Measure valued differentiation for random horizon problems

Heidergott, B.F.; Vazquez-Abad, F.

published in
Markov Processes and Related Fields
2006

document version
Publisher's PDF, also known as Version of record

Markov Processes Relat. Fields **12**, 509–536 (2006)

Markov Processes and Related Fields

Measure-Valued Differentiation for Random Horizon Problems

B. Heidergott[1] **and F.J. Vázquez-Abad**[2]

[1] Vrije Universiteit and Tinbergen Institute De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands. E-mail: bheidergott@feweb.vu.nl
[2] Department of Mathematics and ARC Special Research Centre for Ultra-Broadband Information Networks University of Melbourne, 3010 Australia. E-mail: vazquez@MS.UniMelb.edu.AU

Received July 28, 2004, revised March 18, 2006

Abstract. This paper deals with sensitivity analysis (gradient estimation) of random horizon cost functions of Markov chains. More precisely, we consider general state-space Markov chains and the random horizon is given through a hitting time of the chain onto a predefined set. The "cost" of interest is an expectation of a functional of the stopped process. This encompasses a wide range of models, such as the Gambler's ruin problem and performance evaluation for stationary queueing networks. We work within the framework of measure-valued differentiation and provide a general condition under which the gradient of the random horizon performance can be obtained in a closed form expression. For several scenarios, which occur typically in applications, we subsequently provide sufficient conditions for our general condition to hold. We illustrate our results with a series of examples. Eventually, we discuss unbiased sensitivity estimators and establish a new unbiased estimator for the gradient of stationary Markov chains.

Keywords: weak derivatives, gradient estimation, optimization, Markov chains

AMS Subject Classification: Primary 60J27, Secondary 90C31

1. Introduction

Markov chains are among the most used probabilistic models of discrete event stochastic phenomena that appear, for example, in the areas of manufacturing, transportation, finance and communication. In the past decade such models have been implemented to solve a number of problems related to performance evaluation and optimal design, and new methodologies are being developed in this growing research field. The goal of most research efforts is to find better and more efficient control methods in order to improve the *performance* of the system. In particular, stochastic approximation methods have extended the applicability of gradient search techniques to complex stochastic systems, but their implementation requires the construction of gradient estimators (of the cost function) satisfying certain conditions [15]. Complex stochastic systems, such as Discrete Event and Hybrid Dynamical Systems, can be described mathematically by general state-space Markov chains, such as generalized semi-Markov chains [5, 9]. Cost functions are described either by finite horizon cost functions or so-called "infinite" horizon performances.

Let $\theta \in \Theta$ be a continuous parameter on a compact set, with $\Theta \subset \mathbf{R}$, and let $(\Omega, \mathfrak{F}, \mathsf{P})$ be a probability space independent of θ. Throughout the paper we assume that the system under consideration is modeled as a homogeneous Markov chain $\{X_\theta(n)\}$ on $(\Omega, \mathfrak{F}, \mathsf{P})$ with filtration $\{\mathfrak{F}_\theta(n)\}$ and state-space $(S, \mathcal{S})$. Our notation summarizes the equivalence:

$$\mathsf{E}[g(X_\theta(n), \dots, X_\theta(1))] = \int_\Omega g(X_\theta(n, \omega), \dots, X_\theta(1, \omega))\, \mathsf{P}(d\omega).$$

We use the notation $P_\theta(ds_i; s) = \mathrm{Prob}(X_\theta(i) \in ds_i \mid X_\theta(i-1) = s)$ for the kernel of the Markov chain, so that

$$\mathsf{E}[g(X_\theta(n), \dots, X_\theta(1))] = \int g(s_n, \dots, s_1) \prod_{i=1}^{n} P_\theta(ds_i; s_{i-1})$$

and s_0 represents the initial state. A finite horizon cost function is an expectation of the form $\mathsf{E}[g(X_\theta(n), \dots, X_\theta(1))]$ for a given deterministic integer n. A theory for derivative estimation of finite horizon performances can be found in [14]. A *random horizon cost function* is an expectation of the form:

$$J(\theta) = \mathsf{E}[g_{\tau_\theta}(X_\theta(\tau_\theta), \dots, X_\theta(1))], \tag{1.1}$$

where τ_θ is a stopping time, more precisely, $\{\omega \in \Omega : \tau_\theta(\omega) = n\} \in \mathfrak{F}_\theta(n)$ and $g_n : S^n \to \mathbf{R}$, for $n \in \mathbf{N}$.

This paper deals with sensitivity analysis of random horizon experiments, where we assume throughout the paper that $\tau_\theta = \tau_\theta(B)$ is the first entrance time of the process $\{X_\theta(n)\}$ into some measurable set B. The initial state follows the distribution μ, assumed independent of θ. The sensitivity of a random horizon cost function is then defined as:

$$J'(\theta) = \frac{d}{d\theta} \mathsf{E}[g_{\tau_\theta}(X_\theta(\tau_\theta), \dots, X_\theta(1))]. \tag{1.2}$$

Formulating the stopping criterion via an entrance time of the (possibly enlarged) process into a predefined set B, imposes no severe restriction with respect to generality. However, it is worth emphasizing that the set B itself must not depend on θ. The goal of our analysis is to derive conditions under which the derivative (1.2) with respect to θ can be obtained in a closed form expression for a class of cost functions as large as possible.

Our approach is based on the concept of measure-valued differentiation, as introduced in [14] for finite horizon performances. As will be shown, it turns out that the potential of a Markov kernel is a natural way of expressing derivatives of stopped Markov processes. This has been observed earlier, see [6, 7, 18, 19].

Our results extend the ones known in the literature in the following way.

- The analysis in [6, 7, 18, 19] is restricted to gradients of bounded cost functions of stationary Markov chains. Moreover, in [6] and [7] only finite state-space Markov chains are treated. In contrast, our approach applies to general cost functions of general state-space Markov chains and to the (general) random horizon problem (gradients of stationary Markov chains are addressed in Section 5).

- The theory of weak differentiation as developed in [18, 19] does not cover random horizon problems and the analysis put forward in this paper is the first application of measure-valued derivatives to random horizon experiments.

- The general random horizon problem has been treated in [2] for an admission control problem and a general analysis has been provided in [3]. The key condition needed to estimate the derivative of $\mathsf{E}[\tau_\theta]$ was that τ_θ has finite third moment. We improve upon these results by only requiring the second moment of τ_θ to be finite.

The functional analytic approach presented in this paper is inspired by Pflug's work on weak derivatives, see for example [18, 19]. Moreover, key conditions for differentiability are formulated in terms of Lipschitz conditions for measures in a weak topology, thus emulating the formulation of Glasserman and Yao of the sample path approach of infinitesimal perturbation analysis on a measure theoretic level, see [10].

The paper is organized as follows. Section 2 discusses modeling random horizon experiments via potential kernels. In Section 3 sufficient conditions are established such that the potential kernel associated to a random horizon experiment is differentiable and a closed form expression for its derivative is derived. Moreover, this result is translated into the setting of random variables. It will turn out that the main technical condition is that a certain bound, to be defined later in the text, is finite. For situations that typically occur in applications, sufficient conditions for our bound to be finite are provided in

Section 4. Finally, we discuss gradient estimation for stationary performance characteristics in Section 5.

2. Random horizon experiments and the potential kernel

In this section, the basic notation is introduced. Let $(S, \mathcal{O})$ be a Polish measurable space and let $\mathcal{S}$ be the Borel-field on $(S, \mathcal{O})$. Denote the set of finite (signed) measures on $(S, \mathcal{S})$ by $\mathcal{M}$, and the set of probability measures by $\mathcal{M}_1$. The mapping $P : \mathcal{S} \times S \to [0, 1]$ is a transition kernel, or simply, kernel (on $(S, \mathcal{S})$) if

(a) $P(\cdot; s) \in \mathcal{M}$ for all $s \in S$;

(b) $P(B; \cdot)$ is $\mathcal{S}$-measurable for all $B \in \mathcal{S}$.

If, in condition (a), $\mathcal{M}$ can be replaced by $\mathcal{M}_1$, then we call P a Markov kernel. Denote the set of Markov kernels on $(S, \mathcal{S})$ by $\mathcal{K}_1 \stackrel{\text{def}}{=} \mathcal{K}(S, \mathcal{S})$ and the set of transition kernels by $\mathcal{K} \stackrel{\text{def}}{=} \mathcal{K}(S, \mathcal{S})$. The product of Markov kernels is again a Markov kernel. Specifically, let $P, Q \in \mathcal{K}$, then the product of P and Q is defined as follows. For $s \in S$ and $B \in \mathcal{S}$ set $PQ(B; s) = (P \circ Q)(B, s) = \int_S P(B; z)Q(dz; s)$. Moreover, write $P^n(\cdot; s)$ for the measure obtained by the n fold convolution of P in the above way. For technical reasons, we define P^0 to be the identity mapping.

When an initial distribution $\mu \in \mathcal{M}_1$ is given, P defines a Markov process $\{X_\theta(n)\}$ with state-space $(S, \mathcal{S})$. We write E_μ to indicate that the initial distribution of the Markov chain is μ. For example, one may think of starting $X_\theta(n)$ in an initial state s_0 and take $\mu = \delta_{s_0}$, the Dirac measure in $\{s_0\}$, as initial distribution. For what follows we assume that the initial state $y \in S$ is fixed. Set:

$$\langle g, P \rangle \stackrel{\text{def}}{=} \langle g, P(\cdot; y) \rangle = \int g(x) P(dx; y),$$

where it is understood that the left-hand side depends on the initial value y, and for $\mu \in \mathcal{M}_1$, $\langle g, P\mu \rangle = \int \int g(x) P(dx; y) \mu(dy) \in \mathbf{R}$.

In Section 1 we stated the problem in terms of the estimation of (1.2), where τ_θ is a hitting time adapted to $\{\mathfrak{F}_\theta(n)\}$, that is, $\tau_\theta = \min\{n : X_\theta(n) \in B\}$, $B \in \mathcal{S}$ and B is independent of $\theta \in \Theta$. All infinite horizon problems that use the Renewal Theorem for regenerative estimation fall within this formulation. In addition, models for the estimation of ruin probability and option pricing also stop when a given threshold is reached. These models are used in finance, telecommunications, insurance and optimal repair models. We now introduce a version of $X_\theta(n)$ that freezes as soon as τ_θ is reached. To this end, let $* \notin S$ and define $S_* = S \cup \{*\}$. Correspondingly, denote by $\mathcal{S}_*$ the σ-field generated by $\mathcal{S}$ and $\{*\}$. Cost functions $g\colon S \to \mathbf{R}$ are extended to S_* by setting $g(*) = 0$. The *halted process* evolves according to the Markov kernel $P_{\theta,\tau}$, defined by

$$P_{\theta,\tau}(A;s) = \begin{cases} P_\theta(A;s), & s \notin B \\ \delta_*(A), & s \in B \cup \{*\}. \end{cases} \tag{2.1}$$

In words, once $X_\theta(n)$ hits B, the Markov kernel $P_{\theta,\tau}$ forces the process to jump to the state $*$ and remains there, for ever frozen. The kernel $P_{\theta,\tau}$ is called the *halted Markov kernel*. Note that halted Markov chains are also called *stopped* Markov chains.

Note that $P_{\theta,\tau}$ fails to be a Markov kernel on $(S,\mathcal{S})$ which is due to the fact that $P_{\theta,\tau}(S;s) < 1$ for some $s \in S$. Such Markov kernels are called *deficient* in the literature. Denote the process that evolves according to $P_{\theta,\tau}$ by $X_\theta^\tau(n)$. Note that $X_\theta^\tau(n)$ is constructed in such a way that is resembles $X_\theta(n)$ as long as τ_θ has not occurred. Put another way, for any $n \in \mathbf{N}$ it holds that

$$\mathbf{1}\{\tau_\theta \geq n\} X_\theta^\tau(n) = X_\theta(n), \quad \mathsf{P}\text{-a.s.} \tag{2.2}$$

Let

$$K_\theta = K(P_\theta) \stackrel{\text{def}}{=} \sum_{m=1}^{\infty} P_\theta^m \tag{2.3}$$

denote the *potential kernel* of P_θ, see [17] or [16]. Note that the potential kernel is finite for deficient kernels only. We write $K_{\theta,\tau}$ to indicate the potential kernel of the Markov kernel associated to τ_θ.

Definition 2.1. Let $\Theta \subset \mathbf{R}$ be an open set. For $\theta \in \Theta$, let P_θ be a Markov kernel on some state-space $(S,\mathcal{S})$ and let μ be a distribution on $(S,\mathcal{S})$. The Markov chain with transition probabilities given by P_θ and initial distribution μ is denoted by $\{X_\theta(n)\}$. Let the stopping time $\tau_\theta = \tau_\theta(B)$ be defined through the first entrance time of the Markov chain $\{X_\theta(n)\}$ into $B \in \mathcal{S}$. Let g be a cost function on S. The tuple $(g,\mu,P_\theta,\tau_\theta,\Theta)$ defines a random horizon experiment (r.h.e.) through $J(\theta) = \mathsf{E}_\mu\left[\sum_{n=1}^{\tau_\theta} g(X_\theta(n))\right]$, $\theta \in \Theta$.

The r.h.e. $(g,\mu,P_\theta,\tau_\theta,\Theta)$ is called finite if $J(\theta)$ is finite for all $\theta \in \Theta$.

Let a finite r.h.e. $(g,\mu,P_\theta,\tau_\theta,\Theta)$ be given, then

$$J(\theta) = \langle g, K_{\theta,\tau}\mu \rangle, \quad \theta \in \Theta, \tag{2.4}$$

where $K_{\theta,\tau}$ is the potential kernel associated with the halted Markov kernel $P_{\theta,\tau}$ defined in (2.1). We will always assume that $K_{\theta,\tau}$ is defined by (2.3).

3. Differentiation of the halted Markov chain

In this section, basic definitions and properties of differentiable Markov kernels are established. In particular, the main technical result, providing sufficient conditions for $\mathcal{D}$-differentiability of a potential kernel, is established in Theorem 3.1.

For $\Theta = (a,b) \subset \mathbf{R}$, let $(P_\theta : \theta \in \Theta)$ be a family of Markov kernels on a Polish measurable space $(S,\mathcal{S})$ and assume that the initial distribution is fixed and independent of θ. To simplify the notation, we will suppress the initial distribution when this causes no confusion.

Definition 3.1. Let $\mathcal{D}$ be a set of measurable mappings $g : S \to \mathbf{R}$. Transition kernel P_θ is called $\mathcal{D}$-preserving if

$$\forall g \in \mathcal{D} : \quad \int_S g(u) P_\theta(du; \cdot) \in \mathcal{D}.$$

Denote by $L^1(P_\theta;\Theta) \subset \mathbf{R}^S$ the set of measurable mappings $g : S \to \mathbf{R}$ such that $\int |g(u)| P_\theta(du;s)$ is finite for all $\theta \in \Theta$ and $s \in S$.

Definition 3.2. Let P_θ be a kernel and $\mathcal{D} \subset L^1(P_\theta;\Theta)$. We call P_θ *$\mathcal{D}$-Lipschitz continuous* if for any $g \in \mathcal{D}$ a $K_g \in \mathcal{D}$ exists such that for any $\Delta > 0$ with $\theta + \Delta \in \Theta$

$$\left| \int g(s) P_{\theta+\Delta}(ds;\cdot) - \int g(s) P_\theta(ds;\cdot) \right| \le \Delta K_g(\cdot).$$

Recall that $\mathcal{K}_1$ denotes the set of Markov kernels on $(S,\mathcal{S})$ and $\mathcal{K}$ the set of transition kernels on $(S,\mathcal{S})$.

Definition 3.3. For $\theta \in \Theta$, let $P_\theta \in \mathcal{K}_1$, and let $\mathcal{D} \subset L^1(P_\theta : \Theta)$. We call P_θ *$\mathcal{D}$-differentiable* if a transition kernel $P'_\theta \in \mathcal{K}$ exists such that

$$\frac{d}{d\theta} \int g(u) P_\theta(du;s) = \int g(u) P'_\theta(du;s), \qquad \text{for all} \ \ g \in \mathcal{D}, s \in S,$$

and P'_θ is $\mathcal{D}$-preserving.

In the above definition, if the set $\mathcal{D}$ is taken to be the set of bounded, continuous cost functions, then $\mathcal{D}$-differentiability recovers weak differentiability; see, for example, [19].

Throughout the paper we assume that the set of continuous, bounded mappings are a subset of $\mathcal{D}$, but $\mathcal{D}$ is only required to be in the set of integrable functions: $C^1(P_\theta) \subset \mathcal{D} \subset L^1(P_\theta)$. This implies that $\mathcal{D}$-differentiability extends weak differentiability.

The assumption $C^1(P_\theta) \subset \mathcal{D}$ implies that P'_θ is uniquely defined; see [14] for details. If P_θ is $\mathcal{D}$-differentiable, then, under some mild extra conditions, there exist Markov kernels P_θ^+, P_θ^- and a random variable c_{P_θ} such that for all $g \in \mathcal{D}$

$$\frac{d}{d\theta} \int g(u) P_\theta(du;s) = c_{P_\theta}(s) \Big(\int g(u) P_\theta^+(du;s) - \int g(u) P_\theta^-(du;s) \Big). \tag{3.1}$$

For details we refer to [13].

Definition 3.4. For $\theta \in \Theta$, let $P_\theta \in \mathcal{K}_1$, and let $\mathcal{D} \subset L^1(P_\theta : \Theta)$, such that P_θ is $\mathcal{D}$-differentiable on Θ. Any triple $(c_{P_\theta}(\cdot), P_\theta^+, P_\theta^-)$, with $P_\theta^\pm \in \mathcal{K}_1$ and c_{P_θ} a measurable mapping from S to $\mathbf{R}$, that satisfies (3.1) is called a $\mathcal{D}$-*derivative* of P_θ. The kernel P_θ^+ is called the (normalized) *positive part* of $P_\theta{}'$ and P_θ^- is called the (normalized) *negative part* of $P_\theta{}'$; and $c_{P_\theta}(s)$ is called the *normalizing factor*.

For easy reference we introduce the following definition.

Definition 3.5. We say that $(P_\theta, \mathcal{D})$, with $\theta \in \Theta$, is *Leibnitz* if

- $\mathcal{D} \subset L^1(P_\theta : \Theta)$ and for any $g, f \in \mathcal{D}$ it holds that $g + f \in \mathcal{D}$,
- P_θ is $\mathcal{D}$-preserving and $\mathcal{D}$-Lipschitz,
- P_θ is $\mathcal{D}$-differentiable and P_θ' is $\mathcal{D}$-Lipschitz,

If $(P_\theta, \mathcal{D})$ is Leibnitz for any $\theta \in \Theta$, then we say that $(P_\theta, \mathcal{D})$ is Leibnitz on Θ.

Remark 3.1. If $(P_\theta, \mathcal{D})$ is Leibnitz, then a sufficient condition for $\mathcal{D}$-Lipschitz continuity of P_θ is that for any $g \in \mathcal{D}$ $\sup_{\theta\in\Theta} \left|\langle g, P_\theta`(ds; \cdot)\rangle\right| \in \mathcal{D}$ (for a proof use the Mean Value Theorem). In the same vein, a sufficient condition for $\mathcal{D}$-Lipschitz continuity of P_θ' is that for any $g \in \mathcal{D}$

$$\sup_{\theta\in\Theta} \left| \frac{d^2}{d\theta^2} \langle g, P_\theta(ds; \cdot)\rangle \right| \in \mathcal{D},$$

provided that the second-order derivative exists ($d^2 P_\theta / d\theta^2$ is defined in the obvious way).

Let $\|\cdot\|$ denote the supremum norm on Θ, that is, for $f : \Theta \to \mathbf{R}$ we set $\|f\| = \sup_{\theta\in\Theta} |f_\theta|$.

Definition 3.6. For $\theta \in \Theta$, let $P_\theta \in \mathcal{K}_1$, and let $\mathcal{D} \subset L^1(P_\theta : \Theta)$, such that P_θ is $\mathcal{D}$-differentiable on Θ. For $g \in \mathcal{D}$, we define the bound $H_\mu(P, g)$ by

$$H_\mu(P, g) \stackrel{\text{def}}{=} H_\mu((P_\theta : \theta \in \Theta), g) = \sum_{n=1}^{\infty} \left\| \frac{d}{d\theta} \langle g, P_\theta^n \mu\rangle \right\|.$$

As the next theorem shows, finiteness of $H_\mu(P, g)$ is the key condition for deriving a closed form expression of the gradient of the stopped experiment in (1.2). For various typical scenarios, we will discuss in Section 4 conditions that imply the finiteness of $H_\mu(P, g)$. To simplify the notation, we will often suppress the subscript "θ" whenever this causes no confusion. The proof of the theorem is provided in Section A.1 in the Appendix.

Theorem 3.1. *Let Θ be compact and assume that $(P_\theta, \mathcal{D})$ is Leibnitz on Θ. Denote the potential kernel of P_θ by K_θ and let μ be some initial distribution.*

(i) *Let $g \in \mathcal{D}$. If $\langle g, K_\theta \mu \rangle$ and $H_\mu(P, g)$ are finite, then*

$$\frac{d}{d\theta} \langle g, K_\theta \mu \rangle = \sum_{n=1}^{\infty} \sum_{m=1}^{n} \langle g, P_\theta^{n-m} P_\theta' P_\theta^{m-1} \mu \rangle.$$

(ii) *If $\langle g, K_\theta \mu \rangle$ and $H_\mu(P, g)$ are finite for all $g \in \mathcal{D}$, then $K_\theta \mu$ is $\mathcal{D}$-differentiable with $K_\theta' \mu = K_\theta P_\theta' K_\theta \mu$.*

We now turn to the halted Markov kernel. As a first step, we show that $\mathcal{D}$-differentiability of a transition kernel implies that of the halted kernel $P_{\theta,\tau}$.

Lemma 3.1. *For $\theta \in \Theta$, let $P_\theta \in \mathcal{K}_1$, and let $\mathcal{D} \subset L^1(P_\theta : \Theta)$, such that P_θ is $\mathcal{D}$-differentiable on Θ. Then $P_{\theta,\tau}$ is $\mathcal{D}$-differentiable and if $P_\theta(\cdot; s)$ has $\mathcal{D}$-derivative $(c_{P_\theta}(s), P_\theta^+(\cdot; s), P_\theta^-(\cdot; s))$, then $P_{\theta,\tau}(\cdot; s)$ has $\mathcal{D}$-derivative*

$$\big(c_{P_{\theta,\tau}}(s), P_{\theta,\tau}^+(\cdot; s), P_{\theta,\tau}^-(\cdot; s)\big) = \begin{cases} \big(c_{P_\theta}(s), P_\theta^+(\cdot; s), P_\theta^-(\cdot; s)\big), & s \notin B, \\ \big(0, P_\theta(\cdot; s), P_\theta(\cdot; s)\big), & s \in B \cup \{*\}. \end{cases}$$

Moreover, if $(P_\theta, \mathcal{D})$ is Leibnitz so is $(P_{\theta,\tau}, \mathcal{D})$.

