The Semi-Regenerative Method of Simulation Output Analysis

JAMES M. CALVIN
New Jersey Institute of Technology
PETER W. GLYNN
Stanford University
and
MARVIN K. NAKAYAMA
New Jersey Institute of Technology

We develop a class of techniques for analyzing the output of simulations of a semi-regenerative process. Called the semi-regenerative method, the approach is a generalization of the regenerative method, and it can increase efficiency. We consider the estimation of various performance measures, including steady-state means, expected cumulative reward until hitting a set of states, derivatives of steady-state means, and time-average variance constants. We also discuss importance sampling and a bias-reduction technique. In each case, we develop two estimators: one based on a simulation of a single sample path, and the other a type of stratified estimator in which trajectories are generated in an i.i.d. manner. We establish a central limit theorem for each estimator, so confidence intervals can be constructed.

1. INTRODUCTION

A stochastic process is *regenerative* if, loosely speaking, there exists an infinite sequence of random times, known as regeneration points, at which the process probabilistically restarts. For example, for a positive-recurrent irreducible Markov chain on a discrete state space S, the successive hitting times to a fixed state form one possible regeneration sequence. A sample path of a regenerative process can be divided into independent and identically distributed (i.i.d.) cycles based on the se-

quence of regeneration points. The *regenerative method* (RM) of simulation output analysis [Crane and Iglehart 1975] uses this structure to construct asymptotically valid confidence intervals for the steady-state mean of a regenerative process.

There are several settings in which exploiting regenerative structure and applying the RM lead to improvements over other methods. For example, the only known estimator of the time-average variance constant with convergence rate $t^{-1/2}$, where t is the run length of the simulation, is based on the RM [Henderson and Glynn 2001]. (The time-average variance constant is the variance constant appearing in the central limit theorem for the time average of a process.) Several bias-reduction techniques rely on regenerative structure [Meketon and Heidelberger 1982; Glynn 1994]. Also, it is known that the variance of likelihood ratio derivative estimators (resp., importance-sampling estimators) grows linearly (resp., exponentially) in the length of observations [Reiman and Weiss 1989; Glynn 1990; 1995], so breaking up a sample path into regenerative cycles can be beneficial.

For many regenerative processes there is more than one choice of regeneration sequence to use for the RM. For example, for a Markov chain, returns to any fixed state constitute regenerations. In such settings it would be useful to have methods that exploit multiple regeneration sequences.

In this paper we present a general approach for taking advantage of multiple regeneration sequences. We call it the "semi-regenerative method" (SRM) because of its relationship to the theory of semi-regenerative stochastic processes (Section 10.6 of [Çinlar 1975]). We develop the SRM in the context of Markov chains on a discrete state space S, with the goal of obtaining estimators that have smaller variance than their regenerative counterparts.

The basic idea of the SRM is to fix a set of states $A \subset S$, and we define a *trajectory* as a sample path beginning in a state in A until the first return to A. Then we derive a new representation for the performance measure of interest in terms of expectations of functionals of trajectories. The semi-regenerative estimator results by replacing each expectation with a simulation-based estimator of it.

We develop the SRM for several different classes of performance measures, and in each case, we define two estimators. One is based on simulating a single (long) sample path, which we then divide into trajectories. The other uses a type of stratification, in which trajectory "segments" are sampled in an i.i.d. manner. We establish central limit theorems for each of our semi-regenerative estimators, thus enabling one to construct asymptotically valid confidence intervals.

Other methods for simulating processes with multiple regeneration sequences have been proposed in the literature, including the almost regenerative method (ARM) [Gunther and Wolff 1980], A-segments [Zhang and Ho 1992], semi-stationary processes [Alexopoulos and Shultes 1998], and permuted regenerative estimators [Calvin and Nakayama 1998; 2000b; 2000a]. The other methods, which all result in estimators that differ from semi-regenerative estimators, are based on simulating a single sample path. In addition to our semi-regenerative estimators based on a single sample path, we also consider stratified estimators, which have no analogues with the other methods.

Gunther and Wolff [1980] developed the ARM only for estimating the steady-state mean reward. It fixes two disjoint sets of states, U and V, and divides a sample

path into *almost regenerative cycles* that begin and end with transitions from the set U to the set V. To relate this to the semi-regenerative estimator, let $V = A$ and $U = S \backslash A$, and note that the SRM allows for trajectories consisting of one transition from A back to A, whereas this cannot be an almost regenerative cycle. (Similarly, the approach of [Alexopoulos and Shultes 1998] for semi-stationary processes does not allow transitions from A directly back to A.) If we remove the restriction that the sets U and V are disjoint in the ARM, then the resulting ARM point estimator is the same as the SRM estimator when using one long sample path.

Zhang and Ho [1992] developed the A-segments method to reduce the variance of likelihood ratio derivative estimators. Their technique breaks up a sample path into A-segments determined by returns to the set A, as is done with the SRM, but they construct their estimator using an approach that differs from ours and end up with a different estimator. Also, Zhang and Ho only apply their method to likelihood ratio derivative estimation, and they do not prove a central limit theorem for their estimator, as we do.

Permuted regenerative estimators [Calvin and Nakayama 1998; 2000b; 2000a] are constructed by first running a simulation of a fixed number of cycles from one regeneration sequence. Then for each regeneration sequence, permute the cycles of that sequence along the generated path. Compute an estimate of the performance measure based on this permuted sample path, and averaging over all permuted paths yields the permuted estimator.

Calvin and Nakayama [2002] analyze the difference between semi-regenerative and permuted estimators when estimating the second moment of a cycle reward when $|A| = 2$. They also compare the two estimators to another, a type of V-statistic estimator, which resamples trajectories with replacement. They demonstrate that the three estimators are not the same in general, but they are asymptotically equivalent and satisfy the same central limit theorem. Specifically, they show that for the performance measure considered, the permuted estimator is unbiased and the other two estimators have positive bias, with the bias of the SRM estimator being at least as large of that for the V-statistic estimator.

An alternative approach to using multiple regeneration sequences is to try to increase the frequency of regenerations from a single sequence. Andradóttir et al. [1995] discuss such an approach for simulation of Markov chains. Instead of regenerations occurring with each visit to a fixed state, regenerations may occur (with a certain state-dependent probability) for visits to many states. In the case of the regenerative estimator of the time-average variance constant, basing the estimator on a regenerative subsequence of a regeneration sequence produces an estimator with at least as large variance.

The rest of the paper has the following structure. In Section 2 we develop the mathematical framework for the paper. Throughout this paper we restrict the setting to discrete-time Markov chains on a discrete state space, but the methods also apply to more general semi-regenerative processes. In Section 3 we derive semi-regenerative estimators for steady-state means, and in Section 5 we develop estimators that incorporate importance sampling for estimating steady-state means. We construct estimators for the expected cumulative reward until hitting a set of states, the gradient of a steady-state mean, and the time-average variance constant

in Sections 6, 7, and 4, respectively. In Section 8 we derive a semi-regenerative version of a regenerative low-bias estimator. In Section 9 we consider ratios of steady-state means. We close with some concluding remarks.

Calvin et al. [2001] present (without proofs) some of the results from the current paper. In particular, Calvin et al. [2001] present the semi-regenerative estimator based on a single sample path for the expected cumulative reward until hitting a set, as well as some importance sampling estimators for this measure, which are not in the current paper. Also, Calvin et al. [2001] give some empirical results.

2. MATHEMATICAL FRAMEWORK

Let $X = \{X_j : j = 0, 1, 2, \ldots\}$ be a discrete-time Markov chain (DTMC) on a finite or countably infinite state space S. Let $\Pi = (\Pi(x, y) : x, y \in S)$ be the transition probability matrix of X, and let P_x (resp., E_x, Var_x, and Cov_x) denote the probability measure (resp., expectation, variance, and covariance) given that $X_0 = x$, $x \in S$.

Assumption 2.1. The DTMC X with transition probability matrix Π is irreducible and positive recurrent.

Under Assumption 2.1, X has a unique stationary distribution $\pi = (\pi(x) : x \in S)$, which is the row-vector solution to $\pi = \pi\Pi$ with $\sum_{x \in S} \pi(x) = 1$ and $\pi(x) > 0$ for all $x \in S$.

3. STEADY-STATE MEANS

Let $f : S \to \Re$ be a "reward" function. Our goal is to estimate $\alpha = \pi f \equiv \sum_{x \in S} \pi(x) f(x)$.

Assumption 3.1. The reward function f satisfies

$$\sum_{x \in S} \pi(x) |f(x)| < \infty.$$

3.1 The Regenerative Method

Consider first the regenerative method [Crane and Iglehart 1975]. For $x \in S$, define $\tau_x = \inf\{j \geq 1 : X_j = x\}$. Fix a "return state" $w \in S$. Under Assumptions 2.1 and 3.1, the RM is based on the identity

$$\alpha = \frac{E_w\left[\sum_{j=0}^{\tau_w - 1} f(X_j)\right]}{E_w[\tau_w]}. \tag{1}$$

The moments in (1) are estimated by generating independent copies of

$$\left(\sum_{j=0}^{\tau_w - 1} f(X_j), \tau_w\right) \quad \text{under measure } P_w \tag{2}$$

and forming the sample means. Specifically, let $T_{w,0} = \inf\{j \geq 0 : X_j = w\}$ and $T_{w,k} = \inf\{j > T_{w,k-1} : X_j = w\}$ for $k \geq 1$. Define $\tau_{w,k} = T_{w,k} - T_{w,k-1}$, for $k \geq 1$. Also, define $Y_{w,k} = \sum_{j=T_{w,k-1}}^{T_{w,k}-1} f(X_j)$ for $k \geq 1$. Now fix a large integer n and run a simulation of X up to time $T_{w,n}$, giving a sample path $\{X_j : j = 0, 1, \ldots, T_{w,n}\}$. The

$(Y_{w,k}, \tau_{w,k})$, $k = 1, 2, \ldots, n$, are i.i.d. copies of (2). Set $\bar{Y}_{w,n} = (1/n) \sum_{k=1}^{n} Y_{w,k}$ and $\bar{\tau}_{w,n} = (1/n) \sum_{k=1}^{n} \tau_{w,k}$. Then the regenerative estimator of α is $\widetilde{\alpha}_{w,n} \equiv \bar{Y}_{w,n}/\bar{\tau}_{w,n}$.

Let $\mathcal{N}(\kappa, \Phi)$ denote a normal distribution with mean vector κ and covariance matrix Φ, and let $\xrightarrow{\mathcal{D}}$ denote convergence in distribution. We can form an asymptotically valid confidence interval for α based on the following central limit theorem (e.g., see p. 100 of [Shedler 1993]).

PROPOSITION 3.2. *If Assumption 2.1 holds and if* $E_w\left[\left(\sum_{j=0}^{\tau_w-1}|f(X_j)|\right)^2\right] < \infty$ *and* $E_w[\tau_w^2] < \infty$, *then*

$$n^{1/2}(\widetilde{\alpha}_{w,n} - \alpha) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \widetilde{\sigma}^2)$$

as $n \to \infty$, *where* $\widetilde{\sigma}^2 = (\,Var[Y_{w,k}] - 2\alpha\,Cov(Y_{w,k}, \tau_{w,k}) + \alpha^2\,Var[\tau_{w,k}])/E_w[\tau_{w,k}]$.

3.2 The Semi-Regenerative Estimator for Steady-State Means

We will now develop another estimator for α. Fix a set of states $A \subset S$, $A \neq \emptyset$, and set

$$
\begin{aligned}
T_0 &= \inf\{j \geq 0 : X_j \in A\}, \\
T_k &= \inf\{j > T_{k-1} : X_j \in A\}, \quad k \geq 1, \\
T &= T_1, \\
W_k &= X_{T_k}, \quad k \geq 0.
\end{aligned}
$$

The following result follows from pp. 314–315 of [Çinlar 1975].

PROPOSITION 3.3. *Under Assumption 2.1, $W = \{W_k : k \geq 0\}$ is an irreducible, positive-recurrent discrete-time Markov chain with state space A.*

The process W is sometimes called the "chain on A." Define $R(x,y) = P_x(X_T = y)$ for $x, y \in A$, and let $R = (R(x,y) : x, y \in A)$, which is the transition probability matrix of W. Under Assumption 2.1, Proposition 3.3 implies the existence of a unique stationary distribution $\nu = (\nu(x) : x \in A) \in \Re^{1 \times d}$ for W; i.e., ν is the row vector satisfying $\nu R = \nu$ with $\sum_{x \in A} \nu(x) = 1$ and $\nu(x) > 0$ for all $x \in A$. Let E_ν denote expectation with initial distribution ν. We assume the following:

Assumption 3.4. $|A| = d < \infty$, with $A = \{x_1, x_2, \ldots, x_d\}$.

The SRM is based on the following identity.

PROPOSITION 3.5. *If Assumptions 2.1, 3.1, and 3.4 hold, then*

$$\alpha = \frac{E_\nu\left[\sum_{j=0}^{T-1} f(X_j)\right]}{E_\nu[T]} = \frac{\sum_{i=1}^{d} \nu(x_i) E_{x_i}\left[\sum_{j=0}^{T-1} f(X_j)\right]}{\sum_{i=1}^{d} \nu(x_i) E_{x_i}[T]}.$$

We defer the proof to Remark 3.9 before Theorem 3.10 in Section 3.3. Çinlar [1975], Theorem 10.6.12, provides a proof of this result under different assumptions when the function f is of the form $f(x) = I(x \in B)$ for some set of states $B \subset S$, where $I(\cdot)$ is the indicator function. Also, see [Zhang and Ho 1992].

Using the semi-regenerative identity in Proposition 3.5, we will now develop an estimator for α using a type of stratified sampling. Let

$$Y = \sum_{j=0}^{T-1} f(X_j),$$
$$\tau = T,$$
$$\chi(y) = I(X_T = y), \text{ for } y \in A.$$

Let $p_1, p_2, \ldots, p_d$, be d positive numbers summing to one. Given a "replication budget" n, we will sample $\lfloor p_i n \rfloor$ times from the initial state $x_i \in A$, where for $a \in \Re$, $\lfloor a \rfloor$ is the greatest integer less than or equal to a. Specifically, for each $i = 1, 2, \ldots, d$, let

$$(Y_k(x_i), \tau_k(x_i), \chi_k(x_i, y) : y \in A)$$

for $1 \leq k \leq \lfloor p_i n \rfloor$ be i.i.d. copies of

$$(Y, \tau, \chi(y) : y \in A) \text{ under measure } P_{x_i}.$$

Set

$$R_n(x_i, y) = \frac{1}{\lfloor p_i n \rfloor} \sum_{k=1}^{\lfloor p_i n \rfloor} \chi_k(x_i, y)$$

for $1 \leq i \leq d$ and $y \in A$, and set $R_n = (R_n(x, y) : x, y \in A) \in \Re^{d \times d}$. Clearly,

$$R_n \to R \text{ a.s.} \tag{3}$$

as $n \to \infty$ by the strong law of large numbers. Since R is irreducible and positive recurrent by Proposition 3.3, R_n also is for sufficiently large n, so there exists a unique stationary distribution $\nu_n = (\nu_n(x) : x \in A) \in \Re^{1 \times d}$ for R_n for n sufficiently large; i.e., ν_n satisfies $\nu_n = \nu_n R_n$ with $\sum_{x \in A} \nu_n(x) = 1$ and $\nu_n(x) > 0$ for all $x \in A$ for sufficiently large n by Proposition 3.3. We define the semi-regenerative estimator of α to be

$$\alpha_n = \frac{\sum_{i=1}^{d} \nu_n(x_i) \frac{1}{\lfloor p_i n \rfloor} \sum_{k=1}^{\lfloor p_i n \rfloor} Y_k(x_i)}{\sum_{i=1}^{d} \nu_n(x_i) \frac{1}{\lfloor p_i n \rfloor} \sum_{k=1}^{\lfloor p_i n \rfloor} \tau_k(x_i)}. \tag{4}$$

If $|A| = 1$, then α_n is the standard regenerative estimator of α.

