

On the Existence and Estimation of Performance Measure Derivatives for Stochastic Recursions *

Peter W. Glynn¹ and Pierre L'Ecuyer²

¹ Department of Operations Research, Stanford University
Stanford, CA 94305-4022

² Département d'IRO, Université de Montréal
C.P. 6128, Succ. A, Montréal, H3C 3J7, CANADA

1 Introduction

In this paper, we study a class of Markov chains $X = (X_n : n \geq 0)$ that arise as solutions to stochastic recursions of the form

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

where $Z = (Z_n : n \geq 1)$ is an i.i.d. sequence, called the *innovations* sequence. Our motivation stems from the fact that discrete-event simulation are typically formulated and implemented as stochastic recursions of the form (1). Our aim is to study the behavior of X under perturbations of the distribution that governs Z . Suppose that θ is a real-valued parameter under which the Z_n 's have common distribution K_θ (say). We will provide conditions under which

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

Our conditions are based on stochastic Lyapunov functions and can be expressed in terms of K_θ and the one-step transition function of X . In addition to giving conditions for model "smoothness", our approach also provides derivative estimators which can be computed via simulation. In particular, we develop a likelihood ratio (LR) derivative estimator for the derivative of the steady-state expectation of a functional defined on a Harris recurrent Markov chain.

In §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 LR derivative estimators where the LR can be based on either the filtration associated with the "innovation process" or that associated with the

* The research of the first author was supported by the U. S. Army Research Office under Contract No. DAAL03-91-G-0101 and by the NSF under Contract No. DDM-9101580. The second author's research was supported by NSERC-Canada grant No. OGP0110050 and FCAR-Québec grant No. 93ER1654.

Markov chain itself. In §3, we construct LR's for Harris recurrent Markov chains, while in §4, we study the derivative of such LR's and construct LR derivative estimators for the steady-state performance measures. The main results are stated here without proof, due to space limitations. A complete version of the paper, including all the proofs and some examples, will appear elsewhere.

2 LR's for Finite-Horizon Stochastic Recursions

We consider the recursion (1), where for each n , X_n and Z_n take values in separable metric spaces S_1 and S_2 , respectively, and $f : S_1 \times S_2 \rightarrow S_1$ is jointly measurable. For each $\theta \in \Lambda$, we assume that there exists a family of measurable functions $\{r_n(\theta) : S_1^{n+1} \rightarrow [0, 1], n \geq 0\}$ and a r.v. T such that $P_\theta[T = n | X] = r_n(\theta, X_0, \dots, X_n)$. The random stopping time T is the time horizon, and is determined as $T = \inf\{j : \sum_{n=1}^j r_n(\theta, X_0, \dots, X_n) \geq U\}$, where U is a uniform r.v. over $(0, 1)$. For each $\theta \in \Lambda = (a, b)$, K_θ and μ_θ are probability measures on S_2 and S_1 , that act as the respective distribution of Z_n and X_0 under θ . Let $P_{\theta,x}$ denote the probability law of the process, conditional on $X_0 = x$, and $E_{\theta,x}$ be the corresponding expectation. The sequence X is a time-homogeneous Markov chain having the one-step transition function $P(\theta, x, dy) = P_{\theta,x}[X_1 \in dy]$ for $x, y \in S_1$. Define $P_\theta(d\omega) = \int_{S_1} \mu_\theta(dx) P_{\theta,x}(d\omega)$.

Assumption 1. *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 all $(x_0, \dots, x_n) \in S_1^{n+1}$.

Let $k(\theta, z)$ and $u(\theta, x)$ be densities such 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 $r_T(\theta, X_0, \dots, X_T)/r_T(\theta_0, X_0, \dots, X_T)$ and let $\mathcal{G}_n = \sigma(U, X_0, Z_1, \dots, Z_n)$ for each n .

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

$$E_\theta[Y I(T < \infty)] = E_{\theta_0}[Y \tilde{L}(\theta) I(T < \infty)] \tag{2}$$

for $\theta \in \Lambda_\epsilon$, where I denotes the indicator function and

$$\tilde{L}(\theta) = u(\theta, X_0)\rho(\theta) \prod_{i=1}^T k(\theta, Z_i). \tag{3}$$

Let $p(\theta, x, y)$ be the density of $P(\theta, x, \cdot)$ with respect to $P(\theta_0, x, \cdot)$. Set $\mathcal{F}_n = \sigma(U, X_0, \dots, X_n)$. One obtains an alternative LR representation by conditioning: if Y is a non-negative \mathcal{F}_T -measurable r.v. then

$$E_\theta[Y I(T < \infty)] = E_{\theta_0}[Y L(\theta) I(T < \infty)] \tag{4}$$

where

$$L(\theta) = u(\theta, X_0)\rho(\theta) \prod_{i=1}^T p(\theta, X_{i-1}, X_i). \quad (5)$$

Applying (4) to $L(\theta)$ yields $L(\theta) = E_{\theta_0}[\tilde{L}(\theta) \mid \mathcal{F}_T]$ on the set $\{T < \infty\}$. One can use (2) or (4) to estimate functionals of the measure P_θ , while simulating X under θ_0 . Since $L(\theta)$ is a conditional expectation of $\tilde{L}(\theta)$, the estimation based on (4) is statistically more efficient (or at least as efficient) in terms of variance; 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) 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 costs on the other.

The stopping time T is *non-randomized* if each $r_n(\theta, X_1, \dots, X_n)$ is either 0 or 1. In that case, $\rho(\theta) = 1$ P_{θ_0} -a.s. and the LRs simplify accordingly.

To derive a LR representation for the derivative of P_θ , we shall require that K_θ be suitably smooth in θ . To simplify the notation, we denote P_{θ_0} and E_{θ_0} by P and E , respectively. A “prime” denotes the derivative with respect to θ .

Assumption 2. (i) *There exists $\epsilon > 0$ such that for all $\theta \in \Lambda_\epsilon$, $P_\theta[T < \infty] = 1$;*

(ii) *There exists $\epsilon > 0$ such that for each $x \in S_1$ and $z \in S_2$, $u(\cdot, x_0)$ and $k(\cdot, z)$ are continuously differentiable over Λ_ϵ . Also, $\rho'(\theta_0) = \lim_{h \rightarrow 0} (\rho(\theta_0 + h) - \rho(\theta_0))/h$ exists P -a.s.;*

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

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

To differentiate (2) or (4), and pass the derivative inside the expectation operator, we need to control the behavior of $\tilde{L}(\theta)$ and $L(\theta)$. For that, we will make the following assumption, to control the r.v. T .

Assumption 3. *There exists $z > 1$ such that $E[z^T] < \infty$.*

Theorem 2. *Let Y be a \mathcal{F}_T -measurable r.v. and assume that there exists $\delta > 0$ such that $E[|Y|^{1+\delta}] < \infty$. If Assumptions 1–3 hold, then $E_\theta[Y]$ is differentiable at θ_0 and the derivative is given by $E[Y\tilde{L}'(\theta_0)] = E[YL'(\theta_0)]$, where*

$$\begin{aligned} \tilde{L}'(\theta) &= \tilde{L}(\theta) \left[\frac{u'(\theta, X_0)}{u(\theta, X_0)} + \frac{\rho'(\theta)}{\rho(\theta)} + \sum_{i=1}^T \frac{k'(\theta, Z_i)}{k(\theta, Z_i)} \right]; \\ L'(\theta) &= L(\theta) \left[\frac{u'(\theta, X_0)}{u(\theta, X_0)} + \frac{\rho'(\theta)}{\rho(\theta)} + \sum_{i=1}^T \frac{p'(\theta, X_{i-1}, X_i)}{p(\theta, X_{i-1}, X_i)} \right]. \end{aligned}$$

3 LRs for Harris-Recurrent Stochastic Recursions

We now turn our attention to the construction of LRs and derivative estimators for infinite-horizon (steady-state) systems.

Assumption 4. *There exists $\epsilon > 0$, an integer $m \geq 0$, a (measurable) subset $A \subseteq S_1$, a probability φ on S_1 , and a measurable function $\lambda \geq 0$ for which*

- i) For each $\theta \in \Lambda_\epsilon$, there exists a measurable function $g(\theta, \cdot) \geq 0$ and a constant $\epsilon(\theta) > 0$ such that: (a) $E_{\theta,x}[g(\theta, X_1)] \leq g(\theta, x) - \epsilon(\theta)$ for $x \notin A$; and (b) $\sup_{x \in A} E_{\theta,x}[g(\theta, X_1)] < \infty$.*
- ii) $P_{\theta,x}[X_m \in dy] \geq \lambda(x) \varphi(dy)$ for $x, y \in S_1$, $\theta \in \Lambda_\epsilon$;*
- iii) $\inf \{\lambda(x) : x \in A\} \triangleq \lambda_* > 0$.*

In most applications, A is a compact set, and (ii–iii) follow via a continuity argument. Let $T(A) = \inf \{n \geq 1 : X_n \in A\}$. Then, (i) ensures that $\sup_{x \in A} E_{\theta,x}[T(A)] < \infty$ and $P[X_n \in A \text{ i.o.}] = 1$ (see Nummelin 1984). The function $g(\theta, \cdot)$ is called a “test function” or “stochastic Lyapunov function”.

