VU Research Portal

Measure-valued differentiation for stationary Markov chains

Heidergott, B.F.; Hordijk, A.; Weisshaupt, H.

published in

Mathematics of Operations Research
2006

DOI (link to publisher)

[10.1287/moor.1050.0171](https://doi.org/10.1287/moor.1050.0171)

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

Heidergott, B. F., Hordijk, A., & Weisshaupt, H. (2006). Measure-valued differentiation for stationary Markov chains. *Mathematics of Operations Research*, 31, 154-172. <https://doi.org/10.1287/moor.1050.0171>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

Measure-Valued Differentiation for Stationary Markov Chains

Bernd Heidergott

Vrije Universiteit and Tinbergen Institute, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands,
bheidergott@feweb.vu.nl

Arie Hordijk

Mathematical Institute, Leiden University, P.O.Box 9512, 2300 RA Leiden, The Netherlands,
hordijk@math.leidenuniv.nl

Heinz Weisshaupt

Department of Statistics, University of Vienna, Universitaetsstrasse 5/3, A-1010 Vienna, Austria,
heinz.weisshaupt@univie.ac.at

We study general state-space Markov chains that depend on a parameter, say, θ . Sufficient conditions are established for the stationary performance of such a Markov chain to be differentiable with respect to θ . Specifically, we study the case of unbounded performance functions and thereby extend the result on weak differentiability of stationary distributions of Markov chains to unbounded mappings. First, a closed-form formula for the derivative of the stationary performance of a general state-space Markov chain is given using an operator-theoretic approach. In a second step, we translate the derivative formula into unbiased gradient estimators. Specifically, we establish phantom-type estimators and score function estimators. We illustrate our results with examples from queueing theory.

Key words: measure-valued differentiation; perturbation analysis; normed ergodicity; Markov chains; score function

MSC2000 subject classification: Primary: 52M05; secondary: 78M50

OR/MS subject classification: Primary: probability/markov processes; secondary: simulation

History: Received August 22, 2002; revised June 23, 2003, and December 16, 2004.

1. Introduction. In recent years a great deal of attention has been devoted to the computation of derivatives of performance indicators in stochastic systems. More specifically, suppose that the system can be modelled by a (general state-space) Markov chain $\{X_\theta(n)\}$, depending on a (vector-valued) parameter $\theta \in \Theta$, and assume that the process is ergodic for any $\theta \in \Theta$, that is, $X_\theta(n)$ converges, independent of the initial state, to a steady-state $X_\theta(\infty)$. We would like to compute the gradient of the expected value of a performance function g of the Markov chain in equilibrium, that is, $\nabla \mathbb{E}[g(X_\theta(\infty))]$. A typical example is the GI/G/1 queue, where the distribution of the service times or of the interarrival times depends on a parameter, such as the mean. We may be interested in computing the sensitivity of the expected stationary waiting time $\mathbb{E}[g(W_\theta)]$ with respect to the parameter θ . Moreover, the computation of derivatives allows one to take an additional step and develop optimization procedures for the performance indicator of interest.

The problem dealt with in this paper has a long history. Based on a weak differentiation approach, Pflug derives in [27, 28] for Harris ergodic Markov chains with an atom an asymptotically unbiased estimator for derivatives of continuous bounded stationary performance functions. Kushner and Vázquez-Abad use in [25] a weak differentiation approach for establishing differentiability of the stationary distribution, assuming that the driving Markov kernel has a differentiable density. Based on a likelihood ratio approach, Glynn and L'Ecuyer provide in [11] an asymptotically unbiased single-run estimator for general Harris recurrent chains, where “general” means that the existence of an atom is not assumed. See Glynn et al. [12] for details on the unbiasedness of the estimator.

The philosophy of the present paper differs from the approach predominant in the literature. First, a closed-form formula for the derivative of the stationary performance of a general state-space Markov chain is given using an operator-theoretic approach. Specifically, we will derive sufficient conditions for the existence of the derivative of the stationary distribution of a Markov chain in terms of differentiability and stability properties of the Markov kernel driving the chain. In a second step, we translate the derivative formula into unbiased gradient estimators. A straightforward procedure is to introduce variants of the Markov chain, called *phantoms*, and to evaluate the phantoms over a cycle of the Markov chain. The key point is, however, that our formula can be translated into gradient estimators of various kinds, and we will discuss some (possible) implementations. Specifically, we will recover the estimators of Pflug [27, 28] and will illustrate how the score-function estimator of Glynn and L'Ecuyer [11] emerges from our overall formula if the Markov kernel under consideration possesses a differentiable density satisfying some additional regularity conditions.

The mathematical analysis is performed within the framework of measure-valued differentiation (MVD), see Heidergott and Vázquez-Abad [18]. MVD overcomes the restriction of likelihood ratio approach to differentiable

densities. Moreover, MVD extends the concept of weak differentiability, as introduced in Pflug [27, 28], so that performance measures out of a predefined class $\mathcal{D}$ can be handled, and thereby overcomes the restriction to bounded functions, which is inherent in the results in Pflug [27, 28]. As explained in Heidergott and Vázquez-Abad [18], MVD implies no restriction to a particular estimation method and allows for various interpretation in terms of estimators. We present a new phantom estimator and establish its unbiasedness.

It is our belief that separating the search for closed-form derivative formulae on the one hand and the search for translations of these formulae to efficient algorithms on the other hand will fertilize research on gradient estimation.

The paper is organized as follows. Section 2 introduces MVD. Basic definitions and concepts of MVD are provided in §2.1 and the likelihood ratio counterpart is discussed in §2.2. In §3 the main result of the paper is established, namely, that the stationary distribution of a $\mathcal{D}$ -differentiable Markov chain is $\mathcal{D}$ -differentiable. In §4, we provide—based on ergodicity concepts—a set of sufficient conditions for our main result to hold that can be verified in applications. We then turn to translating the closed-form expression for the derivative of the stationary distribution into gradient estimators in §5. Examples are provided in §6.

2. MVD for Markov chains.

2.1. Background on MVD. Let $(S, \mathcal{T})$ be a Polish measurable space. Let $\mathcal{M}(S, \mathcal{T})$ denote the set of finite (signed) measures on $(S, \mathcal{T})$ and $\mathcal{M}_1(S, \mathcal{T})$ that of probability measures on $(S, \mathcal{T})$.

DEFINITION 2.1. The mapping $P: S \times \mathcal{T} \rightarrow [0, 1]$ is called a (homogeneous) *transition kernel* on $(S, \mathcal{T})$ if

- (a) $P(s; \cdot) \in \mathcal{M}(S, \mathcal{T})$ for all $s \in S$; and
- (b) $P(\cdot; B)$ is $\mathcal{T}$ measurable for all $B \in \mathcal{T}$.

If, in condition (a), $\mathcal{M}(S, \mathcal{T})$ can be replaced by $\mathcal{M}_1(S, \mathcal{T})$, then P is called a *Markov kernel* on $(S, \mathcal{T})$.

Denote the set of transition kernels on $(S, \mathcal{T})$ by $\mathcal{K}(S, \mathcal{T})$ and the set of Markov kernels on $(S, \mathcal{T})$ by $\mathcal{K}_1(S, \mathcal{T})$. Consider a family of Markov kernels $(P_\theta; \theta \in \Theta)$ on $(S, \mathcal{T})$, with $\Theta \subset \mathbb{R}$, and let $L^1(P_\theta; \Theta) \subset \mathbb{R}^S$ denote the set of measurable mappings $g: S \rightarrow \mathbb{R}$, such that $\int_S P_\theta(s; du)|g(u)|$ is finite for all $\theta \in \Theta$ and $s \in S$. We call $P_\theta \in \mathcal{K}(S, \mathcal{T})$ $\mathcal{D}$ -preserving for $\mathcal{D} \subset L^1(P_\theta; \Theta)$ if $\int_S P_\theta(\cdot; du)g(u) \in \mathcal{D}$ for any $g \in \mathcal{D}$. For ease of exposition, we fix θ_0 and let Θ be an open neighborhood of θ_0 .

DEFINITION 2.2. For $\theta \in \Theta$, let $P_\theta \in \mathcal{K}(S, \mathcal{T})$ and $\mathcal{D} \subset L^1(P_\theta; \Theta) \subset \mathbb{R}^S$. We call P_θ $\mathcal{D}$ -continuous at θ_0 if, for any $g \in \mathcal{D}$ and any $s \in S$:

$$\lim_{\Delta \rightarrow 0} \left| \int P_{\theta_0 + \Delta}(s; dz)g(z) - \int P_{\theta_0}(s; dz)g(z) \right| = 0.$$

Furthermore, we call P_θ $\mathcal{D}$ -Lipschitz continuous at θ_0 if, for any $g \in \mathcal{D}$, a $K_g \in \mathcal{D}$ exists such that for any $\Delta \in \mathbb{R}$, with $\theta_0 + \Delta \in \Theta$, and any $s \in S$:

$$\left| \int P_{\theta_0 + \Delta}(s; dz)g(z) - \int P_{\theta_0}(s; dz)g(z) \right| \leq |\Delta|K_g(s).$$

We denote the set of bounded continuous mappings from S to $\mathbb{R}$ by $C_b(S)$. Let $P_\theta \in \mathcal{K}_1(S, \mathcal{T})$, for $\theta \in \Theta$. We call $\mathcal{D} \subset L^1(P_\theta; \Theta)$ a set of *test functions* for $(P_\theta; \theta \in \Theta)$ if $C_b(S) \subset \mathcal{D}$.

DEFINITION 2.3. Let $\mathcal{D} \subset L^1(P_\theta; \Theta)$ be a set of test functions for $(P_\theta; \theta \in \Theta)$. We call $P_\theta \in \mathcal{K}(S, \mathcal{T})$ *differentiable* at θ with respect to $\mathcal{D}$, or $\mathcal{D}$ -differentiable for short, if for any $s \in S$ a $P'_\theta(s; \cdot) \in \mathcal{M}(S, \mathcal{T})$ exists such that for any $g \in \mathcal{D}$:

$$\frac{d}{d\theta} \int_S P_\theta(s; du)g(u) = \int_S P'_\theta(s; du)g(u) \quad (1)$$

and if P'_θ is $\mathcal{D}$ -preserving. If the left-hand side of Equation (1) equals zero for all $g \in \mathcal{D}$, then we say that P'_θ is not significant.

If P_θ is $\mathcal{D}$ -differentiable, then the measure-valued derivative $P'_\theta(\cdot; \cdot)$ is uniquely defined, which stems from the fact that we have assumed that $C_b(S) \subset \mathcal{D}$. It is, however, not clear whether $P'_\theta(\cdot; \cdot)$ is a transition kernel on $(S, \mathcal{T})$. For example, if for any $A \in \mathcal{T}$ its indicator function is in $\mathcal{D}$, then $P'_\theta(\cdot; \cdot)$ is indeed a transition kernel on $(S, \mathcal{T})$, see Heidergott and Vázquez-Abad [18]. As shown in Heidergott et al. [13], if

$$\sup_{\{g \in C_b(S): \sup_{u \in S} |g(u)| \leq 1\}} \left| \int g(z)P'_\theta(s; dz) \right| < \infty, \quad (2)$$

for any $s \in S$, then P'_θ is a transition kernel on $(S, \mathcal{T})$.

REMARK 2.1. In the case that θ_0 is the boundary point of the definition range of P_θ , we take Θ to be a half-open interval with θ_0 as boundary point. In this case, limits have to be understood as one-sided limits.

EXAMPLE 2.1. A typical choice for $\mathcal{D}$ is the set of measurable bounded functions on S , denoted by $\mathcal{D}_b$ (it is easily checked that $\mathcal{D}_b$ is indeed a set of test functions).