The estimator α_n is a type of stratified estimator (Chapter 5 of [Cochran 1977]), in which starting a trajectory from x_i is effectively a sample from stratum i, $1 \leq i \leq d$. In general, for an arbitrary choice of the p_i's, the estimator α_n can have a variance very different from that of the point estimator $(1/T_n) \sum_{j=0}^{T_n-1} f(X_j)$ based on one long sample path of length T_n. Later in this section we discuss methods for determining the weights p_i to minimize the asymptotic variance of α_n. Also, we will carry out an asymptotic comparison of the point estimator of one long sample path and the semi-regenerative estimator in Section 3.3.

We now wish to develop a central limit theorem for the semi-regenerative estimator α_n. For this we will need to make some moment assumptions.

Assumption 3.6. There exists $w \in S$ such that

$$E_w\left[\tau_w^2\right] < \infty, \qquad E_w\left[\left(\sum_{j=0}^{\tau_w-1} |f(X_j)|\right)^2\right] < \infty.$$

PROPOSITION 3.7. *Under Assumptions 2.1 and 3.6, $E_x[Y^2] < \infty$ and $E_x[\tau^2] < \infty$ for all $x \in A$.*

For a proof, see Theorem 4, page 84, of [Chung 1967].

We now state our central limit theorem for α_n. For $i = 1, 2, \ldots, d$, define the matrix $\Psi_i = (\Psi_i(x, y) : x, y \in A)$ with entries

$$\Psi_i(x_j, x_k) = \mathrm{Cov}_{x_i}(\chi(x_j), \chi(x_k)) = -R(x_i, x_j)R(x_i, x_k), \;\; j \neq k, \tag{5}$$

and

$$\Psi_i(x_j, x_j) = \mathrm{Var}_{x_i}(\chi(x_j)) = R(x_i, x_j)(1 - R(x_i, x_j)). \tag{6}$$

For $i, j = 1, 2, \ldots, d$, define $g(x_i, x_j) = \mathrm{Cov}_{x_i}(\chi(x_j), Z)$ and $h(x_i) = \mathrm{Var}_{x_i}(Z)$, where $Z = Y - \alpha\tau$. Let $z = (z(x) : x \in A) \in \Re^{d \times 1}$ with $z(x) = E_x[Z]$. Let e denote the vector of all 1's in $\Re^{d \times 1}$, and let V be the matrix in $\Re^{d \times d}$ in which all rows are equal to ν; i.e., $V = e\nu$. Let $F = (I - R + V)^{-1}$, the fundamental matrix of W, which exists under Assumptions 2.1 and 3.4 by Proposition 3.3 (e.g., see p. 100 of [Kemeny and Snell 1960]). Finally, let $\zeta = (\zeta(x) : x \in A) \in \Re^{d \times 1}$ be defined by $\zeta = Fz$. The following establishes a central limit theorem for α_n; see the appendix for the proof.

THEOREM 3.8. *Under Assumptions 2.1, 3.1, 3.4, and 3.6,*

$$n^{1/2}(\alpha_n - \alpha) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma^2)$$

as $n \to \infty$, where

$$\sigma^2 = \frac{1}{(E_\nu[T])^2} \sum_{i=1}^{d} \frac{\nu^2(x_i)\eta_i}{p_i}, \tag{7}$$

with

$$\eta_i = h(x_i) + 2\sum_{j=1}^{d} g(x_i, x_j)\zeta(x_j) + \sum_{j=1}^{d}\sum_{l=1}^{d} \zeta(x_j)\zeta(x_l)\Psi_i(x_j, x_l). \tag{8}$$

We can consistently estimate σ^2 by estimating each of the quantities in (7) as follows:

—Estimate $E_\nu[T]$ by $\nu_n \bar{\tau}_n$.
—Estimate $\nu(x_i)$ by $\nu_n(x_i)$.
—Estimate $h(x_i)$ by

$$h_n(x_i) = \frac{1}{\lfloor p_i n \rfloor - 1} \sum_{k=1}^{\lfloor p_i n \rfloor} (Y_k(x_i) - \bar{Y}_n(x_i))^2 + \frac{\alpha_n^2}{\lfloor p_i n \rfloor - 1} \sum_{k=1}^{\lfloor p_i n \rfloor} (\tau_k(x_i) - \bar{\tau}_n(x_i))^2$$

$$- \frac{2\alpha_n}{\lfloor p_i n \rfloor - 1} \sum_{k=1}^{\lfloor p_i n \rfloor} (Y_k(x_i) - \bar{Y}_n(x_i))(\tau_k(x_i) - \bar{\tau}_n(x_i)),$$

where $\bar{Y}_n(x_i) = (1/\lfloor p_i n \rfloor) \sum_{k=1}^{\lfloor p_i n \rfloor} Y_k(x_i)$ and $\bar{\tau}_n(x_i) = \frac{1}{\lfloor p_i n \rfloor} \sum_{k=1}^{\lfloor p_i n \rfloor} \tau_k(x_i)$.

—Estimate ζ by $\zeta_n = (I - R_n + V_n)^{-1}(\bar{Y}_n - \alpha_n \bar{\tau}_n)$, where $V_n = e\nu_n$ and $\bar{Y}_n = (\bar{Y}_n(x) : x \in A)$.

—Estimate $g(x_i, x_j)$ by

$$g_n(x_i, x_j)$$
$$= \frac{1}{\lfloor p_i n \rfloor - 1} \sum_{k=1}^{\lfloor p_i n \rfloor} [\chi_k(x_i, x_j) - R_n(x_i, x_j)][(Y_k(x_i) - \alpha_n \tau_k(x_i)) - \widetilde{Z}_n(x_i)],$$

where $\widetilde{Z}_n(x_i) = \bar{Y}_n(x_i) - \alpha_n \bar{\tau}_n(x_i)$.

—Estimate $\Psi_i(x_j, x_l)$ by

$$\Psi_{i,n}(x_j, x_l) = \begin{cases} -R_n(x_i, x_j)R_n(x_i, x_l) & \text{if } j \neq l, \\ R_n(x_i, x_j)(1 - R_n(x_i, x_j)) & \text{if } j = l. \end{cases}$$

The resulting estimator $\widehat{\sigma}_n^2$ of σ^2 is then strongly consistent and can be used to construct an asymptotically valid (as $n \to \infty$) $100(1-\delta)\%$ confidence interval for α given by $(\alpha_n - a_\delta \widehat{\sigma}_n/\sqrt{n}, \ \alpha_n + a_\delta \widehat{\sigma}_n/\sqrt{n})$, where a_δ is chosen so that $P(\mathcal{N}(0,1) \geq a_\delta) = \delta/2$.

The semi-regenerative approach opens up the possibility of stratification (i.e., choosing the $\{p_i\}$) in a way that is impossible to implement in the regenerative context. This is a potentially important "additional degree of freedom" that does not have a regenerative analogue.

For a fixed subset A of states, we now consider the problem of choosing the optimal $\{p_i\}$ to minimize σ^2. By (7) we can write

$$\sigma^2 \equiv \sigma^2(p_1, p_2, \ldots, p_d) = \sum_{i=1}^{d} \frac{c_i}{p_i},$$

where $c_i = \eta_i(\nu(x_i)/E_\nu[T])^2$. It can be shown (see (55) in the appendix) that $c_i \geq 0$. Minimizing σ^2 subject to $\sum_{i=1}^{d} p_i = 1$ and $p_i \geq 0$ yields the optimal $\{p_i^*\}$ given by

$$p_i^* = \frac{\sqrt{c_i}}{\sum_{j=1}^{d} \sqrt{c_j}}; \qquad (9)$$

see Chapter 5 of [Cochran 1977]. Since the $\{c_i\}$ are typically unknown, one approach to apply this result in practice is to use a two-stage procedure. In the first stage, simulate a pilot run to estimate the $\{c_i\}$; one approach for implementing the pilot run is to fix a number n of trajectories, and for each $i = 1, 2, \ldots, d$, and start $n/|A|$ of them from state x_i. In the second stage, using the estimates of the $\{c_i\}$, simulate the production runs with estimates for $\{p_i^*\}$ determined by substituting the estimates of the $\{c_i\}$ into (9).

(Rather than choosing the p_i to minimize variance, we could instead maximize *efficiency*, which is often defined as the inverse of the product of the variance and work; see [Glynn and Whitt 1992]. To define the work in this case, we need to take into account the cost associated with sampling from each stratum i, and a reasonable measure of that cost is $E_{x_i}[\tau]$, the expected number of transitions in a

trajectory starting from x_i. Thus, maximizing efficiency corresponds to minimizing $\sigma^2(p_1, \ldots, p_d) \sum_{i=1}^{d} p_i E_{x_i}[\tau]$, subject to $\sum_{i=1}^{d} p_i = 1$ and $p_i \geq 0$, which in general cannot be solved in closed form, so that numerical methods need to be used.)

We performed numerical experiments to explore the possible benefit of choosing approximately optimal stratification weights to minimize variance using the two-stage procedure described above. The model is a Markov chain on state space $S = \{0, 1, \ldots, 11\}$ that models the maintenance of a machine subject to periodic breakdowns followed by repair intervals. We model the time until a new machine breaks down, and the time to return it to service as discrete phase-type distributions, as shown in Figure 1. States 0 and 1 are working states, and when the chain enters state 2 the machine is out of service. There are three types of repair, with respective probabilities 0.7, 0.27, and 0.03. After repair the system returns to state 0. The cost function f is given by $f(0) = f(1) = f(2) = f(3) = f(4) = 0$, $f(5) = f(6) = f(7) = f(8) = 1$, and $f(9) = f(10) = f(11) = 24$. The steady-state mean cost is approximately $\pi f \approx 0.7037$.

Table I presents simulation results with different choices of A, where we used pilot runs to approximate the weights in (9) that minimize variance. Each experiment consisted of 10^3 independent replications, with each replication simulating a production run of 2×10^6 transitions of the Markov chain, from which a point estimate of the steady-state mean was computed using (4). When we applied the SRM with approximately optimal weights, each replication started with an initial pilot run of 2×10^4 transitions, which was then followed by a production run; in this case, we used the data from the pilot run only to estimate the optimal weights and not to compute the final point estimates. For each choice of A, we computed the sample variance of the point estimates across the 10^3 replications, as well as the bias and mean-square error (MSE). The column labeled "time" provides the total CPU times (in minutes) required to run all 10^3 replications, which includes the time for all the pilot runs for the cases when they are used.

The first row of Table I shows the results for the standard regenerative estimator, with return state 0, and this required no pilot run. The next five rows show the results for the SRM with pilot runs in each replication to estimate the optimal stratification weights in (9). With A containing four states the simulation took about 50% longer than the standard regenerative estimator. The last column in Table I gives the relative efficiency compared to the regenerative method, where the efficiency is the reciprocal of the product of variance and CPU time. The most efficient choice is $A = \{0, 3, 6\}$, which is about 28% more efficient than the standard regenerative estimator. As $|A|$ increases the variance decreases, but the computational cost also increases. The stationary distribution on $A = \{0, 3, 6, 9\}$ is approximately $\nu = (0.662, 0.233, 0.086, 0.019)$, and the optimal weights are approximately $p^* = (0.836, 0.010, 0.027, 0.127)$.

3.3 Comparison of the Stratified Semi-Regenerative Estimator and the Estimator Based on One Long Path

We now compare the asymptotic behaviors of the semi-regenerative estimator α_n and the point estimator based on a simulation of one (long) sample path, which we define as follows. Fix an initial state $X_0 = x_0 \in A$. (We can also define an initial distribution μ on A or S to select X_0, but for simplicity, we just fix $X_0 = x_0$.)

Table I.   Experimental results.

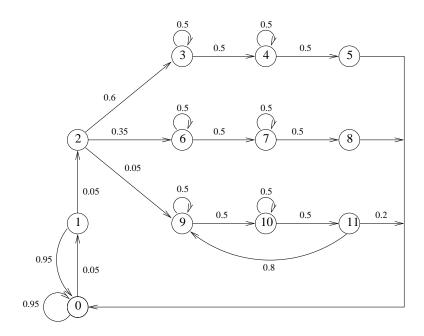| A	variance ($\times 10^{-5}$)	bias ($\times 10^{-4}$)	MSE ($\times 10^{-5}$)	time	efficiency
{0}	10.55	-0.05	10.6	21.30	1.0
{0,3}	7.65	1.45	7.7	24.52	1.19
{0,6}	8.76	-4.11	8.8	23.49	1.09
{0,9}	7.87	1.33	7.9	24.78	1.15
{0,3,6}	6.53	-0.42	6.5	27.19	1.28
{0,3,6,9}	5.76	-1.66	5.8	31.34	1.24

Fig. 1.   Transition probabilities of a machine-repair Markov chain.

Then simulate X up to time T_n, for n fixed (and large), giving us a sample path $\{X_j : j = 0, 1, 2, \ldots, T_n\}$, from which we obtain $\{W_k : k = 0, 1, 2, \ldots, n\}$ with $W_k = X_{T_k}$. Then the point estimator of α based on one long sample path is defined as

$$\alpha'_n = \frac{\sum_{j=0}^{T_n-1} f(X_j)}{T_n}. \tag{10}$$

Note that α'_n is identical to the regenerative estimator when $T_0 = 0$ and T_n are regeneration points.

To facilitate our asymptotic comparison of α'_n and α_n, we will re-express α'_n in a form similar to (4). For $x \in A$, define $H_n(x) = \sum_{k=0}^{n-1} I(W_k = x)$. For $x \in A$, define $T'_1(x) = \inf\{j \geq 0 : X_j = x\}$, and for $k \geq 2$, define $T'_k(x) = \inf\{j > T'_{k-1}(x) : X_j = x\}$. Also, define $\widetilde{T}'_k(x) = \inf\{j > T'_k(x) : X_j \in A\}$, which is the first time after

$T_k'(x)$ that X enters A again. For $x \in A$ and $k = 1, 2, \ldots,$ define

$$Y_k'(x) = \sum_{j=T_k'(x)}^{\widetilde{T}_k'(x)-1} f(X_j), \qquad \tau_k'(x) = \widetilde{T}_k'(x) - T_k'(x).$$

Also, for $x \in A$, define $\widehat{\nu}_n(x) = H_n(x)/n$. Then observe that

$$\alpha_n' = \frac{\sum_{i=1}^d \sum_{k=1}^{H_n(x_i)} Y_k'(x_i)}{\sum_{i=1}^d \sum_{k=1}^{H_n(x_i)} \tau_k'(x_i)} = \frac{\sum_{i=1}^d \widehat{\nu}_n(x_i) \frac{1}{H_n(x_i)} \sum_{k=1}^{H_n(x_i)} Y_k'(x_i)}{\sum_{i=1}^d \widehat{\nu}_n(x_i) \frac{1}{H_n(x_i)} \sum_{k=1}^{H_n(x_i)} \tau_k'(x_i)}. \tag{11}$$

Remark 3.9. Because of Assumption 3.1, all of the $Y_1'(x_i)$ have finite first moments. Also, Assumption 1 implies that each $H_n(x_i) \to \infty$ a.s., so it follows from (11) that $\alpha_n' \to E_\nu[\sum_{j=0}^{T-1} f(X_j)]/E_\nu[T]$ a.s. by the strong law of large numbers. In addition, (10) implies $\lim_{n\to\infty} \alpha_n' = \lim_{m\to\infty}(1/m)\sum_{j=0}^{m-1} f(X_j)$ a.s. since $T_n \to \infty$ a.s. Now $\lim_{m\to\infty}(1/m)\sum_{j=0}^{m-1} f(X_j) = \alpha$ a.s. by the strong law of large numbers for regenerative processes (e.g., see Theorem 2.2 on p. 74 of [Shedler 1993]), thereby establishing Proposition 3.5.

Note the similarity of α_n' in (11) and α_n in (4). We now establish the following central limit theorem for α_n'; see the appendix for the proof.

THEOREM 3.10. *Under Assumptions 2.1, 3.1, 3.4, and 3.6,*

$$n^{1/2}\left(\alpha_n' - \alpha\right) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_1^2)$$

as $n \to \infty$, where

$$\sigma_1^2 = \frac{1}{(E_\nu[T])^2} \sum_{i=1}^d \nu(x_i)\eta_i \tag{12}$$

with η_i defined in (8).

If we set $p_i = \nu(x_i)$, $i = 1, 2, \ldots, d$, in (7), then $\sigma^2 = \sigma_1^2$. Therefore, our semi-regenerative estimator α_n in (4) with $p_i = \nu(x_i)$ has the same asymptotic efficiency as the point estimator α_n' based on one long sample path (if we ignore the added computational cost of constructing α_n). However, choosing the $\{p_i\}$ according to (9) may lead to the semi-regenerative estimator having an asymptotic variance that is strictly smaller than that of α_n'.

As stated before, the regenerative estimator is equivalent to α_n' (when $T_0 = 0$ and T_n are regeneration points). We ran experiments using the steady-state probabilities as the stratification weights p_i in (4) and (7), and the results (not included in Table I) for the variance were consistent with the regenerative method but with increased computational time for the semi-regenerative estimator.

4. TIME-AVERAGE VARIANCE CONSTANT

Suppose that Assumptions 2.1, 3.1, 3.4, and 3.6 hold. Then

$$n^{1/2}\left(\frac{1}{n}\sum_{j=0}^{n-1} f(X_j) - \alpha\right) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \bar{\sigma}^2)$$

as $n \to \infty$; e.g., see p. 74 of [Shedler 1993]. Our goal is now to estimate $\bar{\sigma}^2$, which is known as the time-average variance constant. Assumption 2.1 implies that $T_n/n \to E_\nu[T]$ a.s. as $n \to \infty$ with $0 < E_\nu[T] < \infty$, so

$$T_n^{1/2} \left(\frac{1}{T_n} \sum_{j=0}^{T_n-1} f(X_j) - \alpha \right) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \bar{\sigma}^2)$$

by the random-time-change central limit theorem. But it follows from Theorem 3.10 that

$$\bar{\sigma}^2 = \frac{1}{E_\nu[T]} \sum_{i=1}^{d} \nu(x_i)\eta_i,$$

where η_i is defined in (8).

We now describe an unstratified semi-regenerative estimator of $\bar{\sigma}^2$ based on a single simulated sample path up to time T_n. To do this, we will express $\bar{\sigma}^2$ as a function of expectations of functionals of trajectories. We first define

$$y_1(x) = E_x[Y], \qquad y_2(x) = E_x[Y^2],$$
$$t_1(x) = E_x[\tau], \qquad t_2(x) = E_x[\tau^2], \qquad \upsilon(x) = E_x[Y\tau],$$

and let $y_1 = (y_1(x) : x \in A) \in \Re^{d \times 1}$, $y_2 = (y_2(x) : x \in A) \in \Re^{d \times 1}$, $t_1 = (t_1(x) : x \in A) \in \Re^{d \times 1}$, $t_2 = (t_2(x) : x \in A) \in \Re^{d \times 1}$, $\upsilon = (\upsilon(x) : x \in A) \in \Re^{d \times 1}$, $g = (g(x, y) : x, y \in A) \in \Re^{d \times d}$, and $\Psi = (\Psi_i(x_j, x_k) : i, j, k = 1, 2, \ldots, d) \in \Re^{d \times d \times d}$. Now define the function $r_\sigma : \Re^{d \times 1} \times \Re^{d \times 1} \times \Re^{d \times 1} \times \Re^{d \times 1} \times \Re^{d \times 1} \times \Re^{d \times d} \times \Re^{d \times d} \times \Re^{d \times d \times d} \times \Re^{1 \times d} \times \Re^{d \times d} \to \Re$ given by

$$r_\sigma(w_1, w_2, \ldots, w_{10})$$

$$= \frac{1}{w_9 w_3} \left\{ 2 w_9 w_6 w_{10} \left(w_1 - \frac{w_9 w_1}{w_9 w_3} w_3 \right) + \sum_{i=1}^{d} w_9(x_i) \left[w_2(x_i) - 2 \frac{w_9 w_1}{w_9 w_3} w_5(x_i) \right. \right.$$

$$+ \left(\frac{w_9 w_1}{w_9 w_3} \right)^2 w_4(x_i) - \left(w_1(x_i) - \frac{w_9 w_1}{w_9 w_3} w_3(x_i) \right)^2$$

$$\left. \left. + \left(w_{10} \left(w_1 - \frac{w_9 w_1}{w_9 w_3} w_3 \right) \right)^\top w_8(x_i, \cdot, \cdot) \left(w_{10} \left(w_1 - \frac{w_9 w_1}{w_9 w_3} w_3 \right) \right) \right] \right\},$$

and it is easy to show that $\bar{\sigma}^2 = r_\sigma(\mu_\sigma)$, where $\mu_\sigma = (y_1, y_2, t_1, t_2, \upsilon, g, R, \Psi, \nu, F)$. Recall from Section 3.3 our definitions of $H_n(x)$, $Y_k'(x)$, $\tau_k'(x)$, $\widehat{\nu}_n(x)$, and α_n'. Define $R_n' = (R_n'(x, y) : x, y \in A)$ with

$$R_n'(x, y) = \frac{\sum_{k=0}^{n-1} I(W_k = x, W_{k+1} = y)}{\sum_{k=0}^{n-1} I(W_k = x)}. \tag{13}$$

Also, define $\chi_k'(x, y) = I(X_{\widetilde{T}_k'(x)} = y)$, so we can re-express $R_n'(x, y)$ as

$$R_n'(x, y) = \frac{\sum_{k=1}^{H_n(x)} \chi_k'(x, y)}{H_n(x)}.$$

Our unstratified estimator $\bar{\sigma}_{n'}^2$ of $\bar{\sigma}^2$ is then

$$\bar{\sigma}_{n'}^2 = r_\sigma(\bar{Y}_n', \bar{Y}_{2,n}', \bar{\tau}_n', \bar{\tau}_{2,n}', \upsilon_n', g_n', R_n', \Psi_n', \widehat{\nu}_n, \widehat{F}_n'),$$

where $\bar{Y}'_n = (\bar{Y}'_n(x_i) : i = 1, 2, \ldots, d)$, $\bar{Y}'_{2,n} = (\bar{Y}'_{2,n}(x_i) : i = 1, 2, \ldots, d)$, $\bar{\tau}'_n = (\bar{\tau}'_n(x_i) : i = 1, 2, \ldots, d)$, $\bar{\tau}'_{2,n} = (\bar{\tau}'_{2,n}(x_i) : i = 1, 2, \ldots, d)$, $v'_n = (v'_n(x_i) : i = 1, 2, \ldots, d)$, $g'_n = (g'_n(x_i) : i = 1, 2, \ldots, d)$, $\Psi'_n = (\Psi'_{i,n}(x_j, x_k) : i, j, k = 1, 2, \ldots, d)$, and $\widehat{F}'_n = (I - R'_n - e\widehat{\nu}_n)^{-1}$ with

$$\bar{Y}'_n(x_i) = \frac{1}{H_n(x_i)} \sum_{k=1}^{H_n(x_i)} Y'_k(x_i), \qquad \bar{Y}'_{2,n}(x_i) = \frac{1}{H_n(x_i)} \sum_{k=1}^{H_n(x_i)} Y'^2_{k'}(x_i),$$

$$\bar{\tau}'_n(x_i) = \frac{1}{H_n(x_i)} \sum_{k=1}^{H_n(x_i)} \tau'_k(x_i), \qquad \bar{\tau}'_{2,n}(x_i) = \frac{1}{H_n(x_i)} \sum_{k=1}^{H_n(x_i)} \tau'^2_{k'}(x_i),$$

$$v'_n(x_i) = \frac{1}{H_n(x_i)} \sum_{k=1}^{H_n(x_i)} Y'_k(x_i)\, \tau'_k(x_i),$$

$$g'_n(x_i, x_j) = \frac{1}{H_n(x_i) - 1} \sum_{k=1}^{H_n(x_i)} [\chi'_k(x_i, x_j) - R'_n(x_i, x_j)][Z'_{n,k}(x_i) - \widetilde{Z}'_n(x_i)],$$

$$Z'_{n,k}(x_i) = Y'_k(x_i) - \alpha'_n \tau'_k(x_i), \qquad \widetilde{Z}'_n(x_i) = \bar{Y}'_n(x_i) - \alpha'_n \bar{\tau}'_n(x_i),$$

$$\Psi'_{i,n}(x_j, x_l) = \begin{cases} -R'_n(x_i, x_j)R'_n(x_i, x_l) & \text{if } j \neq l, \\ R'_n(x_i, x_j)(1 - R'_n(x_i, x_j)) & \text{if } j = l. \end{cases}$$

Using similar arguments as in the proof of Theorem 3.10, we can show that

$$n^{1/2}\left[(\bar{Y}'_n, \bar{Y}'_{2,n}, \bar{\tau}'_n, \bar{\tau}'_{2,n}, v'_n, g'_n, R'_n, \Psi'_n, \widehat{\nu}_n, \widehat{F}'_n) - \mu_\sigma\right]$$
$$\xrightarrow{\mathcal{D}} (\bar{N}'_1, \bar{N}'_2, \ldots, \bar{N}'_{10}) \overset{\mathcal{D}}{=} \mathcal{N}(0, \Sigma'_\sigma),$$

for some normal random elements $(\bar{N}'_1, \bar{N}'_2, \ldots, \bar{N}'_{10})$, having a covariance matrix Σ'_σ. As in the case of the covariance matrix Σ' in the proof of Theorem 3.10, many of the entries in Σ'_σ are zero. We will not give all of the non-zero entries of Σ'_σ, but arguing as in the proof of Theorem 3.10, we can show, for example, that

$$\text{cov}(N'_1(x_i), N'_2(x_j)) = \begin{cases} \frac{1}{\nu(x_i)}\left(E_{x_i}[Y^3] - y_1(x_i)y_2(x_i)\right) & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases}.$$

Now we can show (see (57) and (60) of the appendix) that as $n \to \infty$,

$$n^{1/2}(\widehat{\nu}_n - \nu) = n^{1/2}\widehat{\nu}_n(R'_n - R)F + n^{1/2}(\widehat{\nu}_n - \nu'_n)(I - R'_n)F \xrightarrow{\mathcal{D}} \bar{N}'_9 = \nu \bar{N}'_7 F. \quad (14)$$

Also, let $\widehat{V}_n = e\widehat{\nu}_n$, so for sufficiently large n,

$$\begin{aligned}
(R'_n - R) &+ (\widehat{V}_n - V) \\
&= (I - R - V) - (I - R'_n - \widehat{V}_n) \\
&= \left[I - (I - R'_n - \widehat{V}_n)(I - R - V)^{-1}\right](I - R - V) \\
&= (I - R'_n - \widehat{V}_n)\left[(I - R'_n - \widehat{V}_n)^{-1} - (I - R - V)^{-1}\right](I - R - V) \\
&= (I - R'_n - \widehat{V}_n)\left[\widehat{F}'_n - F\right](I - R - V),
\end{aligned}$$

and it follows that

$$n^{1/2}\left(\widehat{F}_n' - F\right) = \widehat{F}_n'\left[n^{1/2}(R_n' - R) + n^{1/2}(\widehat{V}_n - V)\right]F.$$

Consequently, $\bar{N}_{10}' = F(\bar{N}_7' - e\nu\bar{N}_7'F)F$ by (14). Now applying the delta method (e.g., Theorem A, p. 122 of [Serfling 1980]) results in the following central limit theorem for the estimator $\bar{\sigma}_{n'}^2$.

THEOREM 4.1. *Suppose Assumptions 2.1, 3.1, and 3.4 hold, and also that there exists $w \in S$ such that $E_w[\tau_w^4] < \infty$ and $E_w[(\sum_{j=0}^{T-1}|f(X_j)|)^4] < \infty$. Then*

$$n^{1/2}(\bar{\sigma}_{n'}^2 - \bar{\sigma}^2) \xrightarrow{\mathcal{D}} \mathcal{N}(0, D_\sigma^\top \Sigma_\sigma D_\sigma),$$

where D_σ is the vector of partial derivatives of the function r_σ evaluated at μ_σ.

It is straightforward to compute the entries of D_σ, which we largely omit. For example, letting $\frac{\partial}{\partial w_2(x_j)}$ denote the partial derivative with respect to $w_2(x_j)$, we see that

$$\left.\frac{\partial}{\partial w_2(x_j)}r_\sigma(w_1, \ldots, w_{10})\right|_{(w_1,\ldots,w_{10})=\mu_\sigma} = w_9(x_j).$$

Also, we can similarly define a stratified estimator for $\bar{\sigma}^2$, but we omit this.

5. IMPORTANCE SAMPLING FOR STEADY-STATE MEANS

Importance sampling is a variance-reduction technique that can lead to dramatic decreases in variance (when applied appropriately), especially when used in rare-event simulations; see [Glynn and Iglehart 1989] for an overview of importance sampling. We now show how to combine importance sampling with the SRM to estimate steady-state means.

Let $\mathcal{F}_{x,T}$ denote the filtration of the process X up to time T with $X_0 = x$. For $x \in A$, define $P_{x,T}$ to be the probability measure on $\mathcal{F}_{x,T}$ for the process X under the transition probability matrix Π given $X_0 = x$. Now suppose that for each $x \in A$, we define another probability measure $P_{x,T}^*$ (not necessarily Markovian) on $\mathcal{F}_{x,T}$ for X conditional on $X_0 = x$, and let $E_{x,T}^*$ be the corresponding expectation. Also, let P_x^* (resp., E_x^*) be the probability measure (resp., expectation operator) for X induced by the collection of measures $(P_{y,T}^* : y \in A)$, given $X_0 = x$. We need to assume the following.

Assumption 5.1. For each $x \in A$, $P_{x,T}$ is absolutely continuous with respect to $P_{x,T}^*$.

By the Radon-Nikodym theorem (Theorem 32.2 of [Billingsley 1995]), Assumption 5.1 guarantees the existence of a non-negative random variable $L \equiv L(x)$ for which

$$P_{x,T}(C) = E_{x,T}^*[I(C)\,L], \quad C \in \mathcal{F}_{x,T}. \tag{15}$$

The random variable $L = dP_{x,T}/dP_{x,T}^*$ is known as the likelihood ratio (or Radon-Nikodym derivative) of $P_{x,T}$ with respect to $P_{x,T}^*$ (given $X_0 = x$). For example, if the measure $P_{x,T}^*$ is induced by a transition probability matrix $\Pi_x^* = (\Pi_x^*(w,y) : w,y \in S)$, then Assumption 5.1 will hold if $\Pi_x^*(w,y) = 0$ implies $\Pi(w,y) = 0$ for all

$w, y \in S$, and the likelihood ratio for the sample-path trajectory $X_0, X_1, X_2, \ldots, X_T$, given $X_0 = x$, is $L = \prod_{j=0}^{T-1} \Pi(X_j, X_{j+1})/\Pi_x^*(X_j, X_{j+1})$. Thus, L is the likelihood of the observed path $(X_0, X_1, \ldots, X_T)$ under the original measure over the likelihood of the path under the new measure.

We use the importance-sampling measure $P_{x_0}^*$, $x_0 \in A$, to generate a sample path $\{X_j : j \geq 0\}$ of the process X as follows. Set $X_0 = x_0$, so $T_0 = 0$. Then using measure $P_{X_0,T}^*$, generate a sequence of states until set A is hit again, thereby yielding the trajectory $X_1, X_2, \ldots, X_{T_1}$. Now from state X_{T_1}, use measure $P_{X_{T_1},T}^*$ to generate a sequence of states until A is hit again, yielding $X_{T_1+1}, X_{T_1+2}, \ldots, X_{T_2}$. In general, at the kth hit to set A, the process is in state X_{T_k}, and we use measure $P_{X_{T_k},T}^*$ to generate a sequence of states until A is hit again, yielding $X_{T_k+1}, X_{T_k+2}, \ldots, X_{T_{k+1}}$. We define the process $W = \{W_k : k \geq 0\}$ by letting $W_k = X_{T_k}$.

