

NOTICE: This material may be protected by Copyright Law
(Title 17 U.S. code)

9/7/2010

Borrower: ST2
Lending String: *NJR,NJR

Patron: Rhee, Changan

Journal Title: Transactions of the Fourth Army
Conference on Applied Mathematics and Computing /
ISSN:

Volume: (missing) **Issue:**
Month/Year: 1986**Pages:** 917-932

Article Author:

Article Title: Peter Glynn; Sensitivity Analysis for
Stationary Probabilities of a Markov Chain

Imprint: Research Triangle Park, NC ; U.S. Army R

ILL Number: 68643839



Call #: QA9.59.A74 1986 --- COPY of pages

Location: MATH RESV-MATH2
Date: 9/7/2010

Charge
Maxcost: 501FM

Shipping Address:
Stanford University
ILL
Green Library, Room 161
1 Crothers Way
Stanford, CA 94305-6004

Fax: 650 725-6874
Ariel: 171.67.128.63
Odyssey: 171.64.10.44

Rutgers The State University of New Jersey ILL



ILLiad TN: 198384

**SEND TO
ILS -ALEX
LENDING**

SENSITIVITY ANALYSIS
FOR
STATIONARY PROBABILITIES OF MARKOV CHAINS

Peter W. Glynn¹
University of Wisconsin

NOTICE
This material may be protected by
Copyright Law (Title 17 U.S. Code)

ABSTRACT

This paper considers the problem of evaluating the sensitivity of a steady-state cost $\alpha(\theta)$ to underlying uncertainty in a parameter vector θ governing the probabilistic dynamics of the system under consideration. We show that the gradient $\nabla\alpha(\theta)$ plays a fundamental role in the parametric statistical theory for Markov processes. We then survey numerical methods available for evaluating $\nabla\alpha(\theta)$ and introduce a new Monte Carlo estimator for $\nabla\alpha(\theta)$, which is applicable to Markov processes of substantial generality.

¹Research supported by the United States Army under Contract No. DAAG29-84-K-0030 and by NSF Grant ECS-840-4809.

1. INTRODUCTION

Let $X = \{X_n : n \geq 0\}$ be an irreducible positive recurrent Markov chain governed by a transition kernel $P(\theta)$, where θ is a parameter vector taking values in \mathbb{R}^d . If $\pi(\theta)$ is the stationary measure of $P(\theta)$ and $f(\theta, x)$ is the cost of running the chain while in state x , then $\alpha(\theta) \equiv \int f(\theta, x) \pi(\theta, dx)$ is the long-run average cost of running X under parameter choice θ . In many applications settings, it is of interest to compute the sensitivity of α to (infinitesimal) changes in the parameter θ . Specifically, it is frequently useful to be able to evaluate $\nabla \alpha(\theta)$, the gradient of $\alpha(\cdot)$ evaluated at $\theta \in \mathbb{R}^d$. Since it is generally impossible to analytically evaluate $\nabla \alpha(\theta)$ (except for simple models), this paper will concentrate on numerical methods for determining $\nabla \alpha(\theta)$.

This paper is organized as follows. In Section 2, we introduce an important statistical application for these methods. We show that the numerical methods discussed here offer the opportunity to do statistical point, variance, and interval estimation for highly complex functionals of analytically intractable Markov processes. Section 3 is devoted to the formal derivation of an expression for $\nabla \alpha(\theta)$ and describes, for finite state Markov chains, a set of linear equations which characterizes $\nabla \alpha(\theta)$. For complicated stochastic processes, the corresponding linear systems are too complex to solve via standard numerical methods, and Monte Carlo techniques therefore become relevant. Thus, Section 4 provides a (new) Monte Carlo estimator for $\nabla \alpha(\theta)$, which is applicable to Markov chains of substantial generality. Finally, Section 5 offers a brief summary of the paper.

2. STATISTICAL RELEVANCE OF THE GRADIENT

Suppose that the transition kernel P governing the Markov chain X is determined by a finite family of distributions (F_1, \dots, F_m) $(F_1(\gamma_1), \dots, F_m(\gamma_m))$, where each $F_i(\gamma_i)$ is a probability distribution associated with a known parametric family in which $\gamma_i \in \mathbb{R}^{d_i}$. If $\theta = (\gamma_1, \dots, \gamma_m)$, then P can be viewed as a function of θ , namely $P = P(\theta)$.

In statistical contexts, the vector $\theta \in \mathbb{R}^d$ ($d = d_1 + \dots + d_m$) is, in general, unknown. Most of the literature on statistical inference for Markov processes has concentrated on estimation of the "true" parameter θ^* (i.e., estimation of θ when the observed chain X is governed by $P(\theta^*)$) and on related issues such as production of variance estimates and confidence intervals. However, in many applications settings, it is of more practical importance to estimate not θ^* but some associated steady-state cost $\alpha(\theta^*)$.

(2.1) EXAMPLE. Let $X = \{X_n : n \geq 1\}$ be the Markov chain consisting of waiting times of consecutive customers in the $M/M/1/\infty$ queue. (See HEYMAN and SOBEL (1982) for a description.) Arrivals follow an $\exp(\gamma_1)$ distribution, whereas service times are distributed $\exp(\gamma_2)$. Suppose that the long-run customer waiting time $\alpha(\theta)$ is of importance, when $\theta = (\gamma_1, \gamma_2)$. The objective is to produce estimates for $\alpha(\theta^*)$, as well as variance and interval estimates, from observed inter-arrival times Y_{11}, \dots, Y_{n1} as well as observed service times Y_{21}, \dots, Y_{2n2} . Note that in certain settings, the inter-arrival times and service times may have been collected

NOTICE

This material may be protected by
Copyright Law (Title 17 U.S. Code)

from two independent sources, so that no waiting times for the system are available. For example, the queue might correspond to a telephone switching system being designed, in which historical inter-arrival data exists and service time data for the proposed switching device is available.

(2.2) EXAMPLE. Virtually any general discrete-event stochastic system can be formulated as a generalized semi-Markov process (GSMP). A GSMP can, in turn, be viewed as a Markov chain $X = \{X_n : n \geq 0\}$, where $X_n = (S_n, C_n)$ records the "physical state" S_n (e.g., configuration of customers in a queue) and clock readings C_n (e.g., remaining service times for each of the customers in the system) at the n^{th} transition of the GSMP. (For further details, see GLYNN (1983).) GSMP's are characterized probabilistically by certain distributions F_1, \dots, F_λ governing the way clocks are reset (e.g., service times in a queue) and by routing probabilities p_1, \dots, p_k (e.g., the proportion of customers who visit station j after receiving service at station i).

In many applications environments, the distributions $F_1, F_2, \dots, F_\lambda$ and routing probabilities p_1, \dots, p_k are unknown and must be estimated via statistical methods. If one models the distributions F_1, \dots, F_λ as belonging to parametric families (i.e., $F_i = F_i(\gamma_i)$), then the transition function P governing X can be viewed as $P = P(\theta)$, where $\theta = (\gamma_1, \dots, \gamma_\lambda, p_1, \dots, p_k)$. The performance of a stochastic system is often assessed by considering a long-run average cost α for the system which, in this context, can be regarded as a function $\alpha = \alpha(\theta)$ of the unknown parameter θ associated with P . Consequently, an important

statistical objective involves point and interval estimation of $\alpha(\theta^*)$, where θ^* is the "true" parameter governing the system.

We will now outline a method for obtaining point and interval estimates for $\alpha(\theta^*)$, which is applicable to very general stochastic systems. Let $\hat{\theta} = (\hat{\theta}_1(n_1), \dots, \hat{\theta}_d(n_d))$ be an estimator for $\theta^* = (\theta_1^*, \dots, \theta_d^*)$ (n_i is the sample size associated with estimation of θ_i^* .) Such estimators $\hat{\theta}$ are frequently available for complex systems. In particular, one can often appeal to maximum likelihood estimation (MLE) methods for estimating θ^* . Under very general conditions, $\hat{\theta}$ will be asymptotically normal, in the sense that

$$(2.3) \quad \hat{\theta} \stackrel{\mathcal{D}}{\approx} N(\theta^*, C(n_1, \dots, n_d))$$

where $N(\theta^*, C(n_1, \dots, n_d))$ is a multivariate normal r.v. with mean θ^* and covariance matrix $C(n_1, \dots, n_d)$. ($\stackrel{\mathcal{D}}{\approx}$ denotes "has approximately the distribution of".) In certain design settings (see Example 2.1), the data for each of the different components θ_i^* is gathered from independent sources. In this case, $C(n_1, \dots, n_d)$ takes the diagonal form

$$(2.4) \quad C(n_1, \dots, n_d) = \begin{pmatrix} \sigma_1^2/n_1 & & 0 \\ & \ddots & \\ 0 & & \sigma_d^2/n_d \end{pmatrix}.$$

