Measure-Valued Differentiation for Stochastic Processes:
The Finite Horizon Case

Bernd Heidergott

Vrije Universiteit Amsterdam and Tinbergen Institute
Department of Econometrics and Operations Research
De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands
email:bheidergott@feweb.vu.nl

Felisa J. Vázquez-Abad, member *GERAD*
Department of Mathematics and
ARC Special Research Centre for Ultra-Broadband Information Networks,
University of Melbourne, 3010 AUSTRALIA
email: vazquez@iro.umontreal.ca

Abstract

This paper addresses the problem of sensitivity analysis for finite horizon performance measures of general Markov chains. We derive closed form expressions and associated unbiased gradient estimators for derivatives of finite products of Markov kernels by measure-valued differentiation (MVD). In the MVD setting, derivatives of Markov kernels, called $\mathcal{D}$-derivatives, are defined with respect to an appropriately defined class of performance functions $\mathcal{D}$, such that for any performance measure $g \in \mathcal{D}$ the derivative of the integral of g with respect to the one step transition probability of the Markov chain exists. The MVD approach (1) yields results that that can be applied to performance functions out of a predefined class, (2) allows for a product rule of differentiation, that is, analyzing the derivative of the transition kernel immediately yields finite horizon results, (3) provides an operator language approach to differentiation of Markov chains and (4) clearly identifies the trade-off between the generality of performance classes that can be analyzed and the generality of the classes of measures (Markov kernels). The $\mathcal{D}$-derivative of a measure can be interpreted in terms of various (unbiased) gradient estimators and the product rule for $\mathcal{D}$-differentiation yields a product-rule for various gradient estimators.

1 Introduction

Many real-world systems in manufacturing, transportation, communication networks, or finance can be modeled by general state-space Markov chains, such as generalized semi-Markov processes. The past two decades have witnessed an

increased attention for the study of discrete event driven systems (see [5, 6, 18]), with the aim of finding better and more efficient control methods. 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 satisfying certain conditions [14]. This paper addresses this problem using a general probabilistic framework to study sensitivity analysis of finite horizon performance measures for Markov chains.

The motivation for the present paper is to establish an unified mathematical framework for gradient estimation of Markov chains. Our approach summarizes the proof techniques and ideas that are known in the literature in order to establish a meta theory of gradient estimation. In our view, such a meta theory has to meet the following prerequisites:

(i) a general and meaningful concept of differentiation is provided,

(ii) a product rule of differentiation for this concept of differentiation holds,

(iii) statements obtained within the meta theory can be translated into unbiased gradient estimators, and finally

(iv) random horizon problems can be dealt with.

The focus of this paper is on measure-valued differentiation (MVD) for general state-space Markov chains. We will show that MVD provides the means to establish a meta theory of gradient estimation. In particular, we address here topics *(i)* to *(iii)*. The fact the MVD is an operator language approach will prove most helpful when going from *(i)* to *(ii)*. Topic *(iv)* is beyond the scope of this paper and its analysis is postponed to a follow-up paper ([11]).

The paper is organized as follows. Section 2 discusses various approaches to the gradient estimation problem. We illustrate to what extent these methods already have features of the intended meta theory. In Section 3, we introduce measure-valued derivatives and we establish the key technical result, which is the product rule of measure valued differentiation. In Section 4, we show how the conditions of the product rule can be verified in various scenarios that are of importance in applications. For example, when only bounded functions are considered, then the conditions for the product rule can be expressed in a very simple manner. While Section 4 prepares the 'input' of the product rule, Section 5 is concerned with the 'output' of the product rule: This last section shows how the expressions produced by the product rule (and containing signed measures) can be turned into various types of gradient estimators leading to estimators typically obtained from SPA, the Score function method or weak derivatives.

2   Background and Motivation

Let $\{X_\theta(n)\}$, with $\theta \in \Theta \subset \mathbb{R}$, be a Markov Chain with (arbitrary) state space S defined on a common probability space $(\Omega, \mathfrak{F}, \mathbb{P})$, where Θ is the set of parameters (generally these are the control parameters) , such that $(\Omega, \mathfrak{F}, \mathbb{P})$ is independent of θ and $X_\theta(n)$ is well defined on Θ. The problem of sensitivity

analysis can be phrased as follows: For performance functions $g : S \to \mathbb{R}$ find conditions, such that for $n \in \mathbb{N}$

$$\frac{d}{d\theta}\mathbb{E}\Big[g(X_\theta(n),\ldots,X_\theta(1))\Big] \qquad (1)$$

exists and can be obtained in a closed form expression. When derivatives are defined, it is sufficient that Θ be a neighborhood of the point θ of interest. The term *sensitivity analysis* is often used to refer to (unbiased) estimators, called gradient estimators, of this derivative. The last two decades have witnessed a great interest in the problem of finding unbiased estimators for the expression in (1), see for example [3, 5, 6, 13, 15, 17, 18]. The methods available are legion and even experts find it difficult to oversee the various methods and estimators. However, the following three major approaches can be identified: *smoothed perturbation analysis* (SPA) (we remark to the expert reader that we consider infinitesimal perturbation analysis (IPA) as a special case of SPA), *score function* and *weak derivatives*, which will be described in what follows in more detail.

2.1 Smoothed Perturbation Analysis

In the sample-path analysis setting, the dependency of the expectation in the parameter θ is expressed entirely through the performance. If the sample performance is almost surely Lipschitz continuous in θ, then the sample path derivative $dg(X_\theta(n,\omega))/d\theta$ is unbiased for the gradient (1), yielding the so-called infinitesimal perturbation analysis, see [6]. In the presