Editorial Office: *Probability in the Engineering and Informational Sciences*, College of Engineering, Department of Industrial Engineering and Operations Research, University of California, Berkeley, CA 94720, U.S.A. Tel. (415) 642-3615.

Publishing, Production and Advertising Offices: Cambridge University Press, 32 East 57th Street, New York, NY 10022, U.S.A.

Subscription Offices: Cambridge University Press, 110 Midland Avenue, Port Chester, NY 10573, U.S.A. (for US/Canada); or Cambridge University Press, The Edinburgh Building, Shaftesbury Road, Cambridge CB2 2RU, England (for UK/rest of world).

1989 Subscription Information: *Probability in the Engineering and Informational Sciences* (ISSN 0269-9648) is published quarterly. Institutional rate: US $118.00 in the U.S.A. and Canada; UK £87.00 in the U.K. and Eire; UK £93.00 elsewhere. Individual rates (prepayment by personal check or credit card only): US $58.00 in the U.S.A. and Canada, UK £45.00 elsewhere. Prices include postage and insurance; air mail or registered mail extra.

Article Reprints are available from the authors. For orders of 50 or more reprints, contact the Journals Production Office.

Application to mail at second-class rates is pending in New York, NY and additional mailing offices. Postmaster: Send address changes in the U.S.A. and Canada to *Probability in the Engineering and Informational Sciences*, Subscription Department, Cambridge University Press, 110 Midland Avenue, Port Chester, NY 10573.

Probability in the Engineering and Informational Sciences, 3, 1989, 299–318. Printed in the U.S.A.

REPLICATION SCHEMES FOR LIMITING EXPECTATIONS

BENNETT L. FOX[1]
Department of Mathematics
University of Colorado
Denver, Colorado 80204

PETER W. GLYNN[2]
Department of Operations Research
Stanford University
Stanford, California 94305-4022

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.

1. 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 independent and identically distributed (i.i.d.) random vari-

[1]Part of this research done while on leave at the School of Operations Research and Industrial Engineering, Upson Hall, Cornell University, Ithaca, New York 14853-7501, from the University of Montreal. Research supported by the Natural Sciences and Engineering Research Council of Canada.

[2]Research supported by the National Science Foundation under Grant ECS-8404809 and the U.S. Army under Contract No. DAAG29-80-C-0041.

ables. 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 [19] and Frolov and Chentsov [8] showed that, when the derivative $\alpha'(\theta_0)$ of a function $\alpha(\cdot) = EZ(\cdot)$ 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 [5]. The problem there involved estimating $\alpha = \int_0^\infty e^{-t}EZ(\cdot)dt$; 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

$$\alpha = \lim_{\tau\to\infty} \alpha(\tau), \tag{1.1}$$

where $\alpha(\tau) = EX(\tau)$. We do not assume that there is a limiting random variable (r.v.) $X(\infty)$. In fact, the parameterized family $X(\tau)$ may simply be contrived so that (1.1) holds. Our goal here is to estimate the limiting expectation α via simulation. Section 2 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 Eq. (1.1) 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

$$a(t) = \begin{cases} \dfrac{1}{n(t)}\,[X_1(\beta(t)) + \cdots + X_{n(t)}(\beta(t))] & n(t) \geq 1 \\ 0 & n(t) = 0 \end{cases} \tag{1.2}$$

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

$$n(t) = \lfloor t/w(\beta(t)) \rfloor. \tag{1.3}$$

We consider what happens when $EX(\tau) \neq \alpha$ and $X(\tau) \not\Rightarrow \alpha$ but $EX(\tau) \to \alpha$ as $\tau \to \infty$. The sampling plan determines the trade-off 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) [rough] $Z(t)$ converges weakly to α at rate $t^{-1/2}$
[precise] the family $\{t^{1/2}[Z(t) - \alpha]: t > 0\}$ is stochastically bounded, i.e., $\forall\epsilon > 0$ there exists $K(\epsilon)$ such that $P\{t^{1/2}|Z(t) - \alpha| > K(\epsilon)\} < \epsilon$ uniformly in t.

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 3 precisely states and refines the following assertion: under mild hypotheses, all estimators of the form $a(t)$ are inefficient.

Since

$$E[a(t) - \alpha]^2 = \frac{\sigma^2(\beta(t))}{n(t)} + b^2(\beta(t)), \tag{1.4}$$

and

$$n(t) \leq t/w(\beta(t)), \tag{1.5}$$

we get

$$tE[a(t) - \alpha]^2 \geq w(\beta(t))\sigma^2(\beta(t)). \tag{1.6}$$

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 Eq. (3.2.1), we can get a nontrivial limit and specifies this limit distribution. Theorem 3 refines Theorem 1 by showing that all of our hypotheses together imply distribution inefficiency.