If α is continuously differentiable in a neighborhood of θ^* , then a Taylor expansion of α around θ^* shows that (2.3) yields

$$(2.5) \quad \alpha(\hat{\theta}) \stackrel{\mathcal{D}}{\approx} N(\alpha(\theta^*), \nabla\alpha(\theta^*)^t C(n_1, \dots, n_d) \nabla\alpha(\theta^*))$$

where $\nabla\alpha(\theta^*)$ is the (column) gradient of α evaluated at θ^* . (This is the so-called "delta method" of statistics.)

Relation (2.5) shows that if n_1, \dots, n_d are large, then $\alpha(\hat{\theta})$ is a good point estimator for $\alpha(\theta^*)$. Let $\hat{C}(n_1, \dots, n_d)$ be an estimator for $C(n_1, \dots, n_d)$ (such variance estimators are commonly available for MLE point estimates $\hat{\theta}$). Then, (2.5) proves that

$$(2.6) \quad \hat{v} \equiv \nabla\alpha(\hat{\theta})^t \hat{C}(n_1, \dots, n_d) \nabla\alpha(\hat{\theta})$$

is an estimator for the variance of $\alpha(\hat{\theta})$ and,

$$(2.7) \quad [\alpha(\hat{\theta}) - z(\delta)\hat{v}^{1/2}, \alpha(\hat{\theta}) + z(\delta)\hat{v}^{1/2}]$$

is an approximate $100(1-\delta)\%$ confidence interval for $\alpha(\theta^*)$, where $z(\delta)$ is the solution of $P\{N(0,1) \leq z(\delta)\} = 1 - \delta/2$. Thus, provided that $\alpha(\hat{\theta})$ and $\nabla\alpha(\hat{\theta})$ can be evaluated (either analytically or numerically), (2.6) and (2.7) provide a solution to the variance and interval estimation problems discussed above.

In the case that the covariance matrix $C(n_1, \dots, n_d)$ takes the form (2.4), \hat{v} can be expressed as

$$(2.8) \quad \hat{v} = \sum_{i=1}^d \left(\frac{\partial}{\partial \theta_i} \alpha(\hat{\theta}) \right)^2 \hat{\sigma}_i^2 / n_i.$$

Relation (2.8) shows that the contribution of uncertainty in θ_i^* to the variance of $\alpha(\hat{\theta})$ is given by $(\partial\alpha(\hat{\theta})/\partial\theta_i)^2 \hat{\sigma}_i^2 / n_i$. This can be used to determine which component to additionally sample if the current estimator of $\alpha(\theta^*)$ is too "noisy."

NOTICE
This material may be protected by
Copyright Law (Title 17 U.S. Code)

(2.1) EXAMPLE (continued). Because of the simplicity of the M/M/1/ ∞ queue, α can be analytically determined in closed form, namely $\alpha(\gamma_1, \gamma_2) = \gamma_2(\gamma_2 - \gamma_1)^{-1}$ for $\gamma_1 < \gamma_2$ (∞ for $\gamma_1 \geq \gamma_2$). If $\hat{\gamma}_1 < \hat{\gamma}_2$, (2.8) reduces to

$$\hat{v} = \frac{\hat{\gamma}_2^2}{(\hat{\gamma}_2 - \hat{\gamma}_1)^4} \hat{\sigma}_1^2/n_1 + \frac{\hat{\gamma}_1^2}{(\hat{\gamma}_2 - \hat{\gamma}_1)^4} \hat{\sigma}_2^2/n_2,$$

where $\hat{\sigma}_1^2(\hat{\sigma}_2^2)$ is a variance estimate for $\gamma_1(\gamma_2)$ formed from $Y_{11}, \dots, Y_{1n_1}(Y_{21}, \dots, Y_{2n_2})$.

For more complicated systems, such as that described in Example 2.2, $\alpha(\cdot)$ cannot be determined analytically, and so one must turn to numerical algorithms. These algorithms will be described in the remaining sections of this paper.