To assume that the sample performance is bounded ($g \in \mathcal{D}_b$) is often too restrictive in applications. A more convenient set of performance functions is the set $\mathcal{D}^p$ of polynomially bounded performance functions defined by

$$\mathcal{D}^p = \left\{ g: S \rightarrow \mathbb{R} \mid g(x) \leq \sum_{i=0}^p \kappa_i \|x\|^i, \kappa_i \in \mathbb{R}, 0 \leq i \leq p \right\},$$

for some $p \in \mathbb{N}$, where $\|\cdot\|$ denotes a norm on S (assuming that S is indeed equipped with a norm). Most cases that are of interest in applications fall into this setting. The set $\mathcal{D}^p$ is a set of test functions if and only if $\mathcal{D}^p \subset L^1(P_\theta; \Theta)$, or, equivalently, if $\int_S P_\theta(s; du) \|u\|^p$ is finite for any $s \in S$ and $\theta \in \Theta$.

If P'_θ exists, then the fact that $P'_\theta(s; \cdot)$ fails to be a probability measure poses the problem of sampling from $P'_\theta(s; \cdot)$. For $s \in S$ fixed, we can represent $P'_\theta(s; \cdot)$ by its Jordan decomposition as a difference between two positive finite measures. Appropriate rescaling of these measures leads to a decomposition into two probability measures. More precisely, for $s \in S$, let $([P'_\theta]^+(s; \cdot), [P'_\theta]^-(s; \cdot))$ denote the Jordan decomposition of $P'_\theta(s; \cdot)$ and set

$$c_{P_\theta}(s) = [P'_\theta]^+(s; S) = [P'_\theta]^-(s; S) \quad (3)$$

and

$$P_\theta^+(s; \cdot) = \frac{[P'_\theta]^+(s; \cdot)}{c_{P_\theta}(s)}, \quad P_\theta^-(s; \cdot) = \frac{[P'_\theta]^-(s; \cdot)}{c_{P_\theta}(s)}.$$

Then it holds, for all $g \in \mathcal{D}$, that

$$\int_S P'_\theta(s; du) g(u) = c_{P_\theta}(s) \left(\int_S P_\theta^+(s; du) g(u) - \int_S P_\theta^-(s; du) g(u) \right). \quad (4)$$

For the above line of argument we fixed s . For P_θ^+ and P_θ^- to be Markov kernels, we have to consider P_θ^+ and P_θ^- as functions in s and have to establish measurability of $P_\theta^+(\cdot; A)$ and $P_\theta^-(\cdot; A)$ for any $A \in \mathcal{T}$. This problem is equivalent to showing that $c_{P_\theta}(\cdot)$ in (3) is measurable as a mapping from S to $\mathbb{R}$. As shown in Heidergott et al. [13], a sufficient condition for this is (2). Another situation where it can be shown that c_{P_θ} is measurable is when P_θ possesses a Radon-Nikodym derivative with respect to some Markov kernel Q that is independent of θ ; for details we refer to the next section. In applications c_{P_θ} is calculated explicitly and its measurability is therefore established case by case. Specifically, in many examples that are of interest in applications, c_{P_θ} turns out to be a constant and measurability is thus guaranteed; see Heidergott and Vázquez-Abad [18] for more details.

We now introduce the notion of $\mathcal{D}$ -derivative, which extends the concept of a weak derivative.

DEFINITION 2.4. Let P_θ be $\mathcal{D}$ -differentiable at θ . Any triple $(c_{P_\theta}(\cdot), P_\theta^+, P_\theta^-)$, with $P_\theta^\pm \in \mathcal{K}_1(S, \mathcal{T})$ and c_{P_θ} a measurable mapping from S to $\mathbb{R}$ that satisfies (4) is called a $\mathcal{D}$ -derivative of P_θ . The kernel P_θ^+ is called the (normalized) *positive part* of P'_θ and P_θ^- is called the (normalized) *negative part* of P'_θ ; and $c_{P_\theta}(\cdot)$ is called the *normalizing factor*.

We illustrate the concepts introduced above with a simple example.

EXAMPLE 2.2. Let $P, Q \in \mathcal{K}_1(S, \mathcal{T})$ and set

$$P_\theta = \theta P + (1 - \theta)Q, \quad \theta \in [0, 1].$$

Note that $P_\theta \in \mathcal{K}_1(S, \mathcal{T})$ for $\theta \in [0, 1]$, and that $P_0 = Q$ and $P_1 = P$. Specifically, let $\mathcal{D}(P, Q) \triangleq L^1(P_\theta; \Theta)$ denote the set of measurable mappings $g: S \rightarrow \mathbb{R}$ such that both $\int_S P(s; du) |g(u)|$ and $\int_S Q(s; du) |g(u)|$ exist and are finite for any $s \in S$. Moreover, assume that $C_b \subset \mathcal{D}(P, Q)$ and that P and Q , respectively, are $\mathcal{D}(P, Q)$ -preserving. For any $g \in \mathcal{D}(P, Q)$ and any $s \in S$, we now compute

$$\begin{aligned} \frac{d}{d\theta} \int_S P_\theta(s; du) g(u) &= \frac{d}{d\theta} \left(\theta \int_S P(s; du) g(u) + (1 - \theta) \int_S Q(s; du) g(u) \right) \\ &= \int_S P(s; du) g(u) - \int_S Q(s; du) g(u). \end{aligned}$$

Hence, P_θ is $\mathcal{D}(P, Q)$ -differentiable and it is easily seen that P'_θ is $\mathcal{D}(P, Q)$ -preserving. Specifically, $(1, P, Q)$ is an instance of a $\mathcal{D}(P, Q)$ -derivative, for any $\theta \in [0, 1]$, where the derivative at the boundary points has to be understood as one-sided derivatives.

A $\mathcal{D}$ -derivative is not unique. To see this, let $(c_{P_\theta}, P_\theta^+, P_\theta^-)$ be a $\mathcal{D}$ -derivative of P_θ and choose $Q \in \mathcal{K}_1(S; \mathcal{T})$ so that $\int_S g(u) Q(s; du)$ is finite for any $g \in \mathcal{D}$ and $s \in S$. Set

$$\tilde{P}_\theta^+ = \frac{1}{2} P_\theta^+ + \frac{1}{2} Q, \quad \tilde{P}_\theta^- = \frac{1}{2} P_\theta^- + \frac{1}{2} Q.$$

It is easily seen that $(2c_{P_\theta}, \tilde{P}_\theta^+, \tilde{P}_\theta^-)$ satisfies Equation (4) for all $g \in \mathcal{D}$ and is thus a $\mathcal{D}$ -derivative of P_θ .

2.2. The likelihood ratio approach. In the previous section, we have dealt with the problem that $P'_\theta(s; \cdot)$, for $s \in S$, fails to be a probability measure through considering a Jordan-type decomposition of P'_θ . This has led to the concept of a $\mathcal{D}$ -derivative. An alternative approach is to apply a change of measure so that P'_θ becomes $(dP'_\theta/dQ)Q$ for an appropriately chosen Markov kernel Q , where dP'_θ/dQ notably is the Radon-Nikodym derivative of P'_θ with respect to Q . As noted in Heidergott and Hordijk [17], provided that P'_θ has $\mathcal{D}$ -derivative $(c_{P_\theta}, P_\theta^+, P_\theta^-)$, such a kernel Q can always be found. Indeed, fix θ and let

$$Q = \frac{1}{2}P_\theta^+ + \frac{1}{2}P_\theta^-,$$

then Q dominates P'_θ and thus the Radon-Nikodym derivative of P'_θ with respect to Q exists.

The Markov kernel Q defined above is usually too complex to be of use in applications. A particularly nice situation is when Q can be taken as P_θ . As will be explained in the following, dP'_θ/dP_θ then becomes the score function.

Suppose that a Markov kernel $Q \in \mathcal{H}_1$ exists such that P_θ has a Radon-Nikodym derivative with respect to Q and denote the Radon-Nikodym derivative by f_θ , that is, assume that it holds for any $s \in S$:

$$\int P_\theta(s; dz)g(z) = \int Q(s; dz)f_\theta(s; z)g(z),$$

for any $P_\theta(s; \cdot)$ integrable mapping g . In symbols, $f_\theta = dP_\theta/dQ$. In order to avoid confusing Radon-Nikodym “derivatives” with measure-valued “derivatives,” we will refer to f_θ as Q -density in the following.

Assume that $f_\theta(s, z)$ is differentiable with respect to θ for any $s, z \in S$ and let $\mathcal{D}$ denote the set of mappings such that for any $g \in \mathcal{D}$:

$$\frac{d}{d\theta} \int Q(s; dz)f_\theta(s; z)g(z) = \int Q(s; dz) \frac{d}{d\theta} (f_\theta(s; z))g(z). \quad (5)$$

Then, P_θ is $\mathcal{D}$ -differentiable, and it holds that for any $g \in \mathcal{D}$,

$$\int P'_\theta(s; dz)g(z) = \int Q(s; dz) \frac{d}{d\theta} (f_\theta(s; z))g(z). \quad (6)$$

In words, $df_\theta(s; z)/d\theta$ serves as a Q -density of P'_θ (that $df_\theta(s; z)/d\theta$ is indeed measurable in s and z follows from the fact that it is the limit of measurable mappings).

EXAMPLE 2.3. Let $S = \mathbb{R}$. A sufficient condition for (5) to hold for any $g \in \mathcal{D}^p$ (see Example 1 for the definition of $\mathcal{D}^p$) is:

$$\forall s \in S: \int Q(s; dz)|z|^p \sup_{\theta \in \Theta} \left| \frac{d}{d\theta} \Big|_{\theta=\hat{\theta}} f_\theta(s; z) \right| < \infty.$$

Provided that $Q(s; \cdot)$ has finite $(2p)$ th moment, the above follows by the Cauchy-Schwartz inequality from

$$\forall s \in S: \int Q(s; dz) \sup_{\theta \in \Theta} \left| \frac{d}{d\theta} \Big|_{\theta=\hat{\theta}} f_\theta(s; z) \right|^2 < \infty,$$

which is Glynn [10, Assumption (A2)] for $p = 2$.

For easy reference we introduce the following condition.

(SF) For $\theta \in \Theta$, the Markov kernel P_θ has P_θ -density f_θ . In symbols: $P_{\theta_0}(s; dz)f_\theta(s; z) = P_\theta(s; dz)$, for $s \in S$. Moreover, $f_\theta(s, z)$ is differentiable with respect to θ for any $s, z \in S$.

Note that Condition (SF) is implied in Glynn and L'Ecuyer [11, Assumptions (A1) and (A2)]. Under Condition (SF), Equation (6) reads

$$\int P'_\theta(s; dz)g(z) = \int P_\theta(s; dz) \frac{d}{d\theta} f_\theta(s; z)g(z) = \int P_\theta(s; dz) \frac{(d/d\theta)f_\theta(s; z)}{f_\theta(s; z)}g(z). \quad (7)$$

The mapping

$$\text{SF}_\theta(s; z) \triangleq \frac{(d/d\theta)f_\theta(s; z)}{f_\theta(s; z)}$$

is called *score function* in the literature (Glynn [10] and Rubinstein and Shapiro [30]).

REMARK 2.2. Under Condition (SF), the score function is a P_θ -density of P'_θ .