Proof. We only prove the first part of lemma since the proof of the second part follows from the same line of argument. For any $g \in \mathcal{D}$ it holds that

$$\begin{aligned} \frac{d}{d\theta} \int_S g(z) P_{\theta,\tau}(z; s) &= \mathbf{1}\{s \notin B\} \frac{d}{d\theta} \int_S g(z) P_\theta(z; s) \\ &= \mathbf{1}\{s \notin B\} c_{P_\theta}(s) \Big(\int_S g(z) P_\theta^+(z; s) - \int_S g(z) P_\theta^-(z; s) \Big) \\ &= c_{P_{\theta,\tau}}(s) \Big(\int_S g(z) P_{\theta,\tau}^+(z; s) - \int_S g(z) P_{\theta,\tau}^-(z; s) \Big), \end{aligned}$$

for $s \in S$ which concludes the proof. □

We now introduce the "plus" and "minus" processes $X_m^\pm(\cdot)$ to express the derivative in Theorem 3.1 as an expectation. The result will then be summarized in Theorem 3.2. For $m \neq n$, let the transition from $X_m^\pm(n-1)$ to $X_m^\pm(n)$ be governed by $P_{\theta,\tau}$, whereas the transition from $X_m^+(m-1)$ to $X_m^+(m)$ is governed by $P_{\theta,\tau}^+$ and that from $X_m^-(m-1)$ to $X_m^-(m)$ by $P_{\theta,\tau}^-$, respectively. Hence, while $n \neq m$, both $X_m^+(n)$ and $X_m^-(n)$ are driven by the same Markov kernel. Consequently, for $n \geq m \geq 1$ it holds that

$$\mathsf{P}\left(X_m^\pm(n) \in A \mid X_m^\pm(0) = s\right) = P_{\theta,\tau}^{n-m} P_{\theta,\tau}^\pm P_{\theta,\tau}^{m-1}(A; s),$$

and for $1 \leq n < m$ it holds that

$$\mathsf{P}\left(X_m^{\pm}(n) \in A \mid X_m^{\pm}(0) = s\right) = P_{\theta,\tau}^n(A; s),$$

for $A \in \mathcal{S}$ and $s \in S$. We remark that the dependence of $X_m^{\pm}(n)$ on θ is not expressed in the notation in order to avoid too heavy a notation.

Throughout the paper we assume that for each $m < \tau_\theta$, the processes $\{X_m^{\pm}(n)\}$ are constructed from $P_{\theta,\tau}$ using common random numbers (CRN) for all transitions $n < m$.

Remark 3.2. Given m, the processes $X_\theta(n), X_m^+(n)$ and $X_m^-(n)$ have the same distribution for $n < m$. Taking CRN, we choose a particular version of the processes that actually defines the plus and minus processes to have an a.s.-identical trajectory to the so-called “nominal” process $X_\theta(n)$ up to transition $m-1$, that is:

$$X_\theta(n) = X_m^+(n) = X_m^-(n) \quad \text{a.s.}, \quad n < m.$$

Then, at the transition from $m-1$ to m the nominal process “splits” in three different trajcetories, each trajectory governed by a different transition distribution. After this splitting, the transition kernels are again equal for the three processes. Depending on the application, we may choose to continue using either independent or correlated random variables to generate the future transitions. This is called the *coupling* scheme. Which coupling scheme to use highly depends on the particular context of the application and may enlarge the σ-field for derivative estimation.

Denote the first entrance times of $X_m^{\pm}(n)$ into B by $\tau_m^{\pm}$. Once $X_m^{\pm}(n)$ hits B, $X_m^{\pm}(n)$ jumps to $*$ and stays there forever. Hence, for any $n, m \in \mathbf{N}$ it holds that

$$\mathbf{1}\{\tau_m^{\pm} \geq n\} X_m^{\pm}(n) = X_m^{\pm}(n) \quad \mathsf{P}\text{-a.s.} \tag{3.2}$$

Condition $H_\mu(P, g)$ in Theorem 3.1 has no immediate interpretation in terms of a stochastic experiment. For this reason, we introduce in the following definition a new bound that can be expressed as a stochastic experiment in terms of the Markov chains $\{X_\theta(n)\}$, $\{X_m^{\pm}(n)\}$.

Definition 3.7. For $\mathcal{D} \subset L^1(P_\theta : \Theta)$ let a r.h.e. $(g, \mu, P_\theta, \tau_\theta, \Theta)$ with $g \in \mathcal{D}$ be given. Assume that P_θ is $\mathcal{D}$-differentiable on Θ with $\mathcal{D}$-derivative $(c_{P_\theta}, P_\theta^+, P_\theta^-)$. Provided that $g \in \mathcal{D}$, we define the *cumulative bound* $\mathbf{H}_\mu(P, g)$ by

$$\begin{aligned}\mathbf{H}_\mu(P, g) &\stackrel{\text{def}}{=} \mathbf{H}_\mu((P_\theta : \theta \in \Theta), g) \\ &= \mathsf{E}_\mu\Bigg[\sum_{m=1}^{\|\tau_\theta\|} \|c_{P_\theta}(X_\theta(m-1))\| \Bigg(\sum_{n=1}^{\|\tau_m^+\|} \|g(X_m^+(n))\| + \sum_{m=1}^{\|\tau_m^-\|} \|g(X_m^-(n))\|\Bigg)\Bigg].\end{aligned}$$

Note that the above definition implies that $\mathbf{H}_\mu(P,g) = \mathbf{H}_\mu(P_\tau, g)$ for any $g \in \mathcal{D}$. The proof of the following theorem is given in Section A.2 of the Appendix.

Theorem 3.2. *For $\mathcal{D} \subset L^1(P_\theta : \Theta)$ let a r.h.e. $(g, \mu, P_\theta, \tau_\theta, \Theta)$ with $g \in \mathcal{D}$ be given. Assume that $(P_\theta, \mathcal{D})$ is Leibnitz on Θ. If $\mathbf{H}_\mu(P,g)$ is finite, then*

$$\frac{d}{d\theta}\, \mathsf{E}_\mu\left[\sum_{n=1}^{\tau_\theta} g(X_\theta(n))\right]$$
$$= \mathsf{E}_\mu\left[\sum_{m=1}^{\tau_\theta} c_{P_\theta}(X_\theta(m-1))\left(\sum_{n=m}^{\tau_m^+} g(X_m^+(n)) - \sum_{n=m}^{\tau_m^-} g(X_m^-(n))\right)\right].$$

The key conditions for the formula in Theorem 3.2 for the sensitivity of a r.h.e. to hold are that $(P_\theta, \mathcal{D})$ is Leibnitz, that the particular g for which the r.h.e. is carried out lies in $\mathcal{D}$, and that the cumulative bound $\mathbf{H}_\mu(P,g)$ is finite. In applications, P'_θ is typically of rather complex structure and checking whether $(P_\theta, \mathcal{D})$ is Leibnitz leads to cumbersome calculations. However, as we will illustrate with examples, using a conditioning type of approach, the fact that $(P_\theta, \mathcal{D})$ is Leibnitz can be deduced from more elementary considerations.

4. Finiteness of the cumulative bound

Finiteness of the cumulative bound $\mathbf{H}_\mu$ is essential for the existence of the gradient. In this section, sufficient conditions for finiteness of $\mathbf{H}_\mu$ are established.

For the sake of simplicity, we introduce the following assumption.

(A) For $\theta \in \Theta$, let $P_\theta \in \mathcal{K}_1$, and let $\mathcal{D} \subset L^1(P_\theta : \Theta)$ be such that P_θ is $\mathcal{D}$-differentiable on Θ with $\mathcal{D}$-derivative $(c_{P_\theta}, P_\theta^+, P_\theta^-)$. A number $c \in \mathbf{R}$ exists such that, for any $\theta \in \Theta$, $\sup_{s\in S} \|c_{P_\theta}(s)\| \le c$.

Condition (A) imposes no severe restriction. Indeed, in many situations which are of importance in applications, condition (A) is satisfied, see [14]. Let $\mathcal{D}^b$ denote the set of bounded measurable mappings from S onto $\mathbf{R}$.

Remark 4.1. For $g \in \mathcal{D}^b$, the r.h.e. $(g, \mu, P_\theta, \tau_\theta, \Theta)$ is finite if $\mathsf{E}_\mu[\tau_\theta] < \infty$ for any $\theta \in \Theta$.

4.1. Stochastic ordering

Call a "cycle" the process $\{X_\theta(n), n \le \tau_\theta\}$. This terminology is common in the context of regenerative simulation where τ_θ is the time until return to the initial (regenerative) state. The main difficulty in showing that the cumulative bound is finite lies in the fact that the perturbed processes $X_m^{\pm}(n)$ may have longer cycles than the nominal process, that is, $\tau_m^{\pm} > \tau_\theta$. However, for certain processes a dominating stopping time, say T, can be identified, such that $T \geq \max(\tau_\theta, \tau_m^{\pm})$ and the cycles of the nominal and perturbed processes can be controlled via cycles of length T.

Lemma 4.1. *For $\mathcal{D} \subset L^1(P_\theta, \theta : \Theta)$ let a finite r.h.e. $(g, \mu, P_\theta, \tau_\theta, \Theta)$ be given with $g \in \mathcal{D}$. Let $(P_\theta, \mathcal{D})$ be Leibnitz on Θ with $g \in \mathcal{D}$. Suppose that $\{X_\theta(n)\}$, $\{X_m^{\pm}(n)\}$, $m \in \mathbf{N}$, and a random variable $T \in \mathbf{N}$ can be constructed on a common probability space $(\Omega, \mathfrak{F}, \mathsf{P})$ such that with probability one*

$$T \geq \max\big(\|\tau_\theta\|, \|\tau_m^{\pm}\|\big),$$

for any m. Under condition (A) it then holds that

$$\mathbf{H}_\mu(P, g) \leq c\, \mathsf{E}_\mu \left[\sum_{n=1}^{\|T\|} \sum_{m=1}^{\|T\|} \big(\|g(X_n^+(m))\| + \|g(X_n^-(m))\| \big) \right].$$

In particular, if $\mathcal{D} = \mathcal{D}^b$, then

$$\mathsf{E}_\mu \left[\|T\|^2 \right] < \infty \quad \Longrightarrow \quad \mathbf{H}_\mu(P_\tau, g) < \infty.$$

Proof. The statement of the lemma is an immediate consequence of the definition of $\mathbf{H}_\mu(P, g)$ and the stochastic ordering assumption. □

As the following example shows, $T = \tau_\theta$ is sometimes an appropriate choice.

Example 4.1. In [4] a thinning of a Point process was used to build the "Rare Perturbation Analysis" (RPA) derivative estimator. To see how this method is a particular case of the $\mathcal{D}$-derivatives, consider a $GI/G/1$ queuing system under admission control as follows. Customers arrive according to a marked Point process $\{(A_i, Z_i), i = 1, 2, \ldots\}$ where $T_i = A_i - A_{i-1}$ is the interarrival time between customers $i-1$ and i ($A_0 \equiv 0$), and Z_i the service requirement of customer i. It is assumed that $\{T_i\}$ and $\{Z_i\}$ are independent sequences of i.i.d. nonnegative random variables. Each arriving customer is admitted into the queue with probability θ, and it is rejected if not admitted (no feedback), as in [12]. To establish the mathematical framework, one defines a sequence $\{\eta_i(\theta), i = 1, 2, \ldots\}$ of i.i.d. random variables with Bernoulli distribution of parameter θ, that is, $\mathsf{P}(\eta_n(\theta) = 1) = \theta = 1 - \mathsf{P}(\eta_n(\theta) = 0)$, for any n. The waiting times of successive customers follow Lindley's recursion:

$$X_\theta(n+1) = \max\big((X_\theta(n) + Z_n \eta_n(\theta)) - T_n, 0\big), \quad n \geq 1,$$

with initial value $X_\theta(0) = 0$, where a customer with $\eta_n(\theta) = 0$ has "instantaneous service" $Z_n \eta_n(\theta) = 0$ and the effect is as if it disappeared from the queue.