The process X defined in this way may no longer be a Markov chain since we did not assume any particular structure (other than Assumption 5.1) for the measure P_x^*. On the other hand, no matter how the $P_{y,T}^*$, $y \in A$, are defined, the embedded process W is always a Markov chain.

PROPOSITION 5.2. *If Assumptions 2.1, 3.4, and 5.1 hold, then for all $x \in A$, W under measure P_x^* is an irreducible, positive-recurrent discrete-time Markov chain on A.*

PROOF. It is clear that W is a Markov chain. Assumptions 2.1 and 5.1 ensure that W is irreducible since any sample path of X having positive probability under the original measure P_x also has positive probability under the importance-sampling measure P_x^*. Thus, W is positive recurrent by Assumption 3.4. □

Define matrix $R^* = (R^*(x, y) : x, y \in A)$ with elements $R^*(x, y) = P_x^*(X_T = y)$, and note that R^* is the transition probability matrix of W under importance sampling. As shown in Proposition 5.2, Assumptions 2.1, 3.4, and 5.1 ensure that R^* is irreducible and positive recurrent, so R^* has a stationary distribution $\rho = (\rho(x) : x \in A)$.

We can write $\alpha = \pi f$ in Proposition 3.5 as

$$\alpha = \frac{\sum_{i=1}^d \nu(x_i) E_{x_i,T}^* [YL]}{\sum_{i=1}^d \nu(x_i) E_{x_i,T}^* [\tau L]} \tag{16}$$

by (15), where ν is the stationary distribution for the R matrix under the original measure, as before. Expression (16) forms the basis for some semi-regenerative approaches using importance sampling, which we will describe below. For more details on importance sampling in general, see [Glynn and Iglehart 1989].

An advantage of applying importance sampling in a semi-regenerative setting rather than using the RM is that even if all of the $P_{x,T}^*$, $x \in A$, correspond to the same underlying change of measure, the trajectories simulated using importance sampling are shorter (fewer transitions) in the SRM than in the regenerative setting. This suggests that the semi-regenerative estimator will have smaller variance since Glynn [1995] showed that the variance of the likelihood ratio grows approximately exponentially with the number of transitions. Moreover, the SRM has the additional benefit of allowing the $P_{x,T}^*$ to correspond to different underlying changes of measure

for the different $x \in A$, thereby allowing one to tailor the importance sampling for each $x \in A$.

5.1 Stratified Estimation

We start by describing two stratified sampling methods based on (16). For each $x_i \in A$, let

$$(L_k(x_i), Y_k(x_i), \tau_k(x_i), \chi_k(x_i, y) : y \in A)$$

for $1 \le k \le \lfloor p_i n \rfloor$ be i.i.d. copies of

$$(L, Y, \tau, \chi(y) : y \in A) \quad \text{under measure } P^*_{x_i, T}.$$

Set

$$\bar{R}_n(x_i, y) = \frac{1}{\lfloor p_i n \rfloor} \sum_{k=1}^{\lfloor p_i n \rfloor} \chi_k(x_i, y) L_k(x_i)$$

for $1 \le i \le d$ and $y \in A$, and let $\bar{R}_n = (\bar{R}_n(x, y) : x, y \in A)$. Since $E^*_x[\chi(y)L] = E_x[\chi(y)] = R(x, y)$ for all $x, y \in A$ by (15), we have that $\bar{R}_n \to R$ a.s. as $n \to \infty$. Using the fact that R is irreducible and positive recurrent by Proposition 3.3, we can show that $\bar{R}_n$ also is for sufficiently large n under Assumption 5.1. Hence, there exists a stationary distribution $\bar{\nu}_n = (\bar{\nu}_n(x) : x \in A)$ for $\bar{R}_n$. We define our first stratified semi-regenerative importance-sampling estimator of α to be

$$\alpha_n^* = \frac{\sum_{i=1}^d \bar{\nu}_n(x_i) \frac{1}{\lfloor p_i n \rfloor} \sum_{k=1}^{\lfloor p_i n \rfloor} Y_k(x_i) L_k(x_i)}{\sum_{i=1}^d \bar{\nu}_n(x_i) \frac{1}{\lfloor p_i n \rfloor} \sum_{k=1}^{\lfloor p_i n \rfloor} \tau_k(x_i) L_k(x_i)}. \tag{17}$$

To establish a central limit theorem for α_n^*, we need to assume the following.

Assumption 5.3. $E^*_x[Y^2 L^2] < \infty$ and $E^*_x[\tau^2 L^2] < \infty$ for all $x \in A$.

Note that Assumption 5.3 ensures that $E^*_x[\chi(y)L^2] < \infty$ for all $x, y \in A$ since $0 \le \chi(y) \le 1 \le \tau$. Let Var^*_x and Cov^*_x denote variance and covariance, respectively, under P^*_x. Using essentially the same argument that we applied to establish Theorem 3.8, we can prove the following central limit theorem for α_n^*.

THEOREM 5.4. *Under Assumptions 2.1, 3.1, 3.4, 5.1, and 5.3,*

$$n^{1/2} (\alpha_n^* - \alpha) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_*^2)$$

as $n \to \infty$, where

$$\sigma_*^2 = \frac{1}{(E_\nu[T])^2} \sum_{i=1}^d \frac{\nu^2(x_i)\eta_{i*}}{p_i}, \tag{18}$$

with

$$\eta_{i*} = h^*(x_i) + 2 \sum_{j=1}^d g^*(x_i, x_j)\zeta(x_j) + \sum_{j=1}^d \sum_{l=1}^d \zeta(x_j)\zeta(x_l)\Psi_i^*(x_j, x_l), \tag{19}$$

$h^*(x_i) = \text{Var}^*_{x_i}(YL - \alpha\tau L)$, $g^*(x_i, x_j) = \text{Cov}^*_{x_i}(\chi(x_j)L, ZL)$, *and* $\Psi_i^*(x_j, x_l) = \text{Cov}^*_{x_i}(\chi(x_j)L, \chi(x_l)L)$.

Observe that in (17), we use importance sampling to estimate $\nu(x_i)$. However, in some situations we might obtain a better (lower variance) estimate of $\nu(x_i)$ by instead using standard simulation (i.e., without importance sampling). (We will still use importance sampling to estimate $E^*_{x_i,T}[YL]$ and $E^*_{x_i,T}[\tau L]$.) To implement this idea, for each state $x_i \in A$, we will now generate two sets of trajectories starting from x_i, where some of the trajectories will be under the original measure $P_{x_i,T}$ and the others will be generated using the importance-sampling measure $P^*_{x_i,T}$, with all trajectories being mutually independent. Specifically, let $q_1, q_2, \ldots, q_d, r_1, r_2, \ldots, r_d$ be positive numbers such that $\sum_{i=1}^d (q_i + r_i) = 1$. Given a replication budget n, we will sample $\lfloor q_i n \rfloor$ (resp., $\lfloor r_i n \rfloor$) times from initial state $x_i \in A$ using the original measure $P_{x_i,T}$ (resp., importance-sampling measure $P^*_{x_i,T}$). Let

$$(\chi_k(x_i, y) : y \in A)$$

for $1 \le k \le \lfloor q_i n \rfloor$ be i.i.d. copies of

$$(\chi(y) : y \in A) \quad \text{under measure } P_{x_i,T},$$

and let

$$(L_k(x_i), Y_k(x_i), \tau_k(x_i))$$

for $1 \le k \le \lfloor r_i n \rfloor$ be i.i.d. copies of

$$(L, Y, \tau) \quad \text{under measure } P^*_{x_i,T},$$

where $(\chi_k(x_i, y) : y \in A)$ and $(L_k(x_i), Y_k(x_i), \tau_k(x_i))$ are generated independently. Define

$$\widetilde{R}_n(x_i, y) = \frac{1}{\lfloor q_i n \rfloor} \sum_{k=1}^{\lfloor q_i n \rfloor} \chi_k(x_i, y)$$

for $1 \le i \le d$ and $y \in A$, and define $\widetilde{\nu}_n = (\widetilde{\nu}_n(x) : x \in A)$ such that $\widetilde{\nu}_n = \widetilde{\nu}_n \widetilde{R}_n$ with $\widetilde{\nu}_n \ge 0$ and $\widetilde{\nu}_n e = 1$. Then we define another stratified semi-regenerative importance-sampling estimator of α to be

$$\alpha_n^{**} = \frac{\sum_{i=1}^d \widetilde{\nu}_n(x_i) \frac{1}{\lfloor r_i n \rfloor} \sum_{k=1}^{\lfloor r_i n \rfloor} Y_k(x_i) L_k(x_i)}{\sum_{i=1}^d \widetilde{\nu}_n(x_i) \frac{1}{\lfloor r_i n \rfloor} \sum_{k=1}^{\lfloor r_i n \rfloor} \tau_k(x_i) L_k(x_i)}, \tag{20}$$

which satisfies the following central limit theorem.

THEOREM 5.5. *Under Assumptions 2.1, 3.1, 3.4, 5.1, and 5.3,*

$$n^{1/2} \left(\alpha_n^{**} - \alpha \right) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_{**}^2)$$

as $n \to \infty$, where

$$\sigma_{**}^2 = \frac{1}{(E_\nu[T])^2} \sum_{i=1}^d \nu^2(x_i) \left(\frac{h^*(x_i)}{r_i} + \sum_{j=1}^d \sum_{l=1}^d \frac{\zeta(x_j)\zeta(x_l)\Psi_i^*(x_j, x_l)}{q_i} \right), \tag{21}$$

and $h^*(x_i)$ and $\Psi_i^*(x_j, x_l)$ are defined as in Theorem 5.4.

For the estimators in (17) and (20), we are using the same measure $P^*_{x_i,T}$ in the estimation of both $E^*_{x_i,T}[YL]$ and $E^*_{x_i,T}[\tau L]$ in (16). However, in certain contexts,

such as for rare-event simulations (e.g., see [Heidelberger 1995]), it might be more efficient to use different measures for estimating the (conditional) expectations in the numerator and denominator of (16). Thus, suppose $\widetilde{P}^*_{x_i,T}$ is a measure on $\mathcal{F}_{x_i,T}$ such that $P_{x_i,T}$ is absolutely continuous with respect to $\widetilde{P}^*_{x_i,T}$. Let $\widetilde{E}^*_{x_i,T}$ be expectation under measure $\widetilde{P}^*_{x_i,T}$, and let $\widetilde{L} \equiv \widetilde{L}(x_i)$ be the likelihood ratio of $P_{x,T}$ with respect to $\widetilde{P}^*_{x_i,T}$ up to time T. Then we can rewrite (16) as

$$\alpha = \frac{\sum_{i=1}^d \nu(x_i) E^*_{x_i,T}[YL]}{\sum_{i=1}^d \nu(x_i) \widetilde{E}^*_{x_i,T}\left[\tau \widetilde{L}\right]}. \tag{22}$$

We can use (22) as the basis for developing importance-sampling estimators analogous to (17) and (20), but in which different importance-sampling measures are used to estimate the expectations in the numerator and denominator. This idea generalizes a method known as measure-specific importance sampling discussed in [Goyal et al. 1992]. One possibility is to let $\widetilde{P}^*_{x,T} = P_{x,T}$ for all $x \in A$, in which case $\widetilde{E}^*_{x_i,T}\left[\tau \widetilde{L}\right] = E_{x_i}[\tau]$; i.e., we use standard simulation for estimating the expectations in the denominator of (22). This is the analogue to what is suggested in [Goyal et al. 1992], and we might implement this by modifying the estimator in (20) to estimate $\nu(x_i)$ and $E_{x_i}[\tau]$ using the same samples generated under the (original) measure $P_{x_i,T}$.

5.2 Unstratified Estimation

We now develop the estimator corresponding to (17) for when we run a simulation of a single sample path rather than using stratification. To do this, we apply the method described at the beginning of Section 5 for using the importance-sampling measure $P^*_{x_0}$ to generate a sample path $\{X_j : j = 0, 1, 2, \ldots, T_n\}$, from which we get $\{W_k : k = 0, 1, 2, \ldots, n\}$ with $W_k = X_{T_k}$.

To state our new estimator, define $T'_k(x)$, $\widetilde{T}'_k(x)$, $Y'_k(x)$, $\tau'_k(x)$, and $H_n(x)$, for $x \in A$, as in Section 3.3, but now these quantities are under measure $P^*_{x,T}$. Also, for $x \in A$ and $k \geq 1$, define $L'_k(x)$ to be the likelihood ratio of the sample-path trajectory $\{X_j : j = T'_k(x), T'_k(x) + 1, \ldots, \widetilde{T}'_k(x)\}$ conditional on $X_{T'_k(x)}$. Define $\bar{R}'_n = (\bar{R}'_n(x,y) : x, y \in A)$ with

$$\bar{R}'_n(x,y) = \frac{1}{H_n(x)} \sum_{k=1}^{H_n(x)} I(X_{\widetilde{T}'_k(x)} = y) L'_k(x),$$

and let $\bar{\nu}'_n = (\bar{\nu}'_n(x) : x \in A)$ be the stationary distribution of $\bar{R}'_n$. Then we define the analogue of (17) for one sample path as

$$\alpha_n^{*\,'} = \frac{\sum_{i=1}^d \bar{\nu}'_n(x_i) \frac{1}{H_n(x_i)} \sum_{k=1}^{H_n(x_i)} Y'_k(x_i) L'_k(x_i)}{\sum_{i=1}^d \bar{\nu}'_n(x_i) \frac{1}{H_n(x_i)} \sum_{k=1}^{H_n(x_i)} \tau'_k(x_i) L'_k(x_i)}. \tag{23}$$

We then have the following central limit theorem, which can be established using arguments similar to those applied in the proof of Theorem 3.10.

THEOREM 5.6. *Under Assumptions 2.1, 3.1, 3.4, 5.1, and 5.3,*

$$n^{1/2}\left(\alpha_n^{*\prime} - \alpha\right) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_*^{\prime\,2})$$

as $n \to \infty$, where

$$\sigma_*^{\prime\,2} = \frac{1}{\left(E_\nu[T]\right)^2} \sum_{i=1}^{d} \frac{\nu^2(x_i)\eta_{i*}}{\rho(x_i)} \tag{24}$$

with η_{i*} defined as in (19) and ρ is the stationary distribution of R^*.

The reason the $\rho(x_i)$, $i = 1, 2, \ldots, d$, appear in the denominator in (24) is that in (23) we are computing sample averages over $H_n(x_i)$ observations. Note that $H_n(x_i)/n \to \rho(x_i)$ a.s. under measure $P_{x_0}^*$, for any $x_0 \in S$, with $\rho(x_i) > 0$ since R^* is positive recurrent. Thus, application of the random-time-change central limit theorem results in the appearance of the $\rho(x_i)$. For example,

$$n^{1/2}\left(\frac{1}{H_n(x_i)} \sum_{k=1}^{H_n(x_i)} Z_k'(x_i)L_k'(x_i) - z(x_i)\right) \xrightarrow{\mathcal{D}} \mathcal{N}\left(0, \frac{h^*(x_i)}{\rho(x_i)}\right),$$

as $n \to \infty$, where $Z_k'(x_i) = Y_k'(x_i) - \alpha\tau_k'(x_i)$.

We now develop the analogue of (20) for when two independent sample paths of X are generated. Fix x_0 and $x_0^* \in A$. We generate one of the paths using the original measure P_{x_0}, and we use this path to estimate the $\nu(x)$, $x \in A$. The other path is generated under the importance-sampling measure $P_{x_0^*}^*$ and is used to estimate the (conditional) expectations in (16). Specifically, fix $0 < q < 1$, and let $r = 1 - q$. Set $X_0 = x_0$, and use the original measure P_{x_0} to generate a sample path $\{X_j : j = 0, 1, 2, \ldots, T_{\lfloor qn \rfloor}\}$, from which we get $\{W_k : k = 0, 1, 2, \ldots, \lfloor qn \rfloor\}$ with $W_k = X_{T_k}$. Independently of how we generate $\{X_j : j = 0, 1, 2, \ldots, T_{\lfloor qn \rfloor}\}$, fix $X_0^* = x_0^*$ and use the measure $P_{x_0}^*$ to generate a sample path $\{X_j^* : j = 0, 1, 2, \ldots, T_{\lfloor rn \rfloor}^*\}$ in the manner described in Section 5, and this yields $\{W_k^* : k = 0, 1, 2, \ldots, \lfloor rn \rfloor\}$ with $W_k^* = X_{T_k^*}^*$. Here, the T_k^*, $k \geq 0$, are the hitting times of the X^* process to the set A.

For $x \in A$, define $\widetilde{\nu}_n'(x) = \sum_{k=0}^{\lfloor qn \rfloor - 1} I(W_k = x)/\lfloor qn \rfloor$, which is based on the sample path generated using the original measure. Now we define some notation for quantities that are computed based on the sample path generated under importance sampling. For $x \in A$, define $H_n^*(x) = \sum_{k=0}^{\lfloor rn \rfloor - 1} I(W_k^* = x)$. For $x \in A$, define $T_1^*(x) = \inf\{j \geq 0 : X_j^* = x\}$, and for $k \geq 2$, let $T_k^*(x) = \inf\{j > T_{k-1}^*(x) : X_j^* = x\}$. Also, define $\widetilde{T}_k^*(x) = \inf\{j > T_k^*(x) : X_j^* \in A\}$. For $x \in A$ and $k \geq 1$, define

$$Y_k^*(x) = \sum_{j=T_k^*(x)}^{\widetilde{T}_k^*(x)-1} f(X_j^*), \qquad \tau_k^*(x) = \widetilde{T}_k^*(x) - T_k^*(x).$$

Finally, for $x \in A$ and $k \geq 1$, define $L_k^*(x)$ to be the likelihood ratio corresponding to the sample path $\{X_j^* : T_k^*(x) \leq j \leq \widetilde{T}_k^*(x)\}$ given $X_{T_k^*(x)}^*$. Then we define the

analogue of (20) for two sample paths to be

$$\alpha_n^{**\prime} = \frac{\sum_{i=1}^{d} \widetilde{\nu}_n'(x_i) \frac{1}{H_n^*(x_i)} \sum_{k=1}^{H_n^*(x_i)} Y_k^*(x_i) L_k^*(x_i)}{\sum_{i=1}^{d} \widetilde{\nu}_n'(x_i) \frac{1}{H_n^*(x_i)} \sum_{k=1}^{H_n^*(x_i)} \tau_k^*(x_i) L_k^*(x_i)}, \tag{25}$$

which obeys the following central limit theorem.

THEOREM 5.7. *Under Assumptions 2.1, 3.1, 3.4, 5.1, and 5.3,*

$$n^{1/2} \left(\alpha_n^{**\prime} - \alpha \right) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_{**}'^2)$$

as $n \to \infty$, where

$$\sigma_{**}'^2 = \frac{1}{(E_\nu[T])^2} \sum_{i=1}^{d} \nu^2(x_i) \left(\frac{h^*(x_i)}{r} + \sum_{j=1}^{d} \sum_{l=1}^{d} \frac{\zeta(x_j)\zeta(x_l)\Psi_i(x_j, x_l)}{q} \right),$$

with $h^*(x_i)$ defined as in Theorem 5.4 and $\Psi_i(x_j, x_l)$ defined in (5) and (6).

We could also develop (but omit since it is straightforward) an estimator suggested by (22) based on 3 long sample paths. The first path is generated under the original measure P_{x_0} and is used to estimate the $\nu(x_i)$. The second path is generated using the measure $P_{x_0}^*$ and is used to estimate the $E_{x_i,T}^*[YL]$. The third path is generated using measure $\widetilde{P}_{\widetilde{x}_0}^*$ and is used to estimate the $\widetilde{E}_{x_i,T}^*[\tau\widetilde{L}]$.

6. EXPECTED CUMULATIVE REWARD UNTIL A HITTING TIME

Fix a nonempty set $S_0 \subset S$, and let $\Gamma = \inf\{n \geq 0 : X_n \in S_0\}$. For $x \in A$, put

$$\lambda(x) = E_x \left[\sum_{j=1}^{\Gamma} f(X_j) \right],$$

which is the expected cumulative reward up to hitting the set S_0 given that the chain starts in state x. The measure $\lambda(x)$ arises in many contexts. For example, it can be the mean time to failure of a reliability system, or the expected time to buffer overflow in a queueing network. We want to develop semi-regenerative estimators for $\lambda(x)$.

Throughout this section, unless stated otherwise, we no longer assume that Assumptions 2.1, 3.1, 3.6, 5.1, or 5.3 hold. Assume that Assumption 3.4 and the following hold.

Assumption 6.1. For each recurrence class C of states in A, there exists some state $x \in C$ such that $P_x(\Gamma > T) < 1$.

Note that

$$\lambda(x) = E_x \left[\sum_{j=1}^{T \wedge \Gamma} f(X_j) \right] + \sum_{y \in A} E_x \left[I(X_T = y, \Gamma > T) \right] \lambda(y),$$

where $a_1 \wedge a_2 = \min(a_1, a_2)$ for $a_1, a_2 \in \Re$. For $x, y \in A$, put

$$b(x) = E_x \left[\sum_{j=1}^{T \wedge \Gamma} f(X_j) \right], \qquad K(x, y) = E_x \left[I(X_T = y, \Gamma > T) \right].$$

Let $\lambda = (\lambda(x) : x \in A)$, $b = (b(x) : x \in A)$, and $K = (K(x, y) : x, y \in A)$, and note that $\lambda = b + K\lambda$.

PROPOSITION 6.2. *If $|b| < \infty$ and if Assumptions 3.4 and 6.1 hold, then*

$$\lambda = \sum_{m=0}^{\infty} K^m b = (I - K)^{-1} b.$$

Assumption 6.1 ensures that $(I-K)^{-1}$ exists and equals $\sum_{m=0}^{\infty} K^m$. Without this assumption, it is possible that $\sum_{m=0}^{\infty} K^m$ diverges, in which case $(I-K)^{-1}$ need not equal $\sum_{m=0}^{\infty} K^m$. Also, note that Proposition 6.2 does not require irreducibility or recurrence. Finally, the representation of λ in Proposition 6.2 generalizes the well-known result (e.g., [Goyal et al. 1992]) that $\lambda(x) = E_x[\sum_{j=1}^{(\tau_x \wedge \Gamma)} f(X_j)]/E_x[I(\Gamma < \tau_x)]$.

6.1 Stratified Estimation

We now present a stratified semi-regenerative estimator for λ based on Proposition 6.2. Let

$$B = \sum_{j=1}^{T \wedge \Gamma} f(X_j) \quad \text{and} \quad \phi(y) = I(X_T = y, \Gamma > T), \text{ for } y \in A.$$

Let

$$(B_k(x_i), \phi_k(x_i, y) : y \in A)$$

for $1 \leq k \leq \lfloor p_i n \rfloor$ be i.i.d. copies of

$$(B, \phi(y) : y \in A) \text{ under measure } P_{x_i}.$$

Set

$$b_n(x_i) = \frac{1}{\lfloor p_i n \rfloor} \sum_{k=1}^{\lfloor p_i n \rfloor} B_k(x_i), \qquad K_n(x_i, y) = \frac{1}{\lfloor p_i n \rfloor} \sum_{k=1}^{\lfloor p_i n \rfloor} \phi_k(x_i, y)$$

for $1 \leq i \leq d$ and $y \in A$, and let $b_n = (b_n(x) : x \in A)$ and $K_n = (K_n(x, y) : x, y \in A)$. We define the stratified semi-regenerative estimator for λ to be

$$\lambda_n = (I - K_n)^{-1} b_n.$$

Under Assumption 6.1, $(I - K)^{-1}$ exists. Since $K_n \to K$ a.s., evidently $I - K_n$ is non-singular for n sufficiently large, and

$$(I - K_n)^{-1} \to (I - K)^{-1} \quad \text{a.s.} \tag{26}$$

as $n \to \infty$ by the continuity of the inverse mapping at $I - K$. To establish a central limit theorem for λ_n, we will assume the following.

Assumption 6.3. For each $x \in A$, $E_x[B^2] < \infty$.

To prove our central limit theorem for λ_n, we need to get a handle on $(I - K_n)^{-1} - (I - K)^{-1}$ and $b_n - b$. Note that

$$\begin{aligned}
K_n - K &= (I - K) - (I - K_n) \\
&= [I - (I - K_n)(I - K)^{-1}](I - K) \\
&= (I - K_n)[(I - K_n)^{-1} - (I - K)^{-1}](I - K).
\end{aligned}$$

Consequently,

$$(I - K_n)^{-1} - (I - K)^{-1} = (I - K_n)^{-1}(K_n - K)(I - K)^{-1}, \qquad (27)$$

so

$$\begin{aligned}
\lambda_n - \lambda &= (I - K_n)^{-1}b_n - (I - K)^{-1}b \\
&= ((I - K_n)^{-1} - (I - K)^{-1})b_n + (I - K)^{-1}(b_n - b) \\
&= (I - K_n)^{-1}(K_n - K)(I - K)^{-1}b_n + (I - K)^{-1}(b_n - b).
\end{aligned}$$

Under Assumption 6.3 we have that

$$n^{1/2}(K_n - K, b_n - b) \xrightarrow{\mathcal{D}} (\widetilde{N}_1, \widetilde{N}_2) \overset{\mathcal{D}}{=} \mathcal{N}(0, \widetilde{\Sigma}) \qquad (28)$$

as $n \to \infty$, where $\widetilde{\Sigma}$ is some covariance matrix. Therefore, the continuous mapping theorem implies that

$$n^{1/2}(\lambda_n - \lambda) \xrightarrow{\mathcal{D}} (I - K)^{-1}\widetilde{N}_1(I - K)^{-1}b + (I - K)^{-1}\widetilde{N}_2$$

by (26) and since $b_n \to b$ a.s. Finally, because $\lambda = (I - K)^{-1}b$, we obtain the following central limit theorem for λ_n.

THEOREM 6.4. *If $|b| < \infty$ and if Assumptions 3.4, 6.1, and 6.3 hold, then*

$$n^{1/2}(\lambda_n - \lambda) \xrightarrow{\mathcal{D}} (I - K)^{-1}\widetilde{N}_1\lambda + (I - K)^{-1}\widetilde{N}_2$$

as $n \to \infty$, where $(\widetilde{N}_1, \widetilde{N}_2)$ is defined in (28). In particular, for each $k = 1, 2, \ldots, d$,

$$n^{1/2}(\lambda_n(x_k) - \lambda(x_k)) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \widetilde{\sigma}_k^2)$$

as $n \to \infty$, where

$$\widetilde{\sigma}_k^2 = \sum_{i=1}^{d} \frac{J(x_k, x_i)^2}{p_i}\left[v_i + 2\sum_{j=1}^{d}\lambda(x_j)s_{ij} + \sum_{j=1}^{d}\sum_{l=1}^{d}\lambda(x_j)\lambda(x_l)\Delta_i(x_j, x_l)\right], \qquad (29)$$

$J = (J(x, y) : x, y \in A)$ with $J = (I - K)^{-1}$, $v_i = Var_{x_i}(B)$, $s_{ij} = Cov_{x_i}(\phi(x_j), B)$, and $\Delta_i(x_j, x_l) = Cov_{x_i}(\phi(x_j), \phi(x_l))$.

6.2 Unstratified Estimation

We now present a semi-regenerative estimator for λ based on Proposition 6.2 when simulating one sample path. We now assume that Assumption 2.1 holds. Define $H_n(x)$, $T'_k(x)$, and $\widetilde{T}'_k(x)$ as in Section 3.3. Also, for $k \geq 1$, define $\Gamma_k(x) = \inf\{j > T'_k(x) : X_j \in S_0\}$. For $x, y \in A$, let

$$B'_k(x) = \sum_{j=T'_k(x)+1}^{\widetilde{T}'_k(x)\wedge\Gamma_k(x)} f(X_j), \qquad \phi'_k(x, y) = I\left(X_{\widetilde{T}'_k(x)} = y, \Gamma_k(x) > \widetilde{T}'_k(x)\right).$$

Then define the estimators of b and K to be $b'_n = (b'_n(x) : x \in A)$ and $K'_n = (K_n(x,y) : x,y \in A)$, respectively, with

$$b'_n(x) \; = \; \frac{\sum_{k=0}^{n-1} \sum_{j=T_k}^{\Gamma_k \wedge T_{k+1}} f(X_j)\, I(W_k = x)}{\sum_{k=0}^{n-1} I(W_k = x)} \; = \; \frac{1}{H_n(x)} \sum_{k=1}^{H_n(x)} B'_k(x),$$

$$K'_n(x,y) \; = \; \frac{\sum_{k=0}^{n-1} I(W_k = x,\, W_{k+1} = y,\, \Gamma_k > T_{k+1})}{\sum_{k=0}^{n-1} I(W_k = x)} \; = \; \frac{1}{H_n(x)} \sum_{k=1}^{H_n(x)} \phi'_k(x,y),$$

where $\Gamma_k = \inf\{j > T_k : X_j \in S_0\}$. Then we define our semi-regenerative estimator of λ based on one simulation to be $\lambda'_n = (I - K'_n)^{-1} b'_n$, where $\lambda'_n = (\lambda'_n(x) : x \in A)$. Using the techniques we developed in the proof of Theorem 3.10, we can prove that

$$n^{1/2}(\lambda'_n(x_k) - \lambda(x_k)) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \widetilde{\sigma}'^2_k)$$

as $n \to \infty$, where $\widetilde{\sigma}'^2_k$ is the same as $\widetilde{\sigma}^2_k$ in (29) except that each p_i, $i = 1, 2, \ldots, d$, appearing in the denominators in (29) is replaced by $\nu(x_i)$. As shown in Table I, the freedom to choose the $\{p_i\}$ different from the $\{\nu(x_i)\}$ can result in a significant efficiency improvement.

7. DERIVATIVE OF STEADY-STATE REWARD

We now discuss the estimation of derivatives of a performance measure with respect to a model parameter. For example, for a reliability system, we may be interested in computing the derivative of the steady-state availability with respect to the failure rate of one component.

We now assume that the transition probability matrix of X depends on some real-valued parameter θ, where we allow θ to vary in an open interval Θ. Thus, we write $\Pi(\theta) = (\Pi(\theta, x, y) : x, y \in S)$ to emphasize the dependence on θ. Our goal is to compute the derivative of the steady-state mean reward $\alpha = \alpha(\theta)$ with respect to θ, and evaluate this when θ takes on some fixed value $\theta_0 \in \Theta$. We assume the following:

Assumption 7.1. $|S| < \infty$. Also, the family $(\Pi(\theta) : \theta \in \Theta)$ is continuously differentiable in θ, and $\Pi(\theta)$ is irreducible for all $\theta \in \Theta$, with $\{(x,y) \in S \times S : \Pi(\theta, x, y) > 0\}$ independent of $\theta \in \Theta$.

For each $\theta \in \Theta$, the finiteness of S and the irreducibility of $\Pi(\theta)$ imply that X is positive recurrent. Thus, for each $\theta \in \Theta$, there exists a unique stationary distribution $\pi(\theta) = (\pi(\theta, x) : x \in S)$ for X.

Let P^θ_x denote the probability measure of the process X induced by the transition matrix $\Pi(\theta)$ given $X_0 = x$, and let E^θ_x be the corresponding expectation operator.

