SCHOOL OF OPERATIONS RESEARCH
AND INDUSTRIAL ENGINEERING
COLLEGE OF ENGINEERING
CORNELL UNIVERSITY
ITHACA, NEW YORK 14853

TECHNICAL REPORT NO. 778

February, 1988

REPLICATION SCHEMES
FOR LIMITING EXPECTATIONS

by

Bennett L. Fox¹

Peter W. Glynn²

- ¹ Department of Computer Science and Operations Research, University of Montreal. On leave at: School of Operations Research and Industrial Engineering, Upson Hall, Cornell University, Ithaca, New York 14853-7501. The research of this author was supported by the Natural Sciences and Engineering Research Council of Canada.
- ² Department of Operations Research, Stanford University. The research of this author was supported by the National Science Foundation under Grant ECS-8404809 and the U.S. Army under Contract No. DAAG29-80-C-0041.

Replication Schemes for Limiting Expectations

Bennett L. Fox¹

Peter W. Glynn²

February 1988

- 1 Department of Computer Science and Operations Research, University of Montreal. On leave at: School of Operations Research and Industrial Engineering, Upson Hall, Cornell University, Ithaca, New York 14853-7501. The research of this author was supported by the Natural Sciences and Engineering Research Council of Canada.
- 2 Department of Operations Research, Stanford University. The research of this author was supported by the National Science Foundation under Grant ECS-8404809 and the U.S. Army under Contract No. DAAG29-80-C-0041.

Abstract

We show that natural estimators occurring in certain simulation settings have convergence rates less than the canonical rate usually associated with simulation. These natural estimators use replication schemes that attenuate bias. For some important examples we find alternative estimators that converge at the canonical rate. The implications of these asymptotic comparisons for choosing good strategies when the computer-time budget is modest are discussed.

Key words:

simulation, replication schemes, convergence rates, efficiency, Markov chains, gradient estimation.

AMS (MOS) subject classification:

90B99

0. Introduction

Typically, application of the Monte Carlo method leads to estimators which converge at rate $t^{-1/2}$ in the sampling effort t . This convergence rate arises as a result of the $t^{1/2}$ scaling factor which appears in the central limit theorem for sample means of iid random variables. Recently, however, several problems have been analyzed in which the *natural* Monte Carlo sampling scheme has led to a subcanonical convergence rate (i.e, slower than $t^{-1/2}$). Zazanis and Suri (1986) and Frolov and Chentsov (1963) showed that when the derivative $\alpha'(\theta_0)$ of a function $\alpha(\bullet) = EZ(\bullet)$ is evaluated via Monte Carlo simulation of a forward difference approximation to $\alpha'(\theta_0)$, the root mean square error (rmse) converges at rate $t^{-1/4}$. A second subcanonical rate appeared in recent work of Fox and Glynn (1987). The problem there involved estimating $\alpha = \int_0^\infty e^{-t} EZ(\bullet) dt$; if α is estimated by truncating the integral and using Monte Carlo on the resulting quantity, the best possible convergence rate is of order $((\log t)/t)^{1/2}$.

In this paper, we develop a systematic framework, in which the above subcanonical convergence rates are just two manifestations of our general theory. Thus, a major contribution of this paper is to show that the above disparate problems are just two sides of the same coin.

The generic problem that we shall consider goes as follows. Let

$$(0.1) \quad \alpha = \lim_{\tau \rightarrow \infty} \alpha(\tau)$$

where $\alpha(\tau) = EX(\tau)$. We do not assume that there is a limiting r.v. $X(\infty)$. In fact, the parameterized family $X(\tau)$ may simply be contrived so that holds. Our goal here is to estimate the limiting expectation α via simulation. Section 1 gives important examples of this class of problems, and shows that the finite-difference and discounting problems described above are special cases.

A straightforward class of *replication schemes* for (0.1) follows. Given a computer-time budget t , we decide how much of this time to allot per run via a *sampling plan* $\beta(t)$ mapping t into τ . We assume that each run takes $w(\beta(t))$ time. Our estimator of α is

$$(0.2) \quad a(t) = \begin{cases} \frac{1}{n(t)} [X_1(\beta(t)) + \cdots + X_{n(t)}(\beta(t))]; & n(t) \geq 1 \\ 0; & n(t) = 0 \end{cases}$$

where $X_i(\beta(t))$ is the output on run i and the number $n(t)$ of (iid) runs completed by time t is clearly

$$(0.3) \quad n(t) = \lfloor t/w(\beta(t)) \rfloor.$$

We consider what happens when $EX(\tau) \neq \alpha$ and $X(\tau) \not\Rightarrow \alpha$ but $EX(\tau) \rightarrow \alpha$ as $\tau \rightarrow \infty$. The sampling plan determines the tradeoff between the bias of $a(t)$ and its variance.

How good is this class of estimators? To answer this question we define the *canonical* convergence rate of an estimator $Z(t)$ of α based on t units of simulation effort:

- (i) the root mean square error (*rmse*) of $Z(t)$ is $O(t^{-1/2})$
- (ii) $\left\{ \begin{array}{l} \text{[rough]} Z(t) \text{ converges weakly to } \alpha \text{ at rate } t^{-1/2} \\ \text{[precise]} \text{ the family } \{t^{1/2}[Z(t) - \alpha] : t > 0\} \text{ is stochastically bounded, i.e., } \forall \varepsilon > 0 \text{ there} \\ \text{exists } K(\varepsilon) \text{ s.t. } P\{t^{1/2} |Z(t) - \alpha| > K(\varepsilon)\} < \varepsilon \text{ uniformly in } t. \end{array} \right.$

An estimator is *rmse efficient* if it converges at the canonical rate in the sense of (i) above. It is *distribution efficient* if it converges at the canonical rate in the sense of (ii) above. By Chebyshev's inequality, rmse efficiency implies distribution efficiency. If an estimator converges strictly slower than the canonical rate, it is *inefficient*. Section 2 precisely states and refines the following assertion: under mild hypotheses, all estimators of the form $a(t)$ are inefficient.

Since

$$(0.4) \quad E[a(t) - \alpha]^2 = \frac{\sigma^2(\beta(t))}{n(t)} + b^2(\beta(t))$$

and

$$(0.5) \quad n(t) \leq t/w(\beta(t)),$$

we get

$$(0.6) \quad tE[a(t) - \alpha]^2 \geq w(\beta(t))\sigma^2(\beta(t)).$$

In Theorem 1, we show that a subset of our hypotheses implies that $a(t) \Rightarrow \alpha$ forces $w(\beta(t))\sigma^2(\beta(t))$ to infinity and therefore $a(t)$ is rmse inefficient. Theorem 2 shows that, with a different normalization $g(t) = o(t^{1/2})$ specified by equation (2.2.1), we can get a nontrivial limit and specifies this limit distribution. Theorem 3 refines Theorem 1 by showing that all our hypotheses together imply distribution inefficiency.