When the arrival process is a renewal process, the queue is regenerative, and stationary expectations can be obtained via the Renewal Theorem using expectations over one cycle of the queue. More specifically, take as initial distribution the point measure in state 0, (in formula $\mu(\cdot) = \delta_0(\cdot)$), representing an arrival to an idle server. Hence, state $X_\theta(0) = 0$ is the initial state and

$$\tau_\theta = \min\{n \geq 1 : X_\theta(n) = 0\} - 1$$

is the index of the final customer of the first busy cycle; indeed $\min\{n \geq 1 : X_\theta(n) = 0\}$ is the index of the customer that starts the second cycle.

It is not hard to evaluate the derivative of the transition kernel, because in this case $(T_n, Z_n, \eta_n(\theta))$ is independent of $X_\theta(n)$. In words, the time for the next arrival, the current service requirement and the admission variable are independent of the time that the current customer waits in the queue. Moreover, the dynamics of the Markov chain depends on θ only through the Bernoulli distributed random variable $\eta_n(\theta)$. This implies:

$$\begin{aligned}
&\frac{d}{d\theta}\, \mathsf{E}[g(X_\theta(n+1)) \mid X_\theta(n) = s] \\
&\quad = \frac{d}{d\theta}\, \mathsf{E}\left[(1-\theta) g(\max(s - T_n, 0)) + \theta g(\max(s + Z_n - T_n, 0))\right] \\
&\quad = \mathsf{E}\left[g(\max(s + Z_n \eta^+ - T_n, 0)) - g(\max(s + Z_n \eta^- - T_n, 0))\right],
\end{aligned}$$

where the last expectation is w.r.t. (Z_n, T_n), and $\eta^\pm$ are degenerate random variables, $\eta^+ \equiv 1$ and $\eta^- \equiv 0$. Setting $\theta = 1$, yields $\eta_n(\theta) \equiv \eta^+$ and, in the same vein, setting $\theta = 0$ yields $\eta_n(\theta) \equiv \eta^-$. Hence, the difference on the right-hand side of the above equation can be obtained from the nominal transition kernel through setting either $\theta = 1$ or $\theta = 0$. Let $\mathcal{D}$ be the set of mappings $g : \mathbf{R}^+ \to \mathbf{R}$, such that for any $s \in \mathbf{R}^+$ it holds that

$$\mathsf{E}\left[g(X_1(n+1)) \mid X_1(n) = s\right] < \infty,$$

then the $\mathcal{D}$-derivative of the kernel is the triple $(1, P_1(\cdot; s), P_0(\cdot; s))$. Observe that the $\mathcal{D}$-derivative is independent of θ and thus both the kernel and its $\mathcal{D}$-derivative are $\mathcal{D}$-Lipschitz, for a proof use Remark 3.1. Specifically, $(P_\theta, \mathcal{D})$ is Leibnitz on Θ. Hence, the halted chain can be used to estimate the desired gradients using the positive and negative kernels as follows: the original queue is simulated (or observed) and for each customer $m \leq \tau_\theta$, a positive (negative) version of the process $X_m^+(n)$ ($X_m^-(n)$) is obtained by accepting (rejecting) the mth customer, independently of the decision $\eta_m(\theta)$ made for the so-called "nominal" process $X_\theta^\tau(n)$. In this example, the coupling scheme uses CRN throughout, where transitions beyond m for the perturbed processes $X_m^\pm(\cdot)$ are obtained using the same values of $T_n, Z_n, \eta_n(\theta)$ as the nominal process.

We now discuss the Phantom RPA method introduced in [4]. For that problem, arrivals follow a homogeneous Poisson process of rate $\bar{\lambda}$ and there is no admission control. In the admission setting, because of the random admission policy, the admitted arrivals also follow a Poisson process, of rate $\lambda = \theta\bar{\lambda}$. Using the chain rule, and the fact that $\bar{\lambda}$ is independent of θ,

$$\frac{\partial}{\partial\theta} = \Big(\frac{d\lambda}{d\theta}\Big)\frac{\partial}{\partial\lambda} = \Big(\frac{\lambda}{\theta}\Big)\frac{\partial}{\partial\lambda}.$$

In this setting, the positive and negative processes are defined using the same arrival and service requirements $((T_n, Z_n); n \in \mathbf{N})$. Therefore $(X_m^+, m = 1, \ldots, \tau)$ are equal P-a.s. to X_θ, where all customers are accepted. For each customer $m < \tau_\theta$ a negative process X_m^- is obtained by rejecting this customer, which necessarily implies stochastic ordering and $\tau_m^- < \tau_m^+ = \tau_\theta \le \tau_1 \stackrel{\text{def}}{=} T$, P-a.s. Let $S_\tau(\theta) = \sum_{n=1}^{\tau_\theta} Z_n$ denote the total length of the busy cycle. Stochastic ordering implies

$$S_\tau \stackrel{\text{def}}{=} S_\tau(1) \ge S_\tau(\theta),$$

for any $\theta \in \Theta = [0, 1]$. Let

$$|g(X_\theta(n))| \le (S_\tau)^p, \tag{4.1}$$

for $p \in \mathbf{N}$. For example, the pth moment of any sojourn time of the busy cycle under consideration is bounded by $(S_\tau)^p$. Note that (4.1) implies that

$$\|g(X_m^\pm(n))\| \le (S_\tau)^p, \tag{4.2}$$

for any m, n.

Assuming that $\mathsf{E}\,[\tau_1^2 S_\tau^p]$ is finite, it follows by (4.2) from Lemma 4.1 that the cumulative bound is finite (take $T = \tau_1$). Moreover, the integrability condition on the cycle performance in Theorem 3.2 is satisfied. Since $(P_\theta, \mathcal{D})$ is Leibnitz on Θ, we obtain for any $g \in \mathcal{D}$

$$\frac{\partial}{\partial\lambda}\,\mathsf{E}_\mu\left[\sum_{n=1}^{\tau_\theta} g(X_\theta(n))\right] = \frac{1}{\lambda}\,\mathsf{E}\left[\tau\bigg(\sum_{n=\sigma}^{\tau_\theta} g(X_\theta(n)) - \sum_{n=\sigma}^{\tau_\sigma^-} g(X_\sigma^-(n))\bigg)\right], \tag{4.3}$$

where σ is a random index, uniformly chosen from $\{1, \ldots, \tau_\theta\}$. This improves upon the analysis in [4], where the key condition for unbiasedness of the above estimator is $\mathsf{E}\,[\tau_1^3 S_\tau^p] < \infty$.

In the foregoing example $\tau_1 = \tau_m^+$ serves as dominating random variable T in Lemma 4.1, and the condition in Lemma 4.1 reduces to $\mathsf{E}\,[\tau_1^2] < \infty$, which can be guaranteed through appropriate conditions on the nominal model. In the following example, the dominating random variable will be τ_m^-, for $m \in \mathbf{N}$. We will illustrate how finiteness of the second moment of τ_m^- can be deduced from properties of the nominal model, using a change of measure idea first explained in [22].

Example 4.2. The following ruin problem is of importance in Risk Theory. An insurance company receives premiums at some rate depending on the current client portfolio. Claims arrive according to a Poisson process of intensity θ, with θ out of a compact set $\Theta \subset \mathbf{R}$. Let $T_\theta(i)$ be the interarrival times of the Poisson process and $A_\theta(i) = \sum_{k=1}^{i} T_\theta(i)$ the arrival epoch of the ith claim. Consecutive claim amounts form a sequence $\{Y_n\}$ of i.i.d. random variables with density f, and call $S(i) = \sum_{k=1}^{i} Y_k$. The insurance company receives payments according to a premium of c dollars per unit time, so the surplus of the company at claim arrival times evolves according to the recursion:

$$X_\theta(n+1) = X_\theta(n) + cT_\theta(n) - Y_{n+1}; \quad X(0) = u, \tag{4.4}$$

where u is the initial endowment. Ruin occurs when the surplus becomes negative, that is,

$$\tau_\theta = \min\{n \in \mathbf{N} : X_\theta(n) < 0\},$$

and we want to estimate the sensitivity of the *probability of ruin*:

$$J(\theta) = \mathsf{P}(\tau_\theta < \infty).$$

It is well known (see [1,23]) that direct simulation of the surplus process is inefficient: it is not clear when to terminate simulation of sample paths where τ_θ is not finite. As a function of u, ruin is a *rare event*: the probability decreases to zero as u increases. For this setting, Importance Sampling has been successfully used to improve the efficiency in the estimation. Suppose that the claim distribution has a moment generating function $M_Y(R) = \mathsf{E}[\exp RY]$. The idea in [1] is to use a change of measure for the claim distribution, simulating claims from the so-called exponentially tilted density (or Escher transform): $f_0(y) = f(y)\exp\{Ry\}/M_Y(R)$, and using a Poisson process with rate $\vartheta = \theta + cR$. Call E_0 the expectation with respect to this new process. Evaluating the Randon−Nikodym derivative for this change of measure, it follows that for any R,

$$J(\theta) = \mathsf{E}_0\left[\left(\frac{M_Y(R)\theta}{(\theta+Rc)}\right)^{\tau_\theta} \exp\{R(cA(\tau_\theta) - S(\tau_\theta))\}\right].$$

Call

$$K(\theta) = \frac{M_Y(R)\theta}{(\theta+Rc)},$$

and notice that it is a bounded, deterministic and continuously differentiable function of θ. As explained in [1,24], it is possible to choose R so that ruin is certain under the new measure, that is: $\mathsf{P}_0(\tau_\theta < \infty) = 1$.

In order to write the model as a homogeneous Markov chain and $J(\theta)$ in the form (1.2) one enlarges the state space as: $\tilde{X}_\theta(n) = (X_\theta(n), n, I_\theta(n))$, where

$$I_\theta(n+1) = \begin{cases} 0, & \text{if } I_\theta(n) = 0 \text{ and } X_\theta(n) > 0, \\ 1, & \text{if } I_\theta(n) = 0 \text{ and } X_\theta(n) \le 0, \\ -1, & \text{if } I_\theta(n) = 1 \text{ or } -1. \end{cases}$$

In words, $I_\theta(n) = 1$ if and only if ruin occurs at transition n. Note that $\tilde{X}_\theta(n)$ is Markovian. The cost function is here:

$$g_\theta(x, n, i) = \begin{cases} (K(\theta))^n e^{R(x-u)}, & i = 1, \\ 0 & \text{otherwise,} \end{cases}$$

and one has (using (4.4))

$$\begin{aligned}\sum_{n=1}^{\infty} g_\theta(\tilde{X}_\theta(n)) &= (K(\theta))^{\tau_\theta} \exp\{R(X_\theta(\tau_\theta) - u)\} \\ &= (K(\theta))^{\tau_\theta} \exp\{R(cA(\tau_\theta) - S(\tau_\theta))\}.\end{aligned}$$

Let P_θ denote the Markov kernel of $\{\tilde{X}_\theta(n)\}$ and let $\mathcal{K}_\theta$ denote the potential kernel of P_θ. Then it holds for any bounded measurable g that

$$\langle g, \mathcal{K}_\theta \mu \rangle = \mathsf{E}_0 \left[\sum_{n=1}^{\infty} g_\theta(\tilde{X}_\theta(n)) \right] = \mathsf{E}\left[(K(\theta))^{\tau_\theta} \exp\{R(X_\theta(\tau_\theta) - u)\} \right] = J(\theta).$$

We now turn to the derivative of P_θ (and $\mathcal{K}_\theta$, respectively). From the definition of the enlarged state, the derivative of the transition kernel of $\tilde{X}_\theta(n)$, that is P_θ, can be obtained in terms of the derivative of the transition kernel of $X_\theta(n)$, defined in (4.4), because the transition probabilities of the second and the third components are independent of θ. Let $\mathcal{D}^b$ denote the set of bounded measurable mappings from S onto $\mathbf{R}$. For any $h \in \mathcal{D}^b$ and $x > 0$, one has:

$$\begin{aligned}&\frac{d}{d\theta}\ \mathsf{E}_0[h(X_\theta(n+1)) \mid X_\theta(n) = x] \\ &\quad = \frac{d}{d\vartheta} \int\int h(x + ct - y)\big(\vartheta e^{-\vartheta t}\big) f_0(y)\, dt\, dy \\ &\quad = \int \Big(\int h(x + ct - y) \frac{d}{d\vartheta}\big(\vartheta e^{-\vartheta t}\big)\, dt \Big) f_0(y)\, dy \\ &\quad = \int\int h(x + ct - y) \frac{1}{\vartheta}\big(\vartheta e^{-\vartheta t} - t\vartheta^2 e^{-\vartheta t}\big)\, dt\, f_0(y)\, dy \\ &\quad = \frac{1}{\vartheta}\, \mathsf{E}_0\left[h(X_\theta(n+1)) - h(X^-(n+1)) \mid X_\theta(n) = x\right],\end{aligned}$$

where $X^-(n+1) = X_\theta(n) + c\Gamma_n - Y_{n+1}$, for Γ_n a random variable with Gamma distribution of parameters $(2, \vartheta)$, and $Y_n \sim F_0$. Since Γ_n has the same distribution as the sum of two i.i.d. exponential random variables with intensity ϑ, one can couple the processes $X_\theta(n)$, $X^-(n)$ generating an exponential $\hat{T}_\theta$ and using $\Gamma_n = T_\theta(n) + \hat{T}_\theta$. Then necessarily $X^-(n+1) \geq X_\theta(n)$ a.s. and stochastic domination follows for the ensuing processes. Following Remark 3.1, it is easy to see that P_θ is $\mathcal{D}^b$-Lipschitz continuous and calculating the second-order derivative of $\mathsf{E}_0\left[h(X_\theta(n+1)) - h(X^-(n+1)) \mid X_\theta(n) = x\right]$ one checks that P'_θ is $\mathcal{D}^b$-Lipschitz continuous as well. Therefore, $(P_\theta, \mathcal{D}^b)$ is Leibnitz on Θ.

The key condition for expressing the derivative of $J(\theta)$ is finiteness of the cumulative bound. This follows from $\mathsf{E}_0[(\tau_j^-)^2] < \infty$ (for a proof apply Lemma A.4 in the Appendix). The process $\{X_j^-(n), n = 1, \ldots, \tau_j^-\}$ is a surplus process satisfying (4.4), where interarrival times are exponential with mean $1/\vartheta$ except for the jth interarrival, Γ_j, which has a Gamma distribution. Claim amounts are all i.i.d. with distribution F_0. Therefore, using a simple change of measure argument:

$$\begin{aligned}
\mathsf{E}_0[(\tau_j^-)^2] &= \mathsf{E}_0\left[\frac{\vartheta^2 T_\theta(j)\exp\{-\vartheta T_\theta(j)\}}{\vartheta\exp\{-\vartheta T_\theta(j)\}}{\tau_\theta}^2\right] \\
&= \mathsf{E}_0\left[\vartheta T_\theta(j){\tau_\theta}^2\right] \le \sqrt{\vartheta\,\mathsf{E}\left[(T_\theta(1))^2\right]\mathsf{E}_0[{\tau_\theta}^4]},
\end{aligned}$$

for $j \ge 1$, which follows from the Cauchy−Schwartz inequality. Since $T_\theta(1)$ is an exponentially distributed random variable it follows $\mathsf{E}[(T_\theta(1))^2] < \infty$. In [24] it is shown that this surplus process is equivalent to a stable queuing $M/G/1$ model, for which $\mathsf{E}_0[Y_1^4] < \infty$ implies $\mathsf{E}_0[{\tau_\theta}^4] < \infty$. Therefore if the tilted distribution F_0 has a bounded fourth moment, then $\mathbf{H}_\mu(P, g) < \infty$ for any $g \in \mathcal{D}^b$.

The last step in obtaining the derivative of $J(\theta)$ uses an importance sampling type result put forward in Corollary 4.1 in the next section, in view of the fact that g_θ depends explicitly on θ. By calculation, $K'(\theta) = K(\theta)(1/\theta - 1/\vartheta)$, and the constant R chosen for efficient estimation is the so-called *adjustment coefficient* satisfying $\theta + cR = \theta M_y(R)$, so that $K(\theta) = 1$. Therefore Corollary 4.1 allows the application of the chain rule for differentiation, which yields:

$$\begin{aligned}
\frac{d}{d\theta}J(\theta) &= \frac{d}{d\vartheta}\left\langle g_\vartheta, \mathcal{K}_\vartheta\mu\right\rangle \\
&= \left\langle \frac{d}{d\vartheta}g_\vartheta, \mathcal{K}_\vartheta\mu\right\rangle + \left\langle g_\vartheta, \mathcal{K}_\vartheta P'_\vartheta\mathcal{K}_\vartheta\mu\right\rangle \\
&= \mathsf{E}_0\left[\left(\frac{1}{\theta} - \frac{1}{\vartheta}\right)\tau_\theta(K(\theta))^{\tau_\theta}\exp\{R(X_\theta(\tau_\theta) - u)\}\right] \\
&\quad + \frac{1}{\vartheta}\,\mathsf{E}_0\left[\sum_{j=1}^{\tau_\theta}\left(K(\theta)^\tau\exp\{R(X_\theta(\tau_\theta) - u)\}\right.\right. \\
&\quad \left.\left. - (K(\theta))^{\tau_j^-}\exp\{R(X_j^-(\tau_j^-) - u)\}\right)\right] \\
&= \mathsf{E}_0\left[\frac{\tau}{\theta}\exp\{R(X_\theta(\tau_\theta) - u)\} - \frac{1}{\vartheta}\sum_{j=1}^{\tau_\theta}\exp\{R(X_j^-(\tau_j^-) - u)\}\right].
\end{aligned}$$

It is worth noticing that the $\mathcal{D}^b$-derivative w.r.t. the rate of a marked Poisson process can be interpreted (as in the previous example) in terms of the *nominal* process and a *phantom* process where one arrival fails to have a mark: in our queueing example, the arriving "phantom" has null service time, and in the ruin example, an arrival has a null claim. This estimator is equivalent to the phantom RPA estimator given in [23] and also illustrates the case where CRN are used to define the coupling of the plus and minus processes.

4.2. Domination of measures

The analysis of derivatives of stochastic systems simplifies when the distributions involved have densities that are differentiable as functions of θ. In this section, we will illustrate how the conditions for the finiteness of the cumulative bound simplify under the presence of differentiable densities. For $P, Q \in \mathcal{K}$, let P be absolutely continuous with respect to Q. This implies that the Radon−Nikodym derivative of $P(\cdot; s)$ with respect to $Q(\cdot; s)$ exists for all s, and we denote it by $dP/dQ(r; s)$ with $r, s \in S$. If P_θ is absolutely continuous with respect to Q, then the positive and negative part of the $\mathcal{D}$-derivative of P_θ, provided they exist, are given through integrating the positive and negative parts of the derivative of $dP_\theta/dQ(r; s)$. More specifically, it holds

$$\frac{d}{d\theta}\Big(\frac{dP_\theta}{dQ}(r;s)\Big) = \frac{dP'_\theta}{dQ}(r;s) = c_{P_\theta}(s)\Big(\frac{dP^+_\theta}{dQ}(r;s) - \frac{dP^-_\theta}{dQ}(r;s)\Big),$$

for $r, s \in S$. We extend the definition of the Radon−Nikodym derivative of the halted kernel associated with τ_θ by setting $[dP_{\theta,*}/dQ](*; s) = 0$ for any $s \in S_*$.

In the presence of domination, we adapt the definition of $\mathbf{H}_\mu(P, g)$ in the following way.

Definition 4.1. For $\mathcal{D} \subset L^1(P_\theta : \Theta)$ let a r.h.e. $(g, \mu, P_\theta, \tau_\theta, \Theta)$ with $g \in \mathcal{D}$ be given. Let $(P_\theta, \mathcal{D})$ be Leibnitz on Θ and assume that P'_θ is absolutely continuous with respect to P_θ on Θ. For $g \in \mathcal{D}$ we set

$$\begin{aligned}\hat{\mathbf{H}}_\mu(P, g) &\stackrel{\text{def}}{=} \hat{\mathbf{H}}_\mu((P_\theta : \theta \in \Theta), g) \\ &= \mathsf{E}_\mu\left[\sum_{n=1}^{\|\tau_\theta\|} \|g(X_\theta(n))\| \sum_{m=1}^{\|\tau_\theta\|} \Big\|\frac{dP'_\theta}{dP_\theta}(X_\theta(m), X_\theta(m-1))\Big\|\right].\end{aligned}$$

We now present a version of Theorem 3.2 that applies under domination. The proof of the theorem can be found in Section A.3 of the Appendix.

Theorem 4.1. *For $\mathcal{D} \subset L^1(P_\theta : \Theta)$ let a r.h.e. $(g, \mu, P_\theta, \tau_\theta, \Theta)$ with $g \in \mathcal{D}$ be given. Assume that $(P_\theta, \mathcal{D})$ is Leibnitz on Θ and that P'_θ is absolutely continuous* *with respect to* P_θ*. If* $\hat{\mathbf{H}}_\mu(P,g)$ *is finite, then*

$$\frac{d}{d\theta}\ \mathsf{E}_\mu\left[\sum_{n=1}^{\tau_\theta} g(X_\theta(n))\right]$$
$$= \mathsf{E}_\mu\left[\sum_{m=1}^{\tau_\theta} c_{P_\theta}(X_\theta(m-1))\Bigg(\sum_{n=m}^{\tau_m^+} g(X_m^+(n)) - \sum_{n=m}^{\tau_m^-} g(X_m^-(n))\Bigg)\right].$$

Typically, the Radon – Nikodym derivative is locally bounded (which implies that the normalizing constant is bounded too), see [14]. In such cases, the following corollary states a simple condition on the performance that implies finiteness of the cumulative bound.

Corollary 4.1. *For* $\mathcal{D} \subset L^1(P_\theta : \Theta)$ *let a r.h.e.* $(g,\mu,P_\theta,\tau_\theta,\Theta)$ *with* $g \in \mathcal{D}$ *be given. Assume that* $(P_\theta,\mathcal{D})$ *is Leibnitz on* Θ *and that* P'_θ *is absolutely continuous with respect to* P_θ*. If*

$$\sup_{s,r\in S}\left\|\frac{dP'_\theta}{dP_\theta}(s,r)\right\| < \infty,$$

then

$$\mathsf{E}_\mu\left[\|\tau_\theta\|\sum_{m=1}^{\|\tau_\theta\|}\|g(X_\theta(m))\|\right] < \infty$$

implies that $\hat{\mathbf{H}}_\mu(P,g)$ *is finite.*

Example 4.3. Consider again the $GI/G/1$ queuing system under admission control with $\theta \in (0,1)$. In this case τ_m^+ is not stochastically smaller than τ: customers rejected in the nominal process may be accepted in the process $X_m^+(n)$. This example has been considered in [12] and it is often referred to as the "two-sided" phantom RPA method. To the authors' knowledge, no proof exists for unbiasedness of the method in the context of regeneration, unless the system at $\bar\theta = 1$ is stable, in which case stochastic ordering implies finiteness of the pth moment of $\tau_m^\pm$. It is however too restrictive to assume that the queuing system is stable for all $\theta \in [0,1]$, when only a small neighborhood of $\theta \in (0,1)$ is necessary to define the derivatives. To establish boundedness of the cumulative bound, we use now that both $P_\theta^\pm$ are dominated by P_θ, with:

$$\frac{dP_\theta^+}{dP_\theta}(t,z,\eta) = \begin{cases} \dfrac{1}{\theta}, & \text{if } \eta = 1, \\ 0, & \text{if } \eta = 0, \end{cases} \qquad \frac{dP_\theta^-}{dP_\theta}(t,z,\eta) = \begin{cases} 0, & \text{if } \eta = 1, \\ \dfrac{1}{1-\theta}, & \text{if } \eta = 0. \end{cases}$$

Using the notation previously introduced in Example 4.1, we obtain the following conditions for unbiasedness of the RPA estimator, see (4.3). For $\Theta = [a,\theta]$ a subset of $(0,1)$, Corollary 4.1 establishes the desired result: $\mathsf{E}_\mu[\tau_\theta^2(S_\tau(\theta))^p] < \infty$

is a sufficient condition for the RPA gradient estimator to be unbiased for any g satisfying condition (4.1). This thinning of an arrival process of a G/G/1 queuing system has been studied before by Brémaud and Gong, see [2]. However, written in our notation, they require $\mathsf{E}_\mu[\tau_\theta^3(S_\tau(\theta))^p] < \infty$ for the RPA gradient estimator to be unbiased.