Remark. Allowing $m = 0$ in Assumption 4 is non-standard, but permits us to simplify the 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 a.s. in finite time from any other state. Define $A = \{x_*\}$ and $\varphi(dy) = I[x_* \in dy]$. Then, Assumptions 4(ii–iii) hold with $m = 0$, $\lambda(x) = I[x = x_*]$, and $\lambda_* = 1$. In fact, this degenerate case is the only case where Assumption 4 can hold for $m = 0$.

We can now use the so-called “splitting method” (Athreya and Ney 1978, and Nummelin 1978), to show that for each $\theta \in \Lambda_\epsilon$, X possesses a unique σ -finite stationary measure $\pi(\theta)$ having a regenerative representation. Assumption 4 (ii) ensures the existence of a family of transition functions $Q(\theta)$ such that

$$P_{\theta,x}[X_m \in dy] = \lambda(x) \varphi(dy) + (1 - \lambda(x)) Q(\theta, x, dy) \tag{6}$$

for $\theta \in \Lambda_\epsilon$, $x, y \in S_1$. Roughly speaking, if X currently occupies state $x \in S$, then m time units later, with probability $\lambda(x)$, the chain will be distributed according to φ . There is then a random time τ for which X_τ is distributed independently of $X_{\tau-m}$, and the stationary distribution $\pi(\theta)$ can be represented in terms of a ratio formula expressed over the time interval $[0, \tau]$. 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 $m \geq 1$, we will shrink $\lambda(x)$ to make sure that the measures $Q_\beta(\theta, x, \cdot)$ and $P^m(\theta, x, \cdot) \equiv P_{\theta,x}[X_m \in \cdot]$ are equivalent. Fix $\beta \in (0, 1)$ and let

$$\begin{aligned} \varphi_\beta(x, dy) &= \beta \lambda(x) \varphi(dy), \\ Q_\beta(\theta, x, dy) &= (1 - \beta) \lambda(x) \varphi(dy) + (1 - \lambda(x)) Q(\theta, x, dy) \\ &= P_{\theta,x}[X_m \in dy] - \beta \lambda(x) \varphi(dy). \end{aligned} \tag{7}$$

Then, there exist densities $w_i(\theta, x, y)$, $i = 0, 1$, such that

$$\begin{aligned} Q_\beta(\theta, x, dy) &= w_0(\theta, x, y) P^m(\theta, x, dy); \\ \varphi_\beta(x, dy) &= w_1(\theta, x, y) P^m(\theta, x, dy), \end{aligned} \tag{8}$$

and these densities 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 A , spaced at least m steps apart. Define a r.v. τ such that

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

for each $n \geq 1$, and let γ be such that $S_\gamma + m = \tau$. One can identify τ with the finite horizon T of the previous section. Its distribution is concentrated on $\{S_n + m, n \geq 1\}$, and one has

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

and $\mu_\theta \equiv \varphi$. This τ is a time at which the distribution of X is independent of its position at time $\tau - m$. It is the desired "regeneration time" for X under P_θ , and under Assumption 4, there exists $\epsilon > 0$ such that for $\theta \in \Lambda_\epsilon$,

$$\pi_\theta(dx) = \frac{E_\theta \left[\sum_{j=0}^{\tau-1} I(X_j \in dx) \right]}{E_\theta[\tau]}. \quad (9)$$

To construct a LR representation of π_θ in terms of π_{θ_0} , let Assumption 1 be in force and 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 $n \geq 1$ and $\theta \in \Lambda_\epsilon$. Then, it is possible to show that

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

$$\tilde{L}(\theta) = \left(\prod_{i=1}^{\tau} k(\theta, Z_i) \right) \rho(\theta), \quad \text{and} \quad L(\theta) = \left(\prod_{i=1}^{\tau} p(\theta, X_{i-1}, X_i) \right) \rho(\theta).$$

Combining this with (9) and Theorem 1, we obtain:

Corollary 3. *Under Assumptions 4 and 1(i), there exists $\epsilon > 0$ such that for all $\theta \in \Lambda_\epsilon$,*

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

In the degenerate case where $m = 0$, one can take $\beta = 1$, so $\rho(\theta) \equiv 1$.

4 A LR Representation for the Derivative of the Stationary Distribution

A glance at $\tilde{L}(\theta)$ and $L(\theta)$ suggests that a LR derivative formula for the stationary distribution should require differentiability of $p_m(\cdot)$ and $q(\cdot)$ with respect to θ . By recognizing that the derivative of $p_m(\cdot)$ can be defined in terms of the conditional expectation of the derivative of $k(\cdot)$, the required differentiability can be proved by imposing appropriate regularity conditions on $k(\cdot)$. These conditions force $p_m(\cdot, x, y)$ and $q(\cdot, x, y)$ to be well-behaved.

- Assumption 5.** (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 $r \in \mathbb{R}$, $\lim_{\epsilon \rightarrow 0} E \left[\sup_{\theta \in \Lambda_\epsilon} |k(\theta, Z_1)|^r \right] = 1$.*

To show that Assumption 2 (iii) and 3 hold, and apply the results of § 2, we need to control the behavior of τ and γ . Under suitable hypotheses (below), it can be shown that τ and γ have a geometrically dominated tail. We shall assume the existence of the following Lyapunov function:

Assumption 6. *There exists a function $g \geq 0$ defined on S_1 , $\epsilon > 0$, and $r < 1$ such that $E_{\theta_0, x} [g(X_1)] \leq r g(x) - \epsilon$ for $x \notin A$ and $\sup_{x \in A} E_{\theta_0, x} [g(X_1)] < \infty$.*

The set A is then 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 $\sup_{x \in A} E_{\theta_0, x} [z^{T(A)}] < \infty$. (See Nummelin (1984), pp. 90–91 and Chap. 16 of Meyn and Tweedie (1993).)

Proposition 4. *Under Assumptions 1(i), 4, and 5, $P[\gamma > k] \leq (1 - \beta\lambda_*)^k$ for $k \geq 0$. If, in addition, Assumption 6 is in force, then there exists $z > 1$ such that $E[z^\tau] < \infty$.*

It is then possible to show that

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

when these expressions exist. Under Assumptions 1(i) and 4–6, one can prove that $\rho'(\theta_0)$ exists a.s. Verifying that condition does not seem direct, because there is no a priori reason to expect almost sure differentiability, or even continuity over $\theta \in \Lambda_\epsilon$, of the r.v.'s $p_m(\theta, X_{S_i}, X_{S_i+m})$ and $q(\theta, X_{S_i}, X_{S_i+m})$. The proof also turns out to be quite involved.

Then, Assumptions 1–3 hold for our Harris-recurrent setup and Theorem 5 below 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 in Theorem 5 is precisely that of total variation.) 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. For a measure μ on S_1 and a S_1 -measurable function f , we adopt the notation $\mu f = \int_{S_1} f(y) \mu(dy)$. Let $f \geq 0$ be S_1 -measurable. This and our previous assumptions will imply that $\pi_{\theta_0} f^{1+\delta}$ is finite (see Tweedie 1983).

Assumption 7. *There exists a function $g \geq 0$ defined on S_1 and $\epsilon > 0$ such that $E_{\theta_0, x}[g(X_1)] \leq g(x) - \epsilon f(x)^{1+\delta}$ for $x \notin A$, and $\sup_{x \in A} E_{\theta_0, x}[g(X_1)] < \infty$.*

Theorem 5. *Let Assumptions 1(i) and 4–7 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

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

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 5.

Corollary 6. *Under the assumption of Theorem 5, $\pi_\theta g$ is differentiable at $\theta = \theta_0$ for any g satisfying $|g| \leq f$, and*

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

References

- Athneya, K. B. and P. Ney (1978). A new approach to the limit theory of recurrent Markov Chains. *Trans. Amer. Math. Soc.* **245**, 493–501.
- Fox, B. L. and P. W. Glynn (1986). Discrete-time conversion for simulating semi-Markov processes. *Operations Research Letters* **5**, 191–196.
- Meyn, S. P. and R. L. Tweedie (1993). *Markov Chains and Stochastic Stability*. Springer-Verlag, New York.
- Nummelin, E. (1978). A splitting technique for Harris recurrent Markov chains. *Z. Wahrscheinlichkeitstheorie* **43**, 309–318.
- Nummelin, E. (1984). *General Irreducible Markov Chains and Non-negative Operators*, Cambridge University Press, New York.
- Tweedie, R. L. (1983). The existence of moments for stationary Markov chains. *J. Appl. Prob.* **20**, 191–196.
- Vásquez-Abad, F. and H. Kushner (1992). Estimation of the derivative of a stationary measure with respect to a control parameter. *J. Appl. Prob.* **29**, 343–352.