at each such value, and one may select the value of θ which gave the best result. A statistical analysis of such an approach can be performed along the lines of the statistical *ranking and selection* and *multiple comparison* methods [12], [16], [37]. In particular, those values of θ of interest may have been chosen among a much larger (perhaps infinite) set by some *experimental design* (ED) strategy [14].

Suppose now that both parameters are continuous. To optimize w.r.t. θ for fixed v, one can use a stochastic approximation (SA) algorithm (see, e.g., [6], [7], [15], [21], [24], [25], [27], [30] and several further references given there). SA is an iterative procedure which at each step estimates the gradient of the objective function and makes a small step in its opposite direction. The gradient estimator can be based on either SF, PA, finite differences, and so on, and the speed of convergence depends highly on the quality of the gradient estimator that is used (for example, is can be quite slow when using the so-called Kiefer-Wolfowitz SA variant, based on finite differences with independent random numbers). In fact, the SA algorithm can be used as well for optimization w.r.t. v, or w.r.t. both parameters simultaneously. However, it could be in some cases more efficient or convenient to use the SC approach rather than SA for dealing with v. This leads to the following question: Can we design combined (or hybrid) algorithms which use SC for v and SA for θ, while performing optimization w.r.t. both parameters simultaneously ?

The rest of this work is organized as follows. In Section 2, we show how to combine the SF and CMC/ED methods in order to estimate $\ell(v,\theta)$, $\nabla_v\ell(v,\theta)$, and $\nabla_\theta\ell(v,\theta)$ simultaneously for several scenarios (combinations) of (v,θ). Section 3 deals with the minimization of $\ell(v,\theta)$ with respect to both v and θ. We outline three optimization algorithms, each coming in two versions, and provide a convergence proof for the first version of the first algorithm. A numerical example then gives some insight into the behavior of the algorithms and also illustrates some potential difficulties.

2. The "What-if" Problem

In this section, we recall some background material on the SF method and on how a functional estimator w.r.t. v can be obtained. Further details on this are given in [1], [2], [3], [9], [10], [17], [28], [31], [32]. We then explain how to combine SF and CMC in order to estimate $\ell(v,\theta)$ and its gradient *simultaneously* for several values of v and θ. We distinguish the following two cases: (a) θ is *fixed*; (b) θ is *not fixed.*

2.1. The Case of a Fixed θ: a "What-if" Design with Respect to v

Assume that θ is fixed and let $Y_1, Y_2, \ldots,$ be an input sequence of independent identically distributed (iid) random vectors, generated from a density $f(\cdot, v)$, which depends on the parameter vector v. Let $\{L_t : t > 0\}$ be a discrete-time output process driven by $\{Y_t\}$, that is, $L_t = L_t(Y_1, \ldots, Y_t, \theta)$. Assume that $\{L_t\}$ is regenerative with cycle length $\tau = \tau(Y_1, Y_2, \ldots, \theta)$. It is well known [36] that the expected steady-state (average) of $\{L_t\}$ can be written as

$$\ell(v,\theta) = \frac{E_v\left[\sum_{t=1}^{\tau} L_t\right]}{E_v[\tau]} = \frac{\ell_1(v,\theta)}{\ell_2(v,\theta)}, \tag{2}$$

provided that $E_v[\tau] > 0$ and $E_v\left[\left|\sum_{t=1}^{\tau} L_t\right|\right] < \infty$, and similarly when L_t is a continuous-time regenerative process. A finite-horizon model can be viewed as a special case of this: just replace $\ell_2(v,\theta)$ by 1.

Under standard regularity conditions allowing interchangeability of expectation and differentiation (e.g., uniform integrability), one has [1], [9], [17], [32]:

$$\ell_1(v,\theta) = E_g\left[\sum_{t=1}^{\tau} L_t W_t\right]; \tag{3}$$

$$\nabla_v \ell_1(v,\theta) = E_g\left[\sum_{t=1}^{\tau} L_t \nabla_v W_t\right], \tag{4}$$

where $L_t = L_t(Z_1, \ldots, Z_t, \theta)$, $W_t = \prod_{j=1}^{t} f(Z_j, v)/g(Z_j)$, $(Z_1, \ldots, Z_t)$ has density $\prod_{j=1}^{t} g(z_j)$, $g(\cdot)$ is a density that dominates all the $f(\cdot, v)$'s in the sense that $f(z,v) > 0$ for some v implies $g(z) > 0$, and E_g denotes the expectation with respect to g. To obtain similar expressions for ℓ_2, just replace L_t by 1. We call W_t, ∇W_t, $L_t W_t$, and $L_t \nabla W_t$ the *likelihood ratio*, *score function*, *sample performance*, and *sensitivity* processes, respectively. This could also be generalized to larger values of k [17], [32].

If we further assume that $\nabla_\theta L_t(\theta)$ is available from the simulation, that $\nabla_\theta \tau = 0$ (which is typical, since τ is usually piecewise constant as a function of θ), and under a few additional conditions (see [8], [13]), then one also has

$$\nabla_\theta \ell_1(v,\theta) - \ell_1(v,\theta)\nabla_\theta \ell_2(v,\theta) = E_g\left[\sum_{t=1}^{\tau} W_t \nabla_\theta L_t\right]. \tag{5}$$

(Note that the latter bracketted expression is typically *not* an unbiased estimator of $\ell_1(v,\theta)$ when τ depends on θ.) When these conditions are not satisfied, one can still rely to finite differences to estimate the gradients on the left-hand-side of (5), preferably with common random numbers (see [20]).

Remark 1 In this setup, we have implicitly assumed that both v and θ are continous parameters and that the derivatives exist. In the case where either v or θ is discrete, then we just forget about the corresponding derivatives.

Consider now a sample of N iid regenerative cycles, with values τ_i, L_{ti}, and W_{ti} of τ, L_t, and W_t, respectively, for $t \geq 1$ and $1 \leq i \leq N$, again based on the underlying distribution g. Then, under the conditions mentioned above, unbiased estimators of $\ell_1(v,\theta)$ and $\nabla_v \ell_1(v,\theta)$ are given by

$$\begin{aligned}
\overline{\ell}_{1N}(v,\theta) &= \frac{1}{N}\sum_{i=1}^{N}\sum_{t=1}^{\tau_i} L_{ti}W_{ti}, \\
\nabla_v \overline{\ell}_{1N}(v,\theta) &= \frac{1}{N}\sum_{i=1}^{N}\sum_{t=1}^{\tau_i} L_{ti}\nabla_v W_{ti},
\end{aligned}$$

and similarly for ℓ_2 with L_{ti} replaced by 1. Consistent estimators of $\ell(v,\theta) = \ell_1(v,\theta)/\ell_2(v,\theta)$ and $\nabla_v \ell(v,\theta) = (\nabla_v \ell_1(v,\theta) - \ell(v,\theta)\nabla_v \ell_2(v,\theta))/\ell_2(v,\theta)$ are then given by:

$$\overline{\ell}_N(v,\theta) = \frac{\overline{\ell}_{1N}(v,\theta)}{\overline{\ell}_{2N}(v,\theta)} \tag{6}$$

and

$$\nabla_v \overline{\ell}_N(v,\theta) = \frac{\nabla_v \overline{\ell}_{1N}(v,\theta) - \overline{\ell}_N(v,\theta)\nabla_v \overline{\ell}_{2N}(v,\theta)}{\overline{\ell}_{2N}(v,\theta)} \tag{7}$$

respectively. Note that these estimators depend on v only through the W_{ti}'s, which can typically be written explicitly as functions of v for fixed values of the underlying random variables Z_{ji}'s. These estimators then permit one to estimate the function and its gradient in functional form w.r.t. v, from the simulation of N regenerative cycles based on density g. Confidence intervals at any fixed value of v can be computed as explained in [11]. In a similar way, again under the appropriate conditions, a consistent estimator of $\nabla_\theta \ell(v,\theta)$ is given by

$$\nabla_\theta \overline{\ell}_N(v,\theta) = \frac{\nabla_\theta \overline{\ell}_{1N}(v,\theta)}{\overline{\ell}_{2N}(v,\theta)}. \tag{8}$$

2.2. Selecting the Reference Parameter Value

An important question in this context is how to select g. Henceforth, we shall assume that $g(\cdot) = f(\cdot, v_0)$, where v_0 is a fixed value of v called the *reference parameter value*. Now the question is how to select v_0. This has been studied, e.g., in [2], [19], [32]. A good choice of v_0 turns out to be extremely important, because for a given v, the variance of the estimators (6–8) may blow up to a very large value or even become infinite for certain choices of v_0. A good choice of v_0 may reduce the variance at a given v compared to that obtained with $v_0 = v$ (the usual choice in standard simulation). On the other hand, it may happen that for any given value of v_0, the variance blows up for certain values of v.

Solving the problem:

$$\min_{v_0 \in V} \mathrm{Var}_{v_0}[\bar{\ell}_N(v,\theta)], \tag{9}$$

for a given $v \in V$ is very difficult in general. In our context, we are also interested in a value of v_0 that does well for all values of v in a certain region. Asmussen and Rubinstein [2] and Rubinstein and Shapiro [32] have studied the problem (9) in the context where $\ell(v,\theta)$ is the average sojourn time per customer in a single queue. In that context, let $\rho(v)$ denote the traffic intensity (which is assumed to depend on v). Under conditions given in [32], $\mathrm{Var}_{v_0}[\bar{\ell}_N(v,\theta)]$ is strictly convex w.r.t. v_0 and one has $\rho(v_0^*) > \rho(v)$, where v_0^* is the optimal solution of (9). This means that it is best to simulate at a larger traffic intensity than the one at which we want to estimate the performance or its gradient. The trace of the variance of $\nabla_v \bar{\ell}_N(v,\theta)$ has the same property under similar conditions. Under these conditions, for a given v, there exists a traffic intensity $\tilde{\rho} > \rho = \rho(v)$ such that

$$\mathrm{Var}_{v_0}[\bar{\ell}_N(v,\theta)] \le \mathrm{Var}_{v}[\bar{\ell}_N(v,\theta)] \tag{10}$$

if and only if $\rho(v_0) \in [\rho, \tilde{\rho}]$, and similarly for the trace of the variance of $\nabla_v \bar{\ell}_N(v,\theta)$. Conversely, for a fixed $\rho_0 = \rho(v_0)$, there exists an interval $[\hat{\rho}, \rho_0]$, such that

$$\mathrm{Var}_{v_0}[\bar{\ell}_N(v,\theta)] \ge \mathrm{Var}_{v}[\bar{\ell}_N(v,\theta)] \tag{11}$$

for all v such that $\rho(v) \in [\hat{\rho}, \rho_0]$. When going below $\hat{\rho}$ and above ρ_0, the variance of the "what-if" estimator $\bar{\ell}_N(v,\theta)$ typically increases rather slowly and very fast, respectively, w.r.t. ρ. Similar results were obtained for more complex queueing models for which the performance measure is the average sojourn or waiting time. A general recommendation from [2], [32] is: in order to be on the safe side one should choose v_0 such that $\rho(v_0)$ is moderately larger than the nominal value $\rho = \rho(v)$. That does not tell us the precise value of v_0^* in general, but gives us some rough guideline for that particular class of models.

Example 1 Suppose that the performance measure of interest is the average sojourn time in an $M/M/1$ queue with traffic intensity $\rho = \rho(v)$. In this case, $\tilde{\rho}$ can be found analytically as a function of ρ [2]. For example, if $\rho = 0.5$, then $\tilde{\rho} \approx 0.8$, so that the variance is reduced for $\rho_0 \in [\rho, \tilde{\rho}] \approx [0.5, 0.8]$, which is a rather broad interval. Conversely, for $\rho_0 = 0.8$, we obtain variance reduction in the sense of (11) *simultaneously* for all $\rho \in [0.5, 0.8]$.

Example 2 Let $L_t(\theta)$ be the expected waiting time for the t-th customer in a $GI/D/1$ queue and assume that we want to estimate the gradient of the steady-state waiting time, $\nabla_\theta \ell(v,\theta)$, where θ is the (deterministic) service time. To do so, recall first (see [34]) that for a GI/G/1 queue, for $t \le \tau$, one has

$$L_t(\theta) = \sum_{j=1}^{t-1} (Y_j - A_j),$$

where $\tau = \min\{t : L_t(\theta) \leq 0\}$, Y_j and A_j are the service time of customer j and the interarrival time between customers $j-1$ and j, respectively. For the GI/D/1 queue, $L_t(\theta)$ reduces to

$$L_t(\theta) = \sum_{j=1}^{t-1}(\theta - A_j). \tag{12}$$

Differentiating $L_t(\theta)$ with respect to θ, we obtain

$$\nabla_\theta L_t(\theta) = \sum_{j=1}^{t-1}(1 - A_j). \tag{13}$$

Substituting finally $L_t(\theta)$ and $\nabla_\theta L_t(\theta)$ from (12) and (13) into (8), we obtain the estimator $\nabla_\theta \overline{\ell}_N(v,\theta)$ which allows the estimation of $\nabla_\theta \ell(v,\theta)$ from a single simulation, simultaneously for different values v, for a fixed θ.

Now, let θ be fixed, while we are interested in estimating at values $v_1, \ldots, v_{r_1}$ of v in V. In this case we are interested in the following extension of problem (9):

$$\min_{v_0 \in V} \max_{v = v_1, \ldots, v_{r_1}} \mathrm{Var}_{v_0}[\overline{\ell}_N(v,\theta)]. \tag{14}$$

Arguing as before, again in the same queueing context, it seems natural to choose the reference parameter v_0 in such a way that

$$\rho(v_0) \approx \rho^+ = \max_{k=1,\ldots,r_1} \rho(v_k), \tag{15}$$

which means that the reference parameter v_0 should correspond to the *highest traffic intensity* among all traffic intensities associated with the selected values $v_1, \ldots, v_{r_1}$. Of course, it may happen that $\rho^- = \min_{k=1,\ldots,r_1} \rho(v_k) < \hat{\rho}$, in which case the variance is decreased in the sense of (11) whenever $\rho = \rho(v_k)$ is in $[\hat{\rho}, \rho^+]$, and increased otherwise. In the latter case, the variance increase is typically moderate when $\hat{\rho} - \rho$ is not too large and the cycle length τ tends to be small (see [2], [19] for illustrations). If that is not the case, then the set $\{v_1, \ldots, v_{r_1}\}$ should be partitioned into smaller subsets and a different v_0 chosen over each subset.

2.3. Many Values of θ: the "What-if" Design for Both Parameters

Consider now the estimation of $\ell(v,\theta)$ for the *"what-if"* design

$$\{(v,\theta) = (v_k, \theta_j), \quad k = 1, \ldots, r_1, \quad j = 1, \ldots, r_2\}. \tag{16}$$

In this case, the CMC method (based on N regenerative cycles) requires a total of $r_1 r_2 N$ simulations, whereas a straightforward combination of CMC (for θ) with SF/LR method requires only $r_2 N$ simulation runs. The idea is simply to apply the technique used for the case of a fixed θ to each value θ_j of interest, as follows:

a) Select a reference parameter $v_{0,j}$;

b) Perform N simulation runs using density $g(\cdot) = f(\cdot, v_{0,j})$, at $\theta = \theta_j$, and compute L_{ti} and W_{ti} (the latter in functional form) for each run i;

c) Calculate $\overline{\ell}_N(v_k, \theta_j)$ according to (6) for $k = 1, \ldots, r_1$.

A trivial adaptation of the above allows also to compute the gradient estimators (7) and (8) over the "what-if" design.

Remark 2 In typical applications, the traffic intensity is often monotone in each component of v and θ. When this is the case, it is easy to find out the parameter value (v_0, θ_0) that gives the largest traffic intensity ρ_0, and use it as a reference parameter to estimate the performance at *all* other parameter values of interest. On the other hand, using the same v_0 for *all* θ of interest is not really necessary.

Remark 3 It is common practice in simulation to use the *same stream* of random numbers while running different scenarios (see, e.g., [16], [18], [30], [37]), in order to reduce the variance of the differences across scenarios. In the present context, this means that the same stream of random numbers would be used for all values of j, with proper synchronization, so that the differences between the estimates will be due only to the different parameter values, and not to different random numbers.

Remark 4 If θ is a continuous parameter, then the above method can be combined with different standard experimental design (ED) methods, such as the full factorial design, the central composite design, and so on, w.r.t. the parameter θ. Such designs turn out to be particular cases of the above "what-if" design, aimed at fitting a regression curve to the response surface $\ell(v, \theta)$.

Example 3 Suppose that we want to estimate the steady-state expected waiting time in a $GI/G/c/m$ queue with interarrival rate of 1, for all the combinations of r_1 different values of the service rate v and r_2 different values of the buffer size $\theta = m$. To do so, we select a $v_{0,j}$ according to (15) for each buffer size m_j from the set $\{m_1, \ldots, m_{r_2}\}$, then run the corresponding r_2 simulations. Here, the reference parameter $v_{0,j}$ should be the smallest value of v of interest, which (in this case) is the same for all j. In comparison with CMC, the number of runs is reduced from $r_1 r_2$ to r_2. If one also has $\rho \in [\hat{\rho}, \rho_0]$, then a variance reduction is also obtained at the same time. If not, then one may partition the values of v of interest into separate intervals, then select a different v_0 and perform separate simulations for each interval. For a more specific illustration, consider an $M/M/2/m$ queue with $m \in \{5, 10, 15\}$ and $v \in \{1.25, 1.5, 2.0, 5.0\}$. Here, one would choose $v_0 = 1.25$ as the reference parameter value. From numerical experiments with this example, we found that the estimator of $\ell(v, \theta)$ based on the change of measure is more accurate (has less variance) in the sense of (11) than its CMC counterpart for $v \leq 2$, and less accurate for $v = 5.0$, for all values of m considered.

3. Optimization

3.1. Discrete Parameters

Consider the minimization problem:

$$\min_{(v,\theta)\in V\times\Theta} \ell(v,\theta), \tag{17}$$

where $V \times \Theta = \{v_k, \theta_j,\ k = 1, \ldots, r_1,\ j = 1, \ldots, r_2\}$. To estimate the minimizer, one can simply estimate $\ell(v,\theta)$ at all points of $V \times \Theta$ (using perhaps the approach of Section 2.3), and just select the system with the best sample value. Assuming that there is a single best system, the probability of making the correct decision (choosing the truly best system) under that procedure converges to one as $N \to \infty$, under the (weak) conditions that our estimators are consistent. This follows from the strong law of large numbers (see also [32]). Note however that this does not tell us about the probability of making the correct decision for a specific N.

For finite sample sizes N, there exists "ranking and selection" procedures for selecting the best system among a finite number of candidates (here $r_1 \times r_2$ candidates), but these procedures usually assume independence between the performance estimators for the different candidates (see [12], [16]). Such procedures will return one of the candidates, which will be the the best system, i.e., the minimizer of (17), with probability at least p^*, where p^* depends on the difference in performance between the best and second best systems. Similar selection procedures using control variates and common random numbers have been proposed and analyzed recently [37], but the set of assumptions made for the analysis typically do not hold in the context of the methodology outlines in Section 2.3. Developing ranking and selection procedures for that context is a topic for further research.

3.2. Continuous Parameters

We are interested in the minimization problem:

$$\min_{(v,\theta)\in V\times\Theta} \ell(v,\theta), \tag{18}$$

where V and Θ are continuous parameter sets. Before proceeding further, consider also the following two particular cases of (18):

$$\min_{v\in V} \ell(v,\theta), \qquad \text{for fixed } \theta \in \Theta \tag{19}$$

and

$$\min_{\theta\in\Theta} \ell(v,\theta), \qquad \text{for fixed } v \in V. \tag{20}$$

The problems (19) and (20) are well known in the stochastic optimization literature. In particular, we can estimate the optimal solution of (19), say $v^*(\theta)$, by solving its *stochastic counterpart* (SC) (see [32]):

$$\min_{v \in V} \overline{\ell}_N(v, \theta), \tag{21}$$

using a conventional mathematical programming method. The statistical properties of the minimizer of (21), which is taken as an estimator of $v^*(\theta)$, are studied in [32]. Under reasonable conditions, the function $\overline{\ell}_N(\cdot, \theta)$ is twice continuously differentiable, and the minimizer in (21) obeys a central-limit theorem and converges to $v^*(\theta)$ as $O(N^{-1/2})$.

In the second case, if v is fixed and Θ is a compact and convex set, we can estimate the optimal solution of (20), say $\theta^*(v)$, by using a conventional stochastic approximation (SA) algorithm of the following form:

$$\theta_{n+1} \; := \; \pi_\Theta[\theta_n - \gamma_n \psi_n], \tag{22}$$

where π_Θ denotes the projection on the convex set Θ, ψ_n is an estimator of $\nabla_\theta \ell(v, \theta)$ at $\theta = \theta_n$ (computed at iteration n of the SA algorithm), θ_n is the parameter value at the beginning of iteration n, and $\{\gamma_n\}$ is a sequence of gains, decreasing to zero, and such that $\sum_{n=1}^{\infty} \gamma_n = \infty$. A common choice for the sequence of gains is $\gamma_n = \gamma_0/n$, for some appropriate constant γ_0. Under a few additional conditions, the SA algorithm can be shown to converge to the optimizer w.p.1, and convergence rates can also be obtained in several cases; see [6], [7], [15], [23], [27] and other numerous references given there. The use of SA and other similar stochastic iterative methods which use gradient or subgradient estimates in the context of on-line or simulated discrete-event dynamic systems has attracted much attention recently; see, e.g., [7], [21], [22], [26] and the several other references given there.

Remark 5 Here, to simplify the discussion, we have assumed that γ_n is a scalar. However, it can also be a matrix of the same dimension as θ. Indeed, $\gamma_n = \gamma_0/n$, where γ_0 is the inverse of the Hessian at the optimum, is asymptotically optimal under broad conditions [15]. That inverse is of course unknown in practice, but adaptive algorithms have been designed which modify both θ_n and γ_n (adaptively) between iterations. Other techniques (e.g., averaging) can also improve the performance of SA. For further details, see [15], [22], [27], [35] and the references given there.

Let us now turn back to the problem (18). Besides straightforward SA (22), which can be applied to estimate the parameter vector (v^*, θ^*), we shall present three new algorithms based on the programs (19) and (20), which combine SA with the SC method. As we shall see below, those algorithms work iteratively, but differ from each other by the fact that the first algorithm tries to solve the problems (19) and (20) by iterating on v and θ sequentially, in a Gauss-Seidel-like manner, while the other two perform parallel iterations with respect to both groups of variables, in a Jacobi-like manner. The second algorithm is similar to the *algorithm with Relaxation* used in games theory (see, e.g., [4]).

Algorithm 1 : Sequential algorithm

1. *Choose two sequences of positive integers:* $\{M_i,\ i \geq 1\}$ *and* $\{N_i,\ i \geq 1\}$, *and three sequences of positive real numbers:* $\{\beta_i,\ i \geq 1\}$, $\{\epsilon_i,\ i \geq 1\}$, *and* $\{\gamma_n,\ n \geq 1\}$, *following the guidelines given by Assumption 1 below and the remarks that follow. Choose an initial parameter vector* (v_1, θ_1), *which represents our best guess of* (v^*, θ^*). *Let* $i := 1$, $n := 1$, *and* $\overline{\theta}_1 = \theta_1$.

2. *REPEAT*

   (a) *For* v *fixed at* v_i, *perform SA for* M_i *iterations to improve the current value of* θ, *i.e., repeat the following* M_i *times: Compute a gradient estimator* ψ_n *w.r.t.* θ *by simulating at parameter value* (v_i, θ_n), *let*

   $$\theta_{n+1} \ := \ \pi_\Theta[\theta_n - \gamma_n \psi_n], \tag{23}$$

   and increase n *by 1.*

   (b) *Let*

   $$\overline{\theta}_{i+1} := \theta_n. \tag{24}$$

   Simulate the system at some reference parameter value $v_{0,i}$ *(which may depend on* v_i*), with* θ *fixed at* $\overline{\theta}_{i+1}$, *for* N_i *regenerative cycles, and then solve the stochastic counterpart (21). Let* $\tilde{v}_i^*$ *be the solution and* $\tilde{v}_i$ *an approximation of it, such that* $\|\tilde{v}_i - \tilde{v}_i^*\|^2 \leq \epsilon_i$. *Put*

   $$v_{i+1} \ := \ \beta_i \tilde{v}_i + (1 - \beta_i) v_i \ = \ v_i + \beta_i(\tilde{v}_i - v_i) \tag{25}$$

   and increase i *by 1.*

   UNTIL an appropriate stopping criterion is met.

3. *Return* $(v_i, \bar{\theta}_i)$ *as an estimate of the optimal solution* (v^*, θ^*).

Algorithm 2 : Parallel algorithm I

Same as Algorithm 1, except that $\overline{\theta}_{i+1}$ *is replaced by* $\overline{\theta}_i$ *in step 2b. With that modification, steps 2a and 2b can be performed in parallel.*

Algorithm 3 : Parallel algorithm II

Same as Algorithm 1, except that (a–b) in step 2 are replaced by the following. Select a reference parameter value $v_{0,i}$ *and repeat the following* M_i *times: Compute a gradient estimator* ψ_n *by simulating at parameter value* $(v_{0,i}, \theta_n)$, *compute* θ_{n+1} *from (23), and increase* n *by 1. Then, solve the stochastic counterpart (21) built from the data obtained during the last* M_i *SA iterations, assuming (almost correctly) that* θ *was fixed at* $\bar{\theta}_{i+1} = (1/M_i)\sum_{j=n-M_i+1}^{n} \theta_j$. *Let* $\tilde{v}_i$ *be an approximation of the solution* $\tilde{v}_i^*$, *such that* $\|\tilde{v}_i - \tilde{v}_i^*\|^2 \leq \epsilon_i$. *Compute* v_{i+1} *from (25) and increase* i *by 1.*

Algorithms 1 and 2 are stochastic versions of the Gauss-Seidel and Jacobi steepest descent algorithms for nonlinear optimization, respectively. Algorithm 3 is fundamentally different in the sense that the *same* simulations are used for both the SA and SC. Therefore, it could be more economical. However, its analysis and implementation tend to be more difficult. For example, one difficulty could be the choice of $v_{0,i}$, because of the fact that θ_n is not fixed during an SC iteration.

For each of those algorithms, we also consider the following "averaging" versions: replace (24) by

$$\overline{\theta}_{i+1} := \frac{1}{M_i} \sum_{j=n-M_i+1}^{n} \theta_j. \tag{26}$$

We shall call those versions 1', 2', and 3', respectively. In those versions, the value $\overline{\theta}_i$ of θ that is used for the SC is the *average* of all values of θ_n during the last series of SA iterations, instead of just the last θ_n. However, when we go back to SA, we restart from the last θ_n. In our empirical investigations, that kind of averaging gave much better results than just taking the last θ_n in the SC as stated in the "regular" versions of the algorithms. However, the convergence proofs appear technically more difficult in that case, mainly because of the switching back from $\overline{\theta}_i$ to θ_n after the SC.

Remark 6 Under appropriate assumptions, if we suppose that v_i^* converges to some value as $i \to \infty$, then it is not hard to show by standard SA arguments that θ_n must converge w.p.1. Conversely, if θ_n converges to some value, then the arguments of [32] can be used to show that v_i must also converges w.p.1 under appropriate conditions. In both cases, if the function is convex and the optimizer is in the interior of Θ, then the convergence point must be the optimum. However, we want (and we shall) prove convergence without making any a priori assumption about the convergence of one of the two sequences. This entails a little more complication.

We now state a list of *sufficient* conditions and give a convergence proof to the optimum under those conditions. Let $\nabla^2\ell$ denote the Hessian (matrix) of ℓ and $\|\cdot\|^2$ denote the Euclidean norm (or the sum of squares of the elements in the case of a matrix). The vectors are assumed to be column vectors and the "prime" transposes them into line vectors. For $i = 1, 2, \ldots$, define $m_i = 1 + \sum_{j=1}^{i-1} M_j$, let $\sum_{(i)}$ denote $\sum_{n=m_i+1}^{m_i+M_i}$, and let $\tilde{\gamma}_i = \sum_{(i)} \gamma_n$. We can decompose ψ_n, for $n \geq 1$, as

$$\psi_n = \nabla_\theta \ell(v_i, \theta_n) + \zeta_n + \xi_n,$$

where $E[\xi_n \mid v_i, \theta_n] = 0$ and $\zeta_n = E[\psi_n \mid v_i, \theta_n] - \nabla_\theta \ell(v_i, \theta_n)$. The random variable ζ_n represents the conditional bias on the gradient estimator ψ_n at the nth SA iteration, while ξ_n represents the noise.

Assumption 1 *(i) The function $\ell(v, \theta)$ is twice continuously differentiable over $V \times \Theta$, which is a compact and convex subset of the d-dimensional real space for* *some integer d, and there is a unique minimizer* (v^*, θ^*) *which is an interior point of* $V \times \Theta$.

(ii) *The Hessian* $\nabla^2 \ell(v, \theta)$ *is positive definite over* $V \times \Theta$*, with smallest eigenvalue bounded below by* $\lambda_{\min} > 0$ *and largest eigenvalue bounded above by* $\lambda_{\max} < \infty$*, uniformly over* $V \times \Theta$.

(iii) *One has* $\gamma_n \searrow 0$, $0 < \beta_i \leq 1$, $\sum_{i=1}^{\infty} \beta_i^2 / N_i^2 < \infty$, $N_i \to \infty$, $\beta_i / (N_i \tilde{\gamma}_i) \to 0$, $\sum_{i=1}^{\infty} \min(\beta_i,\ \tilde{\gamma}_i) = \infty$, $\tilde{\gamma}_i / \beta_i \leq K_1$, *and* $\epsilon_i \leq K_1 / N_i$ *for some finite constant* K_1.

(iv) *One has* $E[\zeta_n] \to 0$, *and* $\zeta_n \to 0$ *w.p.1 as* $n \to \infty$. *Also,* $\gamma_n E[\|\xi_n\|^2] \to 0$, $\gamma_n E[\|\xi_n\|^2 \mid v_i, \theta_n] \to 0$ *w.p.1, and* $\sum_{n=1}^{\infty} \gamma_n^2 E[\|\xi_n\|^2] < \infty$.

(v) *W.p.1,* $\overline{\ell}_N(v, \theta)$ *is twice continuously differentiable w.r.t. v and there is a finite constant* K_2 *such that*

$$\sup_{(v_0, \theta) \in V \times \Theta} E\left[\|\nabla_v \overline{\ell}_N(v^*(\theta), \theta)\|^4 \mid v_0, \theta\right] \leq K_2^2 / N^2. \tag{27}$$

In many cases, ζ_n is zero and the conditions on it hold trivially. The last condition in Assumption 1(v) is reasonable in view of the fact that $\nabla_v \ell(v^*(\theta), \theta) = 0$.

Observe that the solution v^* of the SC (21) is usually *not* an unbiased estimator of the optimal solution of the original minimization problem (18), but it is a consistent estimator under broad conditions (see [32]). This is why we need to take $N_i \to \infty$. Reasonable choices for the sequences could be $N_i = N_0 + N_1 i$ for fixed constants N_0 and N_1, and $M_i = N_i$, which gives an equal part of the budget to the SA and SC "components" of the algorithm. The role of β_i is to introduce a weighted averaging of the previous values of $\tilde{v}_i$ instead of just taking the last one. The aim of this is mainly to reduce the variance. For example, one can take $\beta_i = N_i / \sum_{j=1}^{i} N_j$, which is equivalent to taking the weighted average:

$$v_{i+1} = \frac{\sum_{j=1}^{i} N_j \tilde{v}_j}{\sum_{j=1}^{i} N_j}.$$

Other possibilities include taking $\beta_i = \beta_0 / (b + i)$ for some positive constants $\beta_0 \leq b$, or β_i equals to a constant. The latter corresponds to exponential smoothing. The standard choice for γ_n is $\gamma_n = \gamma_0 / n$ for some constant γ_0. Finally, one can take $\epsilon_i = K_1 / N_i$ for some constant K_1. We point out that with the above choices of N_i and γ_n, and with β_i equal to a constant, the condition: $\beta_i / (N_i \tilde{\gamma}_i) \to 0$ is not satisfied.Nevertheless, that combination turned out to give the best results in our empirical investigations.

Proposition 1 *Under Assumption 1, one has*

$$\lim_{i \to \infty} \left(\|v_i - v^*\|^2 + \|\overline{\theta}_i - \theta^*\|^2 \right) = 0 \qquad \textit{w.p.1}$$

in Algorithm 1.

Proof: Let $\Delta_i = \|v_i - v^*\|^2 + \|\theta_{m_i} - \theta^*\|^2 = \|v_i - v^*\|^2 + \|\overline{\theta}_i - \theta^*\|^2$ and $\Delta_{i,j} = \|v_i - v^*\|^2 + \|\theta_{m_i+j} - \theta^*\|^2$, $j = 0, \ldots, M_i$. Let $v_i^* = v^*(\overline{\theta}_{i+1})$, the optimal value of v when θ is fixed at $\overline{\theta}_{i+1}$. For $i \geq 1$, $0 \leq j < M_i$, and $n = m_i + j$, define

$$\begin{aligned} D_n &= \nabla_\theta \ell(v_i, \theta_n), \\ S_n &= 2(D_n - \psi_n)'(\theta_n - \theta^*) + \gamma_n \|\psi_n\|^2, \\ T_i &= \|\tilde{v}_i - v_i^*\|^2, \\ s_n &= E[S_n \mid v_i, \theta_n], \\ t_i &= E[T_i \mid v_{i,0}, \overline{\theta}_{i+1}]. \end{aligned}$$

In the remainder of the proof, we will use the following lemmas.

Lemma 1 *There is a constant* $0 < K_3 \leq 1$ *such that for all* $(v, \theta) \in V \times \Theta$,

$$\|v - v^*(\theta)\| + \nabla_\theta \ell(v, \theta)'(\theta - \theta^*) \;\geq\; K_3(\|\theta - \theta^*\|^2 + \|v - v^*\|^2).$$

Proof: First, observe that

$$\begin{aligned} \ell(v,\theta) - \ell(v^*,\theta^*) &= \nabla \ell(v^*,\theta^*)' \begin{pmatrix} \theta - \theta^* \\ v - v^* \end{pmatrix} + \frac{1}{2} \begin{pmatrix} \theta - \theta^* \\ v - v^* \end{pmatrix}' \nabla^2 \ell(\hat{v},\hat{\theta}) \begin{pmatrix} \theta - \theta^* \\ v - v^* \end{pmatrix} \\ &\geq \frac{\lambda_{\min}}{2} (\|\theta - \theta^*\|^2 + \|v - v^*\|^2) \end{aligned}$$

where $(\hat{v}, \hat{\theta})$ lies on the line segment joining (v, θ) to (v^*, θ^*). Since ℓ is convex, one has

$$\begin{aligned} \ell(v,\theta) - \ell(v^*,\theta^*) &\leq \nabla_\theta \ell(v,\theta)'(\theta - \theta^*) + \nabla_v \ell(v,\theta)'(v - v^*) \\ &= \nabla_\theta \ell(v,\theta)'(\theta - \theta^*) + \nabla_v \ell(v^*(\theta),\theta)'(v - v^*) \\ &\quad + (v - v^*(\theta))' \nabla_v^2 \ell(\hat{\hat{v}},\theta)(v - v^*) \\ &\leq \nabla_\theta \ell(v,\theta)'(\theta - \theta^*) + \lambda_{\max} \|v - v^*(\theta)\| \cdot \|v - v^*\|, \end{aligned}$$

where $\hat{\hat{v}}$ lies on the line segment between v and $v^*(\theta)$. Since V is compact, $\|v - v^*\|$ is bounded above, say by K, so that

$$\|\theta - \theta^*\|^2 + \|v - v^*\|^2 \leq \frac{2}{\lambda_{\min}} \nabla_\theta \ell(v,\theta)'(\theta - \theta^*) + \frac{2K\lambda_{\max}}{\lambda_{\min}} \|v - v^*(\theta)\|, \quad (28)$$

and the result follows. ∎

Lemma 2 *There is a finite constant* K_4 *such that for all* $i \geq 1$,

$$t_i \leq K_4/N_i \quad \text{w.p.1 and} \qquad E[T_i^2] \leq K_4^2/N_i^2. \quad (29)$$

Proof: Let $\overline{\ell}_i$ denote the sample function that corresponds to (6) obtained at iteration i with $N = N_i$. From Assumption 1(i,v) and Taylor's expansion, for any $\theta \in \Theta$, one has

$$\overline{\ell}_i(\tilde{v}_i^*, \theta) - \overline{\ell}_i(v_i^*, \theta) = (\tilde{v}_i^* - v_i^*)' \nabla_v \overline{\ell}_i(v_i^*, \theta) + (\tilde{v}_i^* - v_i^*)' \nabla_v^2 \overline{\ell}_i(u_i(\theta), \theta)(\tilde{v}_i^* - v_i^*)/2$$

where $u_i(\theta)$ lies on the line between v_i^* and $\tilde{v}_i^*$. By definition of $\tilde{v}_i^*$, $\overline{\ell}_i(\tilde{v}_i^*, \overline{\theta}_{i+1}) - \overline{\ell}_i(v_i^*, \overline{\theta}_{i+1}) \leq 0$. Therefore,

$$\begin{aligned} 2\|\tilde{v}_i^* - v_i^*\| \cdot \|\nabla_v \overline{\ell}_i(v_i^*, \overline{\theta}_{i+1})\| &\geq (\tilde{v}_i^* - v_i^*)' \nabla_v^2 \overline{\ell}_i(u_i(\overline{\theta}_{i+1}), \overline{\theta}_{i+1})(\tilde{v}_i^* - v_i^*) \\ &\geq \lambda_{\min} \|\tilde{v}_i^* - v_i^*\|^2, \end{aligned}$$

so

$$\|\tilde{v}_i^* - v_i^*\| \leq (2/\lambda_{\min}) \|\nabla_v \overline{\ell}_i(v_i^*, \overline{\theta}_{i+1})\|$$

and, from Assumption 1(v), w.p.1,

$$\begin{aligned} E\left[\|\tilde{v}_i - v_i^*\|^2 \mid v_{i,0}, \overline{\theta}_{i+1}\right] &\leq 2\left(\frac{4}{\lambda_{\min}^2} E\left[\|\nabla_v \overline{\ell}_i(v_i^*, \overline{\theta}_{i+1})\|^2 \mid v_{i,0}, \overline{\theta}_{i+1}\right] + \epsilon_i\right) \\ &\leq \frac{8K_2}{\lambda_{\min}^2 N_i} + \frac{2K_1}{N_i}. \end{aligned}$$

Similarly,

$$\begin{aligned} E[T_i^2] &= E\left[\|\tilde{v}_i - v_i^*\|^4\right] \\ &\leq 8\left(E\left[\|\tilde{v}_i^* - v_i^*\|^4\right] + E\left[\|\tilde{v}_i - \tilde{v}_i^*\|^4\right]\right) \\ &\leq 8\left(\frac{2^4}{\lambda_{\min}^4} E\left[\|\nabla_v \overline{\ell}_i(v_i^*, \overline{\theta}_{i+1})\|^4\right] + \epsilon_i^2\right) \\ &\leq \frac{128K_2^2}{\lambda_{\min}^4 N_i^2} + \frac{8K_1^2}{N_i^2}. \end{aligned}$$

Define $K_4 = \max(8K_2/\lambda_{\min}^2 + 2K_1,\ 128K_2^2/\lambda_{\min}^4 + 8K_1^2)$. ■

We now continue the proof of the proposition. We have:

$$\begin{aligned} \Delta_{i,j+1} &= \|v_i - v^*\|^2 + \|\theta_{n+1} - \theta^*\|^2 \\ &\leq \|v_i - v^*\|^2 + \|\theta_n - \gamma_n \psi_n - \theta^*\|^2 \\ &= \Delta_{i,j} - 2\gamma_n \psi_n'(\theta_n - \theta^*) + \gamma_n^2 \|\psi_n\|^2 \\ &\leq \Delta_{i,j} + \gamma_n S_n - 2\gamma_n D_n'(\theta_n - \theta^*). \end{aligned}$$

Also,

$$\begin{aligned} \Delta_{i+1} &= \|v_{i+1} - v^*\|^2 + \|\overline{\theta}_{i+1} - \theta^*\|^2 \\ &= \|v_i + \beta_i(\tilde{v}_i - v_i) - v^*\|^2 + \|\overline{\theta}_{i+1} - \theta^*\|^2 \\ &\leq (1 - \beta_i)\|v_i - v^*\|^2 + \beta_i \|\tilde{v}_i - v^*\|^2 + \|\overline{\theta}_{i+1} - \theta^*\|^2 \\ &= \|v_i - v^*\|^2 + \|\overline{\theta}_{i+1} - \theta^*\|^2 + \beta_i[\|\tilde{v}_i - v^*\|^2 - \|v_i^* - v^*\|^2] \\ &\quad -\beta_i[\|v_i - v^*\|^2 - \|v_i^* - v^*\|^2] \\ &\leq \Delta_{i,M_i} + \beta_i \|\tilde{v}_i - v_i^*\|^2 - \beta_i \|v_i - v_i^*\|^2. \end{aligned}$$

Combining these inequalities yields

$$\Delta_{i+1} \;\leq\; \Delta_i + \beta_i T_i - \beta_i\|v_i - v_i^*\|^2 + \sum_{(i)} \gamma_n S_n - 2\sum_{(i)} \gamma_n D_n'(\theta_n - \theta^*). \tag{30}$$

Let $\delta_i = E[\Delta_i]$. To complete the proof, we shall show first that $\delta_i \to 0$ as $i \to \infty$, then that $\Delta_i \to 0$ w.p.1. For the former, we will show that for any $\epsilon > 0$, δ_i eventually gets smaller than ϵ for large i, and cannot go over 2ϵ thereafter. We will then do a similar reasoning for Δ_i. We draw some ideas from the proofs of Lemmas 7 and 8 of Ermoliev and Gaivoronski [7].

From Assumptions 1(iv) and the fact that D_n as well as $(\theta_n - \theta^*)$ are bounded, we have that $E[s_n] = E[2\zeta_n'(\theta_n - \theta^*)] + \gamma_n E[\|\psi_n\|^2] \to 0$ as $n \to \infty$. Take an arbitrary $0 < \epsilon < 1$ and define $\delta(\epsilon) = K_3\epsilon$. There is an integer i_0 such that for all $i \geq i_0$ and $n \geq m_i$,

$$\max\left\{\frac{K_4}{N_i},\ \frac{\tilde{\gamma}_i E[s_n]}{\beta_i}\right\} \;\leq\; \frac{\delta^2(\epsilon)}{16} \tag{31}$$

and

$$\max\left\{E[s_n],\ \frac{K_4\beta_i}{N_i\tilde{\gamma}_i}\right\} \;\leq\; \frac{\delta(\epsilon)}{4}. \tag{32}$$

Suppose that

$$\delta_i > \epsilon \text{ for all } i \geq i_0. \tag{33}$$

Then, for each $i \geq i_0$, using Lemma 1 and taking expectations, one has

$$\begin{aligned} E[\|v_i - v_i^*\| + D_n'(\theta_n - \theta^*)] \;&\geq\; K_3 E[\|v_i - v^*\|^2 + \|\theta_n - \theta^*\|^2] \\ &>\; K_3\epsilon \;=\; \delta(\epsilon), \end{aligned} \tag{34}$$

which implies that either

$$E[\|v_i - v_i^*\|] \geq \delta(\epsilon)/2 \tag{35}$$

or

$$E[D_n'(\theta_n - \theta^*)] > \delta(\epsilon)/2 \qquad \text{for all } n \text{ in } \{m_i, \ldots, m_{i+1}-1\}. \tag{36}$$

If $i \geq i_0$ and (35) holds, then $E[\|v_i - v_i^*\|^2] \geq \delta(\epsilon)^2/4$ and, from (30), Lemma 2, and (31),

$$\begin{aligned} \delta_{i+1} - \delta_i \;&\leq\; \beta_i E[t_i - \delta^2(\epsilon)/4)] + \sum_{(i)} \gamma_n E[s_n] \\ &\leq\; \beta_i[K_4/N_i - \delta^2(\epsilon)/4)] + \tilde{\gamma}_i \sup_{m_i \leq n < m_{i+1}} E[s_n] \\ &\leq\; -\beta_i\delta^2(\epsilon)/8. \end{aligned}$$

On the other hand, if (36) holds, from (30), Lemma 2, and (32), one has

$$\begin{aligned}\delta_{i+1}-\delta_i &\leq \beta_i E[t_i]+\sum_{(i)}\gamma_n(E[s_n]-\delta(\epsilon))\\ &\leq \frac{K_4\beta_i}{N_i}+\tilde{\gamma}_i\delta(\epsilon)/4-\tilde{\gamma}_i\delta(\epsilon)\\ &\leq -\tilde{\gamma}_i\delta(\epsilon)/2.\end{aligned}$$

Combining these inequalities yields

$$\delta_{i+1}-\delta_i \leq -\min(\delta(\epsilon)\beta_i,\ \tilde{\gamma}_i)\delta(\epsilon)/8 \tag{37}$$

and

$$\sum_{i=i_0}^{\infty}(\delta_{i+1}-\delta_i) \ \leq\ -\sum_{i=i_0}^{\infty}\min(\delta(\epsilon)\beta_i,\ \tilde{\gamma}_i)\delta(\epsilon)/8 \ =\ -\infty.$$

This implies that $\delta_i \to -\infty$, which is a contradiction because δ_i can never be negative. Therefore, there exists $i_1 \geq i_0$ such that $\delta_{i_1} < \epsilon$. We now claim that $\delta_i < 2\epsilon$ for all $i \geq i_1$. Suppose otherwise, that is, $i_3 = \inf\{i \geq i_1 \mid \delta_i > 2\epsilon\} < \infty$, and let $i_2 = \max\{i < i_3 \mid \delta_i < \epsilon\}$. For $i \geq i_0$, one has

$$\delta_{i+1}-\delta_i \ \leq\ \beta_i E[t_i]+\sum_{(i)}\gamma_n E[s_n] \ \leq\ \frac{K_4\beta_i}{N_i}+\tilde{\gamma}_i \sup_{m_i\leq n<m_{i+1}} E[s_n] \ \leq\ \epsilon/2.$$

Therefore, one must have $i_3 - i_2 > 1$ and $\epsilon < \delta_i < 2\epsilon$ for $i_2 < i < i_3$. Then, by the same reasoning as above, $\delta_{i+1}-\delta_i \leq -\min(\delta(\epsilon)\beta_i,\ \tilde{\gamma}_i)\delta(\epsilon)/8 < 0$ for $i_2 < i < i_3$, which contradicts the definition of i_3. Since ϵ is arbitrary, we have now shown that $\delta_i \to 0$ as $i \to \infty$.

Now, for each $\epsilon > 0$ and integer i_0,

$$\begin{aligned}\epsilon P[\Delta_i \geq \epsilon \text{ for all } i \geq i_0] \ &=\ \epsilon P\left[\inf_{i\geq i_0}\Delta_i \geq \epsilon\right]\\ &\leq\ E\left[\inf_{i\geq i_0}\Delta_i\right] \ \leq\ \inf_{i\geq i_0} E[\Delta_i] \ =\ 0.\end{aligned}$$

Therefore, w.p.1, there exists $i_1 \geq i_0$ such that $\Delta_{i_1} < \epsilon$.

From Lemma 2 and Assumption 1(iii), we have that

$$\sum_{i=1}^{\infty}\beta_i^2 E[(T_i-t_i)^2] < \infty. \tag{38}$$

It then follows from standard martingale theory that

$$\sum_{i=1}^{\infty}\beta_i(T_i-t_i) < \infty. \qquad \text{w.p.1.} \tag{39}$$

We also have

$$\begin{aligned}\sum_{n=1}^{\infty}\gamma_n(S_n - s_n) &= 2\sum_{n=1}^{\infty}\gamma_n\xi_n'(\theta_n-\theta^*) + \sum_{n=1}^{\infty}\gamma_n^2(\|\psi_n\|^2 - E[\|\psi_n\|^2 \mid v_i, \theta_n]) \\ &= 2\sum_{n=1}^{\infty}\gamma_n\xi_n'(\theta_n-\theta^*) + \sum_{n=1}^{\infty}\gamma_n^2\|\xi_n\|^2 + 2\sum_{n=1}^{\infty}\gamma_n^2\xi_n'(D_n+\zeta_n).\end{aligned}$$

Since $E\left[\sum_{n=1}^{\infty}\gamma_n^2\xi_n^2\right] < \infty$ and since $\{\theta_n\}$ and $\{D_n\}$ evolve in compact sets, it follows (again from a martingale argument) that the first and third sums in the last expression are finite w.p.1. The second sum is also finite w.p.1 because all its terms are non-negative and it has finite expectation. We then have

$$\sum_{n=1}^{\infty}\gamma_n(S_n - s_n) < \infty \qquad \text{w.p.1.} \tag{40}$$

It follows from (39) and (40) that, w.p.1, $\beta_i(T_i - t_i) \to 0$, $\gamma_n(S_n - s_n) \to 0$, and so, in view of (30) and since both $\{t_i\}$ and $\{s_n\}$ converge to zero w.p.1, we have that $\max(0,\ \Delta_{i+1} - \Delta_i) \to 0$ w.p.1. Choose i_0 such that for all $i \geq i_0$, $\Delta_{i+1} - \Delta_i < \epsilon/2$ and

$$\sup_{I \geq i_0}\sum_{i=i_0}^{I}\left[\beta_i(T_i - t_i) \ + \ \sum_{(i)}\gamma_n(S_n - s_n)\right] \ \leq \ \epsilon/2. \tag{41}$$

Then, choose $i_1 \geq i_0$ such that $\Delta_{i_1} < \epsilon$. By a similar argument as in the case of the expectation, we now show that $\Delta_i < 2\epsilon$ for all $i \geq i_1$. Suppose that $i_3 = \inf\{i \geq i_1 \mid \Delta_i > 2\epsilon\} < \infty$, and let $i_2 = \max\{i < i_3 \mid \Delta_i < \epsilon\}$. Since $\Delta_{i+1} - \Delta_i < \epsilon/2$, we must have $i_3 - i_2 > 1$. Then, for $i_2 < i < i_3$, we have $\epsilon < \Delta_i < 2\epsilon$ and

$$\begin{aligned}\Delta_{i+1} - \Delta_i &\leq \beta_i(T_i - t_i) + \sum_{(i)}\gamma_n(S_n - s_n) + \beta_i t_i + \sum_{(i)}\gamma_n s_n \\ &\quad -\beta_i\|v_i - v_i^*\|^2 - 2\sum_{(i)}\gamma_n D_n'(\theta_n - \theta^*) \\ &\leq \beta_i(T_i - t_i) + \sum_{(i)}\gamma_n(S_n - s_n) - \min(\delta(\epsilon)\beta_i\ \tilde{\gamma}_i)\delta(\epsilon)/8,\end{aligned}$$

where the last inequality follows from the same arguments that we used to obtain (37), but without the expectation E, and using the fact that $s_n = 2\zeta_n'(\theta_n - \theta^*) + \gamma_n E[\|\psi_n\|^2 \mid v_i, \theta_n] \to 0$ w.p.1 from our assumptions. Combining this with (41), we obtain that $\Delta_{i_3} - \Delta_{i_2+1} < \epsilon/2$. It follows that $\Delta_{i_3} - \Delta_{i_2} < \epsilon$, contradicting the assumption that $\Delta_{i_3} > 2\epsilon$. We have now shown that $\Delta_i \to 0$ w.p.1 as $i \to \infty$. That completes the proof. ∎

3.3. A Numerical Illustration

Example 4 To illustrate those algorithms, we will take a simple example, namely an $M/D/1$ queue, where v is the arrival rate and θ is the (deterministic) service time of each customer. Suppose we want to minimize

$$\alpha(v,\theta) = \ell(v,\theta) + 1/v + 1/\theta, \tag{42}$$

where $\ell(v,\theta) = \ell_1(v,\theta)/\ell_2(v,\theta)$ is the average sojourn time in the system per customer, while $\ell_1(v,\theta)$ and $\ell_2(v,\theta)$ are respectively the expected total sojourn time and the expected number of customers, over one regenerative cycle. We impose the following constraints: $0.1 \le v \le 1.3$ and $0.1 \le \theta \le 0.7$. These constraints will turn out to be inactive at the optimum; however, they make sure that the system will remain stable and that the parameters always take reasonable values all along the optimization process. Indeed, the traffic intensity is bounded as follows: $0.01 \le v\theta \le 0.91$. Minimizing (42) is clearly a rather simple and easy to solve example, but it can nevertheless illustrate quite well our algorithms.

For the present example, one has $\ell(v,\theta) = \theta + v\theta^2/(2(1-v\theta))$; see Wolff [36], p.385. Using this in a deterministic optimization algorithm, one finds that (42) is minimized by taking $(v^*,\theta^*) \approx (1.0824, 0.5412)$. One has $\ell_2(v^*,\theta^*) = 1/(1-v^*\theta^*) \approx 2.414$, $\alpha(v^*,\theta^*) \approx 3.6955$, and the values of the second derivatives of α with respect to v and θ at that point are approximately 3.8076 and 13.437, respectively.

Here, we can use the score function method to estimate the derivative with respect to v, but not the derivative with respect to θ, because the likelihood ratio does not exist (although one could perhaps apply the "push-out" method as in [32], p.229, but we will not do it here). For that second derivative, we will use here perturbation analysis (IPA) [8], [17]. Minimizing (42) is equivalent to finding a zero of the gradient of (42) with respect to (v,θ), or, equivalently, to solving the equations:

$$\ell_2^2(v,\theta)\frac{d}{dv}\alpha(v,\theta) \;=\; \ell_2(v,\theta)\frac{d}{d\theta}\alpha(v,\theta) \;=\; 0, \tag{43}$$

which can also be written as

$$\ell_2(v,\theta)\frac{d}{dv}\ell_1(v,\theta) - \ell_1(v,\theta)\frac{d}{dv}\ell_2(v,\theta) - \ell_2^2(v,\theta)/v^2 \;=\; 0; \tag{44}$$

$$\frac{d}{d\theta}\ell_1(v,\theta) - \ell(v,\theta)\frac{d}{d\theta}\ell_2(v,\theta) - \ell_2(v,\theta)/\theta^2 \;=\; 0. \tag{45}$$

As explained in L'Ecuyer and Glynn [21], one can obtain an unbiased estimator of the left-hand-side of (44) from two independent regenerative cycles and the score function method, and an unbiased estimator of the left-hand-side of (45) from one regenerative cycle with IPA.

The numerical results we present here are for Algorithms 1'–3'. We first tried Algorithms 1–3 and the results were much more noisy. We took sequences of the form $M_i = N_i = N_0 + N_1 i$ and $\gamma_n = \gamma_0/n$, for different values of N_0, N_1, and

γ_0, and tried both $\beta_i = 1/i$ and β_i constant. In each case, the initial parameter value was $(v_1, \theta_1) = (1/2, 1/2)$, and we used $v_{0,i} = 1.3$. To compute $\tilde{v}_i$ in step 2(b) of the algorithm, we used a bisection method and stopped when the size of the interval was smaller than 10^{-4} (so, ϵ_i is negligeably small). The stopping criterion for the "REPEAT…UNTIL" loop was: stop after a total of $\bar{T}$ customers have been simulated, where $\bar{T}$ is a fixed constant.

The function α here satisfies Assumption 1 (i), while (iv) is satisfied since $\zeta_i = 0$ and one can show (much as in L'Ecuyer and Glynn [21]) that $\sup_{(v,\theta)\in V\times\Theta} E[\|\xi_n\|^2 \mid v, \theta] < \infty$. Note that for β_i constant, (iii) does not hold, but that nevertheless gave us the best results empirically.

For each selected set of parameters $(N_0, N_1, \{\beta_i\}, \gamma_0, \bar{T})$, each algorithm was repeated 10 times and we computed the empirical mean, the standard deviation s_d, and the standard error s_e of the 10 final values of v_i and $\bar{\theta}_i$, as in L'Ecuyer, Giroux, and Glynn [22]. If y_k denotes the final value of parameter y for replication k ($y = v_i$ or $\bar{\theta}_i$), the latter quantities are defined by

$$\mu(y) = \frac{1}{10}\sum_{k=1}^{10} y_k; \qquad s_d^2(y) = \frac{1}{9}\sum_{i=1}^{10}(y_i - \mu(y))^2; \qquad s_e^2(y) = \frac{1}{10}\sum_{i=1}^{10}(y_i - \theta^*)^2. \tag{46}$$

A selection of results is given in Tables 1–3 for $\bar{T} = 10^5,\ 10^6$ and 10^7.

From those tables, we can observe the following.

1. When N_i is fixed to a small constant, the algorithm does not converge to the optimizer. This can be seen by looking at the s_d and s_e values when $N_1 = 0$: a small s_d and large s_e indicate a small variance but large bias.

2. All three algorithms appear to converge at the canonical rate of $O(\bar{T}^{-1/2})$; that is, when $\bar{T}$ is multiplied by 100, s_d and s_e are roughly divided by 10.

3. For that particular example, $\beta_i = 0.5$ and $M_i = N_i = 200 + 200i$ appear to work well. It other words, it is better in this case to switch not too frequently between SA and SC, and to insure that a large number of regenerative cycles is used at every SC iteration.

Of course, these observations stand only for this particular example; the algorithms can behave much differently in other situations. Nevertheless, this is a first step towards getting insight about what goes on. We certainly *cannot say* that these algorithms always work well and are easy to implement in general. Among the difficulties that may arise, we mention the following:

1. Implementing the projection on $V \times \Theta$ when it is a non-rectangular set; and deciding what to do if it is non-convex.

2. Choosing appropriate sequences $\{M_i, N_i, \beta_i, \gamma_n\}$ for the problem at hand (the performance of the algorithm is generally quite sensitive to those choices, and our numerical results illustrate that to some extent).

3. Choosing the appropriate $v_{0,i}$ and implementing the SC part of the algorithm, especially for Algorithm 3.

More investigation would be required before making specific recommendations for dealing with those difficulties. Perhaps adaptive heuristic could also be designed. This offers challenging opportunities for further research.

Table 1. Results for Example 3.1, with $\bar{T} = 10^5$.

Algor.	N_0	N_1	β_i	γ_0	$\mu(v_i)$	$\mu(\theta_i)$	$s_d(v)$	$s_d(\theta)$	$s_e(v)$	$s_e(\theta)$
1'	2	0	$1/i$	0.1	1.190	.513	.0045	.0046	.1076	.0290
1'	20	0	$1/i$	0.1	1.216	.507	.0078	.0040	.1337	.0341
1'	20	0	0.5	0.1	1.241	.502	.0620	.0084	.1691	.0400
1'	200	0	0.5	0.1	1.195	.515	.0607	.0115	.1262	.0280
1'	20	2	$1/i$	0.1	1.183	.514	.0211	.0069	.1021	.0276
1'	20	2	$1/i$	0.5	1.180	.516	.0187	.0082	.0993	.0263
1'	20	20	$1/i$	0.1	1.169	.521	.0302	.0084	.0915	.0221
1'	20	20	$1/i$	0.5	1.154	.527	.0341	.0164	.0789	.0210
1'	20	20	0.5	0.5	1.097	.542	.0221	.0119	.0254	.0113
1'	200	200	0.5	0.1	1.068	.545	.0233	.0068	.0262	.0076
1'	200	200	0.5	0.3	1.076	.539	.0197	.0106	.0197	.0103
1'	200	200	0.5	0.5	1.076	.542	.0209	.0155	.0209	.0147
2'	20	2	$1/i$	0.1	1.172	.519	.0226	.0057	.0922	.0227
2'	20	2	$1/i$	0.5	1.178	.516	.0278	.0091	.0993	.0264
2'	20	20	$1/i$	0.1	1.174	.517	.0238	.0067	.0944	.0249
2'	20	20	$1/i$	0.5	1.115	.535	.0387	.0130	.0494	.0137
2'	20	20	0.5	0.5	1.111	.539	.0453	.0100	.0516	.0097
2'	200	200	0.5	0.1	1.040	.554	.0127	.0062	.0444	.0143
2'	200	200	0.5	0.3	1.044	.550	.0147	.0114	.0407	.0138
2'	200	200	0.5	0.5	1.058	.546	.0155	.0095	.0286	.0102
3'	20	2	$1/i$	0.1	1.180	.515	.0130	.0040	.0986	.0270
3'	20	2	$1/i$	0.5	1.133	.525	.0207	.0057	.0545	.0167
3'	20	20	$1/i$	0.1	1.149	.523	.0435	.0125	.0790	.0218
3'	20	20	$1/i$	0.5	1.106	.534	.0300	.0081	.0375	.0106
3'	20	20	0.5	0.5	1.094	.537	.0120	.0020	.0160	.0043
3'	200	200	0.5	0.1	1.052	.553	.0120	.0042	.0325	.0120
3'	200	200	0.5	0.3	1.072	.544	.0097	.0021	.0138	.0034
3'	200	200	0.5	0.5	1.075	.543	.0093	.0029	.0115	.0035

Table 2. Results for Example 1, with $\bar{T} = 10^6$.

Algor.	N_0	N_1	β_i	γ_0	$\mu(v_i)$	$\mu(\theta_i)$	$s_d(v)$	$s_d(\theta)$	$s_e(v)$	$s_e(\theta)$
1'	20	20	$1/i$	0.1	1.145	.524	.0260	.0269	.0670	.0189
1'	20	20	$1/i$	0.5	1.127	.528	.0357	.0103	.0568	.0167
1'	20	20	0.5	0.5	1.085	.540	.0133	.0056	.0128	.0054
1'	200	200	0.5	0.1	1.087	.540	.0078	.0030	.0089	.0033
1'	200	200	0.5	0.3	1.088	.540	.0063	.0028	.0080	.0030
1'	200	200	0.5	0.5	1.085	.542	.0108	.0040	.0104	.0039
2'	20	20	$1/i$	0.1	1.131	.528	.0369	.0099	.0595	.0165
2'	20	20	$1/i$	0.5	1.106	.534	.0197	.0059	.0297	.0092
2'	20	20	0.5	0.5	1.091	.538	.0172	.0064	.0186	.0068
2'	200	200	0.5	0.1	1.080	.542	.0080	.0034	.0080	.0034
2'	200	200	0.5	0.3	1.081	.541	.0116	.0048	.0111	.0045
2'	200	200	0.5	0.5	1.092	.540	.0099	.0050	.0131	.0048
3'	20	20	$1/i$	0.1	1.109	.533	.0221	.0062	.0340	.0098
3'	20	20	$1/i$	0.5	1.075	.543	.0151	.0043	.0164	.0046
3'	20	20	0.5	0.5	1.082	.541	.0061	.0016	.0058	.0016
3'	200	200	0.5	0.1	1.082	.542	.0053	.0009	.0051	.0008
3'	200	200	0.5	0.3	1.083	.541	.0026	.0013	.0025	.0012
3'	200	200	0.5	0.5	1.084	.541	.0029	.0014	.0031	.0013

Table 3. Results for Example 1, with $\bar{T} = 10^7$.

Algor.	N_0	N_1	β_i	γ_0	$\mu(v_i)$	$\mu(\theta_i)$	$s_d(v)$	$s_d(\theta)$	$s_e(v)$	$s_e(\theta)$
1'	200	200	0.5	0.3	1.084	.541	.0041	.0012	.0042	.0013
2'	200	200	0.5	0.3	1.083	.541	.0027	.0014	.0027	.0014
3'	200	200	0.5	0.3	1.083	.541	.0019	.0003	.0019	.0003

Acknowledgments

This work was supported by NSERC-Canada grants no. OGP0110050 and OGP0005491, FCAR-Québec grant no. EQ2831, and the Technion V.P.R. Fund — B.R.L. Bloomfield Industrial Management Research Fund. We wish to thank G. Kochman, B. Polyak, and S. Uryas'ev for valuable comments and suggestions.

References

1. Asmussen, S. and R. Y. Rubinstein (1992a). The efficiency and heavy traffic properties of the score function method for sensitivity analysis of queueing models. *Advances in Applied Probability*, **24**, 172–201.
2. Asmussen, S. and R. Y. Rubinstein (1992b). Performance Evaluation for the Score Function Method in Sensitivity Analysis and Stochastic Optimization. *International Workshop on Computer-Intensive Methods in Discrete Event Systems*, Vienna, 1990, (G. Pflug ed.). Springer-Verlag, 1–12.
3. Asmussen, S., R. Y. Rubinstein, and C. Wang (1994). Estimating Rare Events via Likelihood Ratios: From M/M/1 Queues to Bottleneck Networks, *Journal of Applied Probability*, **31**, 797–815.
4. Basar, T. (1987). Relaxation techniques and asynchronous algorithms for on-line computation of non-cooperative equilibria, *Journal of Economic Dynamics and Control*, **11**, 531–549.
5. Bertsekas, D. P. and J. N. Tsitsiklis (1989). *Parallel and distributed computation: Numerical methods*, Prentice-Hall.
6. Ermoliev, Y. M. (1983). Stochastic Quasigradient Methods and their Application to System Optimization, *Stochastics*, **9**, 1–36.
7. Ermoliev, Y. M. and Gaivoronski, A. A. (1992). Stochastic Quasigradient Methods for Optimization of Discrete Event Systems, *Annals of Operations Research*, **39**, 1–39.
8. Glasserman, P. (1991). *Gradient Estimation via Perturbation Analysis*, Kluwer Academic Press.
9. Glynn, P. W. and D. L. Iglehart (1989). Importance Sampling for Stochastic Simulations, *Management Science*, **35**, 11, 1367–1392.
10. Glynn, P. W. (1990). Likelihood Ratio Gradient Estimation for Stochastic Systems, *Communications of the ACM*, **33**, 10, 75–84.
11. Glynn, P. W., L'Ecuyer, P., and Adès, M. (1991). Gradient Estimation for Ratios, *Proceedings of the 1991 Winter Simulation Conference*, IEEE Press, 986–993.
12. Goldsman, D., Nelson, B., and Schmeiser, B. (1991). Methods for Selecting the best System, *Proceedings of the 1991 Winter Simulation Conference*, IEEE Press, 177–186.
13. Heidelberger, P., X.-R. Cao, M. A. Zazanis, and R. Suri (1988). "Convergence Properties of Infinitesimal Perturbation Analysis Estimates", *Management Science*, **34**, 11, 1281–1302.
14. Kleijnen, J. P. C. and Van Groenendaal, W. (1992). *Simulation: A Statistical Perspective*, Wiley, Chichester.
15. Kushner, H. J. and D. S. Clark (1978). *Stochastic Approximation Methods for Constrained and Unconstrained Systems*, Springer-Verlag, Applied Math. Sciences, Vol. 26.
16. Law, A. M. and Kelton, W. D. (1991). *Simulation Modeling and Analysis*, second edition, McGraw-Hill.
17. L'Ecuyer, P. (1990). A unified view of the IPA, SF, and LR gradient estimation techniques. *Management Science*, **36**, 1364–1384.
18. L'Ecuyer, P. (1992). Convergence rates for steady-state derivative estimator. *Annals of Operations Research*, **39**, 121–136.
19. L'Ecuyer, P. (1993). Two Approaches for Estimating the Gradient in Functional Form, *Proceedings of the 1993 Winter Simulation Conference*, IEEE Press, 338–346.
20. L'Ecuyer, P. and G. Perron (1994). On the Convergence Rates of IPA and FDC Derivative Estimators. *Operations Research* **42**, 643–656.

21. L'Ecuyer, P. and P. W. Glynn (1994). Stochastic Optimization by Simulation: Convergence Proofs for the GI/G/1 Queue in Steady-State. *Management Science* **40**, 1562–1578.
22. L'Ecuyer, P., N. Giroux, and P. W. Glynn (1994). Stochastic Optimization by Simulation: Numerical Experiments for the M/M/1 Queue in Steady-State. *Management Science* **40**, 1245–1261.
23. L'Ecuyer, P. and G. Yin (1994). Budget-Dependent Convergence Rate of Stochastic Approximation. Submitted.
24. Pflug, G. Ch. (1990). On-Line Optimization of Simulated Markov Processes. *Mathematics of Operations Research* **15**, 381–395.
25. Pflug, G. Ch. (1992). Gradient Estimates for the Performance of Markov Chains and Discrete Event Processes. *Annals of Operations Research* **39**, 173–194.
26. Plambeck, E. L., Fu, B.-R., Robinson, S. M., and Suri, R. (1993). Optimizing Performance Functions in Stochastic Systems. Submitted.
27. Polyak, B. T. and Juditsky, A. B. (1992). Acceleration of Stochastic Approximation by Averaging, *SIAM J. on Control and Optimization* **30**, 4, 838–855.
28. Reiman, M. I. and A. Weiss (1989). Sensitivity analysis for simulations via likelihood ratios, *Operations Research* **37**, 830–844.
29. Rubinstein, R. Y. (1976). A Monte Carlo method for estimating the gradient in a stochastic network. Unpublished manuscript, Technion, Haifa, Israel.
30. Rubinstein, R. Y. (1986). *Monte Carlo Optimization Simulation and Sensitivity of Queueing Network*, John Wiley & Sons, Inc., New York.
31. Rubinstein, R. Y. (1992). Monte Carlo Methods for performance evaluation, sensitivity analysis and optimization of stochastic systems, *Encyclopedia of Computer Science and Technology* (Kent ed.), Marcel Deker, Inc., Vol. 25, 211–233.
32. Rubinstein, R. Y. and A. Shapiro (1993). *Discrete Event Systems: Sensitivity Analysis and Stochastic Optimization via the Score Function Method*, John Wiley & Sons.
33. Rubinstein, R. Y. and S. Uryas'ev (1994). On Relaxation Algorithms in Computation of Non-Cooperative Equilibria, *IEEE Transactions on Automatic Control*, **AC-39**, 6, 1263–1268.
34. Suri, R. and M. A. Zazanis (1988). Perturbation analysis gives strongly consistent sensitivity estimates for the M/G/1 queue, *Management Science*, **34**, 1, 39–64.
35. Uryas'ev, S. P. (1992). A Stochastic Quasigradient Algorithm with Variable Metric, *Annals of Operations Research*, **39**, 251–267.
36. Wolff, R. (1989). *Stochastic Modeling and the Theory of Queues*, Prentice-Hall.
37. Yang, W.-N. and Nelson, B. L. (1991). Using Common Random Numbers and Control Variates in Multiple-Comparison Procedures, *Operations Research*, **39**, 4, 583–591.