Remark 4.2. Domination allows for interpreting integration with respect to P'_θ by integration with respect to P_θ provided that the integrand is rescaled by dP'_θ/dP_θ. This approach is called *score function method* in the literature, see for example [20,21]. Specifically, if the conditions in Theorem 4.1 and Corollary 4.1 are satisfied, we obtain the following alternative representation for the derivative in Theorem 4.1:

$$\frac{d}{d\theta}\,\mathsf{E}_\mu\left[\sum_{n=1}^{\tau_\theta} g(X_\theta(n))\right] = \mathsf{E}_\mu\left[\left(\sum_{n=1}^{\tau_\theta} g(X_\theta(n))\right)\left(\sum_{m=1}^{\tau_\theta}\frac{dP'_\theta}{dP_\theta}(X_\theta(m), X_\theta(m-1))\right)\right].$$

5. Gradient estimation for stationary Markov chains

In this section, we address the gradient estimation problem for stationary performance characteristics. Specifically, we use our results to derive a new unbiased gradient estimator for stationary Markov chains. Let $\{X_\theta(n)\}$ be a Harris ergodic Markov chain with atom α, such that α does not depend on θ, see [16] for more details. We will suppress the dependence of the random variables on θ whenever this causes no confusion. Harris ergodicity implies the existence of a unique stationary distribution, which will be denoted by π_θ. For any g absolutely integrable with respect to π_θ, it holds that

$$\langle g, \pi_\theta\rangle = \frac{\mathsf{E}\left[\sum_{n=0}^{\tau_\theta(\alpha)-1} g(X_\theta(n))\right]}{\mathsf{E}\left[\tau_\theta(\alpha)\right]},$$

with $X_\theta(0) \in \alpha$ and $\tau_\theta(\alpha)$ the recurrence time to α. Our analysis so far allows for deriving an unbiased estimator for

$$\mathsf{E}\left[\sum_{n=0}^{\tau_\theta(\alpha)-1} g(X_\theta(n))\right]. \tag{5.1}$$

Note that with $g \equiv 1$ we obtain from this also an unbiased estimator for $\mathsf{E}\left[\tau_\theta(\alpha)\right]$. Incorporating the hitting time into the Markov kernel, we obtain the kernel $P_{\theta,\alpha}$, defined by

$$P_{\theta,\alpha}(A; s) = \begin{cases} P_\theta(A; s), & s \notin \alpha, \\ \delta_*(A), & s \in \alpha \cup \{*\}, \end{cases}$$

for $A \in \mathcal{S}_*$. With the help of $P_{\theta,\alpha}$, the expression in (5.1) reads for any $g \in \mathcal{D}$ absolutely integrable with respect to π_θ

$$\mathsf{E}\left[\sum_{n=0}^{\tau_\theta(\alpha)-1} g(X_\theta(n))\right] = \big\langle g, K_{\theta,\alpha}\mu_\alpha \big\rangle, \tag{5.2}$$

where $K_{\theta,\alpha}$ denotes the potential kernel of $P_{\theta,\alpha}$ and the initial distribution μ_α is a recurrence measure of $\{X_\theta(n)\}$. Note that g is integrable with respect to π if $\langle g, K_{\theta,\alpha}\mu_\alpha\rangle$ is finite. Let $(P_\theta, \mathcal{D})$ be Leibnitz on Θ, then it follows from Lemma 3.1 that $(P_{\theta,\alpha}, \mathcal{D})$ is Leibnitz on Θ. Provided that

$$\langle g, \pi_\theta \rangle < \infty \quad \text{and} \quad \mathbf{H}_{\mu_\alpha}(P_\alpha, g) < \infty,$$

for any $g \in \mathcal{D}$, we obtain from Theorem 3.2 that

$$\frac{d}{d\theta}\big\langle g, K_{\theta,\alpha}\mu_\alpha \big\rangle = \big\langle g, K_{\theta,\alpha}P'_{\theta,\alpha}K_{\theta,\alpha}\mu_\alpha \big\rangle. \tag{5.3}$$

The above formula can be facilitated for estimation in the following way. Let $X^{\pm}_{m,\alpha}(n)$ evolve according to the kernel

$$P^{n-m}_{\theta,\alpha}P^{\pm}_{\theta,\alpha}P^{m-1}_{\theta,\alpha},$$

or, equivalently, for $m \neq n$, let the transition kernel of $X^{\pm}_{m,\alpha}(n)$ be $P_{\theta,\alpha}$, whereas the transition kernel of $X^{+}_{m,\alpha}(m)$ is $P^{+}_{\theta,\alpha}$ and that for $X^{-}_{m,\alpha}(m)$ is $P^{-}_{\theta,\alpha}$, respectively. Without loss of generality, we assume that the processes $\{X^{\pm}_{m,\alpha}(n)\}$ are constructed from $P_{\theta,\alpha}$ using common random numbers. Denote the first entrance times of $X^{\pm}_{m,\alpha}(n)$ into α by $\tau^{\pm}_m(\alpha)$. With this notation, the expression on the right-hand side of equation (5.3) can be written as

$$\mathsf{E}_{\mu_\alpha}\left[\sum_{m=0}^{\tau_\theta(\alpha)-1} c_{P_\theta}(X_\theta(m-1))\left(\sum_{n=m}^{\tau^{+}_m(\alpha)-1} g(X^{+}_{m,\alpha}(n)) - \sum_{n=m}^{\tau^{-}_m(\alpha)-1} g(X^{-}_{m,\alpha}(n))\right)\right].$$

Thus, we obtain the following overall estimator

$$\begin{aligned}\frac{d}{d\theta}\langle g, \pi_\theta\rangle = {} & \frac{1}{\mathsf{E}_{\mu_\alpha}[\tau_\theta(\alpha)]}\,\mathsf{E}_{\mu_\alpha}\left[\sum_{m=1}^{\tau_\theta(\alpha)-1} c_{P_\theta}(X_\theta(m-1))\right.\\ & \left.\times\left(\sum_{n=m}^{\tau^{+}_m(\alpha)-1} g(X^{+}_{m,\alpha}(n)) - \sum_{n=m}^{\tau^{-}_m(\alpha)-1} g(X^{-}_{m,\alpha}(n))\right)\right]\\ & - \frac{1}{\mathsf{E}^2_{\mu_\alpha}[\tau_\theta(\alpha)]}\,\mathsf{E}_{\mu_\alpha}\left[\sum_{m=1}^{\tau_\theta(\alpha)-1} c_{P_\theta}(X_\theta(m-1))\big(\tau^{+}_m(\alpha) - \tau^{-}_m(\alpha)\big)\right]\\ & \times \mathsf{E}_{\mu_\alpha}\left[\sum_{n=1}^{\tau(\alpha)-1} g(X_\theta(n))\right].\end{aligned}$$

Hence, the crucial condition for the above analysis to hold is that the cumulative bound $\mathbf{H}_{\mu_\alpha}(P_\alpha, g)$ is finite for any $g \in \mathcal{D}$, and we refer to Section 4 for sufficient conditions for this. If we replace the stopping times $\tau_\theta(\alpha)$, $\tau_m^+(\alpha)$ and $\tau_m^+(\alpha)$ by a fixed time horizon, say, N, then we recover the (biased) estimator presented by Dai in [7]. Specifically, Dai establishes a central limit theorem (for N towards ∞), and our result can be seen as the sample-path counterpart of this result. With respect to applications, our result is of course stronger since we construct stopping times $\tau_m^\pm(\alpha)$ that produce an unbiased estimator. Furthermore, the analysis in [7] is restricted to bounded cost functions.

Remark 5.1. Under the conditions in Theorem 3.2 and Corollary 4.1, we obtain the following score function estimator for the derivative of the stationary performance

$$\begin{aligned}\frac{d}{d\theta}\langle g, \pi_\theta\rangle = {} & \frac{1}{\mathsf{E}_{\mu_\alpha}[\tau_\theta(\alpha)]}\,\mathsf{E}_{\mu_\alpha}\Bigg[\Bigg(\sum_{n=1}^{\tau_\theta(\alpha)-1} g(X_\theta(n))\Bigg)\\ & \times\Bigg(\sum_{m=1}^{\tau_\theta(\alpha)} \frac{dP'_\theta}{dP_\theta}(X_\theta(m), X_\theta(m-1))\Bigg)\Bigg]\\ & - \frac{1}{\mathsf{E}^2_{\mu_\alpha}[\tau_\theta(\alpha)]}\,\mathsf{E}_{\mu_\alpha}\Bigg[\tau_\theta(\alpha)\sum_{m=1}^{\tau_\theta(\alpha)} \frac{dP'_\theta}{dP_\theta}(X_\theta(m), X_\theta(m-1))\Bigg]\\ & \times \mathsf{E}_{\mu_\alpha}\Bigg[\sum_{n=1}^{\tau_\theta(\alpha)-1} g(X_\theta(n))\Bigg],\end{aligned}$$

for any $g \in \mathcal{D}$. Note that the above estimator is different from the score function estimator established in [11].

Appendix

A.1. Proof of Theorem 3.1

For $N \in \mathbf{N}$, since $(P_\theta, \mathcal{D})$ is Leibnitz, Theorem 1 in [14] yields

$$\sum_{n=1}^{N} \frac{d}{d\theta}\langle g, P_\theta^n \mu\rangle = \sum_{n=1}^{N}\sum_{m=1}^{n} \big\langle g, P_\theta^{n-m} P'_\theta P_\theta^{m-1}\mu\big\rangle. \tag{A.1}$$

Moreover, the first part of Theorem 1 in [14] yields $\mathcal{D}$-Lipschitz continuity on Θ of the expression in (A.1) and $\sum_{n=1}^{N}\langle g, P_\theta^n\mu\rangle$ is thus continuously differentiable on Θ for any $N \in \mathbf{N}$.

Since $H_\mu(P_\theta, g)$ is finite, the limit in (A.1) as $N \to \infty$ exists:

$$\sum_{n=1}^{\infty} \frac{d}{d\theta}\langle g, P_\theta^n \mu\rangle = \sum_{n=1}^{\infty}\sum_{m=1}^{n} \big\langle g, P_\theta^{n-m} P'_\theta P_\theta^{m-1}\mu\big\rangle \tag{A.2}$$

for any $\theta \in \Theta$. By calculation,

$$\lim_{N\to\infty} \left\| \sum_{n=1}^{N} \frac{d}{d\theta} \langle g, P^n \mu \rangle - \sum_{n=1}^{\infty} \frac{d}{d\theta} \langle g, P^n \mu \rangle \right\| \leq \lim_{N\to\infty} \left\| \sum_{n=N+1}^{\infty} \frac{d}{d\theta} \langle g, P^n \mu \rangle \right\|$$
$$\leq \lim_{N\to\infty} \sum_{n=N+1}^{\infty} \left\| \frac{d}{d\theta} \langle g, P^n \mu \rangle \right\| = 0,$$

where the last equality follows from the fact that $H_\mu(P, g)$ is finite. Hence, $\left(\sum_{n=1}^{N} \langle g, P_\theta^n \mu \rangle \right)'$ converges uniformly on Θ. Interchanging differentiation and summation on the left-hand side of (A.2) is thus allowed, which concludes the proof of the first part of the theorem.

