



Likelihood Ratio Gradient Estimation for Stochastic Recursions

Author(s): Peter W. Glynn and Pierre L'Ecuyer

Source: *Advances in Applied Probability*, Vol. 27, No. 4 (Dec., 1995), pp. 1019-1053

Published by: Applied Probability Trust

Stable URL: <http://www.jstor.org/stable/1427933>

Accessed: 21/07/2010 12:40

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=apt>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Applied Probability Trust is collaborating with JSTOR to digitize, preserve and extend access to *Advances in Applied Probability*.

<http://www.jstor.org>

LIKELIHOOD RATIO GRADIENT ESTIMATION FOR STOCHASTIC RECURSIONS

PETER W. GLYNN,* *Stanford University*

PIERRE L'ECUYER,** *Université de Montréal*

Abstract

In this paper, we develop mathematical machinery for verifying that a broad class of general state space Markov chains reacts smoothly to certain types of perturbations in the underlying transition structure. Our main result provides conditions under which the stationary probability measure of an ergodic Harris-recurrent Markov chain is differentiable in a certain strong sense. The approach is based on likelihood ratio 'change-of-measure' arguments, and leads directly to a 'likelihood ratio gradient estimator' that can be computed numerically.

HARRIS-RECURRENT MARKOV CHAIN; REGENERATION

AMS 1991 SUBJECT CLASSIFICATION: PRIMARY 60K05, 60J27, 65C05

SECONDARY 60J10, 60G30, 60G40

1. Introduction

In this paper, we will study the class of Markov chains that arise as solutions to stochastic recursions. Specifically, we shall consider sequences $X = (X_n : n \geq 0)$ that can be represented in the form

$$(1.1) \quad X_{n+1} = h(X_n, Z_{n+1}),$$

where the sequence $Z = (Z_n : n \geq 1)$ is assumed to be i.i.d. (independent and identically distributed). In the case that h is additive, Z is often termed the *innovations sequence*; we shall adopt this terminology for the more general case considered here.

The class of chains that take the form (1.1) is very rich from an applications viewpoint. In fact, Markov chains modeled in discrete time are often formulated as

Received 18 November 1993; revision received 18 May 1994.

* Postal address: Department of Operations Research, Stanford University, Stanford, CA 94305-4022, USA.

The research of this author was supported by the U.S. Army Research Office under Contract No. DAAL03-91-G-0101 and by the National Science Foundation under Contract No. DDM-9101580.

** Postal address: Département d'IRO, Université de Montréal, C.P. 6128, Succ. Centre Ville, Montréal, H3C 3J7, Canada.

This author's research was supported by NSERC-Canada grant No. OGP0110050 and FCAR-Québec Grant No. 93ER-1654.

solutions to stochastic recursions; see Meyn and Tweedie (1993) for examples. Our motivation to study solutions X to (1.1) stems largely from our interest in discrete-event simulation, which is perhaps the most widely used numerical tool for studying stochastic models of production systems, telecommunication networks, and computer systems. Such simulations are typically implemented computationally by recursively updating a certain internal state descriptor that includes information on both the 'physical state' and 'clocks' that govern the behavior of the process. These updates occur at state transition epochs, and take the form (1.1). Consequently, the analysis that we shall pursue in this paper is, at least in principle, applicable to the class of stochastic processes that correspond to discrete-event simulations. (This class can basically be identified with the class of generalized semi-Markov processes studied by König et al. (1967), and by others.)

Our primary goal here is to study the behavior of the Markov chain X under perturbations of the distribution that governs the innovations sequence Z . In particular, suppose that θ is a real-valued parameter under which the Z_n 's have common distribution K_θ (say). In this paper, we will use likelihood ratio 'change-of-measure' arguments to establish conditions under which:

- (i) the expectation of a random variable defined over a randomized time-horizon is differentiable in θ ;
- (ii) the stationary probability measure of X is differentiable in θ (in a sense to be made more precise in Section 4).

We shall also discuss and illustrate how our conditions can be verified by using stochastic Lyapunov functions. These methods permit one, for example, to establish differentiability of the stationary distribution by verifying certain conditions that can be expressed in terms of the distribution K_θ and the one-step transition function of X . These methods are illustrated via applications to the waiting time sequence of the single-server queue and a general class of non-linear storage models. For the single-server queue, our techniques are sharp enough to establish that essentially any functional of the steady-state distribution of the waiting time sequence, having finite mean, is differentiable (see Proposition 6).

The differentiability results that we obtain can be viewed as strengthening the continuity theory for stochastic models studied by, for example, Kennedy (1972), and Whitt (1974), (1980). Of course, it must be added that our theory typically demands more of the underlying perturbation of the process than is the case in existing 'continuity' literature (for example, we basically require some form of differentiability). Derivatives of stationary distributions have also been studied for finite Markov chains. Schweitzer (1968) gives 'close form' expressions (which can be computed by matrix operations) for such derivatives with respect to the transition probabilities of the chain. Golub and Meyer (1986) show how to differentiate the stationary distribution with respect to a parameter θ , assuming that the entries of the transition matrix are differentiable with respect to θ .

In addition to developing theory that can be used to establish model

'smoothness', our approach also provides expressions for the resulting derivatives that can be used to numerically calculate the derivatives via simulation. In particular, we develop a 'likelihood ratio gradient estimator' that can be used to numerically calculate the derivative of the steady-state expectation of a functional defined on a Harris-recurrent Markov chain. This estimator converges at rate $t^{-1/2}$ in the amount of computational effort t , and is the only known estimator having this property that works at the level of generality analyzed here. (For more details on the likelihood ratio gradient estimators in general, see Rubinstein and Shapiro (1993) and the references given there. The method of infinitesimal perturbation analysis is often more efficient but is limited to a much smaller class of models and performance measures than those analysed here; see Glasserman (1991) for details.)

We also consider enhancements to the basic estimator that can improve its numerical efficiency. In particular, we emphasize the fact that the likelihood ratio can be based either directly on the innovations sequence Z or on the chain X itself. We discuss the merits and disadvantages of the two approaches, and offer the results of some numerical computation performed on the waiting time sequence for comparison.

This paper is organized as follows. In Section 2, we consider a finite horizon model where the horizon is a randomized stopping time and provide sufficient conditions under which the expected performance measure is differentiable. We also construct likelihood ratio (LR) derivative estimators where the LR can be based on either the filtration associated with the innovations process or that associated with the Markov chain itself. In Section 3, we construct LRs for Harris-recurrent Markov chains, while in Section 4, we study the derivative of such likelihood ratios and find a LR representation for the derivative of the stationary distribution. From that, we construct LR derivative estimators for the steady-state average cost. The results developed in Sections 3 and 4 build upon those of Section 2. In Section 5, we examine the single-server queue and a storage theory example. The latter is Harris recurrent but has no state that is visited infinitely often with probability 1. For each of these examples, we illustrate how to use our Lyapunov methods to establish smoothness of their corresponding stationary distributions. We also give numerical results for the M/M/1 queue that compare the LR gradient estimator based on the innovations process with that based on the transition probabilities.

2. Likelihood ratios for finite-horizon stochastic recursions

In this section, we shall focus on finite-horizon simulations. We start by formulating the problem more precisely. In particular, we assume that the sequences X and Z take values in separable metric spaces S_1 and S_2 , respectively. Note that \mathbb{R}^d , when equipped with the Euclidean norm, is such a space; see Billingsley (1968). We require that h be a jointly measurable function from $S_1 \times S_2$ into S_1 . We define our basic probability space Ω as $\Omega = (0, 1) \times S_1 \times S_2 \times S_2 \times \dots$. A typical element $\omega \in \Omega$

then takes the form $\omega = (u, x_0, z_1, z_2, \dots)$ where $u \in (0, 1)$, $x_0 \in S_1$, and $z_i \in S_2$ for $i \geq 1$. Then, we can define $U(\omega) = u$, $X_0(\omega) = x_0$, $Z_n(\omega) = z_n$ for $n \geq 1$, and

$$X_{n+1}(\omega) = h(X_n(\omega), Z_{n+1}(\omega))$$

for $n \geq 0$. The random variable U is used to determine a randomized stopping time, as we will see later on. For each $\theta \in \Lambda = (a, b)$, assume that K_θ is a probability measure on S_2 that will act as the distribution of Z_n under θ . We then let $P_{\theta,x}$ be the distribution on Ω under which U has the uniform distribution over $(0, 1)$, $X_0 = x$, and $Z = (Z_n, n \geq 0)$ is an i.i.d. sequence having common distribution K_θ . Specifically,

$$(2.1) \quad P_{\theta,x}(du \times dx_0 \times dz_1 \times \dots \times dz_n) = du \cdot \delta_x(dx_0)K_\theta(dz_1) \cdots K_\theta(dz_n)$$

for $n \geq 1$. With the distributional assumption (2.1), the sequence X is then a (time-homogeneous) Markov chain under $P_{\theta,x}$, having the one-step transition function $P(\theta)$ defined by

$$P(\theta, x, dy) \triangleq P_{\theta,x}[X_1 \in dy]$$

for $x, y \in S_1$.

In a finite-horizon setting, it is natural to permit the initial distribution μ to depend on θ . More precisely, for each $\theta \in \Lambda$, let μ_θ be a probability measure on S_1 . We can then let P_θ be the probability measure on Ω defined by

$$P_\theta(d\omega) = \int_{S_1} \mu_\theta(dx)P_{\theta,x}(d\omega),$$

under which X_0 has distribution μ_θ , and the sequence Z is i.i.d. and independent of X_0 , with common distribution K_θ .

In the most general form of a finite-horizon simulation, the time horizon T is determined by a randomized stopping time. More precisely, for each $\theta \in \Lambda$, we assume that there exists a family of functions $(r_n(\theta): n \geq 0)$, such that for each $n \geq 0$, $r_n(\theta): S_1^{n+1} \rightarrow [0, 1]$ is measurable, and such that

$$P_\theta[T = n \mid X] = r_n(\theta, X_0, \dots, X_n).$$

Demanding this is equivalent to requiring that T be a randomized stopping time with respect to $(\sigma(X_0, \dots, X_n): n \geq 0)$. One can use the random variable U to determine the value of T as follows:

$$T = \inf \left\{ j: \sum_{n=1}^j r_n(\theta, X_0, \dots, X_n) \geq U \right\}.$$

We now turn to the construction of a likelihood ratio (LR) representation of P_θ in terms of P_{θ_0} . We will need to make the following assumption.

A1. There exists $\epsilon > 0$ such that for each $\theta \in \Lambda_\epsilon = (\theta_0 - \epsilon, \theta_0 + \epsilon)$,

- (i) K_θ is absolutely continuous with respect to K_{θ_0} ;
- (ii) μ_θ is absolutely continuous with respect to μ_{θ_0} ;
- (iii) $r_n(\theta, x_0, \dots, x_n) > 0$ implies $r_n(\theta_0, x_0, \dots, x_n) > 0$ for all $n \geq 0$ and $(x_0, \dots, x_n) \in S_1^{n+1}$.

Let $k(\theta, z)$ and $u(\theta, x)$ be the densities of K_θ and μ_θ with respect to K_{θ_0} and μ_{θ_0} , respectively, so that $K_\theta(dz) = k(\theta, z)K_{\theta_0}(dz)$ and $\mu_\theta(dx) = u(\theta, x)\mu_{\theta_0}(dx)$. Let $\rho(\theta)$ denote $r_T(\theta, X_0, \dots, X_T)/r_T(\theta_0, X_0, \dots, X_T)$ on $\{T < \infty\}$ and let $\mathcal{G}_n = \sigma(U, X_0, Z_1, \dots, Z_n)$ for each n . (We omit writing the dependence of $\rho(\theta)$ on X_0, \dots, X_T to simplify the notation.) It is now straightforward to establish the following result, where I denotes the indicator function.

Theorem 1. Let Y be a non-negative \mathcal{G}_T -measurable random variable and let A1 be in force. Then, there exists $\epsilon > 0$ such that

$$(2.2) \quad E_\theta[YI(T < \infty)] = E_{\theta_0}[Y\tilde{L}(\theta)I(T < \infty)]$$

for $\theta \in \Lambda_\epsilon$, where

$$(2.3) \quad \tilde{L}(\theta) = u(\theta, X_0)\rho(\theta)\prod_{i=1}^T k(\theta, Z_i).$$

It turns out that one can obtain an alternative LR representation by conditioning appropriately. Observe that A1 implies that $P(\theta, x, \cdot)$ is absolutely continuous with respect to $P(\theta_0, x, \cdot)$ and let $p(\theta, x, \cdot)$ be the density of $P(\theta, x, \cdot)$ with respect to $P(\theta_0, x, \cdot)$. Set $\mathcal{F}_n = \sigma(U, X_0, \dots, X_n)$. Starting from Theorem 1, it is straightforward to establish that if Y is a non-negative \mathcal{F}_T -measurable random variable, then

$$(2.4) \quad E_\theta[YI(T < \infty)] = E_{\theta_0}[YL(\theta)I(T < \infty)]$$

where

$$(2.5) \quad L(\theta) = u(\theta, X_0)\rho(\theta)\prod_{i=1}^T p(\theta, X_{i-1}, X_i).$$

Since (2.4) holds for any non-negative \mathcal{F}_T -measurable random variable and $L(\theta)$ is itself \mathcal{F}_T -measurable, it follows from the defining property of conditional expectation that

$$L(\theta) = E_{\theta_0}[\tilde{L}(\theta) \mid \mathcal{F}_T]$$

on the set $\{T < \infty\}$. Furthermore, it should be noted that the above analysis establishes that $p(\theta, X_{i-1}, X_i) = E[k(\theta, Z_i) \mid X_{i-1}, X_i]$.

Remark 1. One can use expressions (2.2) and (2.4) to estimate functionals of the measure P_θ , while simulating X under θ_0 . Since $L(\theta)$ is a conditional expectation of $\tilde{L}(\theta)$, it is evident that estimation based on (2.4) is statistically more efficient. Specifically, under mild additional regularity hypotheses, the principle of conditional Monte Carlo asserts that this latter estimator produces smaller confidence intervals for any given number of transitions of X simulated; see Fox and Glynn (1986) for a similar argument. Generally speaking, the more information Z contains relative to X , the greater the gain in statistical efficiency should be. However, (2.2) could be much easier to implement, because the densities $p(\theta, \cdot, \cdot)$ are often rather complicated functions in practice. Therefore, there is typically a trade-off between variance reduction on the one side and ease of implementation and computational cost on the other. As a result, whether (2.4) is preferable to (2.2) or vice versa depends on the problem considered.

Remark 2. The stopping time T is called *non-randomized* if each $r_n(\theta, X_1, \dots, X_n)$ is either 0 or 1. In that case, it follows from A1 (iii) that $\rho(\theta) = 1$ P_{θ_0} -almost surely and the likelihood ratios simplify accordingly.

We will now derive a LR representation for the derivative of P_θ . For that, we shall require that the family of distributions K_θ be suitably smooth in θ . To simplify our notation, let $P(\cdot) \triangleq P_{\theta_0}(\cdot)$ and $E(\cdot) \triangleq E_{\theta_0}(\cdot)$. A prime will denote the derivative with respect to θ .

We shall make the following assumptions.

A2.

- (i) There exists $\epsilon > 0$ such that for each $\theta \in \Lambda_\epsilon$, $P_\theta[T < \infty] = 1$.
- (ii) There exists $\epsilon > 0$ such that for each $x \in S_1$ and $z \in S_2$, $u(\cdot, x)$ and $k(\cdot, z)$ are continuously differentiable on Λ_ϵ .
- (iii) There exists a random variable $\rho'(\theta_0)$ such that

$$\lim_{h \rightarrow 0} E \left[\left| \frac{\rho(\theta_0 + h) - \rho(\theta_0)}{h} - \rho'(\theta_0) \right| \right] = 0.$$

- (iv) For each $p > 0$, there exists $\epsilon = \epsilon(p)$ such that

$$E \left[\sup_{\theta \in \Lambda_\epsilon} |u'(\theta, X_0)|^p \right] < \infty, \quad E \left[\sup_{\theta \in \Lambda_\epsilon} |k'(\theta, Z_1)|^p \right] < \infty, \quad \text{and} \quad \sup_{\theta \in \Lambda_\epsilon} E \left[\left| \frac{\rho(\theta) - 1}{\theta - \theta_0} \right|^p \right] < \infty.$$

Remark 3. Observe that A2 implies that for each $p > 0$, there exists $\epsilon = \epsilon(p)$ such that

$$(2.6) \quad E \left[\sup_{\theta \in \Lambda_\epsilon} |u(\theta, X_0)|^p \right] < \infty \quad \text{and} \quad E \left[\sup_{\theta \in \Lambda_\epsilon} |k(\theta, Z_1)|^p \right] < \infty.$$

Indeed, one can write $u(\theta, X_0) = 1 + (\theta - \theta_0)u'(\xi(\theta), X_0)$ for some $\xi(\theta) \in \Lambda_\epsilon$, and so

$$\begin{aligned} E\left[\sup_{\theta \in \Lambda_\epsilon} |u(\theta, X_0)|^p\right] &\leq E\left[\left(1 + \epsilon \sup_{\theta \in \Lambda_\epsilon} |u'(\theta, X_0)|\right)^p\right] \\ &\leq 2^p E\left[1 + \epsilon^p \sup_{\theta \in \Lambda_\epsilon} |u'(\theta, X_0)|^p\right] < \infty. \end{aligned}$$

The same argument applies to the second expression in (2.6). Assumption A2 also guarantees that

$$(2.7) \quad \lim_{\epsilon \rightarrow 0} E\left[\sup_{\theta \in \Lambda_\epsilon} |k(\theta, Z_1)|^p\right] = E\left[\lim_{\epsilon \rightarrow 0} \sup_{\theta \in \Lambda_\epsilon} |k(\theta, Z_1)|^p\right] = 1.$$

To see this, observe that the uniform integrability of the inside expression on the left permits one to exchange the limit and the expectation, and the inside limit is equal to 1 because $k(\cdot, z)$ is continuous and $k(\theta_0, z) = 1$.

Recall that the transition density $p(\theta, X_i, X_{i+1})$ was constructed using a measure-theoretic argument based on properties of conditional expectation. We will now establish the L^1 convergence of its difference quotient to the random variable $p'(\theta, X_i, X_{i+1}) \triangleq E[k'(\theta, Z_i) | X_{i-1}, X_i]$, again using basic properties of conditional expectation.

Proposition 1. Assume A1 (i) and A2 (ii). Then, there is an $\epsilon > 0$ such that for each $i \geq 1$ and $\theta \in \Lambda_\epsilon$,

$$\lim_{h \rightarrow 0} E\left[\left|\frac{p(\theta + h, X_{i-1}, X_i) - p(\theta, X_{i-1}, X_i)}{h} - p'(\theta, X_{i-1}, X_i)\right|\right] = 0.$$

Furthermore, for each $p > 0$, there exists $\epsilon = \epsilon(p) > 0$ such that

$$\sup_{\theta \in \Lambda_\epsilon} E[|p'(\theta, X_0, X_1)|^p] < \infty \quad \text{and} \quad \sup_{\theta \in \Lambda_\epsilon} E[|p(\theta, X_0, X_1)|^p] < \infty.$$

Proof. Recall that

$$p(\theta, X_{i-1}, X_i) = E[k(\theta, Z_i) | X_{i-1}, X_i].$$

Then, for h such that $\theta + h \in \Lambda_\epsilon$, the continuous differentiability of $k(\cdot, Z_j)$ and the mean value theorem imply the existence of a random variable $\xi \in \Lambda_\epsilon$ such that

$$\begin{aligned} \frac{p(\theta + h, X_{i-1}, X_i) - p(\theta, X_{i-1}, X_i)}{h} &\stackrel{\text{a.s.}}{=} E\left[\frac{k(\theta + h, Z_i) - k(\theta, Z_i)}{h} \mid X_{i-1}, X_i\right] \\ &= E[k'(\xi, Z_i) | X_{i-1}, X_i] \end{aligned}$$

and

$$E[|k'(\xi, Z_i)| | X_{i-1}, X_i] \leq E\left[\sup_{\theta \in \Lambda_\epsilon} |k'(\theta, Z_i)| \mid X_{i-1}, X_i\right].$$

From A2, the latter has finite p th moment for ϵ small enough. From the dominated convergence theorem for conditional expectations, it follows that

$$\begin{aligned} & \lim_{h \rightarrow 0} E \left[\left| \frac{p(\theta + h, X_{i-1}, X_i) - p(\theta, X_{i-1}, X_i)}{h} - p'(\theta, X_{i-1}, X_i) \right| \right] \\ &= E \left[\left| E \left[\frac{k(\theta + h, Z_i) - k(\theta, Z_i)}{h} - k'(\theta, Z_i) \mid X_{i-1}, X_i \right] \right| \right] \\ &= 0. \end{aligned}$$

The finiteness of the two suprema then follows via an application of the conditional Jensen's inequality.

Note that if the derivative of $p(\cdot, X_{i-1}, X_i)$ exists almost surely, then it must be equal to $p'(\cdot, X_{i-1}, X_i)$ a.s. Proposition 1 calculates the limit of the sample path difference quotient. To calculate the limit of the expectation (2.2) or (2.4), we will need to verify that we can pass the derivative inside the expectation operator. An important ingredient in establishing this interchange is to control the behavior of the likelihood ratios $\tilde{L}(\theta)$ and $L(\theta)$. To accomplish this, we will make the following assumption, to control the random variable T :

A3. There exists $z > 1$ such that $E[z^T] < \infty$.

We will also use the following lemma, which will permit us to analyze the difference quotients. This lemma will be used not only in the proof of the next theorem, but also later on, in the proof of Proposition 5, where we will need it to establish the uniform integrability of some difference quotients directly without appealing to the mean value theorem. We denote $\max(x, y)$ by $(x \vee y)$.

Lemma 1. Let z_1, \dots, z_n be non-negative real numbers. Then,

$$\left| \prod_{i=1}^n z_i - 1 \right| \leq \sum_{i=1}^n |z_i - 1| \cdot \prod_{j=1}^n (z_j \vee 1).$$

Proof. This follows by induction on n . The result is obvious for $n = 1$. Assuming that the result holds for $n = k$, note that

$$\begin{aligned} \left| \prod_{i=1}^{k+1} z_i - 1 \right| &\leq z_{k+1} \left| \prod_{i=1}^k z_i - 1 \right| + |z_{k+1} - 1| \\ &\leq (z_{k+1} \vee 1) \left| \prod_{i=1}^k z_i - 1 \right| + |z_{k+1} - 1| \prod_{j=1}^k (z_j \vee 1) \\ &\leq (z_{k+1} \vee 1) \left[\sum_{i=1}^k |z_i - 1| \cdot \prod_{j=1}^k (z_j \vee 1) \right] + |z_{k+1} - 1| \cdot \prod_{j=1}^{k+1} (z_j \vee 1) \\ &= \sum_{i=1}^{k+1} |z_i - 1| \cdot \prod_{j=1}^{k+1} (z_j \vee 1). \end{aligned}$$

We are now ready to state one of our main technical results.

Theorem 2. Assume A1–A3. Then, for each $p > 0$, there exists $\epsilon > 0$ such that

$$\sup_{\theta \in \tilde{\Lambda}_\epsilon} E \left[\left| \frac{L(\theta) - L(\theta_0)}{\theta - \theta_0} \right|^p \right] \leq \sup_{\theta \in \tilde{\Lambda}_\epsilon} E \left[\left| \frac{\tilde{L}(\theta) - \tilde{L}(\theta_0)}{\theta - \theta_0} \right|^p \right] < \infty.$$

Proof. The first inequality follows immediately from the conditional Jensen inequality and the fact that $L(\theta) \stackrel{\text{a.s.}}{=} E[\tilde{L}(\theta) \mid \mathcal{F}]$. So, it remains to prove that the second expression is finite. By Lyapunov's inequality (see, for example, p. 47 of Chung (1974)), it suffices to prove the result for $p > 1$. Noting that $\tilde{L}(\theta_0) = 1$, Lemma 1 yields

$$\begin{aligned} |\tilde{L}(\theta) - \tilde{L}(\theta_0)| &\leq \left\{ |u(\theta, X_0) - 1| + |\rho(\theta) - 1| + \sum_{i=1}^T |k(\theta, Z_i) - 1| \right\} \\ &\quad \cdot (u(\theta, X_0) \vee 1) \cdot (\rho(\theta) \vee 1) \cdot \prod_{i=1}^T (k(\theta, Z_i) \vee 1). \end{aligned}$$

Since we can assume $p > 1$, we may apply Hölder's inequality and then Minkowski's inequality to conclude that

$$\begin{aligned} &\sup_{\theta \in \tilde{\Lambda}_\epsilon} E \left[\left| \frac{\tilde{L}(\theta) - \tilde{L}(\theta_0)}{\theta - \theta_0} \right|^p \right] \\ &\leq \left\{ \sup_{\theta \in \tilde{\Lambda}_\epsilon} E^{1/4p} \left[\left| \frac{u(\theta, X_0) - 1}{\theta - \theta_0} \right|^{4p} \right] + \sup_{\theta \in \tilde{\Lambda}_\epsilon} E^{1/4p} \left[\left| \frac{\rho(\theta) - 1}{\theta - \theta_0} \right|^{4p} \right] \right. \\ (2.8) \quad &+ \left. \sup_{\theta \in \tilde{\Lambda}_\epsilon} E^{1/4p} \left[\sum_{i=1}^T \left| \frac{k(\theta, Z_i) - 1}{\theta - \theta_0} \right|^{4p} \right]^p \right\} \\ &\quad \cdot \left\{ \sup_{\theta \in \tilde{\Lambda}_\epsilon} E[(u(\theta, X_0) \vee 1)^{4p}] \sup_{\theta \in \tilde{\Lambda}_\epsilon} E[(\rho(\theta) \vee 1)^{4p}] \sup_{\theta \in \tilde{\Lambda}_\epsilon} E \left[\prod_{i=1}^T (k(\theta, Z_i) \vee 1)^{4p} \right] \right\}^{1/4} \\ &\triangleq [a_1 + a_2 + a_3]^p (b_1 b_2 b_3)^{1/4}. \end{aligned}$$

We will now show that each quantity in the latter expression is finite. To deal with a_1 , we note that $u(\theta_0, X_0) = 1$. Under A2, the mean value theorem yields the existence of $\xi(\theta) \in \Lambda_\epsilon$, for each $X_0 \in S_1$ and $\theta \in \Lambda_\epsilon$, such that $u(\theta, X_0) - 1 = (\theta - \theta_0)u'(\xi(\theta), X_0)$. Then,

$$(2.9) \quad a_1^{4p} = \sup_{\theta \in \tilde{\Lambda}_\epsilon} E[|u'(\xi(\theta), X_0)|]^{4p} \leq E \left[\sup_{\theta \in \tilde{\Lambda}_\epsilon} |u'(\theta, X_0)|^{4p} \right] < \infty$$

for ϵ small enough, by A2. The finiteness of a_2 is directly guaranteed by A2. For a_3 , we note that $k(\theta_0, Z_i) = 1$ for each i and recall that the Z_i are i.i.d. Under A2, the

mean value theorem can be applied to each of the T summands and for each θ , yielding the existence of $\xi_i(\theta) \in \Lambda_\epsilon$ such that

$$\begin{aligned}
 a_3^{4p} &= \sup_{\theta \in \Lambda_\epsilon} E \left[\left(\sum_{i=1}^T |k'(\xi_i(\theta), Z_i)| \right)^{4p} \right] \\
 &\leq E \left[T^{4p} \max_{1 \leq i \leq T} \sup_{\theta \in \Lambda_\epsilon} |k'(\theta, Z_i)|^{4p} \right] \\
 (2.10) \quad &\leq \left(E[T^{8p}] E \left[\max_{1 \leq i \leq T} \sup_{\theta \in \Lambda_\epsilon} |k'(\theta, Z_i)|^{8p} \right] \right)^{1/2} \\
 &\leq \left(E[T^{8p}] E \left[\sum_{i=1}^T \sup_{\theta \in \Lambda_\epsilon} |k'(\theta, Z_i)|^{8p} \right] \right)^{1/2} \\
 &= \left(E[T^{8p}] E[T] E \left[\sup_{\theta \in \Lambda_\epsilon} |k'(\theta, Z_i)|^{8p} \right] \right)^{1/2},
 \end{aligned}$$

which is again finite, for ϵ small enough, by A2 (iv) and A3. Wald's identity was applied to obtain the final equality.

For b_1 , observe that

$$(2.11) \quad b_1 = \sup_{\theta \in \Lambda_\epsilon} E[(u(\theta, X_0) \vee 1)^{4p}] \leq 1 + \sup_{\theta \in \Lambda_\epsilon} E[u(\theta, X_0)^{4p}] < \infty$$

for ϵ sufficiently small, from Remark 3. For b_2 , observe that

$$\begin{aligned}
 (\rho(\theta) \vee 1)^{4p} &\leq (1 + |\rho(\theta) - 1|)^{4p} \\
 &\leq 2^{4p} (1 + |\rho(\theta) - 1|^{4p}) \\
 &\leq 2^{4p} \left(1 + \epsilon^{4p} \left| \frac{\rho(\theta) - 1}{\theta - \theta_0} \right|^{4p} \right)
 \end{aligned}$$

and so

$$b_2 \leq 2^{4p} \left(1 + \epsilon^{4p} \sup_{\theta \in \Lambda_\epsilon} E \left[\left| \frac{\rho(\theta) - 1}{\theta - \theta_0} \right|^{4p} \right] \right) < \infty$$

by A2 (iv). For b_3 , we have the following inequalities:

$$\begin{aligned}
 (2.12) \quad b_3 &\leq \sup_{\theta \in \Lambda_\epsilon} \sum_{n=1}^{\infty} E \left[\left(\prod_{i=1}^n (k(\theta, Z_i) \vee 1)^{4p}; T \geq n \right) \right] \\
 &\leq \sup_{\theta \in \Lambda_\epsilon} \sum_{n=1}^{\infty} \left(E \left[\prod_{i=1}^n (k(\theta, Z_i)^{8p} \vee 1) \right] \cdot P[T \geq n] \right)^{1/2} \\
 &\leq \sup_{\theta \in \Lambda_\epsilon} \sum_{n=1}^{\infty} (1 + E[|k(\theta, Z_i)^{8p} - 1|])^{n/2} \cdot (P[T \geq n])^{1/2}.
 \end{aligned}$$

From (2.7), for each $\epsilon > 0$, there exists $\delta(\epsilon) > 0$ such that $\sup_{\theta \in \Lambda_\epsilon} E[|k(\theta, Z_i)^{\delta p} - 1|] < \delta(\epsilon)$ and $\lim_{\epsilon \rightarrow 0} \delta(\epsilon) = 0$, while from A3, there exists $\alpha < 1$ and n_0 such that $P[T \geq n] \leq \alpha^n$ for all $n \geq n_0$. Choose ϵ small enough such that $(1 + \delta(\epsilon))\alpha < 1$. Then,

$$b_3 \leq \sum_{n=1}^{n_0-1} (1 + \delta(\epsilon))^{n/2} + \sum_{n=n_0}^{\infty} [(1 + \delta(\epsilon))\alpha]^{n/2} < \infty.$$

This concludes the proof of the theorem.

The derivation of a LR representation for the derivative of P_θ basically reduces to bringing the derivative inside the expectation operator appearing in expressions (2.2) and (2.4). The random variables which then need to be differentiated with respect to θ are $\tilde{L}(\theta)$ and $L(\theta)$.

Proposition 2. Assume A1, A2, and A3. Then,

$$\lim_{h \rightarrow 0} E \left[\left| \frac{\tilde{L}(\theta_0 + h) - \tilde{L}(\theta_0)}{h} - \tilde{L}'(\theta_0) \right| \right] = 0$$

and

$$\lim_{h \rightarrow 0} E \left[\left| \frac{L(\theta_0 + h) - L(\theta_0)}{h} - L'(\theta_0) \right| \right] = 0,$$

where

$$(2.13) \quad \tilde{L}'(\theta_0) = \tilde{L}(\theta_0) \left[\frac{u'(\theta_0, X_0)}{u(\theta_0, X_0)} + \frac{\rho'(\theta_0)}{\rho(\theta_0)} + \sum_{i=1}^T \frac{k'(\theta_0, Z_i)}{k(\theta_0, Z_i)} \right],$$

$$(2.14) \quad L'(\theta_0) = L(\theta_0) \left[\frac{u'(\theta_0, X_0)}{u(\theta_0, X_0)} + \frac{\rho'(\theta_0)}{\rho(\theta_0)} + \sum_{i=1}^T \frac{p'(\theta_0, X_{i-1}, X_i)}{p(\theta_0, X_{i-1}, X_i)} \right] \\ = E[\tilde{L}'(\theta_0) | \mathcal{F}].$$

Proof. Assumption A2 permits us to define the random variable:

$$\tilde{D} = [u'(\theta_0, X_0)\rho(\theta_0) + u(\theta_0, X_0)\rho'(\theta_0)] \prod_{i=1}^T k(\theta_0, Z_i) \\ + u(\theta_0, X_0)\rho(\theta_0) \sum_{i=1}^T k'(\theta_0, Z_i) \prod_{j \neq i} k(\theta_0, Z_j).$$

That assumption also guarantees that we have enough differentiability present for the difference quotient $(\tilde{L}(\theta_0 + h) - \tilde{L}(\theta_0))/h$ to converge in probability to \tilde{D} . In Theorem 2, we have established the uniform integrability of that difference quotient. It then follows that

$$\lim_{h \rightarrow 0} E \left[\left| \frac{\tilde{L}(\theta_0 + h) - \tilde{L}(\theta_0)}{h} - \tilde{D} \right| \right] = 0.$$

To show that \tilde{D} can be written as in (2.13), we need to show that $u(\theta_0, X_0)$ is positive whenever $u'(\theta_0, X_0) \neq 0$ (so that we may divide through by $u(\theta_0, X_0)$), and similarly for k and ρ . Observe, however, that if $\{u(\theta_0, X_0) = 0, u'(\theta_0, X_0) \neq 0\}$ has positive probability, then it follows that $\{u(\theta_0, X_0) < 0\}$ has positive probability as well. This contradiction allows us to divide through by $u(\theta_0, X_0)$, and similarly for k . For ρ , recall that $\rho(\theta_0) = 1$ by definition. Therefore $\tilde{D} = \tilde{L}'(\theta_0)$. The expression (2.14) follows by taking conditional expectations in (2.13) with respect to \mathcal{F}_T and applying Proposition 1.

In Theorem 2 and Proposition 2, we proved that the difference quotients are well-behaved. Those results are the main tools required to establish our next theorem. That theorem provides general conditions under which finite-horizon performance measures are differentiable.

Theorem 3. Let Y be an \mathcal{F}_T -measurable random variable for which there exists $\delta > 0$ such that $E[|Y|^{1+\delta}] < \infty$. If A1–A3 hold, then

$$(2.15) \quad \frac{d}{d\theta} E_\theta[Y] \Big|_{\theta=\theta_0} = E_{\theta_0}[Y\tilde{L}'(\theta_0)].$$

Furthermore, $E_{\theta_0}[Y\tilde{L}'(\theta_0) | \mathcal{F}_T] = YL'(\theta_0)$.

Proof. Given Theorem 1 and Proposition 2, it is sufficient to establish that there exists $p > 1$ and $\epsilon > 0$ such that

$$(2.16) \quad \sup_{\theta_0+h \in \Lambda_\epsilon} E \left[\left| Y \cdot \frac{\tilde{L}(\theta_0+h) - \tilde{L}(\theta_0)}{h} \right|^p \right] < \infty,$$

since this will guarantee the appropriate level of uniform integrability necessary to justify (2.15) (the assertion involving $L'(\theta_0)$ follows by simple conditioning). But (2.16) is an immediate consequence of Theorem 2; just apply Hölder's inequality.

We note that Theorem 3 suggests two different simulation-based estimators for the derivative of a finite-horizon performance measure, one using replicates of $Y\tilde{L}'(\theta_0)$, and the other using replicates of $YL'(\theta_0)$. The principle of conditional Monte Carlo asserts that the estimator based on $YL'(\theta_0)$ has lower variance (but see Remark 1).

3. Likelihood ratios for Harris-recurrent stochastic recursions

We will now turn our attention to the construction of likelihood ratios and gradient estimators for infinite-horizon (steady-state) systems. In order to make the steady-state derivative estimation problem well-defined at $\theta_0 \in \Lambda$, we shall need to require that X possesses a (unique) stationary distribution for each $\theta \in \Lambda_\epsilon$ for some $\epsilon > 0$. Specifically, we shall require the following.

A4. There exists $\epsilon > 0$ such that X is a positive recurrent Harris chain under $P(\theta)$ for each $\theta \in \Lambda_\epsilon$.

It is well known (see Nummelin (1984)) that because S_1 is separable, we can assert that A4 implies that for each $\theta \in \Lambda_\epsilon$, there exists an integer $m(\theta) \geq 0$, a non-negative (measurable) function $\lambda(\theta)$, a (measurable) subset $A(\theta) \subseteq S_1$, and a probability measure $\varphi(\theta)$ on S_1 , such that:

- (i) $P_{\theta,x}[X_n \in A(\theta) \text{ infinitely often}] = 1$ for $x \in S_1$;
- (ii) $P_{\theta,x}[X_{m(\theta)} \in dy] \geq \lambda(\theta, x)\varphi(\theta, dy)$ for $x, y \in S_1$;
- (iii) $\inf\{\lambda(\theta, x): x \in A(\theta)\} > 0$.

In this paper, we shall strengthen these conditions so that they hold uniformly in θ . Specifically, we shall assume that:

A5. There exists $\epsilon > 0$, an integer $m \geq 0$, a (measurable) subset $A \subseteq S_1$, a probability φ on S_1 , and a non-negative (measurable) function λ for which

- (i) $P_{\theta,x}[X_n \in A \text{ infinitely often}] = 1$ for $x \in S_1, \theta \in \Lambda_\epsilon$;
- (ii) $P_{\theta,x}[X_m \in dy] \geq \lambda(x)\varphi(dy)$ for $x, y \in S_1, \theta \in \Lambda_\epsilon$;
- (iii) $\inf\{\lambda(x): x \in A\} \triangleq \lambda_* > 0$.

Remark 4. Allowing $m = 0$ in A4 and A5 is non-standard, but it will permit us to simplify our estimators nicely for systems which have a regenerative state. To be more precise, suppose that there is a specific state $x_* \in S_1$ that is hit in finite time with probability 1 from any other state; that is, $P_{\theta,x}[T < \infty] = 1$ for all $x \in S_1$ and $\theta \in \Lambda_\epsilon$, where $T \triangleq \inf\{n > 0: X_n = x_*\}$. Define $A = \{x_*\}$ and $\varphi(dy) = I[x_* \in dy]$. Then, A5 holds with $m = 0$, $\lambda(x) = I[x = x_*]$, and $\lambda_* = 1$. In fact, this degenerate case is the only case where A5 can hold for $m = 0$.

Remark 5. In most applications, A will be a compact set, and conditions A5 (ii)–(iii) will follow via a continuity argument. To verify A5 (i), let $E_{\theta,x}(\cdot)$ denote the expectation operator corresponding to $P_{\theta,x}(\cdot)$. Suppose that for each $\theta \in \Lambda_\epsilon$, there exists a non-negative (measurable) function $g(\theta, \cdot)$ and a positive constant $\epsilon(\theta)$ such that:

$$(3.1) \quad \begin{aligned} & \text{(i) } E_{\theta,x}[g(\theta, X_1)] \leq g(\theta, x) - \epsilon(\theta) \quad \text{for } x \notin A; \\ & \text{(ii) } \sup_{x \in A} E_{\theta,x}[g(\theta, X_1)] < \infty. \end{aligned}$$

Let $T(A) = \inf\{n \geq 1: X_n \in A\}$. Then, conditions (3.1) ensure that

$$(3.2) \quad \sup_{x \in A} E_{\theta,x}[T(A)] < \infty$$

(see, for example, Nummelin (1984) for details) and hence A5 (i) is satisfied. In fact, (3.2) and A5 together imply A4. (We note that A5 does not guarantee that X is

positive recurrent under $P(\theta)$; the additional hypothesis (3.2) yields this.) The function $g(\theta, \cdot)$ is called a ‘test function’ or ‘stochastic Lyapunov function’ in the literature.

Assumption A5 ensures that for each $\theta \in \Lambda_\epsilon$, the Markov chain X possesses a unique σ -finite stationary measure $\pi(\theta)$ having a regenerative representation. To see this, one uses the so-called ‘splitting method’ due to Athreya and Ney (1978) and Nummelin (1978). This technique consists of observing that A5 (ii) ensures the existence of a family of transition functions $Q(\theta)$ such that

$$(3.3) \quad P_{\theta,x}[X_m \in dy] = \lambda(x)\varphi(dy) + (1 - \lambda(x))Q(\theta, x, dy)$$

for $\theta \in \Lambda_\epsilon$, $x, y \in S_1$. Roughly speaking, (3.3) asserts that if the Markov chain X currently occupies state $x \in S$, then there is a probability $\lambda(x)$ that m time units later, the chain will be distributed according to φ . Because of A5 (i) and (iii), there will therefore be a random time τ at which the state of the chain is distributed independently of the state at time $\tau - m$. The stationary distribution $\pi(\theta)$ can then be represented in terms of a ratio formula expressed over the time interval $[0, \tau]$. Note that if $m = 0$ and $\lambda_* = 1$, then φ is concentrated on a single state x_* and τ is the first hitting time of x_* . For the remainder of this section, we will assume (unless otherwise specified) that $m \geq 1$. For the case where $m = 0$, the development goes through with many simplifications.

To develop likelihood ratio representations for $\pi(\theta)$, we need to make the above discussion more precise. To facilitate this task, we will modify slightly the interpretation of Ω adopted in the previous section. Our interpretation will provide the randomness necessary to ‘split’ $P(\theta)$ and construct the first regeneration time τ , as well as the succeeding regeneration times. Specifically, let $\tilde{\Omega} = S_1 \times S_2^\infty \times \{0, 1\}^\infty$. A typical element $\tilde{\omega} \in \tilde{\Omega}$ then takes the form $(x_0, z_1, z_2, \dots, i_1, i_2, \dots)$. The random variables $(Z_n : n \geq 1)$ and $(X_n : n \geq 0)$ are defined and distributed as before (so we can still denote their probability measure by $P_{\theta,x}$), and we further define the random variables $\eta_n(\tilde{\omega}) = i_n$ for $n \geq 1$.

Before completing the construction of probability measures on $\tilde{\Omega}$, we note that the splitting idea requires the ability to generate variates having distributions given either by φ or by $Q(\theta, x, \cdot)$. We wish to show that such variates can be constructed directly from the simulation of the Markov chain X itself and the 0–1 valued random variables η_n defined above. In other words, no additional randomization will be introduced to generate the appropriate variates. (The details of this type of construction have not previously been explored in the literature on simulation of Harris chains.) To accomplish this task (for $m \geq 1$), we fix $\beta \in (0, 1)$ and let

$$(3.4) \quad \begin{aligned} \varphi_\beta(x, dy) &= \beta\lambda(x)\varphi(dy), \\ Q_\beta(\theta, x, dy) &= (1 - \beta)\lambda(x)\varphi(dy) + (1 - \lambda(x))Q(\theta, x, dy) \\ &= P_{\theta,x}[X_m \in dy] - \beta\lambda(x)\varphi(dy). \end{aligned}$$

Note that introducing β effectively shrinks $\lambda(x)$. The main reason for introducing this shrinkage is to make sure that $Q_\beta(\theta, x, \cdot)$ is equivalent to $P^m(\theta, x, \cdot)$ (in the measure-theoretic sense, i.e. $Q_\beta(\theta, x, dy) = 0$ if and only if $P^m(\theta, x, dy) = 0$) where $P^m(\theta, x, \cdot) \triangleq P_{\theta,x}[X_m \in \cdot]$; this will be used later. Furthermore, $\varphi_\beta(x, \cdot)$ is absolutely continuous with respect to $P^m(\theta, x, \cdot)$. Hence, there exist densities $w_i(\theta, x, y)$, $i = 0, 1$, such that

$$(3.5) \quad \begin{aligned} Q_\beta(\theta, x, dy) &= w_0(\theta, x, y)P^m(\theta, x, dy) \\ \varphi_\beta(x, dy) &= w_1(\theta, x, y)P^m(\theta, x, dy) \end{aligned}$$

Also, these densities are non-negative and satisfy $w_0(\theta, x, y) + w_1(\theta, x, y) = 1$.

Let $S_0 = -m$ and $S_j = \inf \{n \geq S_{j-1} + m : X_n \in A\}$ for $j \geq 1$ be a sequence of hitting times of the set A defined so that at least m time units elapse between such visits to A . We can now define a probability $\tilde{P}_{\theta,x}$ on $\tilde{\Omega}$ as follows:

$$(3.6) \quad \begin{aligned} &\tilde{P}_{\theta,x}(dx_0 \times dz_1 \times \cdots \times dz_n \times \{i_1\} \times \cdots \times \{i_n\}) \\ &= \delta_x(dx_0)K_\theta(dz_1) \cdots K_\theta(dz_n) \prod_{j=1}^n w_{i_j}(\theta, X_{S_j(\omega)}, X_{S_j(\omega)+m}). \end{aligned}$$

Under $\tilde{P}_{\theta,x}$, X and Z are distributed as before. Let

$$P_\theta(d\tilde{\omega}) = \int_{S_1} \varphi(dx) \tilde{P}_{\theta,x}(d\tilde{\omega})$$

be the probability on $\tilde{\Omega}$ under which X_0 has distribution φ (this ensures that X ‘regenerates’ at time 0) and let $E_\theta(\cdot)$ be the corresponding expectation operator. Again, $P(\cdot)$ and $E(\cdot)$ will be a shorthand notation for $P_{\theta_0}(\cdot)$ and $E_{\theta_0}(\cdot)$. In any case, the η_j ’s have conditional distribution given by

$$P_\theta[\eta_j = 1 \mid X_0, Z] = w_1(\theta, X_{S_j}, X_{S_j+m}) = 1 - P_\theta[\eta_j = 0 \mid X_0, Z].$$

With this definition of P_θ , we find that on the event $\{S_l = n\}$,

$$\begin{aligned} &P_\theta[X_{n+m} \in dy \mid X_0, Z_1, \cdots, Z_n, \eta_1, \cdots, \eta_l] \\ &= P_\theta[X_{n+m} \in dy \mid X_n, \eta_l] \\ &= \frac{w_{\eta_l}(\theta, X_n, y)P^m(\theta, X_n, dy)}{\int_{S_1} w_{\eta_l}(\theta, X_n, z)P^m(\theta, X_n, dz)}. \end{aligned}$$

Taking advantage of (3.5) and (3.6), we find that on $\{S_l = n, \eta_l = 1\}$,

$$P_\theta[X_{n+m} \in dy \mid X_0, Z_1, \cdots, Z_n, \eta_1, \cdots, \eta_l] = \varphi(dy).$$

Hence, if we set $\gamma = \inf \{n \geq 1 : \eta_n = 1\}$, we may conclude that $\tau = S_\gamma + m$ is a

randomized stopping time at which the distribution of X is independent of its position at time $\tau - m$. We have

$$P_\theta[\tau = S_n + m \mid X_0, Z] = w_1(\theta, X_{S_n}, X_{S_n+m}) \prod_{j=1}^{n-1} w_0(\theta, X_{S_j}, X_{S_j+m}),$$

which is a function of only X_0, \dots, X_{S_n+m} . As a consequence, τ is the desired 'regeneration time' for X under P_θ , and it follows that under A4 and A5, there exists $\epsilon > 0$ such that

$$(3.7) \quad \pi_\theta(dx) = \frac{E_\theta \left[\sum_{j=0}^{\tau-1} I(X_j \in dx) \right]}{E_\theta[\tau]}$$

for $\theta \in \Lambda_\epsilon$.

Remark 6. The representation (3.7) for π_θ is valid for arbitrary positive recurrent Harris chains. In other words, the construction of τ followed above does not depend on the fact that X is the solution of a stochastic recursion or on the uniformity hypotheses implicit in A5.

We now turn to the construction of a likelihood ratio (LR) representation of π_θ in terms of π_{θ_0} . Assume that:

A6. A1 (i) is in force; i.e. there exists $\epsilon > 0$ such that K_θ is absolutely continuous with respect to K_{θ_0} for $\theta \in \Lambda_\epsilon$.

We note that under A6, $P^n(\theta, x, \cdot) \triangleq P_{\theta,x}[X_n \in \cdot]$ is absolutely continuous with respect to $P^n(\theta_0, x, \cdot)$ for $\theta \in \Lambda_\epsilon$, $x \in S_1$, $n \geq 1$. Let $p_n(\theta, x, y)$ be the density of $P^n(\theta, x, \cdot)$ with respect to $P^n(\theta_0, x, \cdot)$ for $n \geq 1$.

To proceed further, we observe that (3.4) and (3.5) imply that

$$\begin{aligned} \varphi_\beta(x, dy) &= w_1(\theta, x, y)P^m(\theta, x, dy) \\ &= w_1(\theta, x, y)p_m(\theta, x, y)P^m(\theta_0, x, dy) \end{aligned}$$

and

$$\varphi_\beta(x, dy) = w_1(\theta_0, x, y)P^m(\theta_0, x, dy),$$

and hence

$$w_1(\theta, x, y)p_m(\theta, x, y) = w_1(\theta_0, x, y)$$

$P^m(\theta_0, x, \cdot)$ -a.s. Furthermore, if we define $0/0$ to be zero, it is evident that

$$(3.8) \quad \frac{w_1(\theta, x, y)}{w_1(\theta_0, x, y)} = \frac{1}{p_m(\theta, x, y)}$$

$P^m(\theta_0, x, \cdot)$ -almost everywhere. We now take advantage of the fact that because $Q_\beta(\theta, x, \cdot)$ is equivalent to $P^m(\theta, x, \cdot)$, it follows that $Q_\beta(\theta, x, \cdot)$ is absolutely

continuous with respect to $Q_\beta(\theta_0, x, \cdot)$ (this is the prime reason why we took $\beta < 1$). We let $q(\theta, x, y)$ be the corresponding density, and note that (3.4) and (3.5) imply that

$$\begin{aligned} Q_\beta(\theta, x, dy) &= q(\theta, x, y)Q_\beta(\theta_0, x, dy) \\ &= q(\theta, x, y)w_0(\theta_0, x, y)P^m(\theta_0, x, dy). \end{aligned}$$

On the other hand,

$$\begin{aligned} Q_\beta(\theta, x, dy) &= w_0(\theta, x, y)P^m(\theta, x, dy) \\ &= w_0(\theta, x, y)p_m(\theta, x, y)P^m(\theta_0, x, dy) \end{aligned}$$

and thus

$$q(\theta, x, y)w_0(\theta_0, x, y) = w_0(\theta, x, y)p_m(\theta, x, y)$$

$P^m(\theta_0, x, \cdot)$ -almost everywhere. Because of the measure equivalency mentioned above, $p_m(\theta, x, \cdot)$ and $q(\theta, x, \cdot)$ have the same support. Hence, whenever $w_0(\theta_0, x, y) > 0$,

$$(3.9) \quad \frac{w_0(\theta, x, y)}{w_0(\theta_0, x, y)} = \frac{q(\theta, x, y)}{p_m(\theta, x, y)}$$

$P^m(\theta_0, x, \cdot)$ -almost everywhere.

We can now make the connection with the finite-horizon framework of the previous section: take $T = \tau$,

$$r_T(\theta, X_0, \dots, X_T) = w_1(\theta, X_{S_\gamma}, X_{S_\gamma+m}) \prod_{j=1}^{\gamma-1} w_0(\theta, X_{S_j}, X_{S_j+m}),$$

and $\mu_\theta \equiv \varphi$. Using the relation $\tau = S_\gamma + m$, one can construct τ and γ from either Ω or $\tilde{\Omega}$. With (3.8) and (3.9) at our disposal, and since φ does not depend on θ , we obtain that $u(\theta, X_0) \equiv 1$ and

$$\rho(\theta) = \left(\prod_{j=1}^{\gamma-1} \frac{q(\theta, X_{S_j}, X_{S_j+m})}{p_m(\theta, X_{S_j}, X_{S_j+m})} \right) \frac{1}{p_m(\theta, X_{\tau-m}, X_\tau)}.$$

The likelihood ratios $\tilde{L}(\theta)$ and $L(\theta)$ can then be written as

$$(3.10) \quad \tilde{L}(\theta) = \left(\prod_{i=1}^{\tau} k(\theta, Z_i) \right) \left(\prod_{j=1}^{\gamma-1} \frac{q(\theta, X_{S_j}, X_{S_j+m})}{p_m(\theta, X_{S_j}, X_{S_j+m})} \right) \frac{1}{p_m(\theta, X_{\tau-m}, X_\tau)}$$

and

$$(3.11) \quad L(\theta) = \left(\prod_{i=1}^{\tau} p(\theta, X_{i-1}, X_i) \right) \left(\prod_{j=1}^{\gamma-1} \frac{q(\theta, X_{S_j}, X_{S_j+m})}{p_m(\theta, X_{S_j}, X_{S_j+m})} \right) \frac{1}{p_m(\theta, X_{\tau-m}, X_\tau)}.$$

Under A4–A6, A1 holds and Theorem 1 applies, with $\mathcal{G}_\tau = \sigma(X_0, Z_1, \dots, Z_\tau, \eta_1, \dots, \eta_\gamma)$. Combining this with (3.7), we also obtain:

Corollary 1. Under A4–A6, there exists $\epsilon > 0$ such that

$$(3.12) \quad \pi_\theta(dx) = \frac{E_{\theta_0} \left[\sum_{j=0}^{\tau-1} I(X_j \in dx) \tilde{L}(\theta) \right]}{E_{\theta_0}[\tau \tilde{L}(\theta)]} = \frac{E_{\theta_0} \left[\sum_{j=0}^{\tau-1} I(X_j \in dx) L(\theta) \right]}{E_{\theta_0}[\tau L(\theta)]}$$

for $\theta \in \Lambda_\epsilon$.

Remark 7. We note that the last expression in (3.12) is defined only in terms of the chain X and the random variables η_1, \dots, η_τ . This representation for π_θ is in fact valid without any assumption that X be derived from a stochastic recursion, provided that A6 is replaced by an assumption that $P(\theta, x, \cdot)$ is absolutely continuous with respect to $P(\theta_0, x, \cdot)$ for each $\theta \in \Lambda_\epsilon, x \in S_1$.

Remark 8. We must acknowledge that implementing this construction in actual simulations is not easy in general, because w_0 and w_1 may be hard to evaluate. Moreover, when $m > 0$, we must memorize the sequence of states for the last m transitions in order to be able to do the acceptance/rejection test properly. In the degenerate case where $m = 0$, there is no need to shrink $\lambda(x)$: one can take $\beta = 1$. Then, one has $\rho(\theta) \equiv 1$ and the likelihood ratios simplify to

$$\begin{aligned} \tilde{L}(\theta) &= \prod_{i=1}^{\tau} k(\theta, Z_i); \\ L(\theta) &= \prod_{i=1}^{\tau} p(\theta, X_{i-1}, X_i). \end{aligned}$$

4. A likelihood ratio representation for the derivative of the stationary distribution

To obtain a LR representation for the derivative in Section 2, we required the family K_θ to be suitably smooth in θ . One of the major results of this section is that the imposition of appropriate regularity hypotheses on the densities $k(\theta, \cdot)$ in fact forces the densities $p_m(\cdot, x, y)$, and $q(\cdot, x, y)$ appearing on (3.10) and (3.11) to be well-behaved. A similar result for $p(\cdot, x, y)$ was already established in Proposition 1. We shall assume the following conditions.

A7.

- (i) There exists $\epsilon > 0$ such that for each $z \in S_2, k(\cdot, z)$ is continuously differentiable over Λ_ϵ ;
- (ii) For each $p > 0$, there exists $\epsilon = \epsilon(p)$ such that

$$E \left[\sup_{\theta \in \Lambda_\epsilon} |k'(\theta, Z_1)|^p \right] < \infty;$$

(iii) For each $r \in \mathbb{R}$,

$$\lim_{\epsilon \rightarrow 0} E \left[\sup_{\theta \in \Lambda_\epsilon} |k(\theta, Z_1)|^r \right] = 1.$$

Remark 9. Note that in contrast to (2.6), A7 (iii) is assumed to hold not only for positive values of r , but for negative values as well. This will be used in the proof of Proposition 5 (and only there). When the stopping time τ is non-randomized, as when the system is regenerative in the classical sense as indicated in Remarks 2, 4, and 8, then the result of Proposition 5 will hold trivially and A7 (iii) for $r < 0$ is no longer necessary. In the proof of Proposition 3, we will need A7 (iii) for $r > 0$, but that follows from A7 (ii) and the same argument as in Remark 3.

A glance at formulas (3.10) and (3.11) suggests that any LR derivative formula for the stationary distribution will require differentiability of $p_m(\cdot)$ and $q(\cdot)$ in the parameter θ . The next proposition establishes the required differentiability; the key idea in the proof is the recognition that the derivative of $p_m(\cdot)$ can be defined in terms of the conditional expectation of the derivative of $k(\cdot)$.

Proposition 3. Assume A4–A7. Then, for each i and $n \geq 1$, there is an $\epsilon > 0$ such that for each $\theta \in \Lambda_\epsilon$, there exist random variables $p'_n(\theta, X_i, X_{i+n})$ and $q'(\theta, X_{S_i}, X_{S_i+m})$ such that

$$\lim_{h \rightarrow 0} E \left[\left| \frac{p_n(\theta + h, X_i, X_{i+n}) - p_n(\theta, X_i, X_{i+n})}{h} - p'_n(\theta, X_i, X_{i+n}) \right| \right] = 0$$

and

$$\lim_{h \rightarrow 0} E \left[\left| \frac{q(\theta + h, X_{S_i}, X_{S_i+m}) - q(\theta, X_{S_i}, X_{S_i+m})}{h} - q'(\theta, X_{S_i}, X_{S_i+m}) \right| \right] = 0.$$

Proof. The proof for p_n is similar to that of Proposition 1. From the defining relation for a conditional expectation, one has

$$p_n(\theta, X_i, X_{i+n}) = E \left[\prod_{j=i+1}^{i+n} k(\theta, Z_j) \mid X_i, X_{i+n} \right].$$

Hence,

$$\begin{aligned} & \frac{p_n(\theta + h, X_i, X_{i+n}) - p_n(\theta, X_i, X_{i+n})}{h} \\ (4.1) \quad &= E \left[\frac{1}{h} \left(\prod_{j=i+1}^{i+n} k(\theta + h, Z_j) - \prod_{j=i+1}^{i+n} k(\theta, Z_j) \right) \mid X_i, X_{i+n} \right]. \end{aligned}$$

By A7 (i), it follows that for θ sufficiently close to θ_0 ,

$$(4.2) \quad \lim_{h \rightarrow 0} \frac{1}{h} \left(\prod_{j=i+1}^{i+n} k(\theta + h, Z_j) - \prod_{j=i+1}^{i+n} k(\theta, Z_j) \right) = \sum_{j=i+1}^{i+n} k'(\theta, Z_j) \prod_{\substack{l \neq j \\ l=i+1}}^{i+n} k(\theta, Z_l).$$

On the other hand, for h small, the continuous differentiability of $k(\cdot, Z_j)$ and the mean value theorem assert the existence of $\xi \in [\theta - h, \theta + h]$ such that

$$\frac{1}{h} \left(\prod_{j=i+1}^{i+n} k(\theta + h, Z_j) - \prod_{j=i+1}^{i+n} k(\theta, Z_j) \right) = \sum_{j=i+1}^{i+n} k'(\xi, Z_j) \prod_{\substack{l \neq j \\ l=i+1}}^{i+n} k(\xi, Z_l),$$

which is in turn dominated, if $\theta + h \in \Lambda_\epsilon$, by

$$(4.3) \quad \sum_{j=i+1}^{i+n} \sup_{\theta \in \Lambda_\epsilon} |k'(\theta, Z_j)| \prod_{\substack{l \neq j \\ l=i+1}}^{i+n} \sup_{\theta \in \Lambda_\epsilon} k(\theta, Z_l),$$

which has expectation

$$nE \left[\sup_{\theta \in \Lambda_\epsilon} |k'(\theta, Z_1)| \right] \left(E \left[\sup_{\theta \in \Lambda_\epsilon} k(\theta, Z_j) \right] \right)^{n-1}.$$

Assumptions A7 (ii-iii) ensure the finiteness of this expectation. Hence, the dominated convergence theorem for conditional expectations, applied to (4.1) and (4.2), yields

$$\begin{aligned} & \lim_{h \rightarrow 0} E \left[\left| \frac{p_n(\theta + h, X_i, X_{i+n}) - p_n(\theta, X_i, X_{i+n})}{h} - p'_n(\theta, X_i, X_{i+n}) \right| \right] \\ &= \lim_{h \rightarrow 0} E \left[\left| \frac{1}{h} \left(\prod_{j=i+1}^{i+n} k(\theta + h, Z_j) - \prod_{j=i+1}^{i+n} k(\theta, Z_j) \right) - \sum_{j=i+1}^{i+n} k'(\theta, Z_j) \prod_{\substack{l \neq j \\ l=i+1}}^{i+n} k(\theta, Z_l) \right| \right] \\ &= 0 \end{aligned}$$

where

$$p'_n(\theta, X_i, X_{i+n}) = E \left[\sum_{j=i+1}^{i+n} k'(\theta, Z_j) \prod_{\substack{l \neq j \\ l=i+1}}^{i+n} k(\theta, Z_l) \mid X_i, X_{i+n} \right].$$

To handle $q(\cdot)$, we observe that (3.4) and (3.5) imply that

$$\begin{aligned} & P^m(\theta + h, x, dy) - P^m(\theta, x, dy) \\ (4.4) \quad &= Q_\beta(\theta + h, x, dy) - Q_\beta(\theta, x, dy) \\ &= [q(\theta + h, x, y) - q(\theta, x, y)] Q_\beta(\theta_0, x, dy) \\ &= [q(\theta + h, x, y) - q(\theta, x, y)] w_0(\theta_0, x, y) P^m(\theta_0, x, dy). \end{aligned}$$

But

$$(4.5) \quad P^m(\theta + h, x, dy) - P^m(\theta, x, dy) = [p_m(\theta + h, x, y) - p_m(\theta, x, y)]P^m(\theta_0, x, dy)$$

is a (finite) signed measure. Relations (4.4) and (4.5), together with the $P^m(\theta_0, x, \cdot)$ -almost everywhere positivity of $w_0(\theta_0, \cdot)$, therefore yield the equality

$$\frac{q(\theta + h, x, h) - q(\theta, x, y)}{h} = \frac{p_m(\theta + h, x, y) - p_m(\theta, x, y)}{hw_0(\theta_0, x, y)}$$

Moreover, from (3.3)–(3.5),

$$w_0(\theta_0, x, y) = \frac{dQ_\beta(\theta_0, x, \cdot)}{dP_m(\theta_0, x, \cdot)} \geq 1 - \beta.$$

Hence, the second part of the proposition follows again from the dominated convergence theorem.

Remark 10. The proof of Proposition 3 in fact shows that

$$q'(\theta, X_{S_t}, X_{S_t+m}) \stackrel{\text{a.s.}}{=} \frac{p'_m(\theta, X_{S_t}, X_{S_t+m})}{w_0(\theta_0, X_{S_t}, X_{S_t+m})}$$

From that proposition, it also follows that the difference quotient of $\rho(\theta)$ converges in L^1 to $\rho'(\theta)$ at each $\theta \in \Lambda_e$ (although $\rho(\theta)$ is not necessarily continuously differentiable) under A4–A7.

It now remains to show that A2 (iv) and A3 hold, so that all the results of Section 2 apply. To accomplish this, we need to control the random variables τ and γ . In particular, we will show that, under suitable hypotheses on the chain, τ and γ have a geometrically dominated tail. This will require making a further assumption about the set A appearing in A5:

A8. The set A is a Kendall set for the Markov chain having transition function $P(\theta_0)$, i.e. if $T(A) = \inf \{n \geq 1 : X_n \in A\}$, then there exists $z > 1$ such that

$$\sup_{x \in A} E_{\theta_0, x}[z^{T(A)}] < \infty.$$

Remark 11. The verification that a set A is a Kendall set can be implemented via the use of appropriate Lyapunov function methods. In particular, suppose that there exists a non-negative function g defined on S_1 , $\epsilon > 0$, and $r < 1$ such that:

$$(4.6) \quad \begin{aligned} & \text{(i) } E_{\theta_0, x}[g(X_1)] \leq rg(x) - \epsilon \quad \text{for } x \notin A; \\ & \text{(ii) } \sup_{x \in A} E_{\theta_0, x}[g(X_1)] < \infty. \end{aligned}$$

Then, A is a Kendall set; see Nummelin (1984), pp. 90–91 and Chapter 16 of Meyn and Tweedie (1993). Note that such a Lyapunov function automatically implies the

existence of a Lyapunov function satisfying the conditions of Remark 5 at the parameter point $\theta = \theta_0$.

Proposition 4. Under assumptions A4–A7, $P[\gamma > k] \leq (1 - \beta\lambda_*)^k$ for $k \geq 0$. If, in addition, A8 is in force, then there exists $z > 1$ such that $E[z^\tau] < \infty$.

Proof. Relation (3.6) implies that

$$P[\gamma > k] = E[P[\gamma > k \mid X]] = E\left[\prod_{i=1}^k w_0(\theta_0, X_{S_i}, X_{S_i+m})\right].$$

By applying the strong Markov property first at time S_k and then at times $S_{k-1}, S_{k-2}, \dots, S_1$, we conclude that

$$P[\gamma > k] = E\left[\prod_{i=1}^k \bar{w}_0(\theta_0, X_{S_i})\right]$$

where

$$\begin{aligned} \bar{w}_0(\theta_0, X_{S_i}) &= E[w_0(\theta_0, X_{S_i}, X_{S_i+m}) \mid X_{S_i}] \\ &= \int_{S_i} w_0(\theta_0, X_{S_i}, y) P^m(\theta_0, X_{S_i}, dy) \\ &= 1 - \beta\lambda(X_{S_i}) \leq 1 - \beta\lambda_*, \end{aligned}$$

proving the tail bound for γ . To prove the existence of $z > 1$ such that $E[z^\tau] < \infty$, we note that $P[\gamma = k \mid X_j, \eta_i, j \leq S_k, i \leq k] \geq \beta\lambda_*$ on $\{\gamma > k - 1\}$. The result then follows from Lemma 5.6 of Nummelin (1984), p. 88. \square

From (3.5), (3.9), and Remark 10, it is easily verified that

$$\frac{q'(\theta, X_{S_i}, X_{S_i+m})}{q(\theta, X_{S_i}, X_{S_i+m})} \frac{p'_m(\theta, X_{S_i}, X_{S_i+m})}{p_m(\theta, X_{S_i}, X_{S_i+m})} = \frac{p'_m(\theta, X_{S_i}, X_{S_i+m}) w_1(\theta, X_{S_i}, X_{S_i+m})}{p_m(\theta, X_{S_i}, X_{S_i+m}) w_0(\theta, X_{S_i}, X_{S_i+m})}$$

P -a.s., and therefore

$$\begin{aligned} (4.7) \quad & \sum_{i=1}^{\gamma-1} \frac{q'(\theta, X_{S_i}, X_{S_i+m})}{q(\theta, X_{S_i}, X_{S_i+m})} - \sum_{i=1}^{\gamma} \frac{p'_m(\theta, X_{S_i}, X_{S_i+m})}{p_m(\theta, X_{S_i}, X_{S_i+m})} \\ &= \sum_{i=1}^{\gamma-1} \frac{p'_m(\theta, X_{S_i}, X_{S_i+m}) w_1(\theta, X_{S_i}, X_{S_i+m})}{p_m(\theta, X_{S_i}, X_{S_i+m}) w_0(\theta, X_{S_i}, X_{S_i+m})} - \frac{p'_m(\theta, X_{S_\gamma}, X_{S_\gamma+m})}{p_m(\theta, X_{S_\gamma}, X_{S_\gamma+m})}. \end{aligned}$$

If the derivative of $p_m(\cdot, X_{S_i}, X_{S_i+m})$ exists a.s., then it is a.s. equal to $p'_m(\cdot, X_{S_i}, X_{S_i+m})$ and the above expression is a.s. equal to $\rho'(\theta)/\rho(\theta)$.

Recall that the random variables $p_m(\theta, X_{S_i}, X_{S_i+m})$ and $q(\theta, X_{S_i}, X_{S_i+m})$ were constructed indirectly via conditioning arguments. Consequently, when viewing $p_m(\cdot, X_{S_i}, X_{S_i+m})$ and $q(\cdot, X_{S_i}, X_{S_i+m})$ as stochastic processes in θ , there is no *a priori* reason to expect almost sure differentiability or even continuity over $\theta \in \Lambda_\epsilon$. (This can be said even if the derivative exists a.s. at every θ , and this is for the same reason that a Poisson process $N = (N(t) : t \geq 0)$ is discontinuous even though at

every point t , $N(\cdot)$ is continuous a.s. It could happen, for instance, that the set of measure zero on which the derivative fails to exist does depend on θ in such a way that for each ω , there is a value of $\theta \in \Lambda_\epsilon$ where the derivative does not exist.) To proceed further, we will use Lemma 1, which will permit us to analyze the difference quotients directly without appealing to the mean value theorem.

Proposition 5. Assume A4–A8. Then, for each $p > 0$, there exists $\epsilon > 0$ such that

$$\sup_{\theta \in \Lambda_\epsilon} E \left[\left| \frac{\rho(\theta) - \rho(\theta_0)}{\theta - \theta_0} \right|^p \right] < \infty.$$

Proof. Again, by Lyapunov’s inequality, it suffices to prove the result for $p > 1$. To reduce the notational burden, let

$$\begin{aligned} K_i(\theta) &= k(\theta, Z_i), \\ K'_i(\theta) &= k'(\theta, Z_i), \\ P_i(\theta) &= p_m(\theta, X_{S_i}, X_{S_i+m}), \\ Q_i(\theta) &= q(\theta, X_{S_i}, X_{S_i+m}), \\ W_{ij}(\theta) &= w_j(\theta, X_{S_i}, X_{S_i+m}), \quad j = 0, 1. \end{aligned}$$

Noting that $\rho(\theta_0) = 1$, Lemma 1 yields

$$\begin{aligned} |\rho(\theta) - \rho(\theta_0)| &\leq \left\{ \sum_{i=1}^{\gamma} |P_i(\theta)^{-1} - 1| + \sum_{i=1}^{\gamma-1} |Q_i(\theta) - 1| \right\} \\ &\quad \cdot \prod_{i=1}^{\gamma} (P_i(\theta)^{-1} \vee 1) \cdot \prod_{i=1}^{\gamma-1} (Q_i(\theta) \vee 1) \\ &\leq \left\{ \sum_{i=1}^{\gamma} |P_i(\theta) - 1| + \sum_{i=1}^{\gamma-1} |Q_i(\theta) - 1| \right\} \\ &\quad \cdot \prod_{i=1}^{\gamma} (P_i(\theta)^{-2} \vee 1) \cdot \prod_{i=1}^{\gamma-1} (Q_i(\theta) \vee 1). \end{aligned}$$

Now, from Hölder’s and Minkowski’s inequalities,

$$\begin{aligned} &\sup_{\theta \in \Lambda_\epsilon} E \left[\left| \frac{\rho(\theta) - \rho(\theta_0)}{\theta - \theta_0} \right|^p \right] \\ &\leq \left\{ \sup_{\theta \in \Lambda_\epsilon} E^{1/3p} \left[\left(\sum_{i=1}^{\gamma} \left| \frac{P_i(\theta) - 1}{\theta - \theta_0} \right| \right)^{3p} \right] + \sup_{\theta \in \Lambda_\epsilon} E^{1/3p} \left[\left(\sum_{i=1}^{\gamma-1} \left| \frac{Q_i(\theta) - 1}{\theta - \theta_0} \right| \right)^{3p} \right] \right\}^p \\ &\quad \cdot \left(\sup_{\theta \in \Lambda_\epsilon} E \left[\prod_{i=1}^{\gamma} (P_i(\theta)^{-2} \vee 1)^{3p} \right] \sup_{\theta \in \Lambda_\epsilon} E \left[\prod_{i=1}^{\gamma-1} (Q_i(\theta) \vee 1)^{3p} \right] \right)^{1/3} \\ &\triangleq [a_4 + a_5]^p (b_4 b_5)^{1/3}. \end{aligned}$$

By the same argument as in the proof of Proposition 3, one has

$$(4.8) \quad E \left[\left| \frac{P_i(\theta) - P_i(\theta_0)}{\theta - \theta_0} \right|^{6p} \right] \leq m^{6p} E \left[\sup_{\theta \in \Lambda_*} |K'_1(\theta)|^{6p} \right] \cdot \left(E \left[\sup_{\theta \in \Lambda_*} (K_1(\theta))^{6p} \right] \right)^{m-1} \leq C$$

for some finite constant C . Therefore, using Proposition 4, we have

$$\begin{aligned} a_4^{3p} &\leq \sup_{\theta \in \Lambda_*} E \left[\sum_{i=1}^{\infty} I[\gamma \geq i] \left| \frac{P_i(\theta) - P_i(\theta_0)}{\theta - \theta_0} \right|^{3p} \right] \\ &\leq \sup_{\theta \in \Lambda_*} \sum_{i=1}^{\infty} (P[\gamma \geq i])^{1/2} \left(E \left[\left| \frac{P_i(\theta) - P_i(\theta_0)}{\theta - \theta_0} \right|^{6p} \right] \right)^{1/2} \\ &\leq \sum_{i=1}^{\infty} (1 - \beta \lambda_*)^{(i-1)/2} C^{1/2} \\ &< \infty. \end{aligned}$$

For a_5 , we note that the proof of Proposition 3 establishes that

$$(4.9) \quad \frac{Q_i(\theta) - Q_i(\theta_0)}{\theta - \theta_0} = \frac{P_i(\theta) - P_i(\theta_0)}{w_0(\theta_0, X_{S_i}, X_{S_i+m})} \frac{1}{\theta - \theta_0}.$$

Since $(1 - \beta)P(\theta, x, dy) \leq Q_\beta(\theta, x, dy) \leq P(\theta, x, dy)$, it follows that $1 - \beta \leq w_0(\theta_0, X_{S_i}, X_{S_i+m}) \leq 1$ a.s. and consequently

$$(4.10) \quad a_5 \leq (1 - \beta)^{-1} a_4,$$

proving the finiteness of a_5 . Turning now to b_4 , the Cauchy-Schwarz inequality yields

$$(4.11) \quad \begin{aligned} b_4 &= \sup_{\theta \in \Lambda_*} E \left[\prod_{i=1}^{\gamma} (P_i(\theta)^{-2p} \vee 1) \right] \\ &\leq \sup_{\theta \in \Lambda_*} \sum_{n=1}^{\infty} E \left[\prod_{i=1}^n (P_i(\theta)^{-4p} \vee 1) \right] \cdot P[\gamma \geq n]. \end{aligned}$$

Recall that

$$P_i(\theta) \stackrel{\text{a.s.}}{=} E \left[\prod_{j=S_i+1}^{S_i+m} K_j(\theta) \mid X_{S_i}, X_{S_i+m} \right],$$

and note that $(x^{-4p} \vee 1)$ is the maximum of two convex functions and hence convex. So, the conditional Jensen inequality yields

$$\begin{aligned} P_i(\theta)^{-4p} \vee 1 &\leq E \left[\left(\prod_{j=S_i+1}^{S_i+m} K_j(\theta)^{-4p} \right) \vee 1 \mid X_{S_i}, X_{S_i+m} \right] \\ &\leq E \left[\prod_{j=S_i+1}^{S_i+m} (K_j(\theta)^{-4p} \vee 1) \mid X_{S_i}, X_{S_i+m} \right]. \end{aligned}$$

Hence,

$$E \left[\prod_{i=1}^n (P_i(\theta)^{-4p} \vee 1) \right] \leq E \left[\prod_{i=1}^n E \left[\prod_{j=S_i+1}^{S_i+m} (K_j(\theta)^{-4p} \vee 1) \mid X_{S_i}, X_{S_i+m} \right] \right].$$

Applying the strong Markov property at time S_n , we obtain

$$\begin{aligned} & E \left[\prod_{i=1}^n E \left[\prod_{j=S_i+1}^{S_i+m} (K_j(\theta)^{-4p} \vee 1) \mid X_{S_i}, X_{S_i+m} \right] \mid X_0, \dots, X_{S_n} \right] \\ & \cong \prod_{i=1}^{n-1} E \left[\prod_{j=S_i+1}^{S_i+m} (K_j(\theta)^{-4p} \vee 1) \mid X_{S_i}, X_{S_i+m} \right] \cdot E \left[\prod_{j=S_n+1}^{S_n+m} (K_j(\theta)^{-4p} \vee 1) \mid X_{S_n} \right] \\ & = \prod_{i=1}^n E \left[\prod_{j=S_i+1}^{S_i+m} (K_j(\theta)^{-4p} \vee 1) \mid X_{S_i}, X_{S_i+m} \right] \cdot E \left[\prod_{j=1}^n (K_j(\theta)^{-4p} \vee 1) \right] \\ & = \prod_{i=1}^n E \left[\prod_{j=S_i+1}^{S_i+m} (K_j(\theta)^{-4p} \vee 1) \mid X_{S_i}, X_{S_i+m} \right] \cdot (E[(K_1(\theta)^{-4p} \vee 1)])^n. \end{aligned}$$

Successively conditioning at times $S_{n-1}, S_{n-2}, \dots, S_1$, we obtain

$$E \left[\prod_{i=1}^n E \left[\prod_{j=S_i+1}^{S_i+m} (K_j(\theta)^{-4p} \vee 1) \mid X_{S_i}, X_{S_i+m} \right] \right] \cong (E[(K_1(\theta)^{-4p} \vee 1)])^{nm}.$$

So,

$$b_4 \cong \sup_{\theta \in \Lambda_\epsilon} \sum_{n=1}^{\infty} (E[(K_1(\theta)^{-4p} \vee 1)])^{nm} P[\gamma \cong n].$$

Proposition 3 and Assumption A7 can then be exploited, as in (2.12), to obtain that $b_4 < \infty$. For b_5 , we argue as in (4.11) to obtain

$$b_5 \cong \sup_{\theta \in \Lambda_\epsilon} \sum_{n=0}^{\infty} E \left[\prod_{j=0}^{n-1} (Q_j(\theta)^{2p} \vee 1) \right] P[\gamma \cong n].$$

We now apply a conditioning argument similar to that used for b_4 :

$$\begin{aligned} & E \left[\prod_{j=0}^{n-1} (Q_j(\theta)^{2p} \vee 1) \mid X_0, \dots, X_{S_{n-1}} \right] \\ & \cong \prod_{j=0}^{n-2} (Q_j(\theta)^{2p} \vee 1) \cdot E[(Q_{n-1}(\theta)^{2p} \vee 1) \mid X_{S_{n-1}}]. \end{aligned}$$

But

$$\begin{aligned} & E[(Q_{n-1}(\theta)^{2p} \vee 1) \mid X_{S_{n-1}}] \\ & \cong E[(1 + |Q_{n-1}(\theta) - 1|)^{2p} \mid X_{S_{n-1}}] \\ & \cong E[(1 + |Q_{n-1}(\theta) - 1|)^{\lceil 2p \rceil} \mid X_{S_{n-1}}] \\ & = 1 + \sum_{i=1}^{\lceil 2p \rceil} \binom{\lceil 2p \rceil}{i} E[|Q_{n-1}(\theta) - 1|^i \mid X_{S_{n-1}}]. \end{aligned}$$

By (4.9), it is evident that

$$\begin{aligned} & E[|Q_{n-1}(\theta) - 1|^j \mid X_{S_{n-1}}] \\ &= E\left[\left|\frac{P_{n-1}(\theta) - P_{n-1}(\theta_0)}{w_0(\theta_0, X_{S_{n-1}}, X_{S_{n-1}+m})}\right|^j \mid X_{S_{n-1}}\right] \\ &\leq (1 - \beta)^{-j} E[|P_{n-1}(\theta) - P_{n-1}(\theta_0)|^j \mid X_{S_{n-1}}] \\ &\leq (1 - \beta)^{-j} \epsilon^j \sup_{\theta \in \Lambda_\epsilon} E\left[\left|\frac{P_{n-1}(\theta) - P_{n-1}(\theta_0)}{\theta - \theta_0}\right|^j \mid X_{S_{n-1}}\right]. \end{aligned}$$

Arguing as in (4.8), we may therefore conclude that for every $\delta > 0$, there exists $\epsilon > 0$ such that

$$\sup_{\theta \in \Lambda_\epsilon} E[(Q_{n-1}(\theta)^{2p} \vee 1) \mid X_{S_{n-1}}] \leq 1 + \delta.$$

By conditioning on $X_{S_{n-2}}, \dots, X_{S_1}$ and arguing similarly, we obtain the bound

$$\sup_{\theta \in \Lambda_\epsilon} E\left[\prod_{j=0}^{n-1} (Q_j(\theta)^{2p} \vee 1)\right] \leq (1 + \delta)^n,$$

so

$$b_5 \leq \sum_{n=0}^{\infty} (1 + \delta)^n P[\gamma \geq n].$$

By Proposition 4, b_5 is then finite for ϵ small enough, concluding the proof.

Proposition 5 completes the verification of A1–A3 for our Harris-recurrent setup. Theorem 2 shows in that case that the difference quotients are well-behaved. It is the main tool required to establish our next theorem. Theorem 4 shows that the stationary distributions π_θ are in fact differentiable in a very strong sense, namely in an extended version of the total variation norm. (For $f \equiv 1$, the notion of convergence presented will be precisely that of total variation.)

For a measure μ on S_1 and a S_1 -measurable function f , we adopt the notation

$$\mu f \triangleq \int_{S_1} f(y) \mu(dy).$$

We also put

$$Y(g) = \sum_{n=0}^{\tau-1} g(X_n).$$

Theorem 4. Let f be a non-negative S_1 -measurable function and assume that there exists $\delta > 0$ such that $\pi_{\theta_0} f^{1+\delta} < \infty$. If A4–A8 hold, then there exists a finite

signed measure π' such that

$$\limsup_{h \rightarrow 0} \sup_{|g| \leq f} \left| \frac{\pi_{\theta_0+h}g - \pi_{\theta_0}g}{h} - \pi'g \right| = 0$$

and

$$\pi'(\cdot) = \frac{E \left[\sum_{n=0}^{\tau-1} [I(X_n \in \cdot) - \pi_{\theta_0}(\cdot)] \tilde{L}'(\theta_0) \right]}{E[\tau]} = \frac{E[(Y(I(\cdot)) - \tau\pi_{\theta_0}(\cdot)) \tilde{L}'(\theta_0)]}{E[\tau]}.$$

Proof. Assume, to start, that g is non-negative. By (3.12), there exists $\epsilon > 0$ such that for $\theta \in \Lambda_\epsilon$,

$$(4.12) \quad \pi_{\theta}g = u(g; \theta)/l(\theta)$$

where

$$u(g; \theta) = E \left[\sum_{j=0}^{\tau-1} g(X_j) \tilde{L}(\theta) \right],$$

$$l(\theta) = E[\tau \tilde{L}(\theta)].$$

We observe that

$$\begin{aligned} & \sup_{0 \leq g \leq f} \left| \frac{u(g; \theta_0 + h) - u(g; \theta_0)}{h} - E[Y(g) \tilde{L}'(\theta_0)] \right| \\ & \leq E \left[Y(f) \left| \frac{\tilde{L}(\theta_0 + h) - \tilde{L}(\theta_0)}{h} - \tilde{L}'(\theta_0) \right| \right] \\ & \leq E \left[\max_{0 \leq n < \tau} f(X_n) \cdot \tau \cdot \left| \frac{\tilde{L}(\theta_0 + h) - \tilde{L}(\theta_0)}{h} - \tilde{L}'(\theta_0) \right| \right] \\ & \leq E^{1/p} \left[\max_{0 \leq n < \tau} f(X_n)^p \right] \cdot E^{1/q}[\tau^q] \cdot E^{1/r} \left[\left| \frac{\tilde{L}(\theta_0 + h) - \tilde{L}(\theta_0)}{h} - \tilde{L}'(\theta_0) \right|^r \right] \end{aligned}$$

by Hölder's inequality, where $p^{-1} + q^{-1} + r^{-1} = 1$ and $p, q, r > 0$. Choose $p = 1 + \delta$ and use the inequality

$$E \left[\max_{0 \leq n < \tau} f(X_n)^p \right] \leq E \left[\sum_{n=0}^{\tau-1} f(X_n)^p \right] = E[\tau] \pi_{\theta_0} f^{1+\delta}$$

to obtain the finiteness of the first factor. The second factor is finite by Proposition 4, and the proof of Theorem 3 establishes that the third goes to zero. Consequently,

$$(4.13) \quad \limsup_{h \rightarrow 0} \sup_{0 \leq g \leq f} \left| \frac{u(g; \theta_0 + h) - u(g; \theta_0)}{h} - E[Y(g) \tilde{L}'(\theta_0)] \right| = 0.$$

Setting $f \equiv 1$ and noting that $l(\theta) = u(1; \theta)$, we conclude that $l(\cdot)$ is also differentiable at $\theta = \theta_0$, so

$$(4.14) \quad \frac{1}{l(\theta_0 + h)} = \frac{1}{l(\theta_0)} - \frac{l'(\theta_0)}{l^2(\theta_0)h} + o(h).$$

Combining (4.12), (4.13), and (4.14) yields the conclusions of Theorem 4 uniformly in non-negative $g \leq f$. To handle general g , we split g into its positive and negative parts and apply the above argument to the separate pieces.

Remark 12. Theorem 4 requires the hypothesis that $\pi_{\theta_0} f^{1+\delta} < \infty$, where f is a given non-negative S_1 -measurable function. Once again, Lyapunov function methods can be used to verify this condition. In particular, assume that there exists a non-negative function g defined on S_1 and $\epsilon > 0$ such that:

$$(4.15) \quad \begin{aligned} (a) \quad & E_{\theta_0, x}[g(X_1)] \leq g(x) - \epsilon f(x)^{1+\delta} \quad \text{for } x \notin A \\ (b) \quad & \sup_{x \in A} E_{\theta_0, x}[g(X_1)] < \infty. \end{aligned}$$

Then, under A4–A8, Tweedie (1983) has established that finiteness of $\pi_{\theta_0} f^{1+\delta}$ necessarily follows.

Remark 13. Theorem 4 gives general conditions under which the stationary measure of a Markov chain (driven by a stochastic recursion) is differentiable in a strong (total variation-type) sense. Recent work of Vázquez-Abad and Kushner (1992) also addresses this question. The hypotheses given there are quite different and, in particular, are not given in terms of conditions that can be checked directly from the transition function of the chain (unlike, for example, the Lyapunov function criteria used above).

Remark 14. Much of the analysis in this paper is independent of whether the chain is driven by a stochastic recursion of the type described above. Our results could then be generalized. However, because the need for more general results from an applications viewpoint does not seem compelling, we shall not pursue this further.

Noting that $L'(\theta_0) = E[\tilde{L}'(\theta_0) | \mathcal{F}_\tau]$ and that $Y(g)$ and τ are both \mathcal{F}_τ -measurable, we obtain the following corollary to Theorem 4.

Corollary 2. Under the assumptions of Theorem 4, $\pi_{\theta} g$ is differentiable at $\theta = \theta_0$ for any g satisfying $|g| \leq f$, and

$$\begin{aligned} \frac{d}{d\theta} \pi_{\theta} g \Big|_{\theta=\theta_0} &= \frac{E[(Y(g) - (\pi_{\theta_0} g)\tau)\tilde{L}'(\theta_0)]}{E[\tau]} \\ &= \frac{E[(Y(g) - (\pi_{\theta_0} g)\tau)L'(\theta_0)]}{E[\tau]}. \end{aligned}$$

Remark 15. The representation of the derivative of $\pi_{\theta}g$ given in Corollary 2 can be used to construct simulation-based derivative estimators that converge at rate $t^{-1/2}$ in the amount of computational effort t ; see Glynn et al. (1991).

It turns out that because the random variables $Y(g)$ is an additive functional, an alternative representation for the derivative can be constructed. The representation takes advantage of the fact that

$$E[g(X_i)k'(\theta_0, X_j) \mid \mathcal{F}_i] = 0$$

for $i < j$. Consequently, roughly half the cross-product terms appearing in $Y(g)\tilde{L}'(\theta_0)$ (and $Y(g)L'(\theta_0)$) have vanishing expectations. The resulting estimators are called *triangular estimators*.

Corollary 3. Under the assumptions of Theorem 4, one has:

$$\begin{aligned}
 E[\tau] \cdot \frac{d}{d\theta} \pi_{\theta}g \Big|_{\theta=\theta_0} &= E \left[\sum_{j=1}^{\tau} k'(\theta_0, Z_j) \sum_{i=j}^{\tau-1} [g(X_i) - \pi_{\theta_0}g] \right] \\
 &\quad + E \left[\sum_{j=1}^{\gamma-1} q'(\theta_0, X_{S_j}, X_{S_j+m}) \sum_{i=S_j+1}^{\tau-1} [g(X_i) - \pi_{\theta_0}g] \right] \\
 (4.16) \quad &\quad - E \left[\sum_{j=1}^{\gamma} p'_m(\theta_0, X_{S_j}, X_{S_j+m}) \sum_{i=S_j+1}^{\tau-1} [g(X_i) - \pi_{\theta_0}g] \right]
 \end{aligned}$$

$$\begin{aligned}
 &= E \left[\sum_{j=1}^{\tau} p'(\theta_0, X_{j-1}, X_j) \sum_{i=j}^{\tau-1} [g(X_i) - \pi_{\theta_0}g] \right] \\
 &\quad + E \left[\sum_{j=1}^{\gamma-1} q'(\theta_0, X_{S_j}, X_{S_j+m}) \sum_{i=S_j+1}^{\tau-1} [g(X_i) - \pi_{\theta_0}g] \right] \\
 (4.17) \quad &\quad - E \left[\sum_{j=1}^{\gamma} p'_m(\theta_0, X_{S_j}, X_{S_j+m}) \sum_{i=S_j+1}^{\tau-1} [g(X_i) - \pi_{\theta_0}g] \right].
 \end{aligned}$$

5. Examples

The theory that we have developed in the previous sections is well suited to providing sufficient conditions under which steady-state performance measures are differentiable. In particular, suppose that X is the solution to a stochastic recursion for which the measures K_{θ} satisfy A6 and A7. Assume that A5 is satisfied and, for the set A appearing in A5, there exists a non-negative function g , and constants $r < 1$ and $\epsilon > 0$, such that for $\theta \in \Lambda_{\epsilon}$,

$$\begin{aligned}
 (5.1) \quad &(i) \ E_{\theta,x}[g(X_1)] \leq rg(x) - \epsilon \quad \text{for all } x \notin A \\
 &(ii) \ \sup_{x \in A} E_{\theta,x}[g(X_1)] < \infty.
 \end{aligned}$$

Then, Remarks 5, 11, and 12 guarantee that π_θ exists for θ in a neighborhood of θ_0 , and $\pi_\theta f$ is differentiable at θ_0 for each f satisfying the growth condition

$$(5.2) \quad |f(x)| \leq a + b(g(x))^p,$$

where $a, b \geq 0$ and $p < 1$. We will now illustrate these ideas with a couple of examples.

Example 1. Consider the sequence of waiting times in a single FIFO GI/GI/1 queue, $(X_n, n \geq 0)$, with $X_0 = 0$. That sequence follows the well-known recursion

$$X_{n+1} = [X_n + V_n - U_{n+1}]^+,$$

where V_n is the service time of customer n ($n \geq 0$) and U_{n+1} represents the interarrival time between customers n and $n + 1$. This is a special case of (1.1) with $Z_{n+1} = V_n - U_{n+1}$ and $h(x, z) = [x + z]^+$, or with $Z_{n+1} = (V_n, U_{n+1})$ and $h(x, v, u) = [x + v - u]^+$. We will adopt the latter representation, in which Z_{n+1} is a vector of two independent random variables. Let $B(\theta, \cdot)$ and $A(\theta, \cdot)$ be the service time and interarrival time distributions and let $C(\theta, \cdot)$ be the distribution function of $V_n - U_{n+1}$. Assume that over Λ_ϵ , the support of these distributions is independent of θ . Let $c(\theta, \cdot)$ denote the density of $C(\theta, \cdot)$ with respect to $C(\theta_0, \cdot)$, so that $C(\theta, dy) = c(\theta, y)C(\theta_0, dy)$, and similarly for a and b with A and B . This gives $k(\theta, v, u) = b(\theta, v)a(\theta, u)$ and $p(\theta, x, y) = c(\theta, y - x)$.

We assume that k satisfies A6 and A7, and will now examine how to verify A4, A5, and A8 for that example using stochastic Lyapunov functions as suggested in Remarks 5, 11, and 12. For that, we will find a function g that satisfies (5.1)–(5.2). One of our objectives here is to illustrate the use of such functions. There also exist other approaches for verifying A4–A8 for the GI/G/1 queue, based on the fact (for example) that the GI/G/1 queue can be modeled as a random walk (see for example Asmussen (1987) and L'Ecuyer and Glynn (1994)).

To verify A5 (ii)–(iii), take $m = 0$, $A = \{0\}$, and $\varphi(dy) = I[0 \in dy]$. Then, $\rho(\theta) \equiv 1$ and this system is regenerative in the classical sense, with regeneration occurring at each n for which $X_n = 0$.

Now, define $D = V_1 - U_0$. Assume that $E_{\theta_0}[D] < 0$ and that D has a finite and differentiable moment generating function in some neighborhood of zero; that is, there exists $z > 0$ such that $\sup_{\theta \in \Lambda_\epsilon} \varphi_D(\theta, z) < \infty$, where $\varphi_D(\theta, \beta) \triangleq E_\theta[\exp(\beta D)]$. Since $E_{\theta_0}[D] < 0$ and $\varphi_D(\theta, 0) = 1$, it follows that for $\beta > 0$ and $\epsilon > 0$ small enough, one has

$$\bar{r}(\beta, \epsilon) \triangleq \sup_{\theta \in \Lambda_\epsilon} \varphi_D(\theta, \beta) < 1.$$

Define

$$\tilde{\beta} = \inf \{ \beta > 0 : \varphi_D(\theta_0, \beta) \geq 1 \},$$

and let $0 < \beta < \tilde{\beta}$, $r = (\bar{r}(\beta, \epsilon) + 1)/2$, and $g(x) = K \exp(\beta x)I(x > 0)$ for some $K > 0$. Then, for $x > 0$, one has

$$\begin{aligned} E_{\theta,x}[g(X_1)] &= E_{\theta}[K \exp(\beta(x + D))I(D > -x)] \\ &= g(x) \cdot E_{\theta}[\exp(\beta(D))I(D > -x)] \\ &\cong g(x) \cdot \varphi_D(\theta, x) \\ &\leq \bar{r}(\beta, \epsilon)g(x) \leq (2r - 1)g(x) \cong rg(x) - K(1 - r), \end{aligned}$$

which verifies parts (i) of (5.1). For $x = 0$, condition (ii) in (5.1) is trivially verified. This completes the verification of A4, A5, A8, and Remark 12.

As a result, Theorem 4 applies even if f grows exponentially fast in x , provided that it grows no faster than $O(\exp(\beta x))$ for some $\beta < \tilde{\beta}$. This growth rate also typically turns out to be a tight bound, as indicated by the next proposition.

Proposition 6. Suppose that $\varphi_D(\theta_0, \tilde{\beta}) = 1$ for some $\tilde{\beta} > 0$. If $f(x) \sim K \exp(\beta x)$ as $x \rightarrow \infty$, for $K < \infty$, then $\pi_{\theta_0}f < \infty$ if and only if $\beta < \tilde{\beta}$.

Proof. We have just shown the ‘if’ part. Recall that from the Cramèr–Lundberg approximation (Asmussen (1987), p. 269), one has $P_{\theta_0}[X > x] \sim c \exp(-\tilde{\beta}x)$ as $x \rightarrow \infty$, where c is a positive constant. So if X denotes the steady-state waiting time, then

$$\begin{aligned} E[\exp(\beta X) - 1] &= E\left[\beta \int_0^X \exp(\beta x) dx\right] \\ &= \beta E\left[\int_0^{\infty} \exp(\beta x)I(X > x) dx\right] \\ &= \beta \int_0^{\infty} \exp(\beta x)P_{\theta_0}[X > x] dx < \infty \end{aligned}$$

if and only if $\beta < \tilde{\beta}$. So, for $\beta \geq \tilde{\beta}$, $\pi_{\theta_0}f = \infty$.

For a more specific illustration and numerical results comparing the use of $L(\theta)$ with that of $\tilde{L}(\theta)$, consider an M/M/1 queue with arrival rate λ_0 and mean service time θ . Assume that $0 < \theta_0 < 1/\lambda_0$. Details on the specific expressions for k , L , \tilde{L} , and so on, for that case, are given in Glynn and L’Ecuyer (1994), which is a slightly expanded version of this paper. The derivatives $u'(\theta_0)$ and $l'(\theta_0)$ can be estimated by simulating the system at $\theta = \theta_0$ and computing either $\tilde{L}'(\theta_0)Y$ and $\tilde{L}'(\theta_0)\tau$, or $L'(\theta_0)Y$ and $L'(\theta_0)\tau$, where $Y = \sum_{i=1}^{\tau-1} X_i$. The first pair of derivative estimators is based on the innovations process Z (we will denote them by IP) while the second pair is based on the transition probabilities of the Markov chain X (and will be denoted by TP).

To estimate $u'(\theta)$, we can also use the *triangular* LR estimators (4.16)–(4.17), where, roughly speaking, the derivative of each X_i is estimated ‘separately’ using a

likelihood ratio based only on the minimal information required to determine X_i . We will denote them by IP-T and TP-T, respectively.

Suppose that we simulate at $\theta = \theta_0$ for N regenerative cycles and let $\hat{u}_k, \hat{l}_k, \hat{u}'_k$ and \hat{l}'_k denote the (unbiased) estimators of $u(\theta_0), l(\theta_0), u'(\theta_0),$ and $l'(\theta_0),$ respectively, based on cycle i . Unbiased estimators of the latter quantities are obtained by averaging out:

$$\hat{u}(\theta_0) = \frac{1}{N} \sum_{k=1}^N \hat{u}_k$$

and similarly for $\hat{l}(\theta_0), \hat{u}'(\theta_0),$ and $\hat{l}'(\theta_0)$. Then, consistent estimators of $\alpha(\theta_0)$ and $\alpha'(\theta_0)$ are given by:

$$\hat{\alpha}(\theta_0) = \frac{\hat{u}(\theta_0)}{\hat{l}(\theta_0)};$$

$$\hat{\alpha}'(\theta_0) = \frac{\hat{u}'(\theta_0) - \hat{\alpha}(\theta_0)\hat{l}'(\theta_0)}{\hat{l}(\theta_0)}.$$

We performed numerical experiments for this system with $N = 1000, \lambda_0 = 1,$ and different values of θ_0 . Based on these 1000 cycles, we estimated $u(\theta_0), l(\theta_0), \alpha(\theta_0),$ and the derivatives $u'(\theta_0), l'(\theta_0),$ and $\alpha'(\theta_0),$ using IP, TP, IP-T, and TP-T. To estimate the variance of our estimators, we repeated this estimation process $R = 10\,000$ times (that is, $10^4 \times 10^3$ cycles). Table 1 gives the sample variances of those derivative estimators, for $\theta_0 = 0.1, 0.5, 0.8,$ and 0.9 . We also estimated the bias of $\hat{\alpha}'(\theta_0)$ and in all cases, the squared bias was negligible compared to the variance. For this simple case, for $\lambda_0 = 1,$ the exact values are $u(\theta) = \theta^2/(1 - \theta)^2$ and $l(\theta) = 1/(1 - \theta),$ from which one can also derive $u'(\theta), l'(\theta), \alpha(\theta),$ and $\alpha'(\theta)$.

TABLE 1
Experimental results for the $M/M/1$ queue (sample variances)

Derivative	LR approach	$\theta_0 = 0.1$	$\theta_0 = 0.5$	$\theta_0 = 0.8$	$\theta_0 = 0.9$
$u(\theta_0)$		4.59E-6	2.29E-2	1.32E1	6.68E2
$l(\theta_0)$		1.52E-4	5.99E-3	1.82E-1	1.70E0
$\alpha(\theta_0)$		3.36E-6	3.56E-3	2.38E-1	2.64E0
$u'(\theta_0)$	IP	0.0106	11.17	1.76E4	3.57E6
$u'(\theta_0)$	TP	0.0091	10.15	1.62E4	3.16E6
$u'(\theta_0)$	IP-T	0.0088	7.59	1.18E4	2.20E6
$u'(\theta_0)$	TP-T	0.0081	7.30	1.11E4	2.03E6
$l'(\theta_0)$	IP	0.3335	1.33	1.12E2	3.78E3
$l'(\theta_0)$	TP	0.1582	1.12	1.00E2	3.38E3
$\alpha'(\theta_0)$	IP	0.0073	1.50	2.44E2	8.86E3
$\alpha'(\theta_0)$	TP	0.0063	1.34	2.19E2	7.56E3
$\alpha'(\theta_0)$	IP-T	0.0060	0.88	1.30E2	4.02E3
$\alpha'(\theta_0)$	TP-T	0.0055	0.84	1.17E2	3.61E3

Generally speaking, we can see that the triangular estimators have significantly less variance than their ‘more standard’ counterparts (approximately half the variance, in some cases). It turns out that here, the TP estimators do not have much less variance than the IP ones, and this holds for small θ_0 as well as large θ_0 . There is one exception, however, namely the estimation of $l'(\theta_0)$ for small θ_0 .

Example 2. As a second example, we consider the same non-linear storage process as in Example 2 of Glynn (1992). In contrast to Example 1, this chain hits no point infinitely often. Specifically, let X_n represent the volume of water in a reservoir at time n , and $Z_{n+1} \geq 0$ denotes the inflow during period $n + 1$. The model is assumed to satisfy the equation

$$(5.3) \quad X_{n+1} = X_n + Z_{n+1} - aX_{n+1}^b,$$

where $a > 0$ and $b > 0$. This can be rewritten as $X_{n+1} = h(X_n + Z_{n+1})$, where h is the inverse function to $\tilde{h}(x) = x + ax^b$. Let F_θ be the probability distribution function of Z_1 under K_θ . The transition law of the Markov chain is then given by

$$\begin{aligned} P(\theta, x, [0, y]) &= P[h(X_n + Z_{n+1}) \leq y \mid X_n = x] \\ &= P[Z_{n+1} \leq \tilde{h}(y) - x] \\ &= F_\theta[y + ay^b - x], \end{aligned}$$

which can be positive only for $y \geq h(x)$. Then, for $y \geq h(x)$,

$$P(\theta, x, dy) = F'_\theta(y + ay^b - x)(1 + aby^{b-1}) dy.$$

Let us assume that K_θ is such that A6 and A7 are satisfied. For the other conditions, we will use a stochastic Lyapunov function as follows. Let $A = [0, K]$ where $K \geq 0$, $\beta > 0$, and

$$g(x) = \exp[\beta(x + ax^b)].$$

For $x \leq K$, one has $g(x) \leq \exp[\beta(K + aK^b)] < \infty$, and so conditions (ii) in (5.1) holds. For $x > K$, one has

$$\begin{aligned} E_{\theta,x}[g(X_1)] &= E_{\theta,x}[\exp(\beta(X_1 + aX_1^b))] \\ &= E_{\theta,x}[\exp(\beta(X_0 + Z_1))] \\ &= g(x) \exp(-\beta ax^b) \varphi_Z(\theta, \beta), \end{aligned}$$

where $\varphi_Z(\theta, \beta) = E_{\theta,x}[\exp(\beta Z_1)]$ is the moment generating function of Z_1 . We shall assume that

$$(5.4) \quad \sup_{\theta \in \Lambda_\epsilon} \varphi_Z(\theta, \beta) < \infty.$$

Then we can choose K such that $K^b > \sup_{\theta \in \Lambda_\epsilon} \ln(\varphi_Z(\theta, \beta)) / (\beta a)$ (because $a, b > 0$), and so

$$(5.5) \quad \bar{r} \triangleq \sup_{\theta \in \Lambda_\epsilon} \exp(-\beta a K^b) \varphi_Z(\theta, \beta) < 1.$$

Let $r = (1 + \bar{r})/2$. Then,

$$E_{\theta, x}[g(X_1)] \cong \bar{r}g(x) = rg(x) - (1 - r)g(x) \leq rg(x) - (1 - r).$$

This verifies (5.1). It follows that if A5 (ii)–(iii) also hold, then Theorem 4 applies for functions f of the form $f(x) = \exp[\beta'(x + ax^b)]$, for any $\beta' < \beta$, if $\varphi_Z(\cdot, \beta)$ is finite in a neighborhood of θ_0 . Note that if $\sup_{\theta \in \Lambda_\epsilon} \varphi_Z(\theta, \beta) < 1$, then we can take $K = 0$ and so $A = \{0\}$.

To verify A5 (ii)–(iii), we need to make further assumptions on the distribution K_θ . For example, if there exists a lower-bound measure $\tilde{\varphi}$ such that $P(\theta, x, dy) \cong \tilde{\varphi}(dy)$ for all $x \in A$ and $\theta \in \Lambda_\epsilon$, where $\lambda_* \triangleq \int_0^K \tilde{\varphi}(dy) > 0$, then these conditions are verified with $m = 1$ and $\varphi = \tilde{\varphi}/\lambda_*$. Otherwise, the conditions can still hold for larger m , but their actual verification gets more messy.

In Glynn and L'Ecuyer (1994), we verify these conditions and develop specific expressions for a special case of this example in which the distribution of inflows is exponential.

Note that in this model, the strict monotonicity of h guarantees that $\mathcal{F}_n = \mathcal{G}_n$ for each $n \geq 0$, and consequently $L'(\theta_0) = \tilde{L}'(\theta_0)$.

References

- ASMUSSEN, S. (1987) *Applied Probability and Queues*. Wiley, New York.
- ATHREYA, K. B. AND NEY, P. (1978) A new approach to the limit theory of recurrent Markov Chains. *Trans. Amer. Math. Soc.* **245**, 493–501.
- BILLINGSLEY, P. (1968) *Convergence of Probability Measures*. Wiley, New York.
- CHUNG, K. L. (1974) *A Course in Probability Theory*. Academic Press, New York.
- FOX, B. L. AND GLYNN, P. W. (1986). Discrete-time conversion for simulating semi-Markov processes. *Operat. Res. Lett.* **5**, 191–196.
- GLASSERMAN, P. (1991) *Gradient Estimation via Perturbation Analysis*. Kluwer Academic Publishers, Norwell, MA.
- GLYNN, P. W. (1986) Optimization of stochastic systems. *Proc. 1986 Winter Simulation Conf.*, 52–59.
- GLYNN, P. W. (1989) A GSMP formalism for discrete event systems. *Proc. IEEE* **77**, 14–23.
- GLYNN, P. W. (1992) Pathwise convexity and its relation to convergence of time-average derivatives. *Management Sci.* **38**, 1360–1366.
- GLYNN, P. W. AND L'ECUYER, P. (1994) Likelihood ratio gradient estimation for stochastic recursions. Research Report No. G-94, GERAD, University of Montreal.
- GLYNN, P. W., L'ECUYER, P., AND ADÈS, M. (1991) Gradient estimation for ratios. *Proc. 1991 Winter Simulation Conf.* IEEE Press, pp. 986–993.
- 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. Dis. Methods* **7**, 273–281.
- KENNEDY, D. P. (1972) The continuity of the single-server queue. *J. Appl. Prob.* **9**, 370–381.

- KÖNIG, D., MATHES, K. AND NAWROTZKI, K. (1967) *Verallgemeinerungen der Erlangschen und Engsetschen Formeln*. Akademie-Verlag, Berlin.
- L'ECUYER, P. AND GLYNN, P. W. (1994) Stochastic optimization by simulation: Convergence proofs for $GI/G/1$ queues in steady-state. *Management Sci.* **40**, 1562–1578.
- MEYN, S. P. AND TWEEDIE, R. L. (1993) *Markov Chains and Stochastic Stability*. Springer-Verlag, New York.
- NUMMELIN, E. (1978) A splitting technique for Harris recurrent Markov chains. *Z. Wahrscheinlichkeitsthe.* **43**, 309–318.
- NUMMELIN, E. (1984) *General Irreducible Markov Chains and Non-negative Operators*. Cambridge University Press, New York.
- RUBINSTEIN, R. Y. AND SHAPIRO, A. (1993) *Discrete Event Systems: Sensitivity Analysis and Stochastic Optimization by the Score Function Method*. Wiley, New York.
- SCHWEITZER, P. J. (1968) Perturbation theory and finite Markov chains. *J. Appl. Prob.* **5**, 401–413.
- TWEEDIE, R. L. (1983) The existence of moments for stationary Markov chains. *J. Appl. Prob.* **20**, 191–196.
- VÁZQUEZ-ABAD, F. AND KUSHNER, H. (1992) Estimation of the derivative of a stationary measure with respect to a control parameter. *J. Appl. Prob.* **29**, 343–352.
- WHITT, W. (1974) The continuity of queues. *Adv. Appl. Prob.* **6**, 175–183.
- WHITT, W. (1980) Continuity of generalized semi-Markov processes. *Math. Operat. Res.* **5**, 494–501.