In Section 4, we return to the examples of Section 2. For each, we provide mild conditions implying the hypotheses of Section 3. 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 4.2 fall quickly out of our general theory; only (weaker) mean-square-error counterparts were previously known.

Section 5 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.

2. EXAMPLES

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

2.1. Estimation in Finite-Horizon, Continuous-Time Markov Chains

We estimate the 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 [3, p. 236–237].

Let

$$h(Y) = E[f(Z(T))|Y], \quad = \sum_{k=0}^{\infty} f(Y(k))p(k), \tag{2.1.1}$$

where

$$p(k) = e^{-\theta T}(\theta T)^k/k!. \tag{2.1.2}$$

Since taking conditional expectations reduces variance (e.g., see Bratley et al. [1, Section 2.6]), var $h(Y) \leq$ var $f(Z(T))$.

We cannot simulate the infinite summation in Eq. (2.1.1) exactly. So we truncate it to get

$$X(\tau) = \sum_{k=0}^{\lfloor \tau \rfloor} f(Y(k))p(k). \tag{2.1.3}$$

If $E|f(Z(T))| < \infty$, then $\alpha(\tau) = EX(\tau) \to \alpha$ as $\tau \to \infty$. This fits our framework, Eq. (1.1). Here $\beta(t)$ is the truncation point.

Gross et al. [16] studied, among other things, various strategies for implementing Eq. (2.1.3). For improvements in efficiency, see Fox and Glynn [6]. The set of nonnegligible $p(k)$'s needed in Eq. (2.1.3) can be computed efficiently in $O(\sqrt{\theta T})$ time using the scheme of Fox and Glynn [7].

2.2. Gradient Estimation

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

$$\alpha = q'(0). \tag{2.2.1}$$

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:

$$X_1(\tau) = \tau(Z_1(1/\tau) - Z_2(0)) \qquad \text{(forward difference)}, \tag{2.2.2}$$

$$X_2(\tau) = \tau(Z_1(1/2\tau) - Z_2(-1/2\tau)) \qquad \text{(central difference)}, \tag{2.2.3}$$

where $Z_1(\cdot)$ and $Z_2(\cdot)$ 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 $\alpha_i(\tau) = EX_i(\tau) \to \alpha$ as $\tau \to \infty$. 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.

2.3. Discounting Generalized

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

$$\alpha = \int_{[0,\infty)} EZ(t)G(dt). \tag{2.3.1}$$

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 2.1 on the right, we recover $Eh(Y)$ with $h(Y)$ defined by Eq. (2.1.1).

With

$$X(\tau) = \int_{[0,\tau]} Z(s)G(ds) \tag{2.3.2}$$

and $\alpha(\tau) = EX(\tau)$, we again fit our framework, Eq. (1.1).

2.4. 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

$$X(\tau) = [1/(\lfloor \tau \rfloor + 1)] \sum_{k=0}^{\lfloor \tau \rfloor} f(Z(k)). \tag{2.4.1}$$

Under appropriate integrability conditions, there is a random variable $X(\infty)$ such that $X(\tau) \to X(\infty)$ a.s. as $\tau \to \infty$. 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 3.1 requires but which generally does not occur when there is just one recurrence class). Glynn [12] covers the case where $X(\infty)$ is constant except for special f, while Section 4.4 covers the case where $X(\infty)$ is nondegenerate. With $\alpha(\tau) = EX(\tau)$, we again fit our framework.

3. CONVERGENCE RATES

In Section 3.1, 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 3.2. 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 3.3 illustrates how this guides us to good sampling plans. When the bias is exponentially decreasing in τ, the results of Section 3.3 lead to sampling plans giving convergence rates arbitrarily close to the canonical rate as Section 3.4 points out.

3.1. Inefficiency

Let

$$\sigma^2(\tau) = \text{var } X(\tau), \tag{3.1.1}$$

$$b(\tau) = \alpha - EX(\tau). \tag{3.1.2}$$

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 2, all of these hypotheses often hold, as Section 4 shows precisely. With each of our results, we show explicitly the subset of hypotheses used.

H1 (Finite variance): $\sup\{\sigma^2(s): 0 \le s \le \tau\} < \infty$ for all τ.

H2 (Positive variance): there exists τ_0, such that $\inf\{\sigma^2(s): \tau_0 \le s \le \tau\} > 0$ for all $\tau \ge \tau_0$.

H3 (Positive squared bias): $\inf\{b^2(s): 0 \le s \le \tau\} > 0$ for all τ.

H4 (Asymptotic unbiasedness): $EX(\tau) \to \alpha$.

H5 (Nondegeneracy): there exists no sequence $\tau_n \to \infty$ such that $X(\tau_n) \Rightarrow \alpha$ as $n \to \infty$.

H6 (Getting rid of bias by letting $\tau \to \infty$ is costly): $w(\tau)\sigma^2(\tau) \to \infty$ as $\tau \to \infty$.

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 $\inf\{w(\tau): \tau \ge 0\} > 0$.

Remarks on these hypotheses:

R1. H1-H4 and H7 seem reasonable for this problem setting.

R2. H2, H4, and H5 together imply that $\inf\{\sigma^2(s): s \ge \tau_0\} > 0$. This follows from H2 after showing that $\sigma^2(\tau) \not\to 0$ as $\tau \to \infty$. Aiming for a contradiction, suppose that $\sigma^2(\tau_n) \to 0$ along some sequence $\tau_n \to \infty$. Then H4 shows that $E[X(\tau_n) - \alpha]^2 \to 0$. 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 $\tau_n \to \infty$. 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 $\tau_{n_l} \to \tau < \infty$ and $X(\tau_{n_l}) \Rightarrow \alpha$. But by H1, we get $\sup\{EX^2(\tau_{n_l})\rangle \le \sup[EX^2(s): 0 \le s \le \sup \tau_n\} < \infty$, so $\{X(\tau_{n_l})\}$ is uniformly integrable. Hence, $EX(\tau_{n_l}) \to \alpha$, contradicting H3.

R4. To understand H6, observe that R3, H6, and H8 jointly imply that $w(\tau) \to \infty$, or $\sigma^2(\tau) \to \infty$, 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 $\beta(t) \to \infty$, but Proposition 2 says not too fast.

Proposition 1:

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

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

P1.2 (Partial converse): Assume H1, H2, H3, and H8. If $a(t) \Rightarrow \alpha$ and $\alpha \ne 0$, then

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

Proposition 2: Assume H1-H5.

P2.1 (No free lunch): If $a(t) \Rightarrow \alpha$ and $\alpha \ne 0$, then

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

THEOREM 1 (Root-mean-square-error inefficiency): *Assume H1-H3, H6, and H8. If $a(t) \Rightarrow \alpha$ and $\alpha \ne 0$, then $a(t)$ is inefficient:*

$$tE[a(t) - \alpha]^2 \to \infty. \tag{3.1.3}$$

3.2. Extended Weak Limit Points

Let $\{V(t)\colon t \geq 0\}$ 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 $t_n \to \infty$ 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 $P\{V(t) \leq x\} \to 0$ a.s. $t \to \infty$ 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*

$$g(t) = [t/w(\beta(t))\sigma^2(\beta(t))]^{1/2}, \tag{3.2.1}$$

$$\Gamma(t) = g(t)[a(t) - \alpha]. \tag{3.2.2}$$

If $a(t) \Rightarrow \alpha \neq 0$, then the only possible EWLP's of $\{\Gamma(t)\colon t > 0\}$ have the form $N(0,1) + \gamma$ with $\gamma \in [-\infty,\infty]$. The sequence $\{\Gamma(t_n)\colon n \geq 1\}$ converges weakly if and only if

$$g(t_n)b(\beta(t_n)) \to \gamma \tag{3.2.3}$$

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

$$\Gamma(t_n) \Rightarrow N(0,1) + \gamma. \tag{3.2.4}$$

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 EWLP's of $\{t^{1/2}[a(t) - \alpha]\colon t > 0\}$ are $-\infty$ and $+\infty$.*

3.3. Sampling-Plan Guidelines

Our results in Sections 3.1 and 3.2 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 2, we give or cite an efficient alternative estimator in Section 4. 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

$$\tau^{-\delta}\sigma^2(\tau) \to \sigma^2, \tag{3.3.1}$$

$$\tau^{\rho}|b(\tau)| \to b, \tag{3.3.2}$$

$$\tau^{-\lambda}w(\tau) \to w. \tag{3.3.3}$$

By H3, we get $\rho > 0$. The conditions of Eqs. (3.3.1)–(3.3.3) hold in many settings, including several of our examples as we show in Section 4.

In Eq. (3.2.4), we clearly want $|\gamma| < \infty$, so we impose

$$\limsup_{t\to\infty}|g(t)b(\beta(t))| < \infty \tag{3.3.4}$$

to define the class of admissible sampling plans β. Routine calculations show that Eqs. (3.3.1)–(3.3.4) jointly imply

$$\liminf \beta(t)/t^{1/\delta+\lambda+2\rho} > 0. \tag{3.3.5}$$

Sampling plans of the form

$$\beta(t) \sim ct^{1/(\delta+\lambda+2\rho)} \tag{3.3.6}$$

with $0 < c < \infty$ satisfy Eq. (3.3.5) and give

$$g(t) \sim [t^{\rho/(\delta+\lambda+2\rho)}] \cdot (w\sigma^2c^{\delta+\lambda})^{-1/2}, \tag{3.3.7}$$