From (7) an instance of a $\mathcal{D}$ -derivative of P'_θ can be obtained as follows: let

$$c_{P_\theta}(s) = \int Q(s; dz) \max \left(\frac{d}{d\theta} f_\theta(s; z), 0 \right) = \int P_\theta(s; dz) \max(\text{SF}_\theta(s; z), 0)$$

and set

$$\begin{aligned} \text{SF}_\theta^+(s; z) &= \frac{1}{c_{P_\theta}(s)} \max(\text{SF}_\theta(s; z), 0), \\ \text{SF}_\theta^-(s; z) &= \frac{1}{c_{P_\theta}(s)} \max(-\text{SF}_\theta(s; z), 0). \end{aligned}$$

Then it holds for any $g \in \mathcal{D}$ that

$$\begin{aligned} \frac{d}{d\theta} \int P_\theta(s; dz) g(z) &= \int P'_\theta(s; dz) g(z) \\ &= \int P_\theta(s; dz) \text{SF}_\theta(s; z) g(z) \\ &= c_{P_\theta}(s) \left(\int P_\theta(s; dz) \text{SF}_\theta^+(s; z) g(z) - \int P_\theta(s; dz) \text{SF}_\theta^-(s; z) g(z) \right) \\ &= c_{P_\theta}(s) \left(\int P_\theta^+(s; dz) g(z) - \int P_\theta^-(s; dz) g(z) \right), \end{aligned}$$

for Markov kernels $P_\theta^\pm$ having P_θ -densities $\text{SF}_\theta^\pm$. Hence, $(c_{P_\theta}, P_\theta^+, P_\theta^-)$ is a $\mathcal{D}$ -derivative of P_θ . Note that SF_θ is measurable in both arguments, which implies that c_{P_θ} is measurable.

The above analysis shows that the existence of the score function (for a certain set of performance mappings $\mathcal{D}$) implies $\mathcal{D}$ -differentiability of P_θ and yields an instance of a $\mathcal{D}$ -derivative. The converse, however, is not true. A classical example of this kind is the uniform distribution on $[0, \theta]$, denoted by $\mathcal{U}_{[0, \theta]}$. It is easy to see that $\mathcal{U}_{[0, \theta]}$ is $\mathcal{C}$ -differentiable, with $\mathcal{C}$ the set of continuous mappings from S to $\mathbb{R}$, and that any measure-valued derivative of this distribution has a continuous and a discrete part. Obtaining a measure-valued derivative via the score function therefore fails; see Brémaud [3] and Heidergott and Vázquez-Abad [18] for details. Thus, the Markov kernel $P_\theta(s; \cdot) = \mathcal{U}_{[s, s+\theta]}(\cdot)$, with $s \in S = \mathbb{R}$, is $\mathcal{C}$ -differentiable, but the score function does not exist. Another example of this kind is the following.

EXAMPLE 2.4. Let μ be a continuous probability measure on $\mathbb{R}$ with mean x and consider the family of probability measures given by

$$\nu_\theta = \theta\mu + (1 - \theta)\delta_x, \quad \theta \in [0, 1].$$

At $\theta = 0$, ν_θ degenerates to a point mass and, at $\theta = 1$, ν_θ becomes the distribution μ . Hence, letting θ go from 0 to 1, increases the variability of ν_θ and, speaking in terms of random variables, transforms a deterministic variable (which is equal to the first moment of μ) into a random variable with distribution μ . Evaluating the derivative of $\int g(y)\nu_\theta(dy)$ at $\theta = 0$ then measures the sensitivity of the performance function with respect to the simplifying assumption that the considered variable is deterministic. This sensitivity is called *analysis of variability*.

For instance, the inverse throughput of a max-plus linear queueing network can be easily evaluated when all variables are deterministic. Deterministic max-plus models have been successfully applied to railway networks; see, for example, Braker [2] and Heidergott and de Vries [20]. This is in contrast to the situation where (some of the) variables are random. Here, except for special cases, no solutions are known. In order to calculate the sensitivity of the deterministic model with respect to the random nature of the variables, analysis of variability can be applied. Notice that one actually calculates a Taylor series expansion of the inverse throughput with respect to the artificial parameter θ ; see Baccelli and Hong [1] for more on this.

In the following we will show that the score function method is not capable of delivering an unbiased sensitivity estimator for analysis of variability. It is easily seen that ν_θ , with $\theta \in [0, 1]$, is absolutely continuous with respect to ν_{θ_0} , for $\theta_0 \in (0, 1)$, with likelihood ratio

$$\mathbf{f}_{\theta, \theta_0}(y) \triangleq \frac{d\nu_\theta}{d\nu_{\theta_0}}(y) = \frac{\theta}{\theta_0} 1_{\mathbb{R} \setminus \{x\}}(y) + \frac{1 - \theta}{1 - \theta_0} 1_{\{x\}}(y),$$

for any $\theta \in [0, 1]$. Indeed, for $\theta \in [0, 1]$ and $\theta_0 \in (0, 1)$, it holds for any measurable subset A of $\mathbb{R}$:

$$\begin{aligned} \int_A \mathbf{f}_{\theta, \theta_0}(y) \nu_{\theta_0}(dy) &= \int_A \left(\frac{\theta}{\theta_0} 1_{\mathbb{R} \setminus \{x\}}(y) + \frac{1 - \theta}{1 - \theta_0} 1_{\{x\}}(y) \right) (\theta_0\mu + (1 - \theta_0)\delta_x)(dy) \\ &= \theta\mu(A) + (1 - \theta)\delta_x(A), \end{aligned}$$

where the equality follows from the fact that μ is a continuous measure, and set $\{x\}$ therefore has μ -mass zero. Notice that the point $\theta_0 = 0$ is an extreme point: ν_0 does not dominate ν_θ , for $\theta \in (0, 1]$, and $f_{\theta, 0}$ is not defined for any $\theta \in (0, 1]$.

For $\theta, \theta_0 \in (0, 1)$, the pathwise derivative of the likelihood ratio is

$$\mathbf{f}'_{\theta, \theta_0}(y) = \frac{1}{\theta_0} 1_{\mathbb{R} \setminus \{x\}}(y) - \frac{1}{1 - \theta_0} 1_{\{x\}}(y). \quad (8)$$

For $\theta = \theta_0 \in (0, 1)$, the score function reads

$$\text{SF}_\theta(y) = \frac{\mathbf{f}'_{\theta, \theta_0}}{\mathbf{f}_{\theta, \theta_0}} = \frac{1}{\theta} 1_{\mathbb{R} \setminus \{x\}}(y) - \frac{1}{1 - \theta} 1_{\{x\}}(y).$$

It is easily checked that

$$\begin{aligned} \frac{d}{d\theta} \int g(y) \nu_\theta(dy) &= \int g(y) \text{SF}_\theta(y) \nu_\theta(dy) \\ &= \int g(y) \mu(dy) - \int g(y) \delta_x(dy), \end{aligned}$$

for any $\theta \in (0, 1)$. Observe that the score function fails to exist at $\theta = 0$ and $(d/d\theta)|_{\theta=0} \int g(y) \nu_\theta(dy)$ cannot be obtained via the score function. It is worth noting that taking ν_θ , for $\theta > 0$, as dominating measure, a likelihood ratio estimator can be applied.

It is worth noting that Kushner and Vázquez-Abad in [25] use a weak differentiation approach for establishing differentiability of the stationary distribution, assuming that P'_θ has a P_θ -density (this is in Kushner and Vázquez-Abad [25, Condition (A2)]). As we have already pointed out, the approach presented in the present paper allows us to overcome this restriction and avoids any assumption on the existence of an overall dominating kernel.

3. MVD of the stationary distribution. Let P_θ be the transition kernel of an aperiodic Markov chain possessing a unique invariant distribution, denoted by π_θ . See the next section for conditions that are sufficient for the existence of π_θ . Let $L^1(\pi_\theta)$ denote the set of measurable mappings $g: S \rightarrow \mathbb{R}$ such that $\int |g| d\pi_\theta$ is finite. We denote the ergodic projector of π_θ by Π_θ , that is, for any probability measure μ on $(S, \mathcal{F})$ it holds that $\mu \Pi_\theta = \pi_\theta$. To simplify the notation, we set:

$$\mu g = \int g(s) \mu(ds),$$

for $\mu \in \mathcal{M}(S, \mathcal{F})$, and

$$(P_\theta g)(s) = \int P_\theta(s; dz) g(z),$$

for $P_\theta \in \mathcal{H}(S, \mathcal{F})$, provided that the expression exists. Note that $P_\theta g$ is a mapping from S to $\mathbb{R} \cup \{-\infty, \infty\}$. For any $\mu \in \mathcal{M}$ with $\mu(S) = 0$, we have the following rule of computation

$$\mu P_\theta^n g = \mu(P_\theta^n - \Pi_\theta)g, \quad n \in \mathbb{N}, \quad (9)$$

provided that the integrals exist. In what follows, we work locally and fix θ . With slight abuse of notation, we take Θ to be an open neighborhood of θ .

THEOREM 3.1. *If:*

- (i) P is $\mathcal{D}$ -Lipschitz continuous at θ ,
- (ii) for any $\hat{\theta} \in \Theta$, if $h \in \mathcal{D}$ then $P_{\hat{\theta}} h \in \mathcal{D}$,
- (iii) for any $h \in \mathcal{D}$ and any $\Delta \in \mathbb{R}$, with $\theta + \Delta \in \Theta$,

$$\lim_{k \rightarrow \infty} (\pi_{\theta+\Delta} - \pi_\theta) P_\theta^k h = 0,$$

- (iv) for any $h \in \mathcal{D}$,

$$(a) \quad \sum_{n=0}^{\infty} |(P_\theta^n - \Pi_\theta)h| \in \mathcal{D},$$

$$(b) \quad \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta)h \in \mathcal{D},$$

- (v) for any $h \in \mathcal{D}$ a finite number c_h exists such that

$$\pi_{\hat{\theta}}|h| \leq c_h, \quad \forall \hat{\theta} \in \Theta,$$

then π is $\mathcal{D}$ -Lipschitz continuous at θ .

Moreover, if we assume, in addition to the above conditions, that P_θ is $\mathcal{D}$ -differentiable at θ , then π_θ is $\mathcal{D}$ -differentiable at θ with $\mathcal{D}$ -derivative

$$\pi'_\theta = \pi_\theta \sum_{n=0}^{\infty} P'_\theta P_\theta^n,$$

or, equivalently,

$$\pi'_\theta = \pi_\theta P'_\theta \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta).$$

PROOF. Let I denote the unit operator, that is, $\pi_\theta I = \pi_\theta$. Simple algebraic calculation shows that

$$\begin{aligned} (\pi_{\theta+\Delta} - \pi_\theta)(I - P_{\theta+\Delta}) &= \pi_{\theta+\Delta} - \underbrace{\pi_{\theta+\Delta} P_{\theta+\Delta}}_{=\pi_{\theta+\Delta}} - \pi_\theta + \pi_\theta P_{\theta+\Delta} \\ &= \pi_\theta (P_{\theta+\Delta} - P_\theta). \end{aligned}$$

Hence,

$$(\pi_{\theta+\Delta} - \pi_\theta)(I - P_{\theta+\Delta}) \sum_{n=0}^k P_\theta^n = \pi_\theta (P_{\theta+\Delta} - P_\theta) \sum_{n=0}^k P_\theta^n, \quad (10)$$

with $P_\theta^0 = I$. By algebraic calculation,

$$\begin{aligned} (\pi_{\theta+\Delta} - \pi_\theta)(I - P_{\theta+\Delta}) \sum_{n=0}^k P_\theta^n &= (\pi_{\theta+\Delta} - \pi_\theta) \left\{ (I - P_\theta) \sum_{n=0}^k P_\theta^n + (P_\theta - P_{\theta+\Delta}) \sum_{n=0}^k P_\theta^n \right\} \\ &= (\pi_{\theta+\Delta} - \pi_\theta) \left\{ \sum_{n=0}^k P_\theta^n - \sum_{n=1}^{k+1} P_\theta^n + (P_\theta - P_{\theta+\Delta}) \sum_{n=0}^k P_\theta^n \right\} \\ &= (\pi_{\theta+\Delta} - \pi_\theta) \left\{ I - P_\theta^{k+1} + (P_\theta - P_{\theta+\Delta}) \sum_{n=0}^k P_\theta^n \right\} \\ &= \pi_{\theta+\Delta} - \pi_\theta - (\pi_{\theta+\Delta} - \pi_\theta) P_\theta^{k+1} + (\pi_{\theta+\Delta} - \pi_\theta) (P_\theta - P_{\theta+\Delta}) \sum_{n=0}^k P_\theta^n. \end{aligned}$$