Now define the embedded chain W relative to the set A as in Section 2, and let $R(\theta) = (R(\theta, x, y) : x, y \in A)$ be its transition probability matrix with stationary distribution $\nu(\theta) = (\nu(\theta, x) : x \in A)$. Note that $R(\theta, x, y) = P^\theta_x(X_T = y)$, and the set $\{(x,y) \in A \times A : R(\theta, x, y) > 0\}$ is independent of $\theta \in \Theta$ under Assumption 7.1. Also, let $P^{\theta_1}_{\nu(\theta_2)}$ denote the probability measure induced by the transition matrix $\Pi(\theta_1)$ with initial distribution $\nu(\theta_2)$, and let $E^{\theta_1}_{\nu(\theta_2)}$ be the corresponding expectation

operator. According to Proposition 3.5, $\alpha(\theta) = \pi(\theta)f$ can be written as

$$\alpha(\theta) = \frac{E_{\nu(\theta)}^{\theta}\left[\sum_{j=0}^{T-1} f(X_j)\right]}{E_{\nu(\theta)}^{\theta}[T]}.$$

With $\theta_0 \in \Theta$ fixed, we write

$$\alpha(\theta) = \frac{E_{\nu(\theta)}^{\theta_0}\left[\sum_{j=0}^{T-1} f(X_j) \prod_{l=0}^{T-1} \frac{\Pi(\theta,X_l,X_{l+1})}{\Pi(\theta_0,X_l,X_{l+1})}\right]}{E_{\nu(\theta)}^{\theta_0}\left[T \prod_{l=0}^{T-1} \frac{\Pi(\theta,X_l,X_{l+1})}{\Pi(\theta_0,X_l,X_{l+1})}\right]} \equiv \frac{\xi(\theta)}{\kappa(\theta)}.$$

The above change of measure is justified since $P_{\nu(\theta_2)}^{\theta_1}$ is absolutely continuous with respect to $P_{\nu(\theta_4)}^{\theta_3}$ for all $\theta_1, \theta_2, \theta_3, \theta_4 \in \Theta$ by Assumption 7.1. It then follows that

$$\partial\alpha(\theta_0) = \frac{\kappa(\theta_0)\partial\xi(\theta_0) - \xi(\theta_0)\partial\kappa(\theta_0)}{\kappa(\theta_0)^2}, \tag{30}$$

where we use the notation that $\partial g(\theta_0)$ denotes the derivative of $g(\theta)$ taken with respect to θ and evaluated at $\theta = \theta_0$.

We now examine $\partial\xi(\theta_0)$. Observe that

$$\xi(\theta) = \sum_{x \in A} \nu(\theta,x)\, E_x^{\theta_0}\left[\sum_{j=0}^{T-1} f(X_j) \prod_{l=0}^{T-1} \frac{\Pi(\theta,X_l,X_{l+1})}{\Pi(\theta_0,X_l,X_{l+1})}\right].$$

Then Assumption 7.1 ensures that

$$\partial\xi(\theta_0) = \sum_{x \in A} \partial\nu(\theta_0,x)\, E_x^{\theta_0}\left[\sum_{j=0}^{T-1} f(X_j)\right] + \sum_{x \in A} \nu(\theta_0,x)\, E_x^{\theta_0}\left[\sum_{j=0}^{T-1} f(X_j)\,\partial L\right], \tag{31}$$

where

$$\partial L = \sum_{l=0}^{T-1} \frac{\partial\Pi(\theta_0,X_l,X_{l+1})}{\Pi(\theta_0,X_l,X_{l+1})}.$$

Similarly, we can show that

$$\partial\kappa(\theta_0) = \sum_{x \in A} \partial\nu(\theta_0,x)\, E_x^{\theta_0}[T] + \sum_{x \in A} \nu(\theta_0,x)\, E_x^{\theta_0}[T\,\partial L]. \tag{32}$$

These expressions form the basis for applying the likelihood ratio (LR) method for derivative estimation; see, e.g., [Glynn 1990; Reiman and Weiss 1989; Rubinstein 1989] for details on the LR method.

We now need to get a handle on $\partial\nu(\theta_0) = (\partial\nu(\theta_0,x) : x \in A)$. We can show that $R(\theta)$ is continuous and differentiable in θ by using the fact that

$$R(\theta,x,y) = E_x^{\theta_0}\left[I(X_T = y) \prod_{l=0}^{T-1} \frac{\Pi(\theta,X_l,X_{l+1})}{\Pi(\theta_0,X_l,X_{l+1})}\right],$$

so

$$\partial R(\theta_0,x,y) = E_x^{\theta_0}\left[I(X_T = y)\,\partial L\right].$$

Let $\partial R(\theta_0) = (\partial R(\theta_0, x, y) : x, y \in A)$. Glynn [1986] shows that the continuity of $R(\theta)$ in θ implies that $\nu(\theta)$ is also continuous. Then, letting $V(\theta)$ be the matrix with all rows equal to $\nu(\theta)$ (i.e., $V(\theta) = e\nu(\theta)$), Glynn [1986] also establishes that $\partial\nu(\theta_0)$ exists and

$$\partial\nu(\theta_0) = \nu(\theta_0)\partial R(\theta_0)F(\theta_0), \tag{33}$$

where $F(\theta_0) = (I - R(\theta_0) + V(\theta_0))^{-1}$.

For $x \in A$, define

$$y_0(x) = E_x^{\theta_0}\left[\sum_{j=0}^{T-1} f(X_j)\right], \qquad \partial y(x) = E_x^{\theta_0}\left[\sum_{j=0}^{T-1} f(X_j)\,\partial L\right],$$
$$t(x) = E_x^{\theta_0}[T], \qquad \partial t(x) = E_x^{\theta_0}[T\,\partial L],$$

and set $y_0 = (y_0(x) : x \in A)$, $\partial y = (\partial y(x) : x \in A)$, $t = (t(x) : x \in A)$, and $\partial t = (\partial t(x) : x \in A)$. Define the function $r_0 : \Re^{d\times 1} \times \Re^{d\times 1} \times \Re^{d\times 1} \times \Re^{d\times 1} \times \Re^{d\times d} \times \Re^{1\times d} \times \Re^{d\times d} \to \Re$ as

$$r_0(w_1, w_2, \ldots, w_7) = \frac{(w_6 w_5 w_7 w_1 + w_6 w_2)w_6 w_3 - w_6 w_1(w_6 w_5 w_7 w_3 + w_6 w_4)}{(w_6 w_3)^2}. \tag{34}$$

Let $\mu_0 = (y_0, \partial y, t, \partial t, \partial R(\theta_0), \nu(\theta_0), F(\theta_0))$, and note that

$$\partial\alpha(\theta_0) = r_0(\mu_0) \tag{35}$$

by (30), (31), (32), and (33). Equation (35) will be the basis for developing semi-regenerative estimators for $\partial\alpha(\theta_0)$

One advantage of using the SRM rather than the RM to implement the likelihood-ratio derivative method is that the semi-regenerative trajectories are shorter than the regenerative cycles when the return state w of the RM is chosen from the set A. Analyses in [Reiman and Weiss 1989; Glynn 1987] suggest that the variance of likelihood-ratio derivative estimators grows linearly in the length (number of transitions) of the observation, so semi-regenerative derivative estimators should have smaller variance than regenerative derivative estimators. Zhang and Ho [1992] develop a similar idea of dividing regenerative cycles into A-segments, which are the same as trajectories, but they end up with a different estimator than we do.

7.1 Stratified Estimation

We now develop a stratified semi-regenerative estimator based on (35). For $x, y \in A$, define $Y(x)$, $\tau(x)$, and $\chi(x, y)$ as in Section 3.2. Taking $p_1, p_2, \ldots, p_d$, to be d positive numbers summing to one, we let

$$(\partial L_k(x_i), Y_k(x_i), \tau_k(x_i), \chi_k(x_i, y) : y \in A)$$

for $1 \le k \le \lfloor p_i n \rfloor$ be i.i.d. copies of

$$(\partial L, Y, \tau, \chi(y) : y \in A) \quad \text{under measure } P_{x_i}^{\theta_0}.$$

For $i = 1, 2, \ldots, d$, and $y \in A$, set

$$\bar{Y}_n(x_i) = \frac{1}{\lfloor p_i n \rfloor}\sum_{k=1}^{\lfloor p_i n \rfloor} Y_k(x_i), \qquad \partial\bar{Y}_n(x_i) = \frac{1}{\lfloor p_i n \rfloor}\sum_{k=1}^{\lfloor p_i n \rfloor} Y_k(x_i)\,\partial L_k(x_i),$$

$$\bar{\tau}_n(x_i) \;=\; \frac{1}{\lfloor p_i n \rfloor} \sum_{k=1}^{\lfloor p_i n \rfloor} \tau_k(x_i), \qquad \partial \bar{\tau}_n(x_i) \;=\; \frac{1}{\lfloor p_i n \rfloor} \sum_{k=1}^{\lfloor p_i n \rfloor} \tau_k(x_i)\, \partial L_k(x_i),$$

$$R_n(\theta_0, x_i, y) \;=\; \frac{1}{\lfloor p_i n \rfloor} \sum_{k=1}^{\lfloor p_i n \rfloor} \chi_k(x_i, y),$$

$$\partial R_n(\theta_0, x_i, y) \;=\; \frac{1}{\lfloor p_i n \rfloor} \sum_{k=1}^{\lfloor p_i n \rfloor} \chi_k(x_i, y)\, \partial L_k(x_i),$$

and let $\bar{Y}_n = (\bar{Y}_n(x) : x \in A)$, $\partial \bar{Y}_n = (\partial \bar{Y}_n(x) : x \in A)$, $\bar{\tau}_n = (\bar{\tau}_n(x) : x \in A)$, $\partial \bar{\tau}_n = (\partial \bar{\tau}_n(x) : x \in A)$, $R_n(\theta_0) = (R_n(\theta_0, x, y) : x, y \in A)$, and $\partial R_n(\theta_0) = (\partial R_n(\theta_0, x, y) : x \in A$. Let $\nu_n(\theta_0) = (\nu_n(\theta_0, x) : x \in A) \in \Re^{1 \times d}$ be the stationary distribution of $R_n(\theta_0)$. Define $V_n(\theta_0)$ as the matrix with all rows equal to $\nu_n(\theta_0)$; i.e., $V_n(\theta_0) = e\nu_n(\theta_0)$, where e is the vector of all 1's in $\Re^{d \times 1}$. Also, define $F_n(\theta_0) = (I - R_n(\theta_0) + V_n(\theta_0))^{-1}$. Finally, define $\partial \nu_n(\theta_0) = (\partial \nu_n(\theta_0, x) : x \in A)$ as

$$\partial \nu_n(\theta_0) = \nu_n(\theta_0)\, \partial R_n(\theta_0)\, F_n(\theta_0).$$

We then define our stratified semi-regenerative estimator of $\partial \alpha(\theta_0)$ to be

$$\partial \alpha_n(\theta_0) = r_0(\bar{Y}_n, \partial \bar{Y}_n, \bar{\tau}_n, \partial \bar{\tau}_n, R_n(\theta_0), \partial R_n(\theta_0), \nu_n(\theta_0), F_n(\theta_0)),$$

where the function r_0 is defined in (34).

We show in the appendix (see (48)) that

$$\nu_n(\theta_0) - \nu(\theta_0) = \nu_n(\theta_0)(R_n(\theta_0) - R(\theta_0))F(\theta_0), \tag{36}$$

so

$$V_n(\theta_0) - V(\theta_0) = e(\nu_n(\theta_0) - \nu(\theta_0)) = e\nu_n(\theta_0)(R_n(\theta_0) - R(\theta_0))F(\theta_0). \tag{37}$$

We also establish in the appendix (see (49)) that $\nu_n \to \nu$ a.s. as $n \to \infty$, so (3) implies that $V_n(\theta_0) \to V(\theta_0)$ a.s. as $n \to \infty$. Thus,

$$F_n(\theta_0) = (I - R_n(\theta_0) + V_n(\theta_0))^{-1} \to (I - R(\theta_0) + V(\theta_0))^{-1} = F(\theta_0) \;\; \text{a.s.}$$

as $n \to \infty$ by the continuity of the inverse mapping at $I - R(\theta_0) + V(\theta_0)$. Also, in a similar manner to how we established (27), we can also show that

$$\begin{aligned} & F_n(\theta_0) - F(\theta_0) \\ &= F_n(\theta_0)\left[(R_n(\theta_0) - R(\theta_0)) - (V_n(\theta_0) - V(\theta_0))\right]F(\theta_0) \\ &= F_n(\theta_0)\left[(R_n(\theta_0) - R(\theta_0)) - e\nu_n(\theta_0)(R_n(\theta_0) - R(\theta_0))F(\theta_0)\right]F(\theta_0) \end{aligned} \tag{38}$$

by (37).

Note that the finiteness of S and irreducibility of $\Pi(\theta_0)$ by Assumption 7.1 ensure that τ has finite moments of all orders under measure $P_x^{\theta_0}$, $x \in A$. Also, the finiteness of S and the continuous differentiability of $\Pi(\theta)$ imply that f and $\partial \Pi(\theta_0, \cdot, \cdot)/\Pi(\theta_0, \cdot, \cdot)$ are bounded. Thus, Y, $Y \partial L$, τ, $\tau \partial L$, $\chi(y)$, and $\chi(y) \partial L$ all have finite moments of all orders under measure $P_x^{\theta_0}$, $x \in A$, so it follows that

$$n^{1/2}\left[(\bar{Y}_n, \partial \bar{Y}_n, \bar{\tau}_n, \partial \bar{\tau}_n, \partial R_n(\theta_0), \nu_n(\theta_0), F_n(\theta_0)) - \mu_0\right]$$
$$\xrightarrow{\mathcal{D}} (M_1, M_2, \ldots, M_7) \overset{\mathcal{D}}{=} \mathcal{N}(0, \Sigma_0), \tag{39}$$

for some normal random elements $M_1, M_2, \ldots, M_7$, and some covariance matrix Σ_0. Also, note that $n^{1/2}[R_n(\theta_0) - R(\theta_0)] \xrightarrow{\mathcal{D}} M_8$ as $n \to \infty$, where M_8 is a normally distributed matrix for which the ith row of M_8 has covariance matrix $p_i^{-1}\Psi_i$, where Ψ_i is defined in (5) and (6). Moreover, for $i \neq j$, the ith and jth row of M_8 are independent. We can then show that

$$M_6 = \nu(\theta_0)M_8 F(\theta_0), \qquad M_7 = F(\theta_0)(M_8 - e\nu(\theta_0)M_5 F(\theta_0))F(\theta_0),$$

by (36) and (38), respectively. Then applying the delta method (e.g., Theorem A, p. 122 of [Serfling 1980]) results in the following central limit theorem for $\partial \alpha_n(\theta_0)$.

THEOREM 7.2. *Under Assumption 7.1,*

$$n^{1/2}(\partial \alpha_n(\theta_0) - \partial \alpha(\theta_0)) \xrightarrow{\mathcal{D}} \mathcal{N}(0, D_0^\top \Sigma_0 D_0),$$

where Σ_0 is defined in (39) and D_0 is the vector of partial derivatives of the function r_0 evaluated at μ_0, with r_0 defined in (34).

It is straightforward to compute the entries of D_0, which we mostly do not give explicitly. For example, calculating the partial derivative of r_0 with respect to $w_6(x_i)$ and evaluating at μ_0 yields (with a slight abuse of notation)

$$
\begin{aligned}
\frac{\partial r_0(\mu_0)}{\partial w_6(x_i)} = {} & \left\{ \kappa(\theta_0)^2 \left[\kappa(\theta_0) \left(\sum_{l=1}^{d} \partial R(\theta_0, x_i, x_l) \sum_{m=1}^{d} F(\theta_0, x_l, x_m) y_0(x_m) + \partial y(x_i) \right) \right.\right. \\
& + \partial \xi(\theta_0) t(x_i) - y_0(x_i)\, \partial \kappa(\theta_0) \\
& \left.\left. - \xi(\theta_0) \left(\sum_{l=1}^{d} \partial R(\theta_0, x_i, x_l) \sum_{m=1}^{d} F(\theta_0, x_l, x_m) t(x_m) + \partial t(x_i) \right) \right] \right. \\
& \left. - 2\left[\kappa(\theta_0)\partial \xi(\theta_0) - \xi(\theta_0)\partial \kappa(\theta_0) \right] \kappa(\theta_0) t(x_i) \right\} \kappa(\theta_0)^{-4}.
\end{aligned}
$$

Many entries in Σ_0 in (39) are zero because $(\partial L_k(x_i), Y_k(x_i), \tau_k(x_i), \chi_k(x_i, y) : y \in A)$ and $(\partial L_k(x_j), Y_k(x_j), \tau_k(x_j), \chi_k(x_j, y) : y \in A)$ for $i \neq j$ are independent, and the remaining non-zero entries are straightforward to calculate. For example,

$$\text{Cov}_{x_i}^{\theta_0}(M_1(x_i), M_6(x_i, x_j)) = \frac{\text{Cov}_{x_i}^{\theta_0}(Y, \chi(x_j)\, \partial L)}{p_i}, \qquad (40)$$

where Cov^{θ_0} is the covariance operator under parameter value θ_0.