In section 3 we return to the examples of section 1. For each, we provide mild conditions implying the hypotheses of section 2. We also give or cite *efficient* alternative estimators. Typically, these alternatives are also based on multiple replications, but — in contrast to what we assume about $a(t)$ — the output from each run is an unbiased estimator of α . In some, but not all, examples we find sampling plans such that $a(t)$ converges at (subcanonical) rates arbitrarily close to the canonical rate. Some efficient alternatives mentioned are valid only under hypotheses stronger than required for $a(t)$ to converge to α . For finite-difference gradient estimators, the *exact* (subcanonical) convergence rates and the limit distribution in section 3.2 fall quickly out of our general theory; only (weaker) mean-square-error counterparts were previously known.

Section 4 proves our major results.

The theoretical results of this paper *prove* that efficient estimators beat biased estimators of the form $a(t)$ for all large enough t . They suggest that, if an efficient estimator beats $a(t)$ for a given t , it also wins for all larger t . They also show that, if it loses to $a(t)$ for a given t (say in empirical comparison), extrapolation of this ranking to all larger t is erroneous. Dealing systematically with moderate values of t is outside the scope of this paper and would require more sophisticated mathematical tools. Among all efficient estimators, it is generally impossible to prove that any given one is best: their respective variance constants can differ greatly.

1. Examples

In all our examples, one expects that usually the larger $\beta(t)$, the smaller the bias of $a(t)$ but the larger its variance.

1.1. Estimation in finite-horizon, continuous-time Markov chains

We estimate expected terminal reward $\alpha = Ef(Z(T))$, where $Z = \{Z(t): t \geq 0\}$ is a uniformizable, continuous-time Markov chain, f is a real-valued function defined on the state space of Z , and T is a positive constant. Because of the uniformizability, we can represent Z as $Z(t) = Y(N(\theta t))$ where $Y = \{Y(n): n \geq 0\}$ is a discrete-time Markov chain independent of the unit-rate Poisson process N and θ is at least the supremum of the jump rates of Z ; e.g., see Çinlar [(1975), pages 236-237].

Let

$$(1.1.1) \quad h(Y) = E[f(Z(T)) \mid Y] \\ = \sum_{k=0}^{\infty} f(Y(k))p(k)$$

where

$$(1.1.2) \quad p(k) = e^{-\theta T} (\theta T)^k / k! .$$

Since taking conditional expectations reduces variance (e.g., see Bratley, Fox, and Schrage [(1987), §2.6]), $\text{Var } h(Y) \leq \text{Var } f(Z(T))$.

We can't simulate the infinite summation in (1.1.1) exactly. So we truncate it to get

$$(1.1.3) \quad X(\tau) = \sum_{k=0}^{\lfloor \tau \rfloor} f(Y(k))p(k).$$

If $E |f(Z(T))| < \infty$, then $\alpha(\tau) = EX(\tau) \rightarrow \alpha$ as $\tau \rightarrow \infty$. This fits our framework (0.1). Here $\beta(t)$ is the truncation point.

Gross, Miller, and Plastiras (1984) study, among other things, various strategies for implementing (1.1.3). For improvements in efficiency, see Fox and Glynn (1988a). The set of nonnegligible $p(k)$'s needed in (1.1.3) can be computed efficiently in $O(\sqrt{\theta T})$ time using the scheme of Fox and Glynn (1988b).

1.2. Gradient estimation

With $q(\xi) = EZ(\xi)$, we estimate

$$(1.2.1) \quad \alpha = q'(0).$$

If ξ is a controllable input parameter for a simulation, then an estimate of α is useful for sensitivity analysis and for optimizing the expected output $q(\xi)$.

We start by considering two estimators of α :

$$(1.2.2) \quad X_1(\tau) = \tau(Z_1(1/\tau) - Z_2(0)) \quad (\text{forward difference})$$

$$(1.2.3) \quad X_2(\tau) = \tau(Z_1(1/2\tau) - Z_2(-1/2\tau)) \quad (\text{central difference})$$

where $Z_1(\bullet)$ and $Z_2(\bullet)$ are independent copies of Z . If $E|Z(\xi)| < \infty$ for all $\xi \in (-1, 1)$ say and $q(\xi)$ is continuously differentiable at $\xi = 0$, then $\alpha_i(\tau) = EX_i(\tau) \rightarrow \alpha$ as $\tau \rightarrow \infty$. Thus, this problem also fits our framework. Here no limiting random variable $X(\infty)$ exists in general, the case when $Z(\xi)$ is deterministic being an obvious exception. Now $\beta(t)$ is the reciprocal of the increment.

1.3. Discounting generalized

Let Z be a real-valued stochastic process and G be a nondecreasing deterministic function on $[0, \infty)$. Set

$$(1.3.1) \quad \alpha = \int_{[0, \infty)} EZ(t)G(dt).$$

The special case $G(dt) = e^{-\xi t} dt$ corresponds to discounting. We get another special case when $G(dt)$ is Poisson measure with parameter λT ; if in addition $Z(t) = f(Y(\lfloor t \rfloor))$ using the notation of section 1.1 on the right, we recover $Eh(Y)$ with $h(Y)$ defined by (1.1.1).

With

$$(1.3.2) \quad X(\tau) = \int_{[0, \tau]} Z(s)G(ds)$$

and $\alpha(\tau) = EX(\tau)$, we again fit our framework (0.1).

1.4. Steady-state estimation for Markov chains

Let Z be a discrete-time Markov chain with a countable state space. Assume that from any starting state we hit a recurrent state in finite time with probability one. Given a real-valued function f defined on the state space of Z , let

$$(1.4.1) \quad X(\tau) = [1/(\lfloor \tau \rfloor + 1)] \sum_{k=0}^{\lfloor \tau \rfloor} f(Z(k)).$$

Under appropriate integrability conditions, there is a random variable $X(\infty)$ such that $X(\tau) \rightarrow X(\infty)$ *a.s.* as $\tau \rightarrow \infty$. The goal here is to estimate $\alpha = EX(\infty)$.

When we can get to at least two recurrence classes from the starting state, then $X(\infty)$ is nondegenerate (which hypothesis H5 of section 2.1 excludes but which generally occurs when there is just one recurrence class). Glynn (1987b) covers the case where $X(\infty)$ is constant except for special f , while section 3.4 covers the case where $X(\infty)$ is nondegenerate. With $\alpha(\tau) = EX(\tau)$, we again fit our framework.

2. Convergence rates

In section 2.1 we show that biased replication estimators are inefficient in the sense of root mean square error. We extend this to weak-convergence inefficiency in section 2.2. If the respective orders of magnitude of the variance, bias, and work as a function of τ are known, then we can get the (subcanonical) convergence rate as a function of the sampling plan β when the latter has a convenient parametric form. Section 2.3 illustrates how this guides us to good sampling plans. When the bias is exponentially decreasing in τ , the results of section 2.3 lead to sampling plans giving convergence rates arbitrarily close to the canonical rate as section 2.4 points out.

2.1. Inefficiency

Let

$$(2.1.1) \quad \sigma^2(\tau) = \text{Var } X(\tau)$$

$$(2.1.2) \quad b(\tau) = \alpha - EX(\tau).$$

Below we list hypotheses on $\sigma^2(\tau)$, $b(\tau)$, and $w(\tau)$, not all of which need be simultaneously in force. For the (diverse) examples of section 1, all these hypotheses often hold, as section 3 shows precisely. With each of our results, we show explicitly the subset of hypotheses used.

H1 [finite variance]: $\sup\{\sigma^2(s) : 0 \leq s \leq \tau\} < \infty$ for all τ

- H2 [positive variance]: there exists τ_0 , such that $\inf\{\sigma^2(s): \tau_0 \leq s \leq \tau\} > 0$ for all $\tau \geq \tau_0$
- H3 [positive squared bias]: $\inf\{b^2(s): 0 \leq s \leq \tau\} > 0$ for all τ
- H4 [asymptotic unbiasedness]: $EX(\tau) \rightarrow \alpha$
- H5 [nondegeneracy]: there exists no sequence $\tau_n \rightarrow \infty$ such that $X(\tau_n) \Rightarrow \alpha$ as $n \rightarrow \infty$
- H6 [getting rid of bias by letting $\tau \rightarrow \infty$ is costly]: $w(\tau)\sigma^2(\tau) \rightarrow \infty$ as $\tau \rightarrow \infty$
- H7 [uniform integrability]: there exists τ_0 such that $\{X^2(\tau)/\sigma^2(\tau): \tau > \tau_0\}$ is uniformly integrable
- H8 [nonrandom work]: $w(\tau)$ is deterministic with $\inf\{w(\tau): \tau \geq 0\} > 0$.

Remarks on these hypotheses:

- R1. H1-H4 and H7 seem reasonable for this problem setting.
- R2. H2, H4, and H5 together imply that $\inf\{\sigma^2(s): s \geq \tau_0\} > 0$. This follows from H2 after showing that $\sigma^2(\tau) \not\rightarrow 0$ as $\tau \rightarrow \infty$. Aiming for a contradiction, suppose that $\sigma^2(\tau_n) \rightarrow 0$ along some sequence $\tau_n \rightarrow \infty$. Then H4 shows that $E[X(\tau_n) - \alpha]^2 \rightarrow 0$. In turn, this implies that $X(\tau_n) \Rightarrow \alpha$, contradicting H5.
- R3. H1, H3, and H5 jointly imply the nonexistence of a sequence τ_n such that $X(\tau_n) \Rightarrow \alpha$, whether or not $\tau_n \rightarrow \infty$. Aiming for a contradiction, suppose $X(\tau_n) \Rightarrow \alpha$. If $\sup \tau_n = \infty$, we can extract a subsequence τ_{n_l} going to ∞ with $X(\tau_{n_l}) \Rightarrow \alpha$, contradicting H5. If $\sup \tau_n < \infty$, there exists a subsequence τ_{n_l} such that $\tau_{n_l} \rightarrow \tau < \infty$ and $X(\tau_{n_l}) \Rightarrow \alpha$. But by H1 we get $\sup\{EX^2(\tau_{n_l})\} \leq \sup\{EX^2(s): 0 \leq s \leq \sup \tau_n\} < \infty$, so $\{X(\tau_{n_l})\}$ is uniformly integrable. Hence $EX(\tau_{n_l}) \rightarrow \alpha$, contradicting H3.
- R4. To understand H6, observe that R3, H6, and H8 jointly imply that $w(\tau) \rightarrow \infty$ or $\sigma^2(\tau) \rightarrow \infty$ or both.
- R5. Assuming as in H8 that $w(\tau)$ is deterministic oversimplifies most real simulations but significantly simplifies mathematical development. A relaxation of this assumption will be considered elsewhere.

R6. Hypothesis H5 excludes most single-chain steady-state simulations but allows almost all multichain simulations.

Proposition 1 (part 2) says that we want $\beta(t) \rightarrow \infty$, but proposition 2 says not too fast.

Proposition 1.

P1.1 [weak-convergence consistency]: Assume H1, H2, H4, H5, and H8. If $\beta(t) \rightarrow \infty$ and $w(\beta(t))\sigma^2(\beta(t))/t \rightarrow 0$, then

$$a(t) \Rightarrow \alpha$$

P1.2 [partial converse]: Assume H1, H2, H3, and H8. If $a(t) \Rightarrow \alpha$ and $\alpha \neq 0$, then

$$\beta(t) \rightarrow \infty.$$

Proposition 2. Assume H1-H5.

P2.1 [no free lunch]: If $a(t) \Rightarrow \alpha$ and $\alpha \neq 0$, then

$$n(t) \rightarrow \infty.$$

Theorem 1 [root-mean-square-error inefficiency]. Assume H1-H3, H6, and H8. If $a(t) \Rightarrow \alpha$ and $\alpha \neq 0$, then $a(t)$ is inefficient:

$$(2.1.3) \quad tE[a(t) - \alpha]^2 \rightarrow \infty.$$

2.2. Extended weak limit points

Let $\{V(t): t \geq 0\}$ be a family of random variables (not necessarily defined on the same probability space). Let S be a random variable taking values in the extended real line $[-\infty, \infty]$. If there exists a sequence $t_n \rightarrow \infty$ such that $V(t_n) \Rightarrow S$, we call S an EWLP (extended weak limit point) of V . Here weak convergence is relative to the standard topology on the compactified space $[-\infty, \infty]$. Such convergence reduces to standard weak convergence on the real line if S is finite with probability one; if $S = \infty$ a.s., extended weak convergence to S is equivalent to $P\{V(t) \leq x\} \rightarrow 0$ a.s. $t \rightarrow \infty$ for all (finite) x . We require our extended notion of weak convergence because our S may have mass at $-\infty$ or $+\infty$, which standard weak convergence forbids.

Theorem 2. Assume H1-H5 and H7-H8. Let

$$(2.2.1) \quad g(t) = [t/w(\beta(t))\sigma^2(\beta(t))]^{1/2}$$

$$(2.2.2) \quad \Gamma(t) = g(t)[a(t) - \alpha].$$

If $a(t) \Rightarrow \alpha \neq 0$, then the only possible EWLPs of $\{\Gamma(t): t > 0\}$ have the form $N(0, 1) + \gamma$ with $\gamma \in [-\infty, \infty]$. The sequence $\{\Gamma(t_n): n \geq 1\}$ converges weakly if and only if

$$(2.2.3) \quad g(t_n)b(\beta(t_n)) \rightarrow \gamma$$

for some $\gamma \in [-\infty, \infty]$ in which case

$$(2.2.4) \quad \Gamma(t_n) \Rightarrow N(0, 1) + \gamma.$$

We can therefore strengthen the root-mean-square-error-inefficiency conclusion of theorem 1 to distribution inefficiency when H4 and H7 are also assumed:

Theorem 3. Assume H1-H8. The only possible EWLPs of $\{t^{1/2}[a(t) - \alpha]: t > 0\}$ are $-\infty$ and $+\infty$.

2.3. Sampling-plan guidelines

Our results in sections 2.1 and 2.2 show that, when the computer-time budget is *large enough*, efficient alternative estimators (if available) beat all estimators of the form $a(t)$, no matter what the sampling plan. For every problem in section 1, we give or cite an efficient alternative estimator in section 3. Detailed case-by-case studies, theoretical or empirical, are needed to quantify what computer-time budgets are "large enough". Any such theoretical study would have to compare not only the convergence rates of the competing estimators but also the respective variance constants and values of γ , corresponding to the sampling plans considered.

To illustrate how to narrow the scope of sampling plans to consider, suppose that there are finite, nonnegative constants δ , ρ , and λ and positive numbers σ^2 , b , and w such that

$$(2.3.1) \quad \tau^{-\delta}\sigma^2(\tau) \rightarrow \sigma^2$$

$$(2.3.2) \quad \tau^\rho |b(\tau)| \rightarrow b$$

$$(2.3.3) \quad \tau^{-\lambda}w(\tau) \rightarrow w.$$

By H3, we get $\rho > 0$. The conditions (2.3.1)-(2.3.3) hold in many settings, including several of our examples as we show in section 3.

In (2.2.4) we clearly want $|\gamma| < \infty$, so we impose

$$(2.3.4) \quad \limsup_{t \rightarrow \infty} |g(t)b(\beta(t))| < \infty \quad (2.3.4)$$

to define the class of admissible sampling plans β . Routine calculations show that (2.3.1)-(2.3.4) jointly imply

$$(2.3.5) \quad \liminf \beta(t)/t^{1/(\delta+\lambda+2\rho)} > 0.$$

Sampling plans of the form

$$(2.3.6) \quad \beta(t) \sim ct^{1/(\delta+\lambda+2\rho)}$$

with $0 < c < \infty$ satisfy (2.3.5) and give

$$(2.3.7) \quad g(t) \sim [t^{\rho/(\delta+\lambda+2\rho)}] \bullet (w\sigma^2 c^{\delta+\lambda})^{-1/2}$$

$$(2.3.8) \quad \gamma = (w\sigma^2 c^{\delta+\lambda})^{-1/2} bc^{-\rho}$$

where the bracketed factor in (2.3.7) is the maximal (subcanonical) convergence rate obtainable with estimators of the form $a(t)$. Now $[t^{\rho/(\delta+\lambda+2\rho)}] \bullet [a(t) - \alpha]$ converges to a limit random variable $(w\sigma^2 c^{\delta+\lambda})^{1/2} N(0, 1) + bc^{-\rho}$. Minimizing its second moment $w\sigma^2 c^{\delta+\lambda} + b^2 c^{-2\rho}$ over c gives

$$(2.3.9) \quad c^* = \left[\frac{2\rho b^2}{w\sigma^2(\lambda + \delta)} \right]^{1/(\lambda+\delta+2\rho)}$$

as the value of c minimizing the asymptotic mean square error of $a(t)$, under appropriate uniform integrability conditions.

An alternative to minimizing the asymptotic mean square error of $a(t)$ is to minimize the asymptotic mean absolute deviation $E|a(t) - \alpha|$. When $\gamma \neq 0$, we generally get a different value of c^* . This can't happen with unbiased estimators, because γ vanishes in that case. Our (more refined) Theorem 2 allows construction of confidence intervals or determination of estimators minimizing an L_p -norm of the limit distribution for general p . Under uniform-integrability hypotheses, this is equivalent to minimizing the L_p -norm of the estimator's error.

Generally, the parameters w , σ , and b are unknown. While w and σ can easily be estimated consistently, usually only an *a priori* upper bound on b is available and in such cases (2.3.9) has to be adjusted heuristically. We show in section 3.2 that, for finite differences, b can be estimated consistently; so, in that case, (2.3.9) is fully operational.

When b can't be consistently estimated but $\bar{b}$ is a known upper bound, one way to adjust (2.3.8) and (2.3.9) heuristically is to use $\bar{b}$ in place of b .

An alternative is to use a class of sampling plans that gives a limit distribution independent of b , although with a convergence rate epsilon-tilde less. Thus, instead of (2.3.6), we can take

$$(2.3.10) \quad \beta(t) = ct^{1/(\delta+\lambda+2\rho) + 2\varepsilon/(\delta+\lambda)}$$

with $0 < \varepsilon$, $c < \infty$. This gives

$$(2.3.11) \quad t^{[\rho/(\delta+\lambda+2\rho)] - \varepsilon} [a(t) - \alpha] \Rightarrow (w\sigma^2 c^{\delta+\lambda})^{1/2} N(0, 1)$$

where the former term $bc^{-\rho}$ no longer appears. With z_ξ solving $P\{N(0, 1) \leq z_\xi\} = 1 - \xi/2$ and $\hat{\sigma}(t)$ a consistent estimator of σ , we get

$$\left[a(t) - z_\xi \frac{(w(t)\hat{\sigma}^2(t)c^{\delta+\lambda})^{1/2}}{t^{[\rho/(\delta+\lambda+2\rho)] - \varepsilon}}, a(t) + z_\xi \frac{(w(t)\hat{\sigma}^2(t)c^{\delta+\lambda})^{1/2}}{t^{[\rho/(\delta+\lambda+2\rho)] - \varepsilon}} \right]$$

as an asymptotic $100(1 - \xi)\%$ confidence interval for α . In (2.3.11), there is no optimal value for c ; we take c "small". Unless t is very large, a "small" c may lead to an unreasonably small value of $\beta(t)$. Adding a (judiciously-chosen) positive constant to the right side of (2.3.10) has no effect on (2.3.11), but may substantially improve performance for moderate values of t .

2.4. Exponentially-decreasing bias

If $b(\tau) = o(\tau^{-\rho})$ for all $\rho > 0$ (as occurs when $|b(t)|$ decreases exponentially fast), then by picking ρ in (2.3.6) or in (2.3.10) arbitrarily large, with ε arbitrarily small, we can get arbitrarily close to the canonical rate of convergence and get $\gamma = 0$ too. A different functional form of $\beta(t)$, say $\beta(t) = c \log t$, such that $g^2(t)b^2(\beta(t)) \rightarrow \gamma^2 > 0$ sometimes leads to slightly better, though still subcanonical, convergence rates — but with $\gamma \neq 0$. Some of our examples

have exponentially-decreasing bias, as we detail in section 3. A thorough discussion of sampling plans for each of these cases would lead too far afield. For large enough t , the efficient alternative estimators given or cited for each case win. Fine differences in asymptotic growth rates of $\beta(t)$ would be swamped in practice by heuristic adjustments to $\beta(t)$ not affecting asymptotics but needed to make $\beta(t)$ reasonable for small to moderate t , the only case for which biased multiple-replication estimators can be possibly competitive.

3. Examples revisited

We return to the examples considered in section 1 in the corresponding subsections here. In each case we give or cite efficient alternative estimators and find easily-checked conditions implying H1-H8 and $\alpha \neq 0$ for biased multiple-run estimators.

3.1. Estimation in finite-horizon, continuous-time Markov chains

The naive strategy takes the average of $n(t)$ iid copies of $f(Z(T))$, where $n(t)$ is the number of runs completed in computer time t . If W is the (possibly random) computer time to generate a copy, then Glynn and Whitt (1986) and Fox and Glynn (1988a) show (as a special case) that, with i indexing runs,

$$(3.1.1) \quad t^{1/2} \left[\frac{1}{n(t)} \sum_{i=1}^{n(t)} f(Z_i(T)) - \alpha \right] \Rightarrow \left[EW \bullet \text{Var } f(Z(T)) \right]^{1/2} N(0, 1)$$

and thus the canonical rate is obtained. In general, no explicit formula for EW is available but sometimes [see Fox (1987)] the order of magnitude of its worst-case counterpart can be given. Fox and Glynn (1988a) obtain several other efficient estimators, all with lower asymptotic variance constant (= expected work per run times variance per run) than in (3.1.1).

An alternative averages iid copies of (1.1.3) where $\tau = \beta(t)$. In this setting our hypotheses H1-H7 seem reasonable, while the assumed deterministic work in H8 can be taken as a first approximation. Section 4.6 shows that the following (easier-to-check) conditions imply H1-H8:

A1.1: Z is recurrent and uniformizable

A1.2: Y is aperiodic (in particular, not deterministic)

A1.3: $f \geq 0$

A1.4: $0 < \text{Var } f(Z(T)) < \infty$

A1.5: The work to generate $X(\tau)$ is $c(\lfloor \tau \rfloor + 1)$ with the constant c positive.

From A1.5 we get $\lambda = 1$ in (2.3.3). Recalling that in this example τ corresponds to a truncation point, we get $\delta = 0$ in (2.3.1).

Since Y is an aperiodic recurrent chain, generally $Ef(Y(n)) \rightarrow \pi f$ (the steady-state mean with π the stationary probability vector for Y). Using Glynn's (1987c) results on Poisson tail probabilities, we get

$$(3.1.2) \quad b(\tau) \sim \{\pi f e^{-\theta T} (\theta T)^{\lfloor \tau \rfloor} / \lfloor \tau \rfloor!\}$$

where πf often can be computed analytically or numerically. From (3.1.2) we see that the bias goes to zero exponentially fast, so from section 2.4 we can choose a sampling plan to get arbitrarily close to the canonical convergence rate.

3.2. Gradient estimation

For estimating gradients, Zazanis and Suri (1986) point out that "perturbation analysis", *when valid*, has root mean square error $O(t^{-1/2})$ — beating forward differences and central differences. The latter two, however, are valid under much weaker conditions. Heidelberger *et al* (1987) point out certain classes of problems where perturbation analysis gives strongly consistent gradient estimates and other classes where it does not. Glynn's (1986b) method of likelihood ratios for estimating gradients has root mean square error $O(t^{-1/2})$; this method is described in Bratley, Fox, and Schrage [(1987), §2.5]. Neither of the respective domains of validity of likelihood ratios and of perturbation analysis contains the other; they overlap.

For large enough computer-time budgets, finite-difference estimators should be considered only in cases where both perturbation-analysis estimators and likelihood-ratio-based estimators are invalid. The (weak) conditions A2.1-A2.5 below imply that H1-H8 hold, as easily checked, and hence that finite-difference estimators are inefficient:

A2.1: $Z(\xi) \Rightarrow Z(0)$ as $\xi \rightarrow 0$

A2.2: $q(\xi)$ has a continuous third derivative over $[-1, 1]$ [recall: $q(\xi) = EZ(\xi)$]

A2.3: $q^{(2)}(\bullet)$ and $q^{(3)}(\bullet)$ do not vanish anywhere in $[-1, 1]$
[superscript denotes derivative order]

A2.4: $v(\xi) \rightarrow v(0) \neq 0$ as $\xi \rightarrow 0$ where $v(\xi) = \text{Var } Z(\xi)$

A2.5: The work $d(\xi)$ required to generate $Z(\xi)$ is deterministic, positive, and continuous on $[-1, 1]$.

For the forward-difference estimator,

$$(3.2.1) \quad \tau^{-2} \text{Var } X_1(\tau) \rightarrow 2v(0)$$

$$(3.2.2) \quad \tau b_1(\tau) \rightarrow \alpha^{(2)}(0)/2$$

$$(3.2.3) \quad w_1(\tau) \rightarrow 2d(0).$$

For the central-difference estimator,

$$(3.2.4) \quad \tau^{-2} \text{Var } X_2(\tau) \rightarrow 2v(0)$$

$$(3.2.5) \quad \tau^2 b_2(\tau) \rightarrow \alpha^{(3)}(0)/24$$

$$(3.2.6) \quad w_2(\tau) \rightarrow 2d(0).$$

In deterministic settings, the counterparts of (3.2.2) and (3.2.5) are well-known in numerical analysis. Here we obtain them in the same way, via Taylor-like expansions.

In the notation of section 2.3, we have $w = 2d(0)$, $\sigma^2 = 2v(0)$, $\lambda = 0$, and $\delta = 2$. For forward differences, we have $\rho = 1$ and $b = \alpha^{(2)}(0)/2$. For central differences, we have $\rho = 2$ and $b = \alpha^{(3)}(0)/24$. This leads to the (asymptotically) optimal choices

$$(3.2.7) \quad \beta_1(t) = c_1^* t^{1/3}$$

$$(3.2.8) \quad c_1^* = \left[[\alpha^{(2)}(0)]^2 / 16d(0)v(0) \right]^{1/3}$$

$$(3.2.9) \quad \beta_2(t) = c_2^* t^{1/6}$$

$$(3.2.10) \quad c_2^* = \left[[\alpha^{(3)}(0)/24]^2 / 2d(0)v(0) \right]^{1/6}$$

with convergence rates $O(t^{1/4})$ and $O(t^{1/3})$ respectively. By other methods, Zazanis and Suri (1986) and Frolov and Chentsov (1963) find root-mean-square-error counterparts of these rates.

We can estimate the right side of (3.2.2) over τ by the central-difference estimate minus the forward-difference estimate because the bias in the former is an order of magnitude less.

We can estimate the right side of (3.2.5) similarly using an estimate based on Richardson extrapolation with $O(\tau^{-4})$ bias. More generally, we can find a finite-difference formula for gradient estimation with $O(\tau^{-2k})$ bias for $k = 1, 2, 3, \dots$ as shown in Johnson and Riess [(1982), pages 320-322] for example. We can thus get convergence rates arbitrarily close to the canonical rate, but very high-order extrapolation is not practical due to severe roundoff error and to the large number of function evaluations (growing exponentially with the order k).

To estimate an s -dimensional gradient using forward differences requires sets of $s + 1$ simulation runs whereas central differences require sets of $2s$ runs.

3.3. Discounting generalized

Fox and Glynn (1987) and Glynn (1988b) give efficient estimators for generalized discounting.

An alternative uses truncation as in (1.3.2) and averages this (biased) estimator over multiple iid runs. The (weak) conditions

$$\text{A3.1: } \int_0^\infty \int_0^\infty E |Z(s)Z(t)| G(ds)G(dt) < \infty$$

$$\text{A3.2: } \text{Var } X(\infty) > 0$$

$$\text{A3.3: } b(\tau) \neq 0 \text{ for all } \tau$$

$$\text{A3.4: } w(\tau) = c\tau \text{ with } c > 0$$

jointly imply H1-H8, since

$$|X(\tau)| \leq I \stackrel{\text{def.}}{=} \int_0^\infty |Z(s)| G(ds)$$

and $EI^2 < \infty$ (by A3.1) together imply that $\{X^2(\tau); \tau > 0\}$ is uniformly integrable.

In the notation of section 2.3, we get $\delta = 0$, $\sigma^2 = \text{Var } X(\infty)$, $\lambda = 1$, and $w = c$. If (2.3.2) holds with $b \neq 0$, we pick β as in section 2.3. If in (2.3.2) we get $b = 0$ for all $p > 0$, then the bias decreases exponentially fast and we pick β as in section 2.4.

For the case of ordinary continuous discounting, $G(ds) = e^{-\xi s} ds$ with $\xi > 0$, we get

$$(3.3.1) \quad b(t) \sim EZ(\infty)e^{-\xi t}/\xi$$

and so $|b(t)|$ decreases exponentially fast.

3.4. Steady-state estimation for Markov chains

If we know explicitly the recurrence classes, say $C_1, \dots, C_n$, we obtain an efficient estimator as follows. Split the computer-time budget t into $n + 1$ (say equal) parts to estimate respectively the steady-state average reward in $C_1, \dots, C_n$ and $\{P_1, \dots, P_n\}$ where P_i is the probability of hitting C_i starting from the given initial state. Since each part has convergence rate $O(t^{-1/2})$, it follows (easily) that the global simulation also does.

An alternative averages iid copies of the biased estimator $X(\tau)$ given by (1.4.1). The conditions

$$\text{A4.1: } Z(0) = i \text{ a.s.}$$

$$\text{A4.2: } 0 < \text{Var } X(\infty) < \infty$$

$$\text{A4.3: } Ef^2(Z(n)) \rightarrow Ef^2(Z(\infty)) < \infty \text{ as } n \rightarrow \infty$$

$$\text{A4.4: } b(n) \neq 0 \text{ for all } n$$

imply H1-H7. Condition A4.2 excludes (typical) steady-state simulations with just one recurrence class. Uniform integrability is the only non-trivial implication. But Cauchy-Schwartz implies

$$(3.4.1) \quad X^2(\tau) \leq \frac{1}{[\tau] + 1} \sum_{k=0}^{[\tau]} f^2(Z(k))$$

and the right-hand side is uniformly integrable by A4.3. We also assume H8; typically, $\lambda = 1$ in (2.3.3). From A4.3, we get $\delta = 0$ in (2.3.1).

Suppose that convergence to steady state occurs exponentially fast in the sense that

$$(3.4.2) \quad Ef(Z(k)) = Ef(Z(\infty)) + O(q^k)$$

with $0 < q < 1$. Summing from $k = 0$ to n , dividing by $n + 1$, and comparing with (1.4.1), we see that $b(n) = O(1/n)$. This gives $\rho = 1$ in (2.3.2). Glynn's (1986b) results allow estimation of b with low bias. Corresponding to (2.3.11), we get $O(t^{(-1/3)+\epsilon})$ convergence rate with $\gamma = 0$.

4. Proofs

4.1. Proof of Proposition 1

For the first part

$$(4.1.1) \quad E[a(t) - \alpha]^2 = \frac{\sigma^2(\beta(t))}{n(t)} + b^2(\beta(t)).$$

From (0.3) we get

$$(4.1.2) \quad n(t) \geq t/w(\beta(t)) - 1.$$

So

$$(4.1.3) \quad E[a(t) - \alpha]^2 \leq \frac{w(\beta(t))\sigma^2(\beta(t))}{t[1 - w(\beta(t))/t]} + b^2(\beta(t)).$$

By R3, from H1-H5 we get

$$(4.1.4) \quad \inf\{\sigma^2(\tau): \tau > \tau_0\} > 0.$$

From (4.1.4) and (the hypothesis of P1.1) $w(\beta(t))\sigma^2(\beta(t))/t \rightarrow 0$, we get

$$(4.1.5) \quad w(\beta(t))/t \rightarrow 0.$$

From (the hypothesis of P1.1) $\beta(t) \rightarrow 0$ and H4,

$$(4.1.6) \quad b^2(\beta(t)) \rightarrow 0.$$

Combining (4.1.3)-(4.1.6) proves

$$(4.1.7) \quad E[a(t) - \alpha]^2 \rightarrow 0$$

and P1.1 follows.

For the converse, assume that there exists a sequence $t_k \rightarrow \infty$ such that $\beta(t_k) \rightarrow \beta < \infty$. This holds if, contrary to P1.2, $\beta(t) \nrightarrow \infty$ since $\beta(t)$ is nondecreasing. By H1,

$$(4.1.8) \quad \sigma^2(\beta(t_k)) \leq \sup\{\sigma^2(s): 0 \leq s \leq \beta\} < \infty.$$

So if $n(t_k) \rightarrow \infty$, then $\sigma^2(\beta(t_k))/n(t_k) \rightarrow 0$ and hence $a(t_k) - EX(\beta(t_k)) \Rightarrow 0$. Since $b^2(\beta(t_k))$ is uniformly bounded away from zero on $[0, \beta]$ by H3, we get $a(t_k) \Rightarrow \alpha$ — contradicting a hypothesis of P1.2. For the other case $n(t_k) \nrightarrow \infty$, we extract a subsequence t_{k_l} of t_k such that $n(t_{k_l}) \equiv m$. Clearly, $m \geq 1$ since $a(t_{k_l}) \Rightarrow \alpha \neq 0$ by hypotheses of P1.2. Thus, using the definition (0.2),

$$(4.1.9) \quad a(t_{k_l}) = \frac{1}{m} \sum_{i=1}^m X_i(\beta(t_{k_l})) \Rightarrow \alpha.$$

In the next paragraph we show that

$$(4.1.10) \quad \{a(t_{k_l}) \text{ is uniformly integrable}\}.$$

Together (4.1.9) and (4.1.10) imply

$$(4.1.11) \quad E[a(t_{k_l})] \rightarrow \alpha.$$

Since our contradiction argument is assuming that $\beta(t_k) \rightarrow \beta < \infty$, we get

$$(4.1.12) \quad \beta(t_{k_l}) \leq \beta < \infty$$

and so

$$(4.1.13) \quad \inf\{b^2(\beta(t_{k_l}))\} > 0$$

by H3. Now (4.1.13) contradicts (4.1.11).

To complete the proof of P1.2, it remains to show (4.1.10). Observe that

$$\begin{aligned} (4.1.14) \quad E a^2(t_{k_l}) &\leq E \max\{X_i^2(\beta(t_{k_l})): 1 \leq i \leq m\} \\ &\leq E \sum_{i=1}^m X_i^2(\beta(t_{k_l})) \\ &= m EX^2(\beta(t_{k_l})) \end{aligned}$$

$$\leq m \sup\{EX^2(t): 0 \leq t \leq \beta\}$$

$$< \infty$$

by H2. Now (4.1.10) follows from (4.1.12) using Chung [(1974), p. 100, problem 7].

4.2. Proof of Proposition 2

Suppose $n(t) \not\rightarrow \infty$. Then there exists a sequence $t_k \rightarrow \infty$ such that $n(t_k) \equiv m$. If $m = 0$, then $a(t_k) = 0 \neq \alpha$ contradicting the hypothesis $a(t) \Rightarrow \alpha$. Hence $m \geq 1$. Now $a(t) \Rightarrow \alpha$ implies $a(t_k) \Rightarrow \alpha$ which in turn implies

$$(4.2.1) \quad \frac{1}{m} \sum_{i=1}^m X_i(\beta(t_k)) \Rightarrow \alpha$$

as $k \rightarrow \infty$. Taking characteristic functions of both sides of (4.2.1), we get

$$(4.2.2) \quad E[\exp(iuX(\beta(t_k)))] \rightarrow e^{iu\alpha}$$

for all u as $k \rightarrow \infty$. This implies

$$(4.2.3) \quad X(\beta(t_k)) \Rightarrow \alpha.$$

But by P1.2, we get

$$(4.2.4) \quad \beta(t_k) \rightarrow \infty.$$

Now (4.2.3) and (4.2.4) together contradict H5.

4.3. Proof of Theorem 1

From (0.6) we get

$$(4.3.1) \quad tE[a(t) - \alpha]^2 \geq w(\beta(t))\sigma^2(\beta(t)) \rightarrow \infty$$

by P1.2 and H6.

4.4. Proof of Theorem 2

By hypothesis of the theorem, $a(t) \Rightarrow \alpha \neq 0$. Hence $\beta(t) \rightarrow \infty$ and $n(t) \rightarrow \infty$ by Propositions 1 and 2 respectively. From the definition (0.3) of $n(t)$ we therefore get

$$(4.4.1) \quad t/w(\beta(t))n(t) \rightarrow 1$$

and so

$$(4.4.2) \quad g(t) \sim [n(t)/\sigma^2(\beta(t))]^{1/2} = h(t)$$

say. Now

$$(4.4.3) \quad h(t)[a(t) - \alpha] = \sum_{i=1}^{n(t)} \Gamma_i(t) + h(t)b(\beta(t))$$

where

$$(4.4.4) \quad \Gamma_i(t) = h(t)[X_i(\beta(t)) - EX(\beta(t))]/n(t).$$

We see that (2.2.3) holds if and only if $h(t_n)b(\beta(t_n)) \rightarrow \gamma$. Hence, to complete the proof, it suffices to verify the following claim:

$$(4.4.5) \quad \sum_{i=1}^{n(t)} \Gamma_i(t) \Rightarrow N(0, 1).$$

As H1 is not quite strong enough by itself to do the job, our strategy aims to apply the Lindeberg-Feller theorem [Chung (1974), p. 205]. Pick $\eta > 0$. Clearly

$$(4.4.6) \quad \sum_{i=1}^{n(t)} E[\Gamma_i^2(t)]I\{\Gamma_i^2(t) > \eta\} = E[\bar{X}(\beta(t))I\{\bar{X}^2(\beta(t)) > \eta n(t)\}]$$

where

$$(4.4.7) \quad \bar{X}(t) = [X(t) - EX(t)]/n(t).$$

Since $n(t) \rightarrow \infty$, the right side of (4.4.6) converges to zero if $\{\bar{X}^2(\beta(t)): t > t_0\}$ is uniformly integrable. The latter is an easy consequence of H7 and R2. Thus Lindeberg's condition holds and (4.4.5) follows.

4.5. Proof of Theorem 3

By P1.2, we get $\beta(t) \rightarrow \infty$. So by H6 we get

$$(4.5.1) \quad t^{1/2}g^{-1}(t) \rightarrow \infty.$$

Theorem 3 follows immediately from (2.2.2), (2.2.4), and (4.5.1).

4.6. Proof that A1.1-A1.5 imply H1-H8

H8: Follows directly from A1.5.

H7: Note that $0 \leq X(\tau) \leq X(\infty) \leq E\{f(Z(T)) \mid Y\}$. Since $\text{Var } f(Z(T)) < \infty$ implies $EX^2(\infty) < \infty$, clearly $\{X^2(\tau): \tau > 0\}$ is uniformly integrable.

H4 and H1: The argument for H7 also gives $b(\tau) \rightarrow 0$ as $\tau \rightarrow \infty$ and $EX^2(\tau) \rightarrow EX^2(\infty)$.

H3: Since $\text{Var } f(Z(T)) > 0$, there exists a state x such that $f(x) > 0$. By the recurrence of x and the nonnegativity of f , we get

$$\sum_{k=n}^{\infty} f(Y(k))p(k) > 0 \quad a.s.$$

for all $n \geq 1$.