Inserting the right-hand side of the above equation into (10) yields:

$$\pi_\theta (P_{\theta+\Delta} - P_\theta) \sum_{n=0}^k P_\theta^n = \pi_{\theta+\Delta} - \pi_\theta - (\pi_{\theta+\Delta} - \pi_\theta) P_\theta^{k+1} + (\pi_{\theta+\Delta} - \pi_\theta) (P_\theta - P_{\theta+\Delta}) \sum_{n=0}^k P_\theta^n. \quad (11)$$

We now study the limit of the above expression as k tends to ∞ . Firstly, under assumption (iii), we have

$$\lim_{k \rightarrow \infty} (\pi_{\theta+\Delta} - \pi_\theta) P_\theta^{k+1} h = 0. \quad (12)$$

As a second step, we show

$$\lim_{k \rightarrow \infty} \pi_\theta (P_{\theta+\Delta} - P_\theta) \sum_{n=0}^k P_\theta^n h = \pi_\theta (P_{\theta+\Delta} - P_\theta) \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta) h, \quad (13)$$

for $h \in \mathcal{D}$. By assumptions (ii) together with (v),

$$\pi_\theta (P_{\theta+\Delta} - P_\theta) \sum_{n=0}^k P_\theta^n h$$

exists and is finite for any $k \geq 0$ and any $h \in \mathcal{D}$. Elaborating on (9), this implies that

$$\pi_\theta (P_{\theta+\Delta} - P_\theta) \sum_{n=0}^k (P_\theta^n - \Pi_\theta) h$$

exists and is finite for any $k \geq 0$ and any $h \in \mathcal{D}$, where we use the fact that $\pi_\theta(P_{\theta+\Delta} - P_\theta)$ is a signed measure with total mass 0. By assumption (iv)(a), for any $h \in \mathcal{D}$,

$$\left| \sum_{n=0}^k (P_\theta^n - \Pi_\theta)h \right| \leq \sum_{n=0}^{\infty} |(P_\theta^n - \Pi_\theta)h| \in \mathcal{D}.$$

Moreover, by assumption (ii) together with (v),

$$\pi_\theta(P_{\theta+\Delta} - P_\theta) \sum_{n=0}^{\infty} |(P_\theta^n - \Pi_\theta)h|$$

exists and is finite, for any $h \in \mathcal{D}$. Hence, (13) follows from the dominated convergence theorem.

Finally, following the line of argument in the above second step, we show that

$$\begin{aligned} \lim_{k \rightarrow \infty} (\pi_{\theta+\Delta} - \pi_\theta)(P_\theta - P_{\theta+\Delta}) \sum_{n=0}^k P_\theta^n &= \lim_{k \rightarrow \infty} (\pi_{\theta+\Delta} - \pi_\theta)(P_\theta - P_{\theta+\Delta}) \sum_{n=0}^k (P_\theta^n - \Pi_\theta) \\ &= (\pi_{\theta+\Delta} - \pi_\theta)(P_\theta - P_{\theta+\Delta}) \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta). \end{aligned} \quad (14)$$

Taking the limit in (11) as k tends to ∞ , we obtain from (12), (13), and (14), for $h \in \mathcal{D}$:

$$(\pi_{\theta+\Delta} - \pi_\theta)h = \pi_\theta(P_{\theta+\Delta} - P_\theta) \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta)h - (\pi_{\theta+\Delta} - \pi_\theta)(P_\theta - P_{\theta+\Delta}) \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta)h. \quad (15)$$

For $h \in \mathcal{D}$, condition (iv) (b) implies

$$\sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta)h =: \hat{h} \in \mathcal{D},$$

and $\mathcal{D}$ -Lipschitz continuity of P_θ at θ implies

$$\begin{aligned} \left| \pi_\theta(P_{\theta+\Delta} - P_\theta) \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta)h \right| &\leq |\Delta| \pi_\theta K_{\hat{h}} \\ &\leq |\Delta| c_{K_{\hat{h}}} < \infty, \end{aligned}$$

where finiteness of $\pi_\theta K_{\hat{h}}$ is guaranteed by (v). From the same line of argument, we obtain that

$$\begin{aligned} \left| (\pi_{\theta+\Delta} - \pi_\theta)(P_\theta - P_{\theta+\Delta}) \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta)h \right| &\leq (\pi_{\theta+\Delta} + \pi_\theta) \left| (P_\theta - P_{\theta+\Delta}) \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta)h \right| \\ &\leq |\Delta| (\pi_{\theta+\Delta} + \pi_\theta) K_{\hat{h}} \\ &\leq |\Delta| 2c_{K_{\hat{h}}}, \end{aligned}$$

for any $h \in \mathcal{D}$. Because $\mathcal{D}$ is a set of test functions, the constant functions $|\Delta|c_{K_{\hat{h}}}$, $|\Delta|2c_{K_{\hat{h}}}$ as well as their sum, lie in $\mathcal{D}$. From (15) it thus follows that π_θ is $\mathcal{D}$ -Lipschitz continuous at θ . In particular, the Lipschitz factor is the constant function $3c_{K_{\hat{h}}}$.

Starting point for the second part of the theorem is Equation (15). From $\mathcal{D}$ -Lipschitz continuity at θ of both π_θ and P_θ it follows that

$$\lim_{\Delta \rightarrow 0} \frac{1}{\Delta} (\pi_\theta - \pi_{\theta+\Delta})(P_\theta - P_{\theta+\Delta}) \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta)h = 0,$$

for any $h \in \mathcal{D}$. Moreover, because P_θ is $\mathcal{D}$ -Lipschitz at θ , assumption (v) yields

$$\lim_{\Delta \rightarrow 0} \frac{1}{\Delta} \pi_\theta(P_{\theta+\Delta} - P_\theta) \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta)h = \pi_\theta \left(\lim_{\Delta \rightarrow 0} \frac{1}{\Delta} (P_{\theta+\Delta} - P_\theta) \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta)h \right)$$

and, by $\mathcal{D}$ -differentiability of P_θ , this limit equals

$$= \pi_\theta P'_\theta \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta)h.$$

For $h \in \mathcal{D}$, we therefore obtain from (15)

$$\pi_\theta P'_\theta \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta)h = \lim_{\Delta \rightarrow 0} \frac{1}{\Delta} (\pi_\theta - \pi_{\theta+\Delta})h.$$

Using the fact the $P'_\theta(s; \cdot)$ is a signed measure with $P'_\theta(s; S) = 0$, for any $s \in S$, it readily follows from (9) that

$$\pi_\theta \sum_{n=0}^{\infty} P'_\theta P_\theta^n h = \pi_\theta P'_\theta \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta) h,$$

which concludes the proof. $\square$

REMARK 3.1. If $h \in \mathcal{D}$ implies that any $g \in L^1(P_\theta; \Theta)$ with $|g| \leq |h|$ lies in $\mathcal{D}$ as well, then Condition (iv) (a) already implies Condition (iv) (b).

REMARK 3.2. The operator $D_\theta \triangleq \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta)$ is called the *deviation operator* in the theory of Markov chains. Using the deviation operator, the derivative of the stationary distribution can be written in a concise way as $\pi'_\theta = \pi_\theta P'_\theta D_\theta$, provided that the conditions of Theorem 3.1 are satisfied.

4. Ergodicity framework. In this section, we provide sufficient conditions for Conditions (ii) to (iv) to hold. Let $\mathbf{X}(\theta) = \{X_\theta(n)\} = \{X_\theta(s, n)\}$, for $\theta \in \Theta$, be the aperiodic Markov chain with initial state s and transition kernel P_θ , and set, for any $B \in \mathcal{T}$,

$$P_\theta^n(s, B) \triangleq P_\theta(s, n, B) = \mathbb{P}(X_\theta(s, n) \in B).$$

The joint state-space of $\mathbf{X}(\theta)$, $\theta \in \Theta$, is denoted by S . However, for any specific θ the chain $\mathbf{X}(\theta)$ may live on a (measurable) subset S_θ of S . Therefore, ergodicity conditions for $\mathbf{X}(\theta)$ will be expressed with respect to $S_\theta \subset S$. Specifically, we consider $(S_\theta, \mathcal{T}_\theta)$, with $\mathcal{T}_\theta = S_\theta \cap \mathcal{T}$, as a (measurable) space, for $\theta \in \Theta$.

In §4.1, we discuss the general situation. Section 4.2 provides an alternative representation of the $\mathcal{D}$ -derivative of π_θ for the situation where $\mathbf{X}(\theta)$ possesses an atom.

4.1. General chains. The main technical conditions needed for the analysis in this section are introduced subsequently. We will use the following “Lyapunov function” condition:

(C.1) There exists a function $g(s) \geq 0$, $s \in S$, such that any $\theta \in \Theta$,

$$\mathbb{E}[g(X_\theta(s, m_\theta))] - g(s) \leq -\varepsilon + cI_{V_\theta}(s), \quad (16)$$

for some $m_\theta \geq 1$, $\varepsilon > 0$ and $c < \infty$, where for some $d < \infty$

$$V_\theta = \{s \in S_\theta: g(s) \leq d\}$$

and $I_{V_\theta}(s) = 1$ if $s \in V_\theta$ and, otherwise, zero. Note that the function g is the same for each θ .

Furthermore, we need the following Harris-type condition for the set V_θ .

(C.2) For any $\theta \in \Theta$, there exist $n_\theta \geq 0$, $\phi_\theta(\cdot)$ a probability measure on $(S_\theta, \mathcal{T}_\theta)$, and $p_\theta \in (0, 1)$ such that for all $B \in \mathcal{B}$:

$$\inf_{x \in V_\theta} \mathbb{P}(X_\theta(x, n_\theta) \in B) \geq p_\theta \phi_\theta(B). \quad (17)$$

Under Condition (C.1), let

$$\xi_\theta(s) \triangleq g(X_\theta(s, 1)) - g(s), \quad s \in S,$$

and introduce the following condition:

(C.3) The r.v. $\xi_\theta(s)$ is uniformly integrable (in s and θ), and there exist $\lambda > 0$ such that $\xi_\theta(s)e^{\lambda\xi_\theta(s)}$ are uniformly integrable (in s and θ).

Recall that uniform integrability of $\xi_\theta(s)$ in s and θ is defined as

$$\sup_{s, \theta} \int_{t>c} |t| P(|\xi_\theta(s)| \in dt) \searrow 0 \quad \text{for } c \rightarrow \infty,$$

and similarly the uniform integrability of $\xi_\theta(s)e^{\lambda\xi_\theta(s)}$ requires that

$$\sup_{s, \theta} \int_{t>c} |t| P(|\xi_\theta(s)e^{\lambda\xi_\theta(s)}| \in dt) \searrow 0 \quad \text{for } c \rightarrow \infty.$$

Notice that (C.3) is stronger than the assumption that there exists a point at which the corresponding moment-generating functions of all random variables are uniformly bounded; see Borovkov and Hordijk [6] for details.