For the proof of the second part of theorem, note that finiteness of $H_\mu(P, g)$ implies that the expression in (A.2) converges absolutely. Rearranging terms concludes the proof of the theorem.

A.2. Proof of Theorem 3.2

Let P_θ be $\mathcal{D}$-differentiable with $\mathcal{D}$-derivative $(c_{P_\theta}, P_\theta^+, P_\theta^-)$. For $g \in \mathcal{D}$, we set

$$\mathbf{D}_\mu(P_\theta, g) = \mathsf{E}_\mu \Big[\sum_{m=1}^{\tau_\theta} c_{P_\theta}(X_\theta(m-1)) \Big(\sum_{n=1}^{\tau_m^+} \left| g(X_m^+(n)) \right| + \sum_{n=1}^{\tau_m^-} \left| g(X_m^-(n)) \right| \Big) \Big].$$

Lemma A.1. *For $\mathcal{D} \subset L^1(P_\theta : \Theta)$, let a r.h.e. $(g, \mu, P_\theta, \tau_\theta, \Theta)$ with $g \in \mathcal{D}$ be given. Assume that $(P_\theta, \mathcal{D})$ is Leibnitz on Θ. If $\mathbf{D}_\mu(P_\theta, g)$ is finite, then*

$$\sum_{n=1}^{\infty} \sum_{m=1}^{n} \left\langle g, P_{\theta,\tau}^{n-m} P'_{\theta,\tau} P_{\theta,\tau}^{m-1} \mu \right\rangle$$
$$= \mathsf{E}_\mu \Big[\sum_{m=1}^{\tau_\theta} c_{P_\theta}(X_\theta(m-1)) \Big(\sum_{n=1}^{\tau_m^+} g(X_m^+(n)) - \sum_{n=1}^{\tau_m^-} g(X_m^-(n)) \Big) \Big].$$

Proof. To simplify the notation, we suppress the subscript θ for random variables. Observe that

$$\mathsf{E}_\mu \Big[\sum_{m=1}^{\tau} c_{P_\theta}(X(m-1)) \Big(\sum_{n=1}^{\tau_m^+} g(X_m^+(n)) - \sum_{n=1}^{\tau_m^-} g(X_m^-(n)) \Big) \Big]$$
$$= \mathsf{E}_\mu \Big[\sum_{m=1}^{\infty} \mathbf{1}\{\tau \geq m\} c_{P_\theta}(X(m-1))$$
$$\times \sum_{n=1}^{\infty} \left(\mathbf{1}\left\{\tau_m^+ \geq n\right\} g(X_m^+(n)) - \mathbf{1}\left\{\tau_n^- \geq m\right\} g(X_m^-(n)) \right) \Big],$$

finiteness of $\mathbf{D}_\mu(P_\theta, g)$ implies that the sums on the right-hand side of the above equation are absolutely integrable, which yields the right-hand side is equal to

$$\begin{aligned}
&\sum_{m=1}^{\infty}\sum_{n=1}^{\infty} \mathsf{E}_\mu\left[\mathbf{1}\{\tau \geq m\} c_{P_\theta}(X(m-1))\right.\\
&\qquad \left.\times \left(\mathbf{1}\left\{\tau_m^+ \geq n\right\} g(X_m^+(n)) - \mathbf{1}\left\{\tau_m^- \geq n\right\} g(X_m^-(n))\right)\right]\\
&\quad = \sum_{n=1}^{\infty}\sum_{m=1}^{n} \mathsf{E}_\mu\left[\mathbf{1}\{\tau \geq m\} c_{P_\theta}(X(m-1))\right.\\
&\qquad \left.\times \left(\mathbf{1}\left\{\tau_m^+ \geq n\right\} g(X_m^+(n)) - \mathbf{1}\left\{\tau_m^- \geq n\right\} g(X_m^-(n))\right)\right].
\end{aligned}$$

Switching from the process $\{X_\theta(n)\}$ to the halted version, see (2.2), and using (3.2) yields the right-hand side is equal to

$$\begin{aligned}
&\sum_{n=1}^{\infty}\sum_{m=1}^{n} \mathsf{E}_\mu\left[\mathbf{1}\{\tau \geq m\} c_{P_{\theta,\tau}}(X^\tau(m-1))\right.\\
&\qquad\qquad \left.\times \left(\mathbf{1}\left\{\tau_m^+ \geq n\right\} g(X_m^+(n)) - \mathbf{1}\left\{\tau_m^- \geq n\right\} g(X_m^-(n))\right)\right]\\
&\qquad = \sum_{n=1}^{\infty}\sum_{m=1}^{n} \left\langle g, P_{\theta,\tau}^n P_{\theta,\tau}' P_{\theta,\tau}^{m-1}\right\rangle,
\end{aligned}$$

where the last equality follows from Lemma 3.1 (which in particular implies that $c_{P_{\theta,\tau}}(*) = 0$). □

Lemma A.2. *For $\mathcal{D} \subset L^1(P_\theta : \Theta)$, let a r.h.e. $(g, \mu, P_\theta, \tau_\theta, \Theta)$ with $g \in \mathcal{D}$ be given. Assume that $(P_\theta, \mathcal{D})$ is Leibnitz on Θ. For any $\theta \in \Theta$ it holds that*

$$\mathbf{H}_\mu(P, g) \geq H_\mu(P_\tau, g) \quad \text{and} \quad \mathbf{H}_\mu(P, g) \geq \mathbf{D}_\mu(P_\theta, g).$$

Proof. We only prove that $\mathbf{H}_\mu(P, g) \geq H_\mu(P, g)$, since the other relation is obvious. To simplify the notation, we suppress the subscript θ for random variables. For any $g \in \mathcal{D}$ it holds that

$$\begin{aligned}
\mathbf{H}_\mu(P_\tau, g) &= \mathsf{E}_\mu\left[\sum_{m=1}^{\|\tau\|} \|c_{P_\theta}(X(m-1))\| \left(\sum_{n=1}^{\|\tau_m^+\|} \|g(X_m^+(n))\| + \sum_{n=1}^{\|\tau_m^-\|} \|g(X_m^-(n))\|\right)\right]\\
&= \mathsf{E}_\mu\left[\sum_{m=1}^{\infty} \mathbf{1}\{\|\tau\| \geq m\} \|c_{P_\theta}(X(m-1))\|\right.\\
&\quad \left.\times \sum_{n=1}^{\infty}\left(\mathbf{1}\left\{\|\tau_m^+\| \geq n\right\}\|g(X_m^+(n))\| + \mathbf{1}\left\{\|\tau_m^-\| \geq n\right\}\|g(X_m^-(n))\|\right)\right]\\
&= \sum_{n=1}^{\infty}\sum_{m=1}^{\infty} \mathsf{E}_\mu\left[\mathbf{1}\{\|\tau\| \geq m\} \|c_{P_\theta}(X(m-1))\|\right.
\end{aligned}$$

$$
\begin{aligned}
&\quad \times \big(\mathbf{1}\big\{\|\tau_m^+\| \geq n\big\}\|g(X_m^+(n))\| + \mathbf{1}\big\{\|\tau_m^-\| \geq n\big\}\|g(X_m^-(n))\|\big)\big]\\
&\geq \sum_{n=1}^{\infty}\sum_{m=1}^{\infty} \big\| \, \mathsf{E}_\mu \left[\mathbf{1}\{\tau \geq m\} c_{P_\theta}(X(m-1))\right.\\
&\quad \times \big(\mathbf{1}\big\{\tau_m^+ \geq n\big\}|g(X_m^+(n))| + \mathbf{1}\big\{\tau_m^- \geq n\big\}|g(X_m^-(n))|\big)\big]\big\|.
\end{aligned}
$$

Switching from the process $\{X(n)\}$ to the halted versions, see (2.2), and using (3.2) yields the right-hand side is equal to

$$
\begin{aligned}
&\sum_{n=1}^{\infty}\sum_{m=1}^{\infty} \big\| \, \mathsf{E}_\mu \left[\mathbf{1}\{\tau \geq m\} c_{P_{\theta,*}}(X^\tau(m-1))\right.\\
&\qquad \times \big(\mathbf{1}\big\{\tau_m^+ \geq n\big\}|g(X_m^+(n))| + \mathbf{1}\big\{\tau_m^- \geq n\big\}|g(X_m^-(n))|\big)\big]\big\|\\
&\quad = \sum_{n=0}^{\infty}\sum_{m=0}^{\infty} \big\|\big\langle |g|, P_{\theta,\tau}^n |P'_{\theta,\tau}| P_{\theta,\tau}^m \mu\big\rangle\big\|\\
&\quad = \sum_{n=1}^{\infty}\sum_{m=1}^{n} \big\|\big\langle |g|, P_{\theta,\tau}^{n-m} |P'_{\theta,\tau}| P_{\theta,\tau}^{m-1} \mu\big\rangle\big\|\\
&\quad \geq \sum_{n=1}^{\infty} \left\| \sum_{m=1}^{n} \big\langle g, P_{\theta,\tau}^{n-m} P'_{\theta,\tau} P_{\theta,\tau}^{m-1} \mu\big\rangle \right\|\\
&\quad = H_\mu(P_\tau, g),
\end{aligned}
$$

where the last equality follows from the product rule of $\mathcal{D}$-differentiation, see Theorem 1 in [14] □

We turn to the proof of the theorem. We have assumed that $(P_\theta, \mathcal{D})$ is Leibnitz on Θ and thus $(P_{\theta,\tau}, \mathcal{D})$ is Leibnitz on Θ, see Lemma 3.1. By Lemma A.2, $\mathbf{H}_\mu(P) < \infty$ implies $H_\mu(P_\tau) < \infty$. Hence, we may apply Theorem 3.1 to $P_{\theta,\tau}$. Since we have assumed that the r.h.e. is finite, is holds that

$$
\langle g, K_\theta \mu\rangle = \mathsf{E}_\mu \left[\sum_{n=1}^{\tau_\theta} g(X_\theta(n))\right].
$$

Under condition of the theorem, Lemma A.1 yields

$$
\begin{aligned}
&\sum_{n=1}^{\infty}\sum_{m=1}^{n} \big\langle g, P_\theta^{n-m} P'_\theta P_\theta^{m-1}\mu\big\rangle\\
&\quad = \mathsf{E}_\mu \left[\sum_{m=1}^{\tau_\theta} c_{P_\theta}(X_\theta(m-1)) \left(\sum_{n=1}^{\tau_m^+} g(X_m^+(n)) - \sum_{n=1}^{\tau_m^-} g(X_m^-(n))\right)\right]\\
&\quad = \mathsf{E}_\mu \left[\sum_{m=1}^{\tau_\theta} c_{P_\theta}(X_\theta(m-1)) \left(\sum_{n=m}^{\tau_m^+} g(X_m^+(n)) - \sum_{n=m}^{\tau_m^-} g(X_m^-(n))\right)\right],
\end{aligned}
$$

where we use Lemma A.2 to show that $\mathbf{D}_\mu(P_\theta, g)$ is finite, and the last equality follows from the applied coupling scheme, see Remark 3.2. Hence, the proof follows directly from applying Theorem 3.1.

A.3. Proof of Theorem 4.1

Before we can turn to the proof of the theorem, we establish two technical results needed for the proof.