H2, H5, and H6: All follow from

$$(4.6.1) \quad \text{Var } X(\infty) > 0.$$

Before beginning the proof of (4.6.1), we show why A1.2 can't be dropped (though possibly it can be weakened). Without it, Z could alternate between two states 1 and 2 say with the jump rate for $1 \rightarrow 2$ transitions equal to the jump rate for $2 \rightarrow 1$ transitions and $f(1) \neq f(2)$. Then we would get $\text{Var } f(Z(T)) > 0$ but $\text{Var } h(Y) = 0$. It is easy to find intuitively convincing, heuristic arguments that recurrence, aperiodicity, and A1.3 imply (4.6.1), but we could not find a shorter, *rigorous* argument than our proof below.

We argue by contradiction. Suppose that $\text{Var } X(\infty) = 0$. Then

$$(4.6.2) \quad \sum_{k=0}^{\infty} \hat{f}(Y(k))p(k) = 0 \quad a.s.$$

where $\hat{f}(\bullet) = f(\bullet) - Ef(Z(T))$. Let P be the transition matrix of Y and condition on $\{Y(0), \dots, Y(n)\}$. From (4.6.2) we get

$$(4.6.3) \quad \sum_{k=0}^{n-1} \hat{f}(Y(k))p(k) + \sum_{l=0}^{\infty} (P^l \hat{f})(Y(n))p(l+n) = 0 \quad a.s.$$

Likewise

$$(4.6.4) \quad \sum_{k=0}^n \hat{f}(Y(k))p(k) + \sum_{l=0}^{\infty} (P^l \hat{f})(Y(n+1))p(l+n+1) = 0 \quad a.s.$$

Subtract (4.6.3) from (4.6.4):

$$(4.6.5) \quad \sum_{k=0}^{\infty} p(k+n+1)[(P^k \hat{f})(Y(n+1)) - (P^{k+1} \hat{f})(Y(n))] = 0 \quad a.s.$$

Hence, taking the variance of (4.6.5) we get

$$(4.6.6) \quad \sum_{k=0}^{\infty} p^2(k+n+1) \text{Var}[(P^k \hat{f})(Y(n+1)) - (P^{k+1} \hat{f})(Y(n))] = 0$$

noting that the covariance terms vanish. The latter orthogonality follows from

$$(4.6.7) \quad E \{ (P^k \hat{f})(Y(n+1)) - (P^{k+1} \hat{f})(Y(n)) \mid Y(0), \dots, Y(n) \} = 0.$$

From (4.6.6) and $p(l) \neq 0$ for all l we get $(P^k \hat{f})(Y(n+1)) = (P^{k+1} \hat{f})(Y(n)) \quad a.s.$ Hence $\hat{f}(Y(n)) = (P^n \hat{f})(Y(0)) \quad a.s.$ Since P is aperiodic, $(P^n \hat{f})(Y(0)) \rightarrow \pi \hat{f} = 0 \quad a.s.$ as $n \rightarrow \infty$ (see Glynn (1986a)). Let $T_n(s)$ be the n -th time at which state s is visited. By recurrence (A1.1), we get $T_n(s) \rightarrow \infty \quad a.s.$; so $\hat{f}(Y(T_n(s))) \rightarrow 0 \quad a.s.$, i.e., $\hat{f}(s) = 0$.

Because s is arbitrary, we get $\hat{f}(s) = 0$ for all s . Thus, f is constant. This implies $\text{Var } f(Z(T)) = 0$, contradicting A1.4.

Acknowledgment

This research was supported in part by a grant from the Natural Sciences and Engineering Research Council of Canada, by the National Science Foundation under grant ECS-8404809, and by the U.S. Army under Contract No. DAAG29-80-C-0041.

References

- [1] P. BRATLEY, B.L. FOX, and L.E. SCHRAGE, *A Guide to Simulation*. Second Edition. Springer-Verlag, New York, 1987.
- [2] K.L. CHUNG, *A Course in Probability Theory*. Academic Press, New York, 1974.
- [3] E. ÇINLAR, *Introduction to Stochastic Processes*. Prentice-Hall, Englewood Cliffs, New Jersey, 1975.
- [4] B.L. FOX, *Numerical Methods for Transient Markov Chains*. Technical report, Cornell University, 1987.
- [5] B.L. FOX and P.W. GLYNN, *Simulating Discounted Processes*. Technical report, Stanford University, 1987.
- [6] B.L. FOX and P.W. GLYNN, *Discrete-time Conversion for Finite-horizon Markov Processes*. Technical report, Cornell University, 1988a.
- [7] B.L. FOX and P.W. GLYNN, *Computing Poisson Probabilities*. Comm. ACM (to appear), 1988b.
- [8] A.S. FROLOV and N.N. CHENTSOV, *On the Calculation of Definite Integrals Dependent on a Parameter by the Monte Carlo Method*. USSR Comput. Math. and Math. Phys. 4 (1963), 802-807.
- [9] P.W. GLYNN, *Monte Carlo Optimization of Stochastic Systems*. Proceedings Winter Simulation Conference, Washington, D.C., 1986a.
- [10] P.W. GLYNN, *Low-Bias Steady-State Estimation for Equilibrium Processes*. Technical report. University of Wisconsin, 1986b.
- [11] P.W. GLYNN, *Upper Bounds on Poisson Tail Probabilities*. Operations Research Letters 6 (1987a), 9-14.

- [12] P.W. GLYNN, *Limit Theorems for the Method of Replication*. Stochastic Models (to appear), 1987b.
- [13] P.W. GLYNN, *Gradient Estimation with Common Random Numbers*. Technical report in preparation, 1988a.
- [14] P.W. GLYNN, *Randomized Estimators and Time Integrals*. To appear, 1988b.
- [15] P.W. GLYNN and W. WHITT, *The Efficiency of Simulation Estimators*. Technical report. AT&T Bell Laboratories, Holmdel, New Jersey, 1986.
- [16] D. GROSS, D.R. MILLER, and C.G. PLASTIRAS, *Simulation Methodologies for Transient Markov Processes: A Comparative Study Based on Multi-Echelon Repairable Item Inventory Systems*. Proceedings Summer Simulation Conference, Boston, 1984.
- [17] P. HEIDELBERGER, X.R. CAO, M.A. ZAZANIS, and R. SURI, *Convergence Properties of Infinitesimal Perturbation Analysis Estimates*. Technical report, IBM T.J. Watson Research Center, Yorktown Heights, New York, 1987.
- [18] L.W. JOHNSON and R.D. RIESS, *Numerical Analysis*. Addison-Wesley, Reading, Massachusetts, 1982.
- [19] M.A. ZAZANIS and R. SURI, *Comparison of Perturbation Analysis with Conventional Sensitivity Estimates for Stochastic Systems*. Technical report. University of Wisconsin, 1986.