In order to establish conditions (ii) to (iv) in Theorem 3.1, we will work with normed ergodicity. Normed ergodicity dates back to the early eighties; see Hordijk and Dekker [21] and the revised version, which was published as Dekker and Hordijk [8]. It was originally used in analysis of Blackwell optimality; see Dekker and Hordijk [8] and Hordijk and Yushkevich [24] for a recent publication on this topic. Since then, it has been used in various forms under different names in many subsequent papers. In Hordijk and Spieksma [22] it was shown that for a countable Markov chain, which may have one or several classes of essential states (a so-called multichained Markov chain), normed ergodicity is equivalent to geometrical recurrence (for a similar result in Markov decision chains see Dekker et al. [9]). Inspired by this result for a countable Markov chain, a similar

result was proved for a Harris chain in Meyn and Tweedie [26]. In this paper we use the recent results of Borovkov and Hordijk [5], the first part of this technical report has appeared as Borovkov and Hordijk [6].

Let $\mathcal{V}_v$ denote the space of real-valued functions f on S with the finite v -norm

$$\|f\|_v = \sup_{s \in S} \frac{|f(s)|}{v(s)},$$

where $v: S \rightarrow [1, \infty)$. The associated operator norm for a linear operator, say $T: \mathcal{V}_v \rightarrow \mathcal{V}_v$, is defined by

$$\|T\|_v = \sup_{\|f\|_v \leq 1} \|Tf\|_v,$$

and the associated norm of a (signed) measure, say μ , is

$$\|\mu\|_v = \sup_{\|f\|_v \leq 1} |\mu f|.$$

For our analysis, we choose v to be the following mapping:

$$v(s) \triangleq e^{\lambda g(s)}, \quad s \in S, \quad (18)$$

for some positive λ , where g is defined in (C.1).

LEMMA 4.1. *Condition (C.3) implies that for λ small enough*

$$\sup_{\theta \in \Theta} \|P_\theta\|_v < \infty.$$

PROOF. With $\xi_\theta(s) = g(X_\theta(s, 1)) - g(s)$, $s \in S$, and $v(s) = e^{\lambda g(s)}$, $s \in S$, we find that

$$\begin{aligned} \|P_\theta\|_v &= \sup_s \frac{(P_\theta e^{\lambda g})(s)}{e^{\lambda g(s)}} \\ &= \sup_s \mathbb{E}[e^{\lambda(g(X_\theta(s, 1)) - g(s))}] \\ &= \sup_s \mathbb{E}[e^{\lambda \xi_\theta(s)}]. \end{aligned} \quad (19)$$

By (C.3), $\xi_\theta(s)e^{\lambda \xi_\theta(s)}$ is uniformly integrable which implies

$$\sup_{s, \theta} \mathbb{E}[|\xi_\theta(s)e^{\lambda \xi_\theta(s)}|] < \infty.$$

For $s \in S$, we now write

$$\mathbb{E}[e^{\lambda \xi_\theta(s)}] = \mathbb{E}[e^{\lambda \xi_\theta(s)} I(\xi_\theta(s) \leq 1)] + \mathbb{E}[e^{\lambda \xi_\theta(s)} I(\xi_\theta(s) > 1)].$$

The first term on the right-hand side is bounded by e^λ , and the second term is bounded by

$$\sup_{s, \theta} \mathbb{E}[|\xi_\theta(s)e^{\lambda \xi_\theta(s)}|] < \infty.$$

Inserting these bounds into (19) yields

$$\sup_{\theta} \|P_\theta\|_v < \infty. \quad \square$$

Conditions (C.1)–(C.3) imply geometric ergodicity of P_θ for fixed θ as will be shown in the next theorem, which is a direct consequence of Borovkov and Hordijk [6, Theorem 5]. In particular, the type of ergodicity established in the subsequent theorem is called v -geometric ergodicity in Hordijk and Spieksma [22] and v -uniform ergodicity in Meyn and Tweedie [26]. The precise statement is as follows.

THEOREM 4.1. *Conditions (C.1), (C.2), and (C.3) imply that, for any $\theta \in \Theta$, there exist $c_\theta < \infty$ and $0 < \rho_\theta < 1$ such that*

$$\|P_\theta^n - \Pi_\theta\|_v \leq c_\theta \rho_\theta^n. \quad (20)$$

Note that this theorem implies $\|\Pi_\theta\|_v < \infty$ for any θ . Let

$$\mathcal{D}_v = \{g: S \rightarrow \mathbb{R} | \exists r \in \mathbb{R}: |g(s)| \leq r \cdot v(s), s \in S\}.$$

In words, $\mathcal{D}_v$ is the set of mappings g from S to $\mathbb{R}$ that are bounded by $r \cdot v$ (for some finite number r). It is easily seen that Lemma 4.1 implies that any $g \in \mathcal{D}_v$ is integrable with respect to any P_θ , for $\theta \in \Theta$, or, more formally:

$$\mathcal{D}_v \subset L^1(P_\theta, \Theta).$$

THEOREM 4.2. Let Θ denote an open neighborhood of θ . If

- (i) P_θ is $\mathcal{D}_v$ -Lipschitz continuous in θ ,
- (ii) Conditions (C.1) to (C.3) are satisfied,

then π_θ is $\mathcal{D}_v$ -Lipschitz continuous.

Moreover, if we assume, in addition to the above conditions, that P_θ is $\mathcal{D}_v$ -differentiable, then π_θ is $\mathcal{D}_v$ -differentiable with $\mathcal{D}_v$ -derivative

$$\pi'_\theta = \pi_\theta \sum_{n=0}^{\infty} P'_\theta P_\theta^n,$$

or, equivalently,

$$\pi'_\theta = \pi_\theta P'_\theta \sum_{n=0}^{\infty} (P_\theta^n - \Pi_\theta).$$

PROOF. We will show that conditions (ii) to (iv) of Theorem 3.1 hold. Condition (ii) is a straightforward consequence of Lemma 4.1. Note that

$$\|P_\theta^n - \Pi_\theta\|_v \leq c_\theta \rho_\theta^n$$

implies that, for any $h \in \mathcal{D}_v$,

$$|(P_\theta^n - \Pi_\theta)h| \leq \bar{c}_\theta \rho_\theta^n v,$$

with

$$\bar{c}_\theta \triangleq c_\theta \|h\|_v.$$

Hence,

$$\sum_{n=0}^{\infty} |(P_\theta^n - \Pi_\theta)h| =: \hat{h} \leq \frac{\bar{c}_\theta}{1 - \rho_\theta} v.$$

Note that $[\bar{c}_\theta/(1 - \rho_\theta)]v \in \mathcal{D}_v$. Because for any measurable mapping g with $|g| \leq h$ for some $h \in \mathcal{D}_v$ it follows that $g \in \mathcal{D}_v$, we conclude $\hat{h} \in \mathcal{D}_v$ and condition (iv) (a) holds. Moreover, because $|g| \in \mathcal{D}_v$ implies $g \in \mathcal{D}_v$, condition (iv) (b) holds as well. With relation (9) we have that

$$(\Pi_{\theta+\Delta} - \Pi_\theta)P_\theta^k h = (\Pi_{\theta+\Delta} - \Pi_\theta)(P_\theta^k - \Pi_\theta)h.$$

Condition (iii) then follows from (20) of Theorem 4.1. $\square$

REMARK 4.1. In Heidergott and Hordijk [16], $\|\cdot\|_v$ -norm differentiability of π_θ has been established. Generally speaking, $\mathcal{D}$ -differentiability is a weaker condition than $\|\cdot\|_v$ -norm differentiability. Studying the relationship between the two notions is the topic of further research.

REMARK 4.2. Note that Theorem 4.2 remains true for $\mathcal{D}$ a subset of $\mathcal{D}_v$ if it is assumed that (ii) of Theorem 1 holds for $\mathcal{D}$. For example, taking $\mathcal{D}$ to be the subset of all continuous functions in $\mathcal{D}_v$ yields the existence of the derivative of Markov kernels such as $P_\theta(s; \cdot) = \mathcal{U}_{[s, s+\theta]}(\cdot)$, with $s \in \mathbb{R}$.

4.2. Chains with an atom. Throughout this section, we assume that Conditions (C.1)–(C.3) are satisfied. The setup is as in the previous section, with the additional assumption that the chain possesses an atom, say α . The expression of the stationary distribution for a regenerative process is well known (see Ross [29]):

$$\pi_\theta h = \frac{1}{\mathbb{E}[\tau_\theta(s)]} \mathbb{E} \left[\sum_{m=0}^{\tau_\theta(s)-1} h(X_\theta(s, m)) \right], \quad s \in \alpha,$$

where $\tau_\theta(s)$ is the recurrence time to atom α , i.e., $\tau_\theta(s) = \min\{m \geq 1: X_\theta(s, m) \in \alpha\}$ and $\tau_\theta(s) = \infty$ if $X_\theta(s, m) \neq \alpha$ for $m \geq 1$. With the notation $(Q_\theta(s))f = \int_{S \setminus \alpha} P_\theta(s; dy)f(y)$, $s \in S$, we can write

$$\mathbb{E} \left[\sum_{m=0}^{\tau_\theta(s)-1} h(X_\theta(s, m)) \right] = \sum_{m=0}^{\infty} (Q_\theta(s))^m h, \quad (21)$$

and for any $h \in \mathcal{D}_v$ the above expression is finite. Hence, with e the vector with all components equal to 1,

$$\mathbb{E}[\tau_\theta(s)] = \sum_{m=0}^{\infty} (Q_\theta(s))^m e,$$

and thus, provided that $s \in \alpha$,

$$\pi_\theta h = \frac{\sum_{m=0}^{\infty} (Q_\theta(s))^m h}{\sum_{m=0}^{\infty} (Q_\theta(s))^m e}.$$

Using the taboo representation, the deviation operator can be obtained as follows:

$$D_\theta = (I - \Pi_\theta) \sum_{m=0}^{\infty} (Q_\theta(s))^m (I - \Pi_\theta);$$

see Hordijk and Spieksma [23], where a proof can be found for the denumerable state-space. Elaborating on (9), it follows that $P'_\theta \Pi_\theta h = 0$ for any $h \in \mathcal{D}_v$, which gives

$$P'_\theta (I - \Pi_\theta) \sum_{m=0}^{\infty} (Q_\theta(s))^m (I - \Pi_\theta) = P'_\theta \sum_{m=0}^{\infty} (Q_\theta(s))^m (I - \Pi_\theta).$$

We thus obtain the following representation for the derivative

$$\forall h \in \mathcal{D}_v: \quad \pi'_\theta h = \pi_\theta \sum_{m=0}^{\infty} P'_\theta (Q_\theta(s))^m (I - \Pi_\theta) h. \quad (22)$$

5. Gradient estimation. In this section, we illustrate how our result can be applied to gradient estimation. We assume that the conditions put forward in Theorem 3.1 are satisfied for $\mathcal{D}$. The starting point will be our formula $\pi'_\theta = \pi_\theta \sum_{n \geq 0} P'_\theta P_\theta^n$ for the $\mathcal{D}$ -derivative of the stationary distribution. In order to facilitate this formula for gradient estimation, two problems have to be addressed: (1) translating P'_θ into a proper Markov transition kernel, and (2) turning the deviation operator into a proper simulation experiment.

Our approach to (1) is to replace P'_θ by an instance of a $\mathcal{D}$ -derivative. Specifically, we will consider two Markov chains, called *phantom Markov chains* (or *phantoms* for short), one with initial distribution $\pi_\theta P_\theta^+$ and one with initial distribution $\pi_\theta P_\theta^-$. This will be discussed in §5.1. In the fortunate case, however, that condition (SF) holds, we may replace P'_θ by P_θ multiplied by the score function, which simplifies the resulting estimator considerably, as we will show in §5.3. As for (2), Conditions (C.1)–(C.3) imply ergodicity with geometric rate, and the infinite sum can be approximated by a sufficiently large finite sum, which will be explained in §5.1. If the Markov chain is Harris recurrent with an atom, one can construct a coupling and replace the infinite sum by observing the Markov chain over a random number of transitions. This will be addressed in §5.2.

5.1. Phantom estimators for general chains. Throughout this section we assume that the conditions put forward in Theorem 3.1 are in force for $\mathcal{D}$. Let P_θ be $\mathcal{D}$ -differentiable with $\mathcal{D}$ -derivative $(c_{P_\theta}, P_\theta^+, P_\theta^-)$ (see Definition 2.4), and let $X_\theta^\pm(s, n)$, with initial state s , evolve according to the kernel

$$P_\theta^\pm P_\theta^n,$$

or, equivalently, for $n > 0$, let the transition from $X_\theta^\pm(s, n)$ to $X_\theta^\pm(s, n+1)$ be governed by P_θ , whereas the transition from $X_\theta^+(s, 0)$ to $X_\theta^+(s, 1)$ is governed by P_θ^+ and that from $X_\theta^-(s, 0)$ to $X_\theta^-(s, 1)$ by P_θ^- , respectively. Hence, for $n > 0$, both $X_\theta^+(s, n)$ and $X_\theta^-(s, n)$ are driven by the same Markov kernel. The chains $\{X_\theta^+(s, n): n > 0\}$ and $\{X_\theta^-(s, n): n > 0\}$ are called *phantoms*. More specifically, let the phantom chain $Z_\theta(s, n) = (X_\theta^+(s, n), X_\theta^-(s, n))$, for $n \geq 0$, be given as follows: for $A, B \in \mathcal{T}$, $\mathbb{P}(Z_\theta(s, 1) \in A \times B) = P_\theta^+(s; A)P_\theta^-(s; B)$ and, for $n \geq 1$, $\mathbb{P}(Z_\theta(s, n+1) \in A \times B | Z_\theta(s, n) = z) = P_\theta(z; A)P_\theta(z; B)$, for $z \in S$. While this kind of independent coupling is always possible, one can often construct a more efficient coupling through, for example, using common random numbers.

The name “phantom” originates from Suri and Cao [31], where a queueing network was studied, the perturbed version of which had fewer customers than the nominal version. The customers, which are present in the nominal system but not in the perturbed system, act like “phantom customers” in the perturbed system. In a broader sense, the name “phantom” refers to a version of a stochastic process that differs from the nominal one in a certain aspect and was coined in Brémaud and Vázquez-Abad [4].

Let

$$\mathbf{D}_\theta = \sum_{n=0}^{\infty} P'_\theta P_\theta^n = P'_\theta D_\theta,$$

which yields $\pi'_\theta = \pi_\theta \mathbf{D}_\theta$ as the expression for the derivative of π_θ . The above formula extends the result in Heidergott and Cao [15] to Markov chains with general state-space. Elaborating on the phantom chains $X_\theta^\pm(s, m)$, we obtain

$$(\mathbf{D}_\theta g)(s) = \sum_{n=0}^{\infty} \mathbb{E}[c_{P_\theta}(s)(g(X_\theta^+(s, n)) - g(X_\theta^-(s, n)))],$$

for $g \in \mathcal{D}$, and eventually

$$\pi_\theta \sum_{n=0}^{\infty} P'_\theta P_\theta^n g = \mathbb{E}_{\pi_\theta}[(\mathbf{D}_\theta g)(X_\theta)], \quad (23)$$

where X_θ is distributed according to π_θ , and to indicate this fact, we write $\mathbb{E}_{\pi_\theta}$ for the expectation operator on the right-hand side.

The right-hand side of (23) is not suited as gradient estimator because it contains π_θ , which is generally unknown, and it contains an infinite sum. Following the train of thought put forward in Pflug [27], the expression on the right-hand side of (23) can be approximated as follows. Introduce the truncated version of $\mathbf{D}_\theta$ defined as follows. For $N > 0$, set

$$\mathbf{D}_\theta^N = \sum_{n=0}^N P'_\theta P_\theta^n,$$

which is shorthand notation for

$$(\mathbf{D}_\theta^N g)(s) = \sum_{n=0}^N \mathbb{E}[c_{P_\theta}(s)(g(X_\theta^+(s, n)) - g(X_\theta^-(s, n)))],$$

for $s \in S$ and $g \in \mathcal{D}$. Furthermore, for $s \in S$ fixed, let $P_{\theta, s}^m$ denote the distribution of the chain after m transition started in state s . Provided that Conditions (C.1)–(C.3) are in force, it holds

$$\begin{aligned} |P_{\theta, s}^M \mathbf{D}_\theta^N g - \pi_\theta \mathbf{D}_\theta g| &\leq |(P_{\theta, s}^M - \pi_\theta) \mathbf{D}_\theta^N g| + |\pi_\theta (\mathbf{D}_\theta^N g - \mathbf{D}_\theta g)| \\ &\leq \left(\rho_\theta^M \|P'_\theta\|_v \frac{1 - \rho_\theta^{N+1}}{1 - \rho_\theta} + \|\pi_\theta\|_v \|P'_\theta\|_v \frac{\rho_\theta^{N+1}}{1 - \rho_\theta} \right) c_\theta v(s). \end{aligned} \quad (24)$$

Hence, the truncated estimator $P_{\theta, s}^M \mathbf{D}_\theta^N g$ converges as M and N tend to infinity exponentially fast towards the desired derivative. Pflug proposed this type of approximation in [27]. Unfortunately, the numerical values for $\|\pi_\theta\|_v$, $\|P'_\theta\|_v$, c_θ and ρ_θ are only known in special cases.

As we will illustrate in the following section, under the additional assumption that the Markov chain possesses an atom, an exact expression for π'_θ can be derived using a random horizon for the simulation experiment.

5.2. Phantom estimators for chains with an atom. We now turn to the situation where $\mathbf{X}(\theta)$ possesses an atom, denoted by α , such that α is independent of θ . Let $X_\theta(\alpha, n)$ denote the Markov chain started in α and denote the first entrance time after $n = 1$ of the chain into α by τ_θ . Denote by $\tau_\theta^\pm(s)$ the first time that $X_\theta^+(s, m)$ and $X_\theta^-(s, m)$ simultaneously hit α :

$$\tau_\theta^\pm(s) = \inf\{m \geq 1: X_\theta^+(s, m) \in \alpha \text{ and } X_\theta^-(s, m) \in \alpha\}$$

and $\tau_\theta^\pm(s) = \infty$ if the set on the right-hand side is empty.

REMARK 5.1. The stopping time $\tau_\theta^\pm$ can alternatively be defined as

$$\tau_\theta^\pm(s) = \inf\{m \geq 1: X_\theta^+(s, m) = X_\theta^-(s, m)\},$$

and $\tau_\theta^\pm(s) = \infty$ if the set on the right-hand side is empty. This definition will, for example, be preferred for chains with discrete state-space.

For any $g \in \mathcal{D}$, we obtain:

$$\begin{aligned} (\mathbf{D}_\theta g)(s) &= \sum_{n=0}^{\infty} (\mathbb{E}[c_{P_\theta}(s)(g(X_\theta^+(s, n)) - g(X_\theta^-(s, n)))1_{\tau_\theta^\pm(s) > n}]) \\ &= \mathbb{E}\left[c_{P_\theta}(s) \left(\sum_{n=0}^{\tau_\theta^+(s)} g(X_\theta^+(s, n)) - \sum_{n=0}^{\tau_\theta^-(s)} g(X_\theta^-(s, n)) \right)\right]. \end{aligned} \quad (25)$$

If conditions (C.1)–(C.3) hold, then the finiteness of the above expression follows from Theorem 2 and the following inequalities:

$$\mathbb{E}[|c_{P_\theta}(s)(g(X_\theta^+(s, n)) - g(X_\theta^-(s, n)))1_{\tau_\theta^\pm(s) > n}|] \leq \|P'_\theta P_\theta^n\|_v v(s) \|g\|_v.$$

and

$$\|P'_\theta P_\theta^n\|_v = \|P'_\theta P_\theta^n (I - \Pi_\theta)\|_v \leq \|P'_\theta\|_v c_\theta \rho_\theta^n < \infty.$$

The expression for the stationary distribution for a chain with an atom is well known (see Ross [29]) and, provided the conditions put forward in Theorem 4.2 hold, we obtain from (23) for any $g \in \mathcal{D}$:

$$\frac{d}{d\theta} \mathbb{E}_{\pi_\theta}[g(X_\theta)] = \frac{1}{\mathbb{E}[\tau_\theta]} \mathbb{E} \left[\sum_{n=0}^{\tau_\theta-1} (\mathbf{D}_\theta g)(X_\theta(\alpha, n)) \right]. \quad (26)$$

Inserting the expression for $\mathbf{D}_\theta g$ provided in (25), the overall estimator becomes

$$\begin{aligned} \frac{d}{d\theta} \mathbb{E}_{\pi_\theta}[g(X_\theta)] = \frac{1}{\mathbb{E}[\tau_\theta]} \mathbb{E} \left[\sum_{n=0}^{\tau_\theta-1} c_{P_\theta}(X_\theta(\alpha, n)) \left(\sum_{m=0}^{\tau_\theta^+(X_\theta(\alpha, n))-1} g(X_\theta^+(X_\theta(\alpha, n), m)) \right. \right. \\ \left. \left. - \sum_{m=0}^{\tau_\theta^-(X_\theta(\alpha, n))-1} g(X_\theta^-(X_\theta(\alpha, n), m)) \right) \right]; \end{aligned} \quad (27)$$

a formal proof can be found in Heidergott et al. [14], which is the report version of this paper.

5.3. Single-run estimators. Assume that the conditions put forward in Theorem 4.2 are satisfied and that the chain possesses an atom, say, α independent of θ . Under Condition (SF), it holds that $P'_\theta = \text{SF}_\theta P_\theta$ and Equation (22) reads

$$\pi'_\theta h = \pi_\theta \sum_{n=0}^{\infty} (\text{SF}_\theta P_\theta)(Q_\theta)^n (I - \Pi_\theta) h, \quad (28)$$

for $h \in \mathcal{D}_v$, where Q_θ is the taboo version of P_θ with respect to α . Recall that $X_\theta(s, m)$ denotes the m th state of the Markov chain with initial value s , i.e., $X_\theta(s, 0) = s$. Using Equation (28) and the fact that $X_\theta(X_\theta(s, n), m) = X_\theta(s, n+m)$, we obtain for $h \in \mathcal{D}_v$:

$$\begin{aligned} \pi'_\theta h &= \mathbb{E}_{\pi_\theta} \left[\text{SF}_\theta(X_\theta; X_\theta(X_\theta, 1)) \left(\sum_{m=0}^{\tau_\theta(X_\theta(X_\theta, 1))-1} h(X_\theta(X_\theta, m+1)) - \tau_\theta(X_\theta(X_\theta, 1)) \pi_\theta h \right) \right] \\ &= \frac{1}{\mathbb{E}[\tau_\theta(s)]} \mathbb{E} \left[\sum_{n=0}^{\tau_\theta(s)-1} \text{SF}_\theta(X_\theta(s, n); X_\theta(s, n+1)) \left(\sum_{m=0}^{\tau_\theta(X_\theta(s, n+1))-1} h(X_\theta(s, n+m+1)) - \tau_\theta(X_\theta(s, n+1)) \pi_\theta h \right) \right], \end{aligned}$$

for $h \in \mathcal{D}_v$ and $s \in \alpha$; see the previous section for details. At $\tau_\theta(X_\theta(s, \tau_\theta(s)))$, the Markov chain X_θ enters α and a new cycle starts. The expected value of the last factor in the above expression is equal to 0 for initial state $s \in \alpha$. Moreover, for $n < \tau_\theta(s)$ it holds $\tau_\theta(s) = \tau_\theta(X_\theta(s, n))$ and the above expression for the derivative simplifies to

$$\pi'_\theta h = \frac{1}{\mathbb{E}[\tau_\theta(s)]} \mathbb{E} \left[\sum_{n=0}^{\tau_\theta(s)-1} \text{SF}_\theta(X_\theta(s, n); X_\theta(s, n+1)) \left(\sum_{m=0}^{\tau_\theta(s)-1} h(X_\theta(s, m+n+1)) - \tau_\theta(s) \pi_\theta h \right) \right] \quad (29)$$

for $h \in \mathcal{D}_v$ and $s \in \alpha$, which recovers the formula of Glynn and L'Ecuyer [11].