3. A FORMULA FOR THE GRADIENT OF THE STEADY-STATE

Let $P(\theta)$ be the transition function for X under parameter θ , so that $P(\theta, x, A)$ is the corresponding conditional probability that $X_{n+1} \in A$, given that $X_n = x$. For an initial distribution $\mu(\theta)$, let P_θ be the probability measure on the path-space of X associated with $P(\theta)$, namely

$$P_\theta\{X_0 \in A_0, \dots, X_n \in A_n\} = \int_{A_0} \mu(\theta, dx_0) \int_{A_1} P(\theta, x_0, dx_1) \dots \int_{A_n} P(\theta, x_{n-1}, dx_n).$$

If X is Harris recurrent under $P(\theta)$ (see REVUZ (1984)), then there exists a unique probability measure $\pi(\theta)$ such that

$$(3.1) \quad \frac{1}{n} \sum_{k=0}^{n-1} f(\theta, X_k) \rightarrow \int_S f(\theta, x) \pi(\theta, dx) P_\theta \text{ a.s.}$$

as $n \rightarrow \infty$ (for a large class of $f(\theta)$'s). The measure $\pi(\theta)$ is stationary for $P(\theta)$, in the sense that

$$(3.2) \quad \pi(\theta, \cdot) = \int_S P(\theta, x, \cdot) \pi(\theta, dx) .$$

(S is the state space of X .) In fact, $\pi(\theta)$ is the unique probability measure satisfying (3.2). Our goal is to numerically compute $\alpha(\theta)$ and $\nabla \alpha(\theta)$, where $\alpha(\theta)$ is the steady-state limit

$$(3.3) \quad \alpha(\theta) = \int_S f(\theta, x) \pi(\theta, dx) .$$

Since (3.2) only determines $\pi(\theta)$ up to a multiplicative constant, it is necessary to add an additional constraint stating that the total mass $\pi(\theta, S)$ equals 1. The quantity $\alpha(\theta)$ is then the unique solution of the integral equation system

$$(3.4) \quad \begin{aligned} \pi(\theta, \cdot) &= \int_S P(\theta, x, \cdot) \pi(\theta, dx) \\ \pi(\theta, S) &= 1 \\ \alpha(\theta) &= \int_S f(\theta, x) \pi(\theta, dx) . \end{aligned}$$

The system (3.4) is well known and has been extensively studied. If S is finite, then $P(\theta)$ is a finite matrix and (3.4) becomes

$$\begin{aligned}
 (3.5) \quad & \pi(\theta)^t = \pi(\theta)^t P(\theta) \\
 & \pi(\theta)^t e = 1 \\
 & \alpha(\theta) = \pi(\theta)^t f(\theta)
 \end{aligned}$$

(all vectors are column vectors; e is the vector consisting of 1's).

As we shall see, a similar system describes the gradient $\nabla \alpha(\theta)$ of α . Let us formally suppose that the transition function $P(\theta)$ can be expanded as

$$P(\theta + h e_i) = P(\theta) + h Q_i(\theta) + o(h)$$

where e_i is the i^{th} unit vector in \mathbb{R}^d . Assume that $\pi(\theta + h e_i)$ is formally differentiable at $h = 0$, so that there exists a signed measure $\eta_i(\theta)$ such that

$$(3.6) \quad \pi(\theta + h e_i) = \pi(\theta) + h \eta_i(\theta) + o(h) .$$

The stationarity equation (3.2) implies that $\eta_i(\theta)$ must satisfy

$$(3.7) \quad \eta_i(\theta, dx) - \int_S \eta_i(\theta, dx) P(\theta, x, \cdot) = \int_S Q_i(\theta, x, \cdot) \pi(\theta, dx)$$

(formally differentiate both sides of (3.2)). (The equation (3.7) is Poisson's equation for the kernel $P(\theta)$.) These formal calculations can be made rigorous, even in general state space; such arguments will appear elsewhere.

In finite state space, the arguments are more straightforward and have previously appeared in SCHWEITZER (1968), GOLUB and MEYER (1986), and MEYER

and STEWART (1986). We give a very elementary proof in the Appendix to this paper; our argument uses only elementary Markov chain theory. Note that in finite state space, (3.7) becomes $\eta_i(\theta)^t(I-P(\theta)) = \pi(\theta)^t Q_i(\theta)$. This does not uniquely identify $\eta_i(\theta)$, since $\eta_i(\theta) + \delta\pi(\theta)$ also satisfies the equation, for all δ . Note that since $\pi(\theta)^t e = 1$ for all θ , it follows that $\eta_i(\theta)^t e = 0$ (see (3.6)). Let $\Pi(\theta)$ be the matrix in which all rows are identical to $\pi(\theta)$. It is easily verified that since $\eta_i(\theta)^t e = 0$, $\eta_i(\theta)^t \Pi(\theta) = 0$. Consequently, $\eta_i(\theta)$ also satisfies

$$(3.8) \quad \eta_i(\theta)^t(I - P(\theta) + \Pi(\theta)) = \pi(\theta)^t Q_i(\theta) .$$

It is well known (see KEMENY and SNELL (1960), p. 100) that $(I-P(\theta)+\Pi(\theta))$ has an inverse, called the fundamental matrix, which we shall denote $F(\theta)$. Hence, in finite state space, the i^{th} component of $\nabla\alpha(\theta)$ can be computed as the solution of the system

$$(3.9) \quad \begin{aligned} \eta_i(\theta)^t &= \pi(\theta)^t Q_i(\theta) F(\theta) \\ \frac{\partial}{\partial \theta_i} \alpha(\theta) &= \pi(\theta)^t f'_i(\theta) + \eta_i(\theta)^t f(\theta) \end{aligned}$$

where $f'_i(\theta)$ is the vector in which the j^{th} component is $\partial f(\theta, j) / \partial \theta_i$.

Consequently, when S is finite, the systems of linear equations (3.5) and (3.9) may be solved numerically to obtain $\alpha(\theta)$ and $\nabla\alpha(\theta)$. If S is not finite (or if the number of elements in S is large), numerical methods not dependent on explicit solution of linear equations must be considered. In the next section, we show how Monte Carlo methods can be used to advantage here.

NOTICE

This material may be protected by
Copyright Law (Title 17 U.S. Code)

4. MONTE CARLO EVALUATION OF STEADY-STATE GRADIENTS

A critical assumption underlying the analysis of this section is that it is possible to generate sample trajectories of X under the measure P_θ . For the examples that we have in mind (see particularly Example 2.2), this assumption is clearly in force.

Assuming now that X has distribution P_θ , relation (3.1) states that

$$(4.1) \quad \frac{1}{n} \sum_{k=0}^{n-1} f(\theta, X_k) \rightarrow \alpha(\theta) \quad P_\theta \text{ a.s.}$$

as $n \rightarrow \infty$. In other words, rather than solving the integral equation system (3.4), one may numerically approximate $\alpha(\theta)$ by the sample average appearing on the left-hand side of (4.1). The simplicity of this numerical procedure, as well as its broad applicability, is the source of the power of the Monte Carlo method. Our objective here is to obtain a similar Monte Carlo algorithm for evaluation of the gradient $\nabla \alpha(\theta)$.

Observe that (at least formally) we have

$$(4.2) \quad \frac{\partial}{\partial \theta_i} \alpha(\theta) = \int_S \frac{\partial}{\partial \theta_i} f(\theta, x) \pi(\theta, dx) + \int_S f(\theta, x) \eta_i(\theta, dx).$$

A Monte Carlo estimator for the first term appearing on the right-hand side of (4.2) is given by the sample mean

$$(4.3) \quad \frac{1}{n} \sum_{k=0}^{n-1} \frac{\partial}{\partial \theta_i} f(\theta, X_k).$$

It remains to obtain an estimator for the second term.