Lemma A.3. *Let $\mathcal{D} \subset L^1(P_\theta : \Theta)$ and a r.h.e. $(g, \mu, P_\theta, \tau_\theta, \Theta)$ with $g \in \mathcal{D}$ be given. Assume that $(P_\theta, \mathcal{D})$ is Leibnitz on Θ and that P'_θ is absolutely continuous with respect to P_θ. Then,*

$$\mathbf{D}_\mu(P_{\theta,\tau}, g) = \mathsf{E}_\mu \left[\sum_{n=1}^{\tau_\theta} |g(X_\theta(n))| \sum_{m=1}^{\tau_\theta} \left| \frac{dP'_\theta}{dP_\theta}(X(m), X(m-1)) \right| \right].$$

Proof. To simplify the notation, we suppress the subscript θ for random variables. Direct calculation yields

$$\begin{aligned}
\mathbf{D}_\mu(P_\tau, g) &= \mathsf{E}_\mu \left[\sum_{m=1}^{\tau} c_{P_\theta}(X(m-1)) \left(\sum_{n=1}^{\tau_m^+} |g(X_m^+(n))| + \sum_{n=1}^{\tau_m^-} |g(X_m^-(n))| \right) \right] \\
&= \mathsf{E}_\mu \Bigg[\sum_{m=1}^{\infty} \mathbf{1}\{\tau \geq m\} c_{P_\theta}(X(m-1)) \\
&\qquad \times \sum_{n=1}^{\infty} \left(\mathbf{1}\{\tau_m^+ \geq n\} |g(X_m^+(n))| + \mathbf{1}\{\tau_m^- \geq n\} |g(X_m^-(n))| \right) \Bigg].
\end{aligned}$$

Switching from $\{X_\theta(n)\}$ to the halted version, see (2.2), yields the right-hand side is equal to

$$\begin{aligned}
&\mathsf{E}_\mu \Bigg[\sum_{m=1}^{\infty} \mathbf{1}\{\tau \geq m\} c_{P_{\theta,*}}(X^\tau(m-1)) \\
&\qquad \times \sum_{n=1}^{\infty} \left(\mathbf{1}\{\tau_m^+ \geq n\} |g(X_m^+(n))| + \mathbf{1}\{\tau_m^- \geq n\} |g(X_m^-(n))| \right) \Bigg] \\
&= \sum_{m=1}^{\infty} \sum_{n=0}^{\infty} \left\langle |g|, P_{\theta,\tau}^n |P'_{\theta,\tau}| P_{\theta,\tau}^{m-1} \mu \right\rangle \\
&= \sum_{m=1}^{\infty} \sum_{n=0}^{\infty} \left\langle |g| \left| \frac{dP'_{\theta,\tau}}{dP_{\theta,*}} \right|, P_{\theta,\tau}^n P_{\theta,\tau} P_{\theta,\tau}^{m-1} \mu \right\rangle \\
&= \mathsf{E}_\mu \left[\sum_{n=1}^{\infty} \sum_{m=1}^{\infty} \mathbf{1}\{\tau \geq n\} |g(X^\tau(n))| \mathbf{1}\{\tau \geq m\} \left| \frac{dP'_{\theta,*}}{dP_{\theta,\tau}}(X^\tau(m), X^\tau(m-1)) \right| \right]
\end{aligned}$$

$$= \mathsf{E}_\mu \left[\sum_{n=1}^{\tau} |g(X(n))| \sum_{m=1}^{\tau} \left| \frac{dP'_\theta}{dP_\theta}(X(m), X(m-1)) \right| \right],$$

where the last equality follows from (2.2) together with Lemma 3.1. □

Lemma A.4. *Let $\mathcal{D} \subset L^1(P_\theta : \Theta)$ and a r.h.e. $(g, \mu, P_\theta, \tau_\theta, \Theta)$ with $g \in \mathcal{D}$ be given. Assume that $(P_\theta, \mathcal{D})$ is Leibnitz on Θ and that P'_θ is absolutely continuous with respect to P_θ. For any $\theta \in \Theta$ it holds*

$$\hat{\mathbf{H}}_\mu(P, g) \geq H_\mu(P_\tau, g) \quad \text{and} \quad \hat{\mathbf{H}}_\mu(P, g) \geq \mathbf{D}_\mu(P_\theta, g).$$

Proof. We only prove that $\hat{\mathbf{H}}_\mu(P_\tau, g) \geq H_\mu(P_\tau, g)$, since the other relation is obvious (for $\hat{\mathbf{H}}_\mu(P, g) \geq \mathbf{D}_\mu(P_\theta, g)$ use Lemma A.3). To simplify the notation, we suppress the subscript θ for random variables. For any $g \in \mathcal{D}$ it holds that

$$\begin{aligned}
H_\mu(P_\tau, g) &= \sum_{n=1}^{\infty} \left\| \sum_{m=1}^{n} \langle g, P_{\theta,\tau}^{n-m} P'_{\theta,\tau} P_{\theta,\tau}^{m-1} \mu \rangle \right\| \\
&\leq \sum_{n=1}^{\infty} \sum_{m=1}^{n} \| \langle g, P_{\theta,\tau}^{n-m} P'_{\theta,\tau} P_{\theta,\tau}^{m-1} \mu \rangle \| \\
&\leq \mathsf{E}_\mu \Bigg[\sum_{n=1}^{\infty} \sum_{m=1}^{n} \mathbf{1}\{\|\tau\| \geq n\} \|g(X^\tau(n))\| \mathbf{1}\{\|\tau\| \geq m\} \\
&\qquad \times \left\| \frac{dP'_{\theta,*}}{dP_{\theta,*}}(X^\tau(m), X^\tau(m-1)) \right\| \Bigg].
\end{aligned}$$

Switching from the halted process $\{X^\tau(n)\}$ to the original version, see (2.2), and using Lemma 3.1 yields the right-hand side is equal to

$$\begin{aligned}
&\mathsf{E}_\mu \Bigg[\sum_{n=1}^{\infty} \sum_{m=1}^{n} \mathbf{1}\{\|\tau\| \geq n\} \|g(X(n))\| \mathbf{1}\{\|\tau\| \geq m\} \left\| \frac{dP'_\theta}{dP_\theta}(X(m), X(m-1)) \right\| \Bigg] \\
&= \mathsf{E}_\mu \left[\sum_{n=1}^{\|\tau\|} \|g(X(n))\| \sum_{m=1}^{\|\tau\|} \left\| \frac{dP'_\theta}{dP_\theta}(X(m), X(m-1)) \right\| \right].
\end{aligned}$$

□

We now turn to the proof of Theorem 4.1. We have assumed that $(P_\theta, \mathcal{D})$ is Leibnitz on Θ and thus $(P_{\theta,\tau}, \mathcal{D})$ is Leibnitz on Θ, see Lemma 3.1. By Lemma A.4, $\hat{\mathbf{H}}_\mu(P) < \infty$ implies $H_\mu(P_\tau) < \infty$. Hence, we may apply Theorem 3.1 to $P_{\theta,\tau}$. Since we have assumed that the r.h.e. is finite is holds that

$$\langle g, K_\theta \mu \rangle = \mathsf{E}_\mu \left[\sum_{n=1}^{\tau_\theta} g(X_\theta(n)) \right].$$

By Lemma A.4, $\hat{\mathbf{H}}_\mu(P) < \infty$ implies $\mathbf{D}_\mu(P_\tau) < \infty$, and under conditions in Theorem 4.1, Lemma A.1 yields

$$\sum_{n=1}^{\infty} \sum_{m=1}^{n} \langle g, P_\theta^{n-m} P'_\theta P_\theta^{m-1} \mu \rangle$$
$$= \mathsf{E}_\mu \left[\sum_{m=1}^{\tau_\theta} c_{P_\theta}(X_\theta(m-1)) \left(\sum_{n=1}^{\tau_m^+} g(X_m^+(n)) - \sum_{n=1}^{\tau_m^-} g(X_m^-(n)) \right) \right].$$

Hence, the proof follows directly from applying Theorem 3.1.

Acknowledgments

The work of the first author was done while he was with EURANDOM, P.O.Box 513, 5600 MB Eindhoven, the Netherlands, where he was supported by Deutsche Forschungsgemeinschaft under grant He 3139/1-1. The work of the second author was partially supported by NSERC grants of the Government of Canada. We wish to express our gratitude to EURANDOM and GERAD, who sponsored the second author's visit to Eindhoven and the first author's visit to Montreal, respectively.

References

[1] S. Asmussen and H.M. Nielsen (1995) Ruin probabilities via local adjustment coefficients. *J. Appl. Probab.* **32**, 736–755.

[2] P. Brémaud and W. Gong (1993) Derivatives of likelihood rations and smoothed perturbation analysis for the routing problem. *ACM Trans. Modelling and Computer Simulation* **3**, 134–161.

[3] P. Brémaud and L. Massoulié (1995) Maximal coulping rare perturbation analysis with a random horizon. *Discrete Event Systems: Theory and Applications* **5**, 319–342.

[4] P. Brémaud and F. Vázquez-Abad (1992) On the pathwise computation of derivatives with respect to the rate of a point process: The phantom RPA method. *Queueing Systems* **10**, 249–270.

[5] C. Cassandras and S. Lafortune (1999) *Introduction to Discrete Event Systems.* Kluwer Academic, Boston.

[6] X. Cao (1998) The relation among potentials, perturbation analysis, and Markov decision processes. *Discrete Event Systems: Theory and Applications* **8**, 71–87.

[7] L. Dai (2000) Perturbation analysis via coupling. *IEEE Trans. Automatic Control* **45**, 614–628.

[8] M. Fu and J.-Q. Hu (1997) *Conditional Monte Carlo.* Kluwer Academic, Boston.

[9] P. Glasserman (1991) *Gradient Estimation via Perturbation Analysis.* Kluwer Academic Publishers, Boston.

[10] P. Glasserman (1994) Perturbation analysis of production networks. In: *Stochastic Modeling and Analysis of Manufacturing Systems*, D. Yao (ed.), Springer Series in Operations Research, Springer, New York, Chapter 6, 233–280.

[11] P. Glynn and P. L'Ecuyer (1995) Likelihood ratio gradient estimation for stochastic recursions. *Adv. Appl. Probab.* **27**, 1019–1053.

[12] W.-B. Gong and Y.-C. Ho (1990) Smoothed (conditional) perturbation analysis of discrete-event dynamical systems. *IEEE Trans. Automatic Control* **39**, 1218–1230.

[13] B. Heidergott, A. Hordijk and H. Weisshaupt (2003) Derivatives of Markov kernels and their Jordan decomposition. EURANDOM Report 2003-001.

[14] B. Heidergott and F. Vázquez-Abad (2006) Measure-valued differentiation for Markov chains. To appear in *JOTA*.

[15] H. Kushner and G. Yin (1997) *Stochastic Approximation and Applications.* Springer Verlag, New York.

[16] S. Meyn and R. Tweedie (1993) *Markov Chains and Stochastic Stability.* Springer, London.

[17] E. Nummelin (1984) *General Irreducible Markov Chains and Non-negative Operators.* Cambridge University Press, Cambridge.

[18] G. Pflug (1992) Gradient estimates for the performance of Markov chains and discrete event processes. *Ann. Operations Research* **39**, 173-194.

[19] G. Pflug (1996) *Optimisation of Stochastic Models.* Kluwer Academic, Boston.

[20] M. Reiman and A. Weiss (1989) Sensitivity analysis for simulations via likelihood ratios. *Operations Research* **37**, 830-844.

[21] R. Rubinstein and A. Shapiro (1993) *Discrete Event Systems: Sensitivity Analysis and Optimization by the Score Function Method.* Wiley.

[22] F. Vázquez-Abad and K. Davis (1995) Efficient implementation of the phantom RPA method, with an application to a priority queueing system. Rapport technique No. 989, Département d'iro, Université de Montréal, August 1995.

[23] F. Vázquez-Abad (2000) RPA pathwise derivative estimation of ruin probabilities. *Insurance Mathematics and Economics* **26**, 269–288.

[24] F. Vázquez-Abad and P. LeQuoc (2001) Sensitivity analysis for ruin probabilities. *Journal of the Operational Research Society* **52**, 71–81.