$$\gamma = (w\sigma^2c^{\delta+\lambda})^{-1/2}bc^{-\rho} \tag{3.3.8}$$

where the bracketed factor in Eq. (3.3.7) is the maximal (subcanonical) convergence rate obtainable with estimators of the form $a(t)$. Now $[t^{\rho/(\delta+\lambda+2\rho)}] \cdot [a(t) - \alpha]$ converges to a limit random variable $(w\sigma^2c^{\delta+\lambda})^{1/2}N(0,1) + bc^{-\rho}$. Minimizing its second moment $w\sigma^2c^{\delta+\lambda} + b^2c^{-2\rho}$ over c gives

$$c^* = \left[\frac{2\rho b^2}{w\sigma^2(\lambda+\delta)}\right]^{1/(\lambda+\delta+2\rho)} \tag{3.3.9}$$

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 $\gamma \not\equiv 0$, we generally get a different value of c^*. This cannot 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. Although w and σ can easily be estimated consistently, usually only an *a priori* upper bound on b is available, and in such cases Eq. (3.3.9) has to be adjusted heuristically. We show in Section 4.2 that, for finite differences, b *can* be estimated consistently; so, in that case, Eq. (3.3.9) is fully operational.

When b cannot be consistently estimated but $\bar{b}$ is a known upper bound, one way to adjust Eqs. (3.3.8) and (3.3.9) 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 epsilontically less. Thus, instead of Eq. (3.3.6), we can take

$$\beta(t) = ct^{1(\delta+\lambda+2\rho)+2\epsilon/(\delta+\lambda)} \tag{3.3.10}$$

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

$$t^{[\rho/(\delta+\lambda+2\rho)]-\epsilon}[a(t) - \alpha] \Rightarrow (w\sigma^2 c^{\delta+\lambda})^{1/2} N(0,1), \tag{3.3.11}$$

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

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

as an asymptotic $100(1 - \xi)\%$ confidence interval for α. In Eq. (3.3.11), 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 Eq. (3.3.10) has no effect on Eq. (3.3.11), but may substantially improve performance for moderate values of t.

3.4. 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 Eq. (3.3.6) or in Eq. (3.3.10) 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 $g^2(t)b^2(\beta(t)) \to \gamma^2 > 0$ sometimes leads to slightly better, though still subcanonical, convergence rates—but with $\gamma \ne 0$. Some of our examples have exponentially decreasing bias, as we detail in Section 4. 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.

4. EXAMPLES REVISITED

We return to the examples considered in Section 2 in the corresponding subsections here. In each case, we give or cite efficient alternative estimators and find easily checked conditions implying H1–H8 and $\alpha \ne 0$ for biased multiple-run estimators.

4.1. Estimation in Finite-Horizon, Continuous-Time Markov Chains

The naive strategy takes the average of $n(t)$ i.i.d. 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 [15] and Fox and Glynn [6] show (as a special case) that, with i indexing runs,

$$t^{1/2}\left[\frac{1}{n(t)} \sum_{i=1}^{n(t)} f(Z_i(T)) - \alpha\right] \Rightarrow [EW \cdot \operatorname{var} f(Z(T))]^{1/2} N(0,1), \tag{4.1.1}$$

and thus the canonical rate is obtained. In general, no explicit formula for EW is available but sometimes (see Fox [4]) the order of magnitude of its worst-case counterpart can be given. Fox and Glynn [6] obtain several other efficient estimators, all with lower asymptotic variance constant (equal to expected work per run times variance per run) than in Eq. (4.1.1).

An alternative averages i.i.d. copies of Eq. (2.1.3), 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 5.6 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: $f \ge 0$.

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

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 Eq. (3.3.3). Recalling that in this example τ corresponds to a truncation point, we get $\delta = 0$ in Eq. (3.3.1).

Since Y is an aperiodic recurrent chain, generally $Ef(Y(n)) \to \pi f$ (the steady-state mean with π the stationary probability vector for Y). Using Glynn's [11] results on Poisson tail probabilities, we get

$$b(\tau) \sim \{\pi f e^{-\theta T}(\theta T)^{\lfloor \tau \rfloor}\}/\lfloor \tau \rfloor !, \tag{4.1.2}$$

where πf often can be computed analytically or numerically. From Eq. (4.1.2), we see that the bias goes to zero exponentially fast, so from Section 3.4 we can choose a sampling plan to get arbitrarily close to the canonical convergence rate.

4.2. Gradient Estimation

For estimating gradients, Zazanis and Suri [19] 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. [17] point out certain classes of problems where perturbation analysis gives strongly consistent gradient estimates and other classes where it does not. Glynn's [10] method of likelihood ratios for estimating gradients has root-mean-square error $O(t^{-1/2})$; this method is described in Bratley et al. [1, Section 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 $\xi \to 0$.

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

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

A2.4: $\nu(\xi) \to \nu(0) \neq 0$ as $\xi \to 0$, where $\nu(\xi) = \text{var } Z(\xi)$.

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

For the forward-difference estimator,

$$\tau^{-2}\text{var } X_1(\tau) \to 2\nu(0), \tag{4.2.1}$$

$$\tau b_1(\tau) \to \alpha^{(2)}(0)/2, \tag{4.2.2}$$

$$w_1(\tau) \to 2d(0). \tag{4.2.3}$$

For the central-difference estimator,

$$\tau^{-2}\text{var } X_2(\tau) \to 2\nu(0), \tag{4.2.4}$$

$$\tau^2 b_2(\tau) \to \alpha^{(3)}(0)/24, \tag{4.2.5}$$

$$w_2(\tau) \to 2d(0). \tag{4.2.6}$$

In deterministic settings, the counterparts of Eqs. (4.2.2) and (4.2.5) are well-known in numerical analysis. Here we obtain them in the same way, via Taylor-like expansions. Using common random numbers may reduce the mean square errors of both forward and central differences, by an order of magnitude if synchronization is good enough; see Glynn [13].

In the notation of Section 3.3, we have $w = 2d(0)$, $\sigma^2 = 2\nu(0)$, $\lambda = 0$, 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

$$\beta_1(t) = c_1^* t^{1/3}, \tag{4.2.7}$$

$$c_1^* = ([\alpha^{(2)}(0)]^2/16d(0)\nu(0))^{1/3}, \tag{4.2.8}$$

$$\beta_2(t) = c_2^* t^{1/6} \tag{4.2.9}$$

$$c_2^* = ([\alpha^{(3)}(0)/24]^2/2d(0)\nu(0))^{1/6}, \tag{4.2.10}$$

with convergence rates $O(t^{1/4})$ and $O(t^{1/3})$, respectively. By other methods, Zazanis and Suri [19] and Frolov and Chentsov [8] find root-mean-square-error counterparts of these rates.

We can estimate the right side of Eq. (4.2.2) 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 Eq. (4.2.5) 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 $k = 1,2,3,\ldots$ as shown in Johnson and Riess [18, p. 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.

4.3. Discounting Generalized

Fox and Glynn [5] and Glynn [14] give efficient estimators for generalized discounting.

An alternative uses truncation as in Eq. (2.3.2) and averages this (biased) estimator over multiple i.i.d. runs. The (weak) conditions,

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

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

A3.3: $b(\tau) \neq 0$ for all τ,

A3.4: $w(\tau) = c\tau$ 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 $\{X^2(\tau)\colon \tau > 0\}$ is uniformly integrable.

In the notation of Section 3.3, we get $\delta = 0$, $\sigma^2 = \text{var}\, X(\infty)$, $\lambda = 1$, and $w = c$. If Eq. (3.3.2) holds with $b \neq 0$, we pick β as in Section 3.3. If in Eq. (3.3.2), we get $b = 0$ for all $\rho > 0$, then the bias decreases exponentially fast and we pick β as in Section 3.4.

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

$$b(t) \sim EZ(\infty)e^{-\xi t}/\xi \tag{4.3.1}$$

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

4.4. Steady-State Estimation for Markov Chains

If we know explicitly the recurrence classes, say $C_1, \ldots, C_n$, 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 $C_1, \ldots, C_n$ and $\{P_1, \ldots, P_n\}$, 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 i.i.d. copies of the biased estimator $X(\tau)$ given by Eq. (2.4.1). The conditions,

A4.1: $Z(0) = i$ a.s.,

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

A4.3: $Ef^2(Z(n)) \to Ef^2(Z(\infty)) < \infty$ as $n \to \infty$,

A4.4: $b(n) \neq 0$ for all n,

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

$$X^2(\tau) \equiv \frac{1}{\lfloor \tau \rfloor + 1} \sum_{k=0}^{\lfloor \tau \rfloor} f^2(Z(k)), \tag{4.4.1}$$

and the right-hand side is uniformly integrable by A4.3. We also assume H8; typically, $\lambda = 1$ in Eq. (3.3.3). From A4.3, we get $\delta = 0$ in Eq. (3.3.1).

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

$$Ef(Z(k)) = Ef(Z(\infty)) + O(q^k) \tag{4.4.2}$$

with $0 < q < 1$. Summing from $k = 0$ to n, dividing by $n + 1$, and comparing with Eq. (2.4.1), we see that $b(n) = O(1/n)$. This gives $\rho = 1$ in Eq. (3.3.2). Glynn's [10] results allow estimation of b with low bias. Corresponding to Eq. (3.3.11), we get $O(t^{(-1/3)+\epsilon})$ convergence rate with $\gamma = 0$.