As in the finite state space context, one expects that the signed measure $\eta_1(\theta)$ will satisfy $\eta_1(\theta, S) = 0$. As a consequence, it follows from (3.7) that $\eta_1(\theta)$ should satisfy

$$(4.4) \quad \eta_1(\theta, \cdot) - \int_S \eta_1(\theta, dx) P(\theta, x, \cdot) + \int_S \eta_1(\theta, dx) \pi(\theta, \cdot) = \int_S Q_1(\theta, x, \cdot) \pi(\theta, dx) .$$

Letting $\Pi(\theta)$ be the operator $\Pi(\theta, x, \cdot) = \pi(\theta, \cdot)$, one can write (4.4) symbolically as

$$(4.5) \quad \eta_1(\theta)(I - P(\theta) + \Pi(\theta)) = \pi(\theta) Q_1(\theta) .$$

(This is the general state space analogue of (3.8).) The formal inverse of $(I - P(\theta) + \Pi(\theta))$ is given by

$$\sum_{k=0}^{\infty} (P(\theta) - \Pi(\theta))^k .$$

Because of the stationarity of $\pi(\theta)$ and the independence of $\Pi(\theta, x, \cdot)$ from x , it follows that $(P(\theta) - \Pi(\theta))^k = P(\theta)^k - \Pi(\theta)$, for $k \geq 1$. Hence, a formal analysis of (4.5) shows that

$$\eta_1(\theta) = \pi(\theta) Q_1(\theta) + \sum_{k=1}^{\infty} \pi(\theta) Q_1(\theta) (P^k(\theta) - \Pi(\theta)) .$$

For the same reason that $\eta_1(\theta, S) = 0$, $Q_1(\theta, x, S) = 0$ and hence $Q_1(\theta)\Pi(\theta) = 0$. Consequently,

NOTICE

This material may be protected by
Copyright Law (Title 17 U.S. Code)

(4.6)

Su
respect
 $Q_1(\theta, x,$
typical

which c

where
probabi
transit
represe

(4.7)

The for
Ea
hypothe

t the signed
ce, it follows

$$(4.6) \quad \eta_1(\theta)f(\theta) = \sum_{k=0}^{\infty} \pi(\theta) Q_1(\theta) P(\theta)^k f(\theta) .$$

(\theta, \cdot)

Suppose that the measures $P(\cdot, x, dy)$ are absolutely continuous with respect to $P(\theta, x, dy)$ in a neighborhood of θ . Then, one expects that $Q_1(\theta, x, dy)$ has a density with respect to $P(\theta, x, dy)$, call it $q_1(\theta, x, y)$. A typical term on the right-hand side of (4.6) then takes the form

write (4.4)

$$\int_S \pi(\theta, dx) \int_S q_1(\theta, x, y) P(\theta, x, dy) \int_S P^k(\theta, y, dz) f(\theta, z)$$

which can be represented probabilistically as an expectation:

formal inverse of

$$\tilde{E}_\theta [q_1(\theta, X_0, X_1) f(\theta, X_{k+1})]$$

where $\tilde{E}_\theta(\cdot)$ is the expectation corresponding to \tilde{P}_θ , and \tilde{P}_θ is the probability on path-space associated with initial distribution $\pi(\theta)$ and transition function $P(\theta)$. Thus, the second term in (4.2) has the formal representation

of $\Pi(\theta, x, \cdot)$

$k \geq 1$. Hence,

$$(4.7) \quad \sum_{k=0}^{\infty} \tilde{E}_\theta [q_1(\theta, X_0, X_1) f(\theta, X_{k+1})] .$$

The formula (4.7) is the key to the Monte Carlo analysis.

Each term in (4.7) can be consistently estimated (under suitable hypotheses) via

$$\frac{1}{n} \sum_{j=0}^{n-1} q_1(\theta, X_j, X_{j+1}) f(\theta, X_{j+k+1})$$

when X evolves according to transition function $P(\theta)$ (regardless of X_0 initial distribution). In order to estimate the infinite sum, a standard device is to consider an estimator of the form

$$(4.8) \quad \sum_{k=0}^{\lambda(n)} \frac{1}{n-\lambda(n)} \sum_{j=0}^{n-\lambda(n)-1} q_1(\theta, X_j, X_{j+1}) f(\theta, X_{j+k+1})$$

where the truncation point $\lambda(n)$ is keyed to the sample size n in such a way that $\lambda(n) \rightarrow \infty$ with $\lambda(n)/n \rightarrow 0$. The particular choice of $\lambda(n)$ effects a compromise between bias and variance effects in estimating the infinite sum (4.7).

Since $q_1(q, x, y)$ is generally easily computable (for S countable, $q_1(\theta, j, k) = (\partial P(\theta, j, k) / \partial \theta_j) \cdot P(\theta, j, k)^{-1}$), (4.7) provides a Monte Carlo solution to estimating the appropriate gradient.