One can also construct a semi-regenerative estimator based on (35) for one sample path, but we omit this.

8. LOW-BIAS ESTIMATOR

In this section we describe a variation of the semi-regenerative estimator that we expect to have lower bias than the standard semi-regenerative estimator. Put $\widehat{\alpha}_m = m^{-1} \sum_{j=0}^{m-1} f(X_j)$. In great generality, it is known that

$$E_x[\widehat{\alpha}_m] = \pi f + m^{-1}\beta + o(m^{-1})$$

as $m \to \infty$; see [Glynn and Heidelberger 1990]. We wish to have a semi-regenerative estimator β_m for β. This is a semi-regenerative analogue of a regenerative estimator proposed by [Glynn 1994].

Suppose $|S| < \infty$ with X irreducible and aperiodic. For $x \in S$, let $f_c(x) = f(x) - \alpha$. Then we can express the bias of $\widehat{\alpha}_m$ as

$$\frac{1}{m}E_x\left[\sum_{j=0}^{m-1} f(X_j)\right] - \pi f \;=\; \frac{1}{m}E_x\left[\sum_{j=0}^{m-1} f_c(X_j)\right]$$

$$= \frac{1}{m}\sum_{j=0}^{\infty} E_x[f_c(X_j)] - \frac{1}{m}\sum_{j=m}^{\infty} E_x[f_c(X_j)] \;\equiv\; \frac{1}{m}\beta(x) + o\left(\frac{1}{m}\right).$$

Our goal is to reduce the order $1/m$ bias by subtracting an estimate of $\beta(x)/m$.

Let $\mu(x) = E_x\left[\sum_{j=0}^{T-1} f_c(X_j)\right]$ for $x \in A$, and $\mu = (\mu(x) : x \in A)$. Then, it follows that for $x \in A$,

$$\beta(x) \;=\; E_x\left[\sum_{k=0}^{\infty}\sum_{j=T_k}^{T_{k+1}-1} f_c(X_j)\right] \;=\; \sum_{k=0}^{\infty} E_x[\mu(W_k)]$$

$$= \sum_{k=0}^{\infty}\sum_{y \in A} R^k(x,y)\mu(y) \;=\; (F\mu)(x),$$

i.e., $\beta = F\mu$.

We now discuss how to estimate β from a simulation of one sample path. Let $N_m = \sup\{n \geq 0 : T_n \leq m\}$. Recall our definitions of $\widehat{\nu}_n$ and R'_n from Section 3.3, and let $\widehat{V}_n = e\widehat{\nu}_n$. Then $\beta_m = (I - R'_{N_m} - \widehat{V}_{N_m})^{-1}\mu'_m$ estimates β, where $\mu'_m = (\mu'_m(x) : x \in A)$ is an estimate of μ given by

$$\mu'_m(x) = \frac{\frac{1}{N_m}\sum_{k=0}^{N_m-1}\sum_{j=T_k}^{T_{k+1}-1}(f(X_j) - \widehat{\alpha}_m)I(W_k = x)}{\frac{1}{N_m}\sum_{k=0}^{N_m-1} I(W_k = x)}.$$

We then expect $\widehat{\alpha}_m - \beta_m(x)/m$ to have lower bias than $\widehat{\alpha}_m$ does, provided $X_0 = x \in A$.

9. RATIOS OF STEADY-STATE MEANS

In this section we consider a performance measure γ of the form $\gamma = (\pi f_1)/(\pi f_2)$, where f_1 and f_2 are real-valued reward functions on S, and π is the stationary measure of X on S. Note that (1) implies that

$$\gamma = \frac{E_w[\sum_{j=0}^{T-1} f_1(X_j)]/E_w[T]}{E_w[\sum_{j=0}^{T-1} f_2(X_j)]/E_w[T]} = \frac{E_w[\sum_{j=0}^{T-1} f_1(X_j)]}{E_w[\sum_{j=0}^{T-1} f_2(X_j)]} \tag{41}$$

for any state $w \in S$ when Assumption 2.1 holds. Also, we can write

$$\gamma = \frac{E_\nu[\sum_{j=0}^{T-1} f_1(X_j)]/E_\nu[T]}{E_\nu[\sum_{j=0}^{T-1} f_2(X_j)]/E_\nu[T]} = \frac{E_\nu[\sum_{j=0}^{T-1} f_1(X_j)]}{E_\nu[\sum_{j=0}^{T-1} f_2(X_j)]}$$

by Proposition 3.5.

The measure γ is a generalization of the ratio formula for steady-state means in Proposition 3.5. Also, γ arises in practice when applying discrete-time conversion to a continuous-time Markov chain [Hordijk et al. 1976; Fox and Glynn 1986]. The basic idea is to condition on the embedded discrete-time Markov chain, which results in replacing the random exponential holding times in each state with their (conditional) means, and this is guaranteed to reduce variance. Specifically, suppose that $U = \{U(t) : t \geq 0\}$ is a positive-recurrent, irreducible continuous-time Markov chain on state space S with embedded discrete-time Markov chain X. Suppose that $q(x)$ is the total transition rate of U out of state $x \in S$, and we are interested in computing $\gamma = \lim_{t \to \infty}(1/t)\int_0^t f_3(U(s))ds$, where $f_3 : S \to \Re$ is some reward function. Then for any state $w \in S$, the steady-state mean reward of the continuous-time Markov chain can be expressed as

$$\gamma = \frac{E_w[\sum_{j=0}^{T-1} f_3(X_j)/q(X_j)]}{E_w[\sum_{j=0}^{T-1} 1/q(X_j)]},$$

which has exactly the form in (41) by letting $f_1(x) = f_3(x)/q(x)$ and $f_2(x) = 1/q(x)$ for $x \in S$.

The methods described in the previous sections can easily be modified to work with this more general setting of dealing with the performance measure γ rather than α. Thus, we can handle continuous-time Markov chains in our framework of discrete-time Markov chains.

10. CONCLUDING REMARKS

For many of our estimators considered in this paper, confidence intervals are desirable. (But, for example, there is typically no need for a confidence interval for the bias correction of Section 8.) However, the central limit theorem for many of our estimators is complicated. From an implementation standpoint, it may be desirable to produce confidence intervals without having to explicitly work out the corresponding central limit theorem, followed by consistent estimation of the corresponding variance constant.

One way of doing this is by "sectioning." Sectioning works even in the stratified sampling context. Given an integer $\ell > 0$, we section the computer budget c into ℓ different pieces, each of size c/ℓ. We then apply our stratification weights $p_1, \ldots, p_d$ to each of the ℓ sub-budgets; this gives us ℓ i.i.d. estimators. We also get one estimator for the total budget c. Specifically, suppose the goal is to estimate some performance measure α. Let $\widetilde{\alpha}_i$ be the estimator for sub-budget i, $1 \leq i \leq \ell$, where $\widetilde{\alpha}_1, \widetilde{\alpha}_2, \ldots, \widetilde{\alpha}_\ell$ are i.i.d. Also let $\widetilde{\alpha} = \sum_{i=1}^{\ell} \widetilde{\alpha}_i/\ell$ be the estimator associated with the entire budget. Then, under appropriate moment assumptions,

$$\sqrt{c}(\widetilde{\alpha}_1 - \alpha, \widetilde{\alpha}_2 - \alpha, \ldots, \widetilde{\alpha}_\ell - \alpha, \widetilde{\alpha} - \alpha) \xrightarrow{\mathcal{D}} \left(N_1, N_2, \ldots, N_\ell, \frac{1}{\ell}\sum_{i=1}^{\ell} N_i\right),$$

as $c \to \infty$, where $N_1, N_2, \ldots, N_\ell$ are i.i.d. $\mathcal{N}(0, \sigma_\star^2)$ for some $\sigma_\star^2$. Hence,

$$\frac{\sqrt{\ell}(\widetilde{\alpha} - \alpha)}{\sqrt{\frac{1}{\ell-1}\sum_{i=1}^{\ell}(\widetilde{\alpha}_i - \widetilde{\alpha})^2}} \xrightarrow{\mathcal{D}} t_{\ell-1}$$

as $c \to \infty$, where $t_{\ell-1}$ is a Student-t random variable with $\ell - 1$ degrees of freedom. The sectioning approach avoids the complications of deriving an explicit central limit theorem. But this is at the cost of producing a Student-t interval, as opposed to a normal interval. Hence, the intervals are slightly larger and more variable than in the case where we consistently estimate the variance constant explicitly.

APPENDIX

PROOF OF THEOREM 3.8. For each $i = 1, 2, \ldots, d$, let $Z_k(x_i) = Y_k(x_i) - \alpha \tau_k(x_i)$, $k \geq 1$, and set

$$\bar{Z}_n(x_i) = \frac{1}{\lfloor p_i n \rfloor} \sum_{k=1}^{\lfloor p_i n \rfloor} Z_k(x_i).$$

Let $\bar{Z}_n = (\bar{Z}_n(x) : x \in A) \in \Re^{d \times 1}$ and $\bar{\tau}_n = (\bar{\tau}_n(x) : x \in A) \in \Re^{d \times 1}$. Note that

$$n^{1/2}(\alpha_n - \alpha) = n^{1/2} \frac{\sum_{i=1}^{d} \nu_n(x_i) \bar{Z}_n(x_i)}{\sum_{i=1}^{d} \nu_n(x_i) \bar{\tau}_n(x_i)} = n^{1/2} \frac{\nu_n \bar{Z}_n}{\nu_n \bar{\tau}_n} = C_n + D_n, \qquad (42)$$

where

$$C_n = n^{1/2} \frac{\nu_n \bar{Z}_n}{E_\nu[T]}, \qquad D_n = n^{1/2} \nu_n \bar{Z}_n \left(\frac{1}{\nu_n \bar{\tau}_n} - \frac{1}{E_\nu[T]} \right). \qquad (43)$$

We begin by analyzing C_n. The second equality in Proposition 3.5 implies that

$$\nu z = 0. \qquad (44)$$

Hence,

$$C_n = n^{1/2} \frac{(\nu_n - \nu) \bar{Z}_n + \nu(\bar{Z}_n - z)}{E_\nu[T]}. \qquad (45)$$

We now need to analyze $\nu_n - \nu$ and $\bar{Z}_n - z$.

Assumptions 2.1 and 3.4 imply that R is finite and irreducible. It then follows that R_n is finite and irreducible for n sufficiently large by (3). Since $\nu R = \nu$ and $\nu_n = \nu_n R_n$,

$$\nu_n - \nu = \nu_n R_n - \nu R = \nu_n(R_n - R) + (\nu_n - \nu)R, \qquad (46)$$

and so

$$(\nu_n - \nu)(I - R) = \nu_n(R_n - R). \qquad (47)$$

Then because $\nu_n e - \nu e = 1 - 1 = 0$ and $V = e\nu$, we have that $(\nu_n - \nu)V = 0$. Hence, we can rewrite (47) as

$$(\nu_n - \nu)(I - R + V) = \nu_n(R_n - R).$$

Therefore,

$$\nu_n - \nu = \nu_n(R_n - R)F. \qquad (48)$$

Note that $R_n - R \to 0$ a.s. by (3). Also, $|A| < \infty$ by Assumption 3.4, and ν_n is bounded for all n. Thus, (48) implies that

$$\nu_n \to \nu \quad \text{a.s.} \qquad (49)$$

as $n \to \infty$. Also because of (48), (45) becomes

$$C_n = n^{1/2} \frac{\nu_n(R_n - R)F\bar{Z}_n + \nu(\bar{Z}_n - z)}{E_\nu[T]}. \tag{50}$$

Under Assumptions 2.1 and 3.6 (see Proposition 3.7),

$$n^{1/2}(R_n - R, \bar{Z}_n - z) \xrightarrow{\mathcal{D}} (N_1, N_2) \stackrel{\mathcal{D}}{=} \mathcal{N}(0, \Sigma), \tag{51}$$

where $\stackrel{\mathcal{D}}{=}$ denotes equality in distribution. Thus, (N_1, N_2) are jointly normally distributed with some covariance matrix Σ. (We will examine the structure of the matrix Σ below.) Because $\bar{Z}_n \to z$ a.s. and by (49), $(n^{1/2}(R_n - R), n^{1/2}(\bar{Z}_n - z), \nu_n, \bar{Z}_n) \xrightarrow{\mathcal{D}} (N_1, N_2, \nu, z)$ as $n \to \infty$ by Theorem 3.9 of [Billingsley 1999]. Hence, the continuous mapping theorem (e.g., Theorem 2.7 of [Billingsley 1999]) implies that

$$C_n \xrightarrow{\mathcal{D}} \frac{\nu N_1 Fz + \nu N_2}{E_\nu[T]}, \tag{52}$$

as $n \to \infty$.

We now consider D_n from (43). Note that for each $x \in A$, $\bar{\tau}_n(x) \to E_x[\tau]$ a.s. and $E_x[\tau] > 0$. Thus, (49) implies that $(1/(\nu_n\bar{\tau}_n)) - (1/E_\nu[T]) \xrightarrow{\mathcal{D}} 0$ by the continuous mapping theorem. So it follows from (52) that

$$D_n \xrightarrow{\mathcal{D}} 0. \tag{53}$$

Therefore, using (42) and (52) gives

$$n^{1/2}(\alpha_n - \alpha) \xrightarrow{\mathcal{D}} \frac{\nu(N_1\zeta + N_2)}{E_\nu[T]} \tag{54}$$

by the converging-together lemma (see Theorem 25.4 of [Billingsley 1995]), since $\zeta = Fz$.

We now examine the structure of the covariance matrix Σ in (51), which we will need to know to determine the exact form of the variance of the limiting distribution in (54). Many of the entries in Σ are zero since $(R_n(x_i, \cdot), \bar{Z}_n(x_i))$ is independent of $(R_n(x_j, \cdot), \bar{Z}_n(x_j))$ for $i \neq j$, where we use the notation that for $x \in A$, $M(x, \cdot) = (M(x, y) : y \in A)$ for a matrix $M = (M(u, v) : u, v \in A)$. We now consider separately the nonzero components of Σ. Recall we previously defined the matrix Ψ_i in (5) and (6), and note that

$$n^{1/2}\left(R_n(x_i, \cdot) - R(x_i, \cdot)\right) \xrightarrow{\mathcal{D}} \mathcal{N}(0, p_i^{-1}\Psi_i),$$

as $n \to \infty$. Also, $n^{1/2}(R_n(x_i, x_j) - R(x_i, x_j), \bar{Z}_n(x_i) - z(x_i))$ converges in distribution to a normal random vector with mean 0 and covariance matrix

$$p_i^{-1} \begin{pmatrix} \Psi_i(x_j, x_j) & g(x_i, x_j) \\ g(x_i, x_j) & h(x_i) \end{pmatrix}.$$

Now the only thing left to show is that the asymptotic variance σ^2 is given by (7). In (54), $N_1\zeta + N_2$ is a random vector in which each component $(N_1\zeta + N_2)(x_i)$ is normally distributed since it is a linear combination of dependent normals, and

$$\mathrm{Var}((N_1\zeta + N_2)(x_i)) = \frac{\eta_i}{p_i}. \tag{55}$$

For $i \neq j$, $(N_1(x_i, \cdot), N_2(x_i))$ and $(N_1(x_j, \cdot), N_2(x_j))$ are independent because of the independence of $(R_n(x_i, \cdot), \bar{Z}_n(x_i))$ and $(R_n(x_j, \cdot), \bar{Z}_n(x_j))$. Therefore, $N_1\zeta + N_2$ is a normally distributed random vector with independent components, which implies that $\nu(N_1\zeta + N_2)$ is a linear combination of the independent $(N_1\zeta + N_2)(x)$, $x \in A$, and so (7) follows. $\square$

PROOF OF THEOREM 3.10. For $x \in A$, define $Z_k'(x) = Y_k'(x) - \alpha\tau_k'(x)$, $k \geq 1$, and set

$$\bar{Z}_n'(x) = \frac{1}{H_n(x)} \sum_{k=1}^{H_n(x)} Z_k'(x), \qquad \bar{\tau}_n'(x) = \frac{1}{H_n(x)} \sum_{k=1}^{H_n(x)} \tau_k'(x).$$

Let $\bar{Z}_n' = (\bar{Z}_n'(x) : x \in A)$ and $\bar{\tau}_n' = (\bar{\tau}_n'(x) : x \in A)$. Since $\nu z = 0$ by (44), we have that

$$n^{1/2}(\alpha_n' - \alpha) = C_n' + D_n',$$

where

$$C_n' = n^{1/2}\frac{\widehat{\nu}_n\bar{Z}_n'}{E_\nu[T]} = n^{1/2}\frac{(\widehat{\nu}_n - \nu)\bar{Z}_n' + \nu(\bar{Z}_n' - z)}{E_\nu[T]},$$

$$D_n' = n^{1/2}\widehat{\nu}_n\bar{Z}_n'\left(\frac{1}{\widehat{\nu}_n\bar{\tau}_n'} - \frac{1}{E_\nu[T]}\right),$$

and $\widehat{\nu}_n = (\widehat{\nu}_n(x) : x \in A)$.

For $x \in A$, $\nu(x) > 0$ by Proposition 3.3. Also,

$$\widehat{\nu}_n(x) = \frac{H_n(x)}{n} \to \nu(x) > 0 \text{ a.s.} \tag{56}$$

as $n \to \infty$, so $H_n(x) \to \infty$ a.s. Recall the definition of R_n' at (13). By (56), $R_n' \to R$ a.s., and since R is irreducible and positive recurrent by Proposition 3.3, R_n' also is for sufficiently large n. Thus, for n sufficiently large, we can define $\nu_n' = (\nu_n'(x) : x \in A)$ such that $\nu_n' = \nu_n'R_n'$, $\nu_n' \geq 0$, and $e\nu_n' = 1$. Observe that

$$\begin{aligned}\widehat{\nu}_n - \nu &= \widehat{\nu}_nR_n' - \nu R + (\widehat{\nu}_n - \nu_n')(I - R_n') \\ &= \widehat{\nu}_n(R_n' - R) + (\widehat{\nu}_n - \nu)R + (\widehat{\nu}_n - \nu_n')(I - R_n').\end{aligned}$$

Now arguing as we did to go from (46) to (48), we can show that

$$\widehat{\nu}_n - \nu = \widehat{\nu}_n(R_n' - R)F + (\widehat{\nu}_n - \nu_n')(I - R_n')F. \tag{57}$$

Hence,

$$C_n' = n^{1/2}\frac{\widehat{\nu}_n(R_n' - R)F\bar{Z}_n' + \nu(\bar{Z}_n' - z)}{E_\nu[T]} + n^{1/2}\frac{(\widehat{\nu}_n - \nu_n')(I - R_n')F\bar{Z}_n'}{E_\nu[T]}. \tag{58}$$

Note that (56) implies that $\bar{Z}_n'(x) \to z(x)$ a.s. and $\bar{\tau}_n'(x) \to E_x[\tau]$ a.s., and the random-time-change central limit theorem (e.g., p. 32 of [Serfling 1980]) implies that

$$n^{1/2}(R_n' - R, \bar{Z}_n' - z) \xrightarrow{\mathcal{D}} (N_1', N_2') \overset{\mathcal{D}}{=} \mathcal{N}(0, \Sigma'), \tag{59}$$

i.e., (N_1', N_2') are jointly normally distributed with some covariance matrix Σ', for which we will later give the exact form. It then follows that, as in (52),

$$n^{1/2} \frac{\widehat{\nu}_n(R_n' - R)F\bar{Z}_n' + \nu(\bar{Z}_n' - z)}{E_\nu[T]} \xrightarrow{\mathcal{D}} \frac{\nu N_1' \zeta + \nu N_2'}{E_\nu[T]}$$

by (59) since $\widehat{\nu}_n \to \nu$ a.s. and $\bar{Z}_n' \to z$ a.s.

We now show that

$$n^{1/2}(\widehat{\nu}_n - \nu_n') \xrightarrow{\mathcal{D}} 0 \tag{60}$$

as $n \to \infty$. To see this, define $\epsilon_n(x)$ by

$$\epsilon_n(x) = \widehat{\nu}_n(x) - \widehat{\nu}_n R_n'(x) = \frac{H_n(x)}{n} - \sum_{y \in A} \frac{H_n(y)}{n} \frac{\sum_{k=0}^{n-1} I(W_k = y, W_{k+1} = x)}{\sum_{k=0}^{n-1} I(W_k = y)}$$

$$= \frac{1}{n} \left(H_n(x) - \sum_{y \in A} \sum_{k=0}^{n-1} I(W_k = y, W_{k+1} = x) \right)$$

$$= \frac{1}{n} \left(I(W_0 = x) - I(W_n = x) \right).$$

Therefore, $|\epsilon_n(x)| \leq 1/n$ for each x. Let V_n' be the matrix with each row equal to ν_n', and let $F_n' = (I - R_n' + V_n')^{-1}$. For any probability vector $\mu \in \Re^{1 \times d}$, $\mu V_n' = \nu_n'$. Then $\widehat{\nu}_n = \widehat{\nu}_n R_n' + \epsilon_n = \widehat{\nu}_n (R_n' - V_n') + \nu_n' + \epsilon_n$, or

$$\widehat{\nu}_n (I - R_n' + V_n') = \nu_n' + \epsilon_n = \nu_n' (I - R_n' + V_n') + \epsilon_n.$$

Therefore, $\widehat{\nu}_n - \nu_n' = \epsilon_n F_n'$. Since $F_n' \to F$ a.s. and $|\epsilon_n(x)| \leq 1/n$ for all x, (60) follows.

Hence, the second term on the right-hand side of (58) satisfies

$$n^{1/2} \frac{(\widehat{\nu}_n - \nu_n')(I - R_n')F\bar{Z}_n'}{E_\nu[T]} \xrightarrow{\mathcal{D}} 0$$

since $R_n' \to R$ a.s. and $\bar{Z}_n' \to z$ a.s. Also, as in (53) we can similarly show that $D_n' \xrightarrow{\mathcal{D}} 0$. Therefore,

$$n^{1/2}(\alpha_n' - \alpha) \xrightarrow{\mathcal{D}} \frac{\nu N_1' \zeta + \nu N_2'}{E_\nu[T]}.$$

We now examine the structure of the covariance matrix Σ' in (59). Even though $R_n'(x_i, \cdot)$ is not independent of $R_n'(x_j, \cdot)$ for $i \neq j$, it turns out that $N_1'(x_i, \cdot)$ and $N_1'(x_j, \cdot)$ are independent; e.g., see Theorem 3.1 and p. 23 of [Billingsley 1961].

We now show that $\text{Cov}(N_2'(x_i), N_2'(x_j)) = 0$ for $i \neq j$. For $x_i \in A$, define

$$\widehat{Z}_n'(x_i) = \frac{1}{\lfloor \nu(x_i)n \rfloor} \sum_{k=1}^{\lfloor \nu(x_i)n \rfloor} Z_k'(x_i).$$

By slightly modifying the argument on p. 20 of [Billingsley 1961], we can prove that for each $x_i \in A$,

$$n^{1/2}(\bar{Z}_n'(x_i) - z(x_i)) - n^{1/2}(\widehat{Z}_n'(x_i) - z(x_i)) \xrightarrow{\mathcal{D}} 0$$

as $n \to \infty$. Now let $\widehat{Z}'_n = (\widehat{Z}'_n(x_i) : i = 1, 2, \ldots, d)$, and it then follows that

$$n^{1/2}(\bar{Z}'_n - z) - n^{1/2}(\widehat{Z}'_n - z) \xrightarrow{\mathcal{D}} 0$$

as $n \to \infty$. Hence, the converging-together lemma implies that $n^{1/2}(\widehat{Z}'_n - z) \xrightarrow{\mathcal{D}} N'_2$ as $n \to \infty$, where N'_2 is defined in (59); i.e., $n^{1/2}(\widehat{Z}'_n - z)$ and $n^{1/2}(\bar{Z}'_n - z)$ have the same limiting distribution. But the Markov property implies that for each n, $\widehat{Z}'_n(x_i)$ and $\widehat{Z}'_n(x_j)$ are independent for $i \neq j$. Consequently, $N'_2(x_i)$ and $N'_2(x_j)$ are independent for $i \neq j$, so $\mathrm{Cov}(N'_2(x_i), N'_2(x_j)) = 0$. We can similarly show that $N'_1(x_i, \cdot)$ and $N'_2(x_j)$ are independent for $i \neq j$.

Thus, Σ' has non-zero entries in the same places as Σ in (51) does. The difference between Σ' and Σ is that the non-zero entries in Σ are divided by the appropriate p_i, whereas they are divided by the appropriate $\nu(x_i)$ in Σ'. The divisors $\nu(x_i)$ in Σ' arise from the random-time-change central limit theorem because $H_n(x_i)/n \to \nu(x_i)$ a.s. as we showed in (56). For example, note that $n^{1/2}(\bar{Z}'_n(x_i) - z) \xrightarrow{\mathcal{D}} \mathcal{N}(0, h(x_i)/\nu(x_i))$, where $h(x_i) = \mathrm{Var}_{x_i}(Z)$ as before. Thus, as in the proof of Theorem 3.8, we can work out the exact form of the asymptotic variance σ_1^2 to be (12).   □

REFERENCES

ALEXOPOULOS, C. AND SHULTES, B. C. 1998. The balanced likelihood ratio method for estimating performance measures of highly reliable systems. In *Proceedings of the 1998 Winter Simulation Conference*, D. J. Medeiros, E. F. Watson, J. S. Carson, and M. S. Manivannan, Eds. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey, 1479–1486.

ANDRADÓTTIR, S., CALVIN, J. M., AND GLYNN, P. W. 1995. Accelerated regeneration for Markov chain simulations. *Probability in the Engineering and Informational Sciences 9*, 497–523.

BILLINGSLEY, P. 1961. Statistical methods in Markov chains. *Annals of Mathematical Statistics 32*, 12–40.

BILLINGSLEY, P. 1995. *Probability and Measure*, Third ed. John Wiley & Sons, New York.

BILLINGSLEY, P. 1999. *Convergence of Probability Measures*, Second ed. John Wiley & Sons, New York.

CALVIN, J. M., GLYNN, P. W., AND NAKAYAMA, M. K. 2001. Importance sampling using the semi-regenerative method. In *Proceedings of the 2001 Winter Simulation Conference*, B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, Eds. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey, 441–450.

CALVIN, J. M. AND NAKAYAMA, M. K. 1998. Using permutations in regenerative simulations to reduce variance. *ACM Transactions on Modeling and Computer Simulation 8*, 153–193.

CALVIN, J. M. AND NAKAYAMA, M. K. 2000a. Central limit theorems for permuted regenerative estimators. *Operations Research 48*, 776–787.

CALVIN, J. M. AND NAKAYAMA, M. K. 2000b. Simulation of processes with multiple regeneration sequences. *Probability in the Engineering and Informational Sciences 14*, 179–201.

CALVIN, J. M. AND NAKAYAMA, M. K. 2002. A comparison of output-analysis methods for simulations of processes with multiple regeneration sequences. In *Proceedings of the 2002 Winter Simulation Conference*, E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, Eds. Institute of Electrical and Electronics Engineers, Piscataway, New Jersey, 328–335.

ÇINLAR, E. 1975. *Introduction to Stochastic Processes*. Prentice Hall, Englewood Cliffs, New Jersey.

CHUNG, K. L. 1967. *Markov Chains with Stationary Transition Probabilities*. Springer-Verlag, New York.

COCHRAN, W. G. 1977. *Sampling Techniques*, Third ed. Wiley, New York.

CRANE, M. AND IGLEHART, D. L. 1975. Simulating stable stochastic systems, III: Regenerative processes and discrete-event simulations. *Operations Research 23*, 33–45.

FOX, B. L. AND GLYNN, P. W. 1986. Discrete time conversion for simulating semi-Markov processes. *Oper. Res. Lett. 5*, 191–196.

GLYNN, P. AND HEIDELBERGER, P. 1990. Bias properties of budget constrained monte carlo simulations. *Operations Research 38*, 801–814.

GLYNN, P. W. 1986. Sensitivity analysis for stationary probabilities of a Markov chain. In *Proceedings of the 4th Army Conference on Applied Mathematics and Computing*. 917–932.

GLYNN, P. W. 1987. Likelihood ratio derivative estimation: An overview. In *Proceedings of the 1987 Winter Simulation Conference*, A. Thesen, H. Grant, and W. D. Kelton, Eds. 366–375.

GLYNN, P. W. 1990. Likelihood ratio derivative estimators for stochastic systems. *Comm. ACM 33*, 75–84.

GLYNN, P. W. 1994. Some topics in regenerative steady-state simulation. *Acta Applicandae Mathematicae 34*, 225–236.

GLYNN, P. W. 1995. Importance sampling for Markov chains: Asymptotics for the variance. *Stochastic Models 10*, 701–717.

GLYNN, P. W. AND IGLEHART, D. L. 1989. Importance sampling for stochastic systems. *Management Science 35*, 1367–1393.

GLYNN, P. W. AND WHITT, W. 1992. The asymptotic efficiency of simulation estimators. *Operations Research 40*, 505–520.

GOYAL, A., SHAHABUDDIN, P., HEIDELBERGER, P., NICOLA, V., AND GLYNN, P. W. 1992. A unified framework for simulating Markovian models of highly dependable systems. *IEEE Transactions on Computers C-41*, 36–51.

GUNTHER, F. L. AND WOLFF, R. W. 1980. The almost regenerative method. *Operations Research 28*, 375–387.

HEIDELBERGER, P. 1995. Fast simulation of rare events in queueing and reliability models. *ACM Transactions on Modeling and Computer Simulation 5*, 43–85.

HENDERSON, S. AND GLYNN, P. 2001. Regenerative steady-state simulation of discrete-event systems. *ACM Transactions on Modeling and Computer Simulation 11*, 313–345.

HORDIJK, A., IGLEHART, D. L., AND SCHASSBERGER, R. 1976. Discrete-time methods for simulating continuous-time Markov chains. *Adv. in Appl. Probab. 8*, 772–788.

KEMENY, J. G. AND SNELL, J. L. 1960. *Finite Markov Chains*. Van Nostrand, Princeton, NJ.

MEKETON, M. S. AND HEIDELBERGER, P. 1982. A renewal theoretic approach to bias reduction in regenerative simulations. *Management Science 24*, 173–181.

REIMAN, M. I. AND WEISS, A. 1989. Sensitivity analysis for simulations via likelihood ratios. *Operations Research 37*, 830–844.

RUBINSTEIN, R. Y. 1989. Sensitivity analysis and performance extrapolation for computer simulation models. *Operations Research 37*, 72–81.

SERFLING, R. J. 1980. *Approximation Theorems of Mathematical Statistics*. John Wiley & Sons, New York.

SHEDLER, G. S. 1993. *Regenerative Stochastic Simulation*. Academic Press, San Diego.

ZHANG, B. AND HO, Y.-C. 1992. Improvements in the likelihood ratio method for steady-state sensitivity analysis and simulation. *Performance Analysis 15*, 177–195.