5. PROOFS

5.1. Proof of Proposition 1

For the first part,

$$E[a(t) - \alpha]^2 = \frac{\sigma^2(\beta(t))}{n(t)} + b^2(\beta(t)). \tag{5.1.1}$$

From Eq. (1.3), we get

$$n(t) \geq t/w(\beta(t)) - 1. \tag{5.1.2}$$

So

$$E[a(t) - \alpha]^2 \leq \frac{w(\beta(t))\sigma^2(\beta(t))}{t[1 - w(\beta(t))/t]} + b^2(\beta(t)). \tag{5.1.3}$$

By R3, from H1-H5, we get

$$\inf\{\sigma^2(\tau): \tau > \tau_0\} > 0. \tag{5.1.4}$$

From Eq. (5.1.4) and (the hypothesis of P1.1) $w(\beta(t))\sigma^2(\beta(t))/t \to 0$, we get

$$w(\beta(t))/t \to 0. \tag{5.1.5}$$

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

$$b^2(\beta(t)) \to 0. \tag{5.1.6}$$

Combining Eqs. (5.1.3)-(5.1.6) proves

$$E[a(t) - \alpha]^2 \to 0 \tag{5.1.7}$$

and P1.1 follows.

For the converse, assume that there exists a sequence $t_k \to \infty$ such that $\beta(t_k) \to \beta < \infty$. This holds if, contrary to P1.2, $\beta(t) \not\to \infty$ since $\beta(t)$ is nondecreasing. By H1,

$$\sigma^2(\beta(t_k)) \leq \sup\{\sigma^2(s): 0 \leq s \leq \beta\} < \infty. \tag{5.1.8}$$

So if $n(t_k) \to \infty$, then $\sigma^2(\beta(t_k))/n(t_k) \to 0$ 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 $a(t_k) \not\Rightarrow \alpha$—contradicting a hypothesis of P1.2. For the other case $n(t_k) \not\to \infty$, 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 of Eq. (1.2),

$$a(t_{k_l}) = \frac{1}{m} \sum_{i=1}^{m} X_i(\beta(t_{k_l}) \Rightarrow \alpha. \tag{5.1.9}$$

In the next paragraph, we show that

$$\{a(t_{k_l}) \text{ is uniformly integrable}\}. \tag{5.1.10}$$

Together, Eqs. (5.1.9) and (5.1.10) imply

$$E[a(t_{k_l})] \to \alpha. \tag{5.1.11}$$

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

$$\beta(t_{k_l}) \le \beta < \infty \tag{5.1.12}$$

and so

$$\inf\{b^2(\beta(t_{k_l}))\} > 0 \tag{5.1.13}$$

by H3. Now Eq. (5.1.13) contradicts Eq. (5.1.11).

To complete the proof of P1.2, it remains to show Eq. (5.1.10). Observe that

$$\begin{aligned} Ea^2(t_{k_l}) &\le E\max\{X_i^2(\beta(t_{k_l})): 1 \le i \le m\}, \\ &\le E\sum_{i=1}^{m} X_i^2(\beta(t_{k_l})), \\ &= mEX^2(\beta(t_{k_l})), \\ &\le m\sup\{EX^2(t): 0 \le t \le \beta\}, \\ &< \infty \end{aligned} \tag{5.1.14}$$

by H2. Now Eq. (5.1.10) follows from Eq. (5.1.12) using Chung [2, p. 100, problem 7].

5.2. Proof of Proposition 2

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

$$\frac{1}{m}\sum_{i=1}^{m} X_i(\beta(t_k)) \Rightarrow \alpha \tag{5.2.1}$$

as $k \to \infty$. Taking characteristic functions of both sides of Eq. (5.2.1), we get

$$E[\exp(iuX(\beta(t_k)))] \to e^{iu\alpha} \tag{5.2.2}$$

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

$$X(\beta(t_k)) \Rightarrow \alpha. \tag{5.2.3}$$

But by P1.2, we get

$$\beta(t_k) \to \infty. \tag{5.2.4}$$

Now Eqs. (5.2.3) and (5.2.4) together contradict H5.

5.3. Proof of Theorem 1

From Eq. (1.6), we get

$$tE[a(t) - \alpha]^2 \ge w(\beta(t))\sigma^2(\beta(t)) \to \infty \tag{5.3.1}$$

by P1.2 and H6.