A few remarks on the relationship between phantom and single-run formulae via likelihood ratios are in order here. Comparing the single-run formula in (29) with the phantom formula (27), one notices that the single-run formula involves the term $\pi_\theta h$. Because $\pi_\theta h$ is typically not known, it has to be estimated along with estimating the derivative. The resulting estimator is asymptotically unbiased; see Glynn et al. [12]. No type of estimator dominates the other. Single-run estimators are typically easy to implement, but usually tend to suffer from significant variance. A phantom estimator consumes computer storage for keeping track of the parallel phantoms, which makes the estimator usually cumbersome to implement, but typically has a low variance. As a rule of thumb, a single-run estimator should be applied whenever possible. However, counterexamples to this rule can be found in the literature as well. For instance, Pflug discusses in [28, §4.3.2] a Markov chain example for which a phantom estimator outperforms a single-run estimator. Heidergott and Vázquez-Abad present in [19] a public transportation problem where a phantom estimator has considerably less variance than the single-run estimator. A thorough analysis of the relationship between single-run estimators and phantom-type ones is a challenging subject for further research.

We conclude the discussion of estimators by mentioning that Glynn and L'Ecuyer provide in [11] a single-run estimator for the general case of a Harris recurrent chain, where “general” means that the existence of an atom

is not assumed. Moreover, Equation (29) was established in Glynn and L'Ecuyer [11] for performance mappings out of the set

$$\mathcal{D}_\delta = \{g: |g| \leq f\}$$

for some f with $\pi_\theta f^{1+\delta} < \infty$, with $\delta > 0$. Notice that for given $\delta > 0$, a suitable choice of λ in (18) implies $\pi_\theta v^{1+\delta} < \infty$, which gives $\mathcal{D}_\delta \subset \mathcal{D}_v$.

6. Examples. Consider a single-server queue with i.i.d. exponentially distributed service times with rate μ . Service times and interarrival times are independent, and let the interarrival times be a sequence of i.i.d. random variables following a Cox distribution with rates η_j , $j = 1, 2$, and parameter θ , that is, the interarrival times consist with probability $1 - \theta$ of a single exponentially distributed stage with rate η_1 , and a second stage with rate η_2 follows with probability θ . Let h_μ denote the density of the exponential distribution with rate μ , and write $\mathcal{E}_\mu$ for the distribution. Denoting the density of the sum of two independent exponentially distributed random variables with rate η_1 and η_2 by $h_{\eta_1, \eta_2}(x)$ and the corresponding distribution function by $\mathcal{E}_{(\eta_1, \eta_2)}$, the density of the interarrival times is given by

$$h_\theta(x) = (1 - \theta)h_{\eta_1}(x) + \theta h_{\eta_1, \eta_2}(x), \quad x \geq 0. \quad (30)$$

The parameter of interest is θ . Observe that, for $\theta = 1$, $h_\theta(x) = h_{\eta_1, \eta_2}(x)$ and the interarrival times follow a phase-type distribution, whereas for $\theta = 0$, $h_\theta(x) = h_{\eta_1}(x)$ and the interarrival times follow an exponential distribution.

6.1. Discrete state-space. Let $X_\theta(n) = (X_\theta(1, n), X_\theta(2, n))$ be the state of the embedded discrete-time Markov chain of the queueing system, with $X_\theta(1, n) \in \mathbb{N}$ the total number of customers in the system, and $X_\theta(2, n) \in \{1, 2\}$ the stage of the interarrival time. Let

$$P_\theta((k, i); (k', i')) = \mathbb{P}(X_\theta(m+1) = (k', i') | X_\theta(m) = (k, i)),$$

for $(k, i), (k', i') \in \mathbb{N} \times \{1, 2\}$. Then, the probability that an arrival occurs is

$$P_\theta((k, 1); (k+1, 1)) = (1 - \theta) \frac{\eta_1}{\eta_1 + \mu 1_{k>0}},$$

$$P_\theta((k, 2); (k+1, 1)) = \frac{\eta_2}{\eta_2 + \mu 1_{k>0}},$$

the probability that the state of the interarrival time jumps from Stage 1 to 2 is

$$P_\theta((k, 1); (k, 2)) = \theta \frac{\eta_1}{\eta_1 + \mu 1_{k>0}},$$

and the probability that a departure occurs is

$$P_\theta((k, 1); (k-1, 1)) = \frac{\mu}{\eta_1 + \mu},$$

$$P_\theta((k, 2); (k-1, 2)) = \frac{\mu}{\eta_2 + \mu}.$$

Set $P = P_1$ and $Q = P_0$, then

$$P_\theta = \theta P + (1 - \theta)Q. \quad (31)$$

For any θ the process is a discrete-time Markov chain that is irreducible, and hence any state is an atom. In particular, $(0, 1)$ is an atom for each of the processes.

Let us verify Conditions (C.1) to (C.3). We assume that for any θ the process is ergodic. By well-known results for the $G/M/1$ queue (see Cohen [7]), this requires that the mean service time must be smaller than the mean interarrival time. Hence, for all $0 \leq \theta \leq 1$,

$$(1 - \theta) \frac{1}{\eta_1} + \theta \left(\frac{1}{\eta_1} + \frac{1}{\eta_2} \right) > \frac{1}{\mu},$$

or

$$\frac{1}{\eta_1} > \frac{1}{\mu}.$$

For the Lyapunov function, we try a function from $S = \mathbb{N}_0 \times \{1, 2\}$ to $\mathbb{R}_+$ that is linear in the number of customers:

$$g(k, i) \triangleq ck, \quad (32)$$

for some $c > 0$ and $i = 1, 2$.

Easy calculation gives that $P_\theta g(s) - g(s) < 0$ for any $s = (k, i)$, with $k \geq 1$ and $i = 1, 2$, provided that $\mu > \eta_2$ and $\mu > \eta_1$.

We conclude that for any $c > 0$ the function g in (32) is a Lyapunov function with

$$V_\theta = \{(0, 1), (0, 2)\} \quad \text{for } \theta > 0$$

and

$$V_\theta = \{(0, 1)\} \quad \text{for } \theta = 0.$$

Condition (C.1) is thus satisfied with $m_\theta = 1$. Without loss of generality, we take $c = 1$ in the Lyapunov function g .

REMARK 6.1. While the ergodicity condition is

$$\frac{1}{\eta_1} > \frac{1}{\mu}, \quad (33)$$

we also required above that

$$\frac{1}{\eta_2} > \frac{1}{\mu}. \quad (34)$$

As shown in the appendix, we do not need relation (34) for satisfying Condition (C.1) if we take m_θ sufficiently large in relation (16).

For an ergodic Markov chain with atom, the Harris condition is automatically fulfilled, which implies that Condition (C.2) is satisfied. It is straightforward to check Condition (C.3). By Theorem 4.1 together with relation (18), we may choose the bounding function v as $v(s, i) \triangleq e^{\lambda s}$, for $(s, i) \in S$, for sufficiently small λ , which gives

$$\mathcal{D}_v = \{g: S \rightarrow \mathbb{R} \mid \exists r \in \mathbb{R}: |g(s, i)| \leq r \cdot e^{\lambda s}, (s, i) \in S\}.$$

Note that Theorem 4.1 implies that $\mathcal{D}_\nu \subset L^1(P_\theta, \Theta) \cap L^1(\Pi_\theta, \Theta)$. Only the transition out of state $(k, 1)$ depends on θ . Specifically, for $g \in \mathcal{D}_\nu$, it holds:

$$\frac{d}{d\theta} \sum_{r \in S} P_\theta((k, 1); r) g(r) = g(k, 2) \frac{\eta_1}{\eta_1 + \mu 1_{k>0}} - g(k+1, 1) \frac{\eta_1}{\eta_1 + \mu 1_{k>0}},$$

whereas

$$\frac{d}{d\theta} \sum_{r \in S} P_\theta((k, 2); r) g(r) = 0.$$

Set

$$P^+((k, 1); (k, 2)) = \frac{\eta_1}{\eta_1 + \mu 1_{k>0}} = P^-((k, 1); (k+1, 1)),$$

$$P^+((k, 1); (k+1, 1)) = 0 = P^-((k, 1); (k, 2)),$$

and for any other pair of states s, s' set $P_\theta^+(s; s') = P_\theta(s; s') = P_\theta^-(s; s')$. In words, under P^+ there are two possible events that can trigger the Markov chain to leave state $(k, 1)$: A departure takes place, or the phase of the interarrival is increased. Under P^- the two possible events that can trigger the Markov chain to leave state $(k, 1)$ are: A departure takes place, or an arrival occurs. Then, for any $g \in \mathcal{D}_\nu$ it holds that

$$\frac{d}{d\theta} \sum_{r \in S} P_\theta(s; r) g(r) = \sum_{r \in S} P^+(s; r) g(r) - \sum_{r \in S} P^-(s; r) g(r), \quad s \in S,$$

which implies that $(1, P^+, P^-)$ is a $\mathcal{D}_\nu$ -derivative of P_θ . Observe that $P^+ = P_1$ and $P^- = P_0$. Notice that the $\mathcal{D}_\nu$ -derivative is independent of θ , and one-sided derivatives at the boundary points $\theta = 0, 1$ are thus well defined.

A quick way of obtaining the above result is as follows. We revisit the representation of P_θ as the mixture of the kernels P and Q in (31). For any $g \in \mathcal{D}_\nu$, it holds that $\int P(s; du)g(u)$ and $\int Q(s; du)g(u)$ exist and are finite. Hence, by Example 2.2, P_θ is $\mathcal{D}_\nu$ -differentiable with $\mathcal{D}_\nu$ -derivative $(1, P, Q)$, where $P^+ = P$ and $P^- = Q$.

Because the $\mathcal{D}_\nu$ -derivative of P_θ is independent of θ , P_θ is $\mathcal{D}_\nu$ -Lipschitz continuous at any $\theta \in [0, 1]$. Hence, provided that $1/\eta_1 > 1/\mu$, Theorem 4.2 applies and the estimators presented in §5.2 and §5.3 are unbiased (respectively, asymptotically unbiased) for any $g \in \mathcal{D}_\nu$.