It turns out that (4.6) is closely related to a formula which one obtains when one uses likelihood ratio change-of-measure ideas to evaluate gradients. These connections will be explored more fully in a future paper.

5. SUMMARY

We have shown that the gradient $\nabla \alpha(\theta)$ of steady-state quantity α plays a critical role in the variance and interval estimation theory for steady-state estimators $\alpha(\hat{\theta})$ of complex stochastic systems. In some sense, the large-sample variance and interval estimation theory is fully solved given that one can evaluate $\alpha(\hat{\theta})$ and $\nabla \alpha(\hat{\theta})$. Numerical methods for dealing with $\alpha(\hat{\theta})$ when the system is Markov are, of course, well

known. However, numerical algorithms for evaluating $\nabla\alpha(\hat{\theta})$ are a recent development. We have therefore provided a self-contained exposition of the relevant theory, and discuss both Monte Carlo (see (4.3) and (4.8)) and non-Monte Carlo (see (3.9)) approaches to solving the problem.

APPENDIX

Let $P(\cdot)$ be a family of $n \times n$ stochastic matrices which are:

- (i) irreducible in a neighborhood of θ
- (ii) differentiable at θ .

Under (i), $P(\cdot)$ has a unique stationary distribution $\pi(\cdot)$ in a neighborhood of θ . Our goal is to rigorously verify the first equation in (3.9).

Given the existence of the inverse matrix $F(\theta) = (I - P(\theta) + \Pi(\theta))^{-1}$, (3.9) follows immediately once the differentiability of $\pi(\theta)$ is established. Note that for h sufficiently small,

$$\pi(\theta + he_1) - \pi(\theta) = \pi(\theta + he_1)^T [P(\theta) + hQ_1(\theta) + o(h)] - \pi(\theta)^T P(\theta)$$

so

$$[\pi(\theta + he_1) - \pi(\theta)]^T (I - P(\theta)) = h\pi(\theta + he_1)^T Q_1(\theta) + o(h)$$

(note that $o(h)\pi(\theta + he_1) = o(h)$ since all terms in $\pi(\theta + he_1)$ are uniformly (in h) bounded by 1). Since $\Pi(\theta)$ has identical rows and $\pi(\theta + he_1)$ is stochastic for $h \geq 0$, it follows that $[\pi(\theta + he_1) - \pi(\theta)]\Pi(\theta) = 0$. Hence,

$$(A1) \quad [\pi(\theta+he_1) - \pi(\theta)]^t = h\pi(\theta+he_1)^t Q_1(\theta) F(\theta) + o(h) .$$

Again, since $\pi(\theta+he_1)$ is uniformly bounded in h , it is evident from (A1) that $\pi(\theta+he_1)$ is continuous at $h = 0$. Thus, (A1) implies that

$$[\pi(\theta+he_1) - \pi(\theta)]^t = h\pi(\theta)^t Q_1(\theta) F(\theta) + o(h)$$

$$\text{i.e., } \eta_1(\theta)^t = \pi(\theta)^t Q_1(\theta) F(\theta) ,$$

which is the required result.

NOTICE

This material may be protected by
Copyright Law (Title 17 U.S. Code)

REFERENCES

1. GLYNN, P.W. (1983). On the role of generalized semi-Markov processes in simulation output analysis. Proc. of the 1983 Winter Simulation Conference, 38-42.
2. GOLUB, G.H. and MEYER, C.D. (1986). Using the QR factorization and group inversion to compute, differentiate, and estimate the sensitivity of stationary probabilities for Markov chains. SIAM J. Alg. Disc. Meth. 7, 273-281.
3. HEYMAN, D.P. and SOBEL, M.J. (1982). Stochastic Models in Operations Research, Volume 1. McGraw-Hill, New York.
4. KEMENY, J.G. and SNELL, J.L. (1960). Finite Markov Chains. Van Nostrand, Princeton, New Jersey.
5. MEYER, C.D. and STEWART, G.W. (1986). Derivatives and perturbations of eigenvectors. Unpublished manuscript.
6. REVUZ, D. (1984). Markov Chains. North Holland, New York.
7. SCHWEITZER, P.J. (1960). Perturbation theory and finite Markov chains. J. Appl. Prob. 5, 401-413.