5.4. Proof of Theorem 4

By hypothesis of the theorem, $a(t) \Rightarrow \alpha \ne 0$. Hence $\beta(t) \to \infty$ and $n(t) \to \infty$ by Propositions 1 and 2, respectively. From the definition in Eq. (1.3) of $n(t)$, we therefore get

$$t/w(\beta(t))n(t) \to 1, \tag{5.4.1}$$

and so

$$g(t) \sim [n(t)/\sigma^2(\beta(t))]^{1/2} = h(t), \tag{5.4.2}$$

say. Now

$$h(t)[a(t) - \alpha] = \sum_{i=1}^{n(t)} \Gamma_i(t) + h(t)b(\beta(t)), \tag{5.4.3}$$

where

$$\Gamma_i(t) = h(t)[X_i(\beta(t)) - EX(\beta(t))]/n(t). \tag{5.4.4}$$

We see that Eq. (3.2.3) holds if and only if $h(t_n)b(\beta(t_n)) \to \gamma$. Hence, to complete the proof, it suffices to verify the following claim:

$$\sum_{i=1}^{n(t)} \Gamma_i(t) \Rightarrow N(0,1). \tag{5.4.5}$$

As H1 is not quite strong enough by itself to do the job, our strategy aims to apply the Lindeberg–Feller theorem (Chung [2, p. 205]). Pick $\eta > 0$. Clearly,

$$\sum_{i=1}^{n(t)} E[\Gamma_i^2(t)]I\{\Gamma_i^2(t) > \eta\} = E[\bar{X}(\beta(t))I\{\bar{X}^2(\beta(t)) > \eta n(t)\}], \tag{5.4.6}$$

where

$$\bar{X}(t) = [X(t) - EX(t)]/n(t). \tag{5.4.7}$$

Since $n(t) \to \infty$, the right side of Eq. (5.4.6) 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 Eq. (5.4.5) follows.

5.5. Proof of Theorem 3

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

$$t^{1/2}g^{-1}(t) \to \infty. \tag{5.5.1}$$

Theorem 3 follows immediately from Eqs. (3.2.2), (3.2.4), and (5.5.1).

5.6. Proof that A1.1–A1.5 Imply H1–H8

H8: Follows directly from A1.5.

H7: Note that $0 \le X(\tau) \le X(\infty) \le E\{f(Z(T)) \mid Y\}$. Since 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 $b(\tau) \to 0$ as $\tau \to \infty$ and $EX^2(\tau) \to EX^2(\infty)$.

H3: Since 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

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

for all $n \ge 1$.

H2, H5, and H6: All follow from

$$\text{var } X(\infty) > 0. \tag{5.6.1}$$

Before beginning the proof of Eq. (5.6.1), we show why A1.2 cannot be dropped (although possibly it can be weakened). Without it, Z could alternate between two states, 1 and 2 say, with the jump rate for $1 \to 2$ transitions equal to the jump rate for $2 \to 1$ transitions and $f(1) \ne f(2)$. Then we would get var $f(Z(T)) > 0$ but var $h(Y) = 0$. It is easy to find intuitively convincing, heuristic arguments that recurrence, aperiodicity, and A1.3 imply Eq. (5.6.1), but we could not find a shorter, *rigorous* argument than our proof below.

We argue by contradiction. Suppose that var $X(\infty) = 0$. Then

$$\sum_{k=0}^{\infty} \hat{f}(Y(k))p(k) = 0 \qquad \text{a.s.} \tag{5.6.2}$$

where $\hat{f}(\cdot) = f(\cdot) - Ef(Z(T))$. Let P be the transition matrix of Y and condition on $\{Y(0), \ldots, Y(n)\}$. From Eq. (5.6.2), we get

$$\sum_{k=0}^{n-1} \hat{f}(Y(k))p(k) + \sum_{l=0}^{\infty} (P^l\hat{f})(Y(n))p(l+n) = 0, \qquad \text{a.s.} \tag{5.6.3}$$

Likewise,

$$\sum_{k=0}^{n} \hat{f}(Y(k))p(k) + \sum_{l=0}^{\infty} (P^l\hat{f})(Y(n+1))p(l+n+1) = 0, \qquad \text{a.s.} \tag{5.6.4}$$

Subtract Eq. (5.6.3) from Eq. (5.6.4):

$$\sum_{k=0}^{\infty} p(k+n+1)\,[(P^k\hat{f})(Y(n+1)) - (P^{k+1}\hat{f})(Y(n))] = 0. \qquad \text{a.s.} \tag{5.6.5}$$

Hence, taking the variance of Eq. (5.6.5), we get

$$\sum_{k=0}^{\infty} p^2(k+n+1)\text{var}[(P^k\hat{f})(Y(n+1)) - (P^{k+1}\hat{f})(Y(n))] = 0 \tag{5.6.6}$$

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

$$E\{(P^k\hat{f})(Y(n+1)) - (P^{k+1}\hat{f})(Y(n)) \mid Y(0), \ldots, Y(n)\} = 0. \tag{5.6.7}$$

From Eq. (5.6.6) and $p(l) \ne 0$ for all l, we get $(P^k\hat{f})(Y(n+1)) = (P^{k+1}\hat{f})(Y(n))$ a.s. Hence, $\hat{f}(Y(n)) = (P^n\hat{f})(Y(0))$ a.s. Since P is aperiodic, $(P^n\hat{f})(Y(0)) \to \pi\hat{f} = 0$ a.s. as $n \to \infty$ (see Glynn [9]). Let $T_n(s)$ be the nth time at which state s is visited. By recurrence (A1.1), we get $T_n(s) \to \infty$ a.s.; so $\hat{f}(Y(T_n(s))) \to 0$ 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 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. Bratley, P., Fox, B.L., & Schrage, L.E. (1987). *A guide to simulation*, 2nd Ed. New York: Springer-Verlag.
2. Chung, K.L. (1974). *A course in probability theory*. New York: Academic Press.
3. Çinlar, E. (1975). *Introduction to stochastic processes*. Englewood Cliffs, New Jersey: Prentice-Hall.
4. Fox, B.L. (1988). *Numerical methods for transient Markov chains*. Technical Report, Cornell University.
5. Fox, B.L. & Glynn, P.W. (1989). *Simulating discounted processes. Management Science*, to appear.
6. Fox, B.L. & Glynn, P.W. (1988). *Discrete-time conversion for finite-horizon Markov processes*. Technical Report, Cornell University.
7. Fox, B.L. & Glynn, P.W. (1988). Computing Poisson probabilities. *Communications of the ACM* 31: 440–445.
8. Frolov, A.S. & Chentsov, N.N. (1963). On the calculation of definite integrals dependent on a parameter by the Monte Carlo method. *USSR Computational Mathematics and Mathematical Physics* 4: 802–807.
9. Glynn, P.W. (1986). *Monte Carlo optimization of stochastic systems*. Proceedings Winter Simulation Conference, Washington, D.C.
10. Glynn, P.W. (1986). *Low-bias steady-state estimation for equilibrium processes*. Technical Report, University of Wisconsin.
11. Glynn, P.W. (1987). Upper bounds on Poisson tail probabilities. *Operations Research Letters* 6: 9–14.
12. Glynn, P.W. (1987). Limit theorems for the method of replication. *Stochastic Models* 3: 343–350.

13. Glynn, P.W. (1989). *Gradient estimation with common random numbers.* Technical Report in preparation.
14. Glynn, P.W. (1985). *Randomized estimators and time integrals.* Technical Report, Mathematics Research Center, University of Wisconsin–Madison.
15. Glynn, P.W. & Whitt, W. (1986). *The efficiency of simulation estimators.* Technical Report. AT&T Bell Laboratories, Holmdel, New Jersey.
16. Gross, D., Miller, D.R., & Plastiras, C.G. (1984). *Simulation methodologies for transient Markov processes: A comparative study based on multi-echelon reparable item inventory systems.* Proceedings Summer Simulation Conference, Boston, Massachusetts.
17. Heidelberger, P., Cao, X.R., Zazanis, M.A., & Suri, R. (1988). *Convergence properties of infinitesimal perturbation analysis estimates. Management Science* 34: 1281–1302.
18. Johnson, L.W. & Riess, R.D. (1982). *Numerical analysis.* Reading, Massachusetts: Addison-Wesley.
19. Zazanis, M.A. & Suri, R. (1986). *Comparison of perturbation analysis with conventional sensitivity estimates for stochastic systems.* Technical Report, University of Wisconsin.

Probability in the Engineering and Informational Sciences, 3, 1989, 319–321. Printed in the U.S.A.

ESTIMATING THE MEAN NUMBER OF RENEWALS BY SIMULATION

SHELDON M. ROSS
Department of Industrial Engineering and Operations Research
University of California
Berkeley, California 94720

An estimator, based on a simulation, is given for the expected number of events by time t of a renewal process. The estimator is obtained by combining the techniques of control variates and conditional expectation.

ESTIMATING THE RENEWAL FUNCTION

Suppose, possibly as part of a larger study, that we are interested in using simulation to estimate, for a fixed value t, $E[N(t)]$, where $\{N(s), s \geq 0\}$ is a renewal process having interarrival distribution F. Let X_i, $i \geq 1$, denote the interarrival times, so that

$$N(t) = \max\left\{n: \sum_{i=1}^{n} X_i \leq t\right\},$$

and let $\mu = E[X_i]$.

The above process is easily simulated by generating the interarrival times until their sum exceeds t. Letting $N(t)$ – the raw simulation estimator – denote the number of simulated events by time t, then a natural quantity to use as a control is the sequence of $N(t) + 1$ interevent times that were generated. That is, we make use of the fact that

$$E\left[\sum_{i=1}^{N(t)+1} (X_i - \mu)\right] = 0.$$

This research was supported by the Air Force Office of Scientific Research under Grant AFOSR-86-0153 with the University of California.