6.2. Continuous state-space. Let $W_\theta(n)$ be the waiting time of the n th customer arriving to the system. We have $S = \mathbb{R}$ and we take the usual norm on $\mathbb{R}$ for $\|\cdot\|_S$, i.e., for any $x \in \mathbb{R}$ we have $\|x\|_S = |x|$. Let $\{A_\theta(n)\}$ be the i.i.d. sequence of interarrival times and $\{S(n)\}$ the i.i.d. sequence of service times, respectively. Lindley's recursion yields:

$$W_\theta(n+1) = \max(W_\theta(n) + S(n) - A_\theta(n+1), 0), \quad n \geq 1,$$

and $W_\theta(1) = 0$. For $w > 0$, the transition kernel for the waiting times is given by

$$P_\theta(u; (0, w]) = \int_0^\infty \left(\int_{\max(u+s-w, 0)}^{u+s} h_\theta(a) da \right) f(s) ds$$

and

$$P_\theta(u; \{0\}) = \int_0^\infty \left(\int_{u+s}^\infty h_\theta(a) da \right) f(s) ds,$$

with h_θ as in (30) and f the density of an exponential distribution with parameter μ . For the ergodicity of Markov chain $W_\theta(n)$, $0 \leq \theta \leq 1$, we need that (see Cohen [7])

$$\frac{1}{\mu} < \frac{1}{\eta_1}. \quad (35)$$

Stability of the queue on $\Theta = [0, 1]$ is implied by stability of the queue at $\theta = 0$. By (35) the average service time for any $\theta \in [0, 1]$ is larger than the average interarrival time. Hence, for c sufficiently large, it holds

$$\epsilon \triangleq \inf_\theta \int_0^c a h_\theta(a) da - \int_0^\infty s f(s) ds > 0.$$

Choose g such that $g(u) = u$, for $u \geq 0$, and $g(u) = 0$ otherwise. For $u \geq c$, it holds that

$$\begin{aligned} \int P_\theta(u; dy) g(y) &= \int_0^\infty \left(\int_0^{u+s} g(u+s-a) h_\theta(a) da \right) f(s) ds \\ &= \int_0^\infty \left(\int_0^{u+s} -a h_\theta(a) da + u + s \right) f(s) ds \\ &\leq \int_0^\infty \left(\int_0^c -a h_\theta(a) da + u + s \right) f(s) ds \\ &= - \int_0^c a h_\theta(a) da + u + \int_0^\infty s f(s) ds \\ &\leq u - \epsilon. \end{aligned}$$

Hence, g satisfies Condition (C.1) with $V_\theta = [0, c]$. This Markov chain has an atom, and it is straightforward to check Conditions (C.2) and (C.3). The bounding function is $v(s) \triangleq e^{\lambda s}$, $s \in \mathbb{R}_+$, which gives

$$\mathcal{D}_v = \{g: S \rightarrow \mathbb{R} \mid \exists r \in \mathbb{R}: |g(s)| \leq r \cdot e^{\lambda s}, s \in S\}.$$

Note that Theorem 4.1 implies that

$$\mathcal{D}_v \subset L^1(P_\theta, \Theta) \cap L^1(\Pi_\theta, \Theta).$$

As in the previous example, the kernel can be written as a mixture of two kernels $P = P_1$ and $Q = P_0$ as follows: $P_\theta = \theta P + (1 - \theta)Q$, for $\theta \in [0, 1]$. Specifically, under P , the interarrival times are distributed according to $\mathcal{E}_{(\eta_1, \eta_2)}$, and under Q , the interarrival times have distribution $\mathcal{E}_{\eta_1}$. Hence, P_θ falls into the setup of Example 2.2, and it readily follows that P_θ is $\mathcal{D}_v$ -differentiable with $\mathcal{D}_v$ -derivative $(1, P, Q)$. Because the $\mathcal{D}_v$ -derivative of P_θ is independent of θ , P_θ is $\mathcal{D}_v$ -Lipschitz continuous at any $\theta \in [0, 1]$. Hence, if the system is stable at $\theta = 1$ (see (35)), then Theorem 4.2 applies, and the estimators presented in §5.2 and §5.3 are unbiased (respectively, asymptotically unbiased) for any $g \in \mathcal{D}_v$.

Appendix. Let $g(k, i) = k$ for $(k, i) \in S$, and suppose that relation (33) is satisfied, but relation (34) not, then

$$\eta_2 > \mu > \eta_1. \quad (\text{A1})$$

Below we assume that the starting state is large and that the queueing system is never empty on the time interval we consider. Note that in this case the transitions of the phases are not dependent on the state. Hence, the long-run fraction of states of the embedded Markov chain for which the interarrival process is in Phase 1 and Phase 2 can be computed from the two-state Markov chain with transition probabilities

$$p(1, 2) = 1 - p(1, 1) = \frac{\theta\eta_1}{\eta_1 + \mu}$$

and

$$p(2, 1) = 1 - p(2, 2) = \frac{\eta_2}{\eta_2 + \mu}.$$

Calculation gives that the stationary probabilities π_1, π_2 are

$$\pi_1 = \frac{\eta_2(\eta_1 + \mu)}{\eta_2(\eta_1 + \mu) + \theta\eta_1(\eta_2 + \mu)}$$

and

$$\pi_2 = \frac{\theta\eta_1(\eta_2 + \mu)}{\eta_2(\eta_1 + \mu) + \theta\eta_1(\eta_2 + \mu)}.$$

By calculation, for $k \geq 1$,

$$\mathbb{E}[g(X_\theta((k, 1), 1))] - g(k, 1) = \frac{(1 - \theta)\eta_1 - \mu}{\eta_1 + \mu}$$

and

$$\mathbb{E}[g(X_\theta((k, 2), 1))] - g(k, 2) = \frac{\eta_2 - \mu}{\eta_2 + \mu};$$

see Heidergott et al. [14] for details. We have thus shown that

$$\pi_1(\mathbb{E}[g(X_\theta((k, 1), 1))] - g(k, 1)) - \pi_2(\mathbb{E}[g(X_\theta((k, 2), 1))] - g(k, 2)) < 0$$

if and only if

$$\eta_2(\eta_1 - \mu) - \theta\eta_1\mu < 0. \quad (\text{A2})$$

Writing $\mathbb{E}[g(X_\theta(s, t))] - g(s)$ as a telescope sum

$$\mathbb{E}[g(X_\theta(s, t))] - g(s) = \mathbb{E}\left[\sum_{l=1}^t (\mathbb{E}[g(X_\theta(s, l)) | X_\theta(s, l-1)] - g(X_\theta(s, l-1)))\right] \quad (\text{A3})$$

it follows from (A2) together with assumption (A1) that, for t sufficiently large, the right-hand side of (A3) is negative for any state $s = (k, i)$, with $k \geq t$. Take $m_\theta = t$, then relation (16) is satisfied for $d = t$. We conclude that indeed the ergodicity condition is sufficient for Condition (C.1).

Note that we have shown that the linear function g satisfies the Condition (C.1), which we use for our analysis. For proving that the Markov chain is positive recurrent the limiting-average drift argument in Meyn and Tweedie [26, Chapter 19] could also be used.

Acknowledgments. Part of the work was done while the first author was with TU Eindhoven, Department of Mathematics and Computing Science, Division of Mathematics, Eindhoven, The Netherlands. Part of the work of the third author was done while this author was with EURANDOM, Eindhoven, The Netherlands.

References

- [1] Baccelli, F., D. Hong. 2000. Analytic expansions of $(\max, +)$ Lyapunov exponents. *Ann. Appl. Probab.* **10** 779–827.
- [2] Braker, J. 1991. Max-algebra modelling and analysis of time-dependent transportation networks. *Proc. 1st Eur. Control Conf.*, Grenoble, France, 1831–1836. Hermes, Paris, France.
- [3] Brémaud, P. 1992. Maximal coupling and rare perturbation analysis. *Queueing Systems: Theory Appl.* **11** 307–333.
- [4] Brémaud, P., 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: Theory Appl.* **10** 49–270.
- [5] Borovkov, A. A., A. Hordijk. 2000. On normed ergodicity of Markov chains. Technical report MI 2000-40 Leiden University, Leiden, The Netherlands.
- [6] Borovkov, A. A., A. Hordijk. 2004. Characterization and sufficient conditions for normed ergodicity of Markov chains. *Adv. Appl. Probab.* **36** 227–242.
- [7] Cohen, J. W. 1969. *The Single Server Queue*. North-Holland, Amsterdam, The Netherlands.
- [8] Dekker, R., A. Hordijk. 1988. Average, sensitive and Blackwell optimal policies in denumerable Markov decision chains with unbounded rewards. *Math. Oper. Res.* **13** 395–421.

- [9] Dekker, R., A. Hordijk, F. M. Spieksma. 1994. On the relation between recurrence and ergodicity properties in denumerable Markov decision chains. *Math. Oper. Res.* **19** 539–559.
- [10] Glynn, P. 1990. Likelihood ratio gradient estimation for stochastic systems. *Comm. ACM* **33** 75–84.
- [11] Glynn, P., P. L'Ecuyer. 1995. Likelihood ratio gradient estimation for stochastic recursions. *Adv. Appl. Probab.* **27** 1019–1053.
- [12] Glynn, P., P. L'Ecuyer, M. Adès. 1991. Gradient estimation for ratios. B. Nelson, W. Kelton, G. Clark, eds. *Proc. 1991 Winter Simulation Conf.* Winter Simulation Conference, San Diego, CA, 986–993.
- [13] Heidergott, B., A. Hordijk, H. Weisshaupt. 2003. Derivatives of Markov kernels and their Jordan decomposition. EURANDOM Report 2003-001, EURANDOM, Eindhoven, The Netherlands.
- [14] Heidergott, B., A. Hordijk, H. Weisshaupt. 2002. Measure-valued differentiation for stationary Markov chains. EURANDOM Report 2002-027, EURANDOM, Eindhoven, The Netherlands.
- [15] Heidergott, B., X. Cao. 2002. A note on the relation between weak derivatives and perturbation realization. *IEEE Trans. Automatic Control* **47** 1112–1115.
- [16] Heidergott, B., A. Hordijk. 2003. Taylor series expansions for stationary Markov chains. *Adv. Appl. Probab.* **35** 1046–1070.
- [17] Heidergott, B., A. Hordijk. 2004. Single-run gradient estimators via measure-valued differentiation. *IEEE Trans. Automatic Control* **49** 1843–1846.
- [18] Heidergott, B., F. Vázquez-Abad. 2000. Measure-valued differentiation for stochastic processes: The finite horizon case. EURANDOM Report 2000-033, EURANDOM, Eindhoven, The Netherlands.
- [19] Heidergott, B., F. Vázquez-Abad. 2004. Gradient estimation for a problem in public transportation: A comparison of SPA, SF and MVD. International Workshop on DES (WODES'04), Reims, France, 241–246.
- [20] Heidergott, B., R. de Vries. 2001. Towards a control theory for transportation networks. *Discrete Event Dynam. Systems* **11** 371–398.
- [21] Hordijk, A., R. Dekker. 1983. Average, sensitive and Blackwell optimal policies in denumerable Markov decision chains with unbounded rewards. Report No. 83-36, Institute of Applied Mathematics and Computing Science, Leiden University, Leiden, The Netherlands.
- [22] Hordijk, A., F. M. Spieksma. 1992. On ergodicity and recurrence properties of a Markov chain with an application to an open Jackson network. *Adv. Appl. Probab.* **24** 343–376.
- [23] Hordijk, A., F. M. Spieksma. 1994. A new formula for the deviation matrix. F. P. Kelly, ed. Chapter 36 in *Probability, Statistics and Optimization*. Wiley, New York.
- [24] Hordijk, A., A. A. Yushkevich. 1999. Blackwell optimality in the class of all policies in Markov decision chains with a Borel state space and unbounded rewards. *Math. Methods Oper. Res.* **50** 421–448.
- [25] Kushner, H., F. Vázquez-Abad. 1992. Estimation of the derivative of a stationary measure with respect to a control parameter. *J. Appl. Probab.* **29** 343–353.
- [26] Meyn, S. P., R. L. Tweedie. 1993. *Markov Chains and Stochastic Stability*. Springer, London, UK.
- [27] Pflug, G. 1992. Gradient estimates for the performance of Markov chains and discrete event processes. *Ann. Oper. Res.* **39** 173–194.
- [28] Pflug, G. 1996. *Optimisation of Stochastic Models*. Kluwer Academic Publishers, Boston, MA.
- [29] Ross, S. M. 1970. *Applied Probability Models with Optimization Applications*. Holden-Day, San Francisco, CA.
- [30] Rubinstein, R., A. Shapiro. 1993. *Discrete Event Systems: Sensitivity Analysis and Optimization by the Score Function Method*. John Wiley and Sons, New York.
- [31] Suri, R., X. Cao. 1983. The phantom customer and marked customer methods for optimization of closed queueing networks with blocking and general service times. *Performance Evaluation Rev.* **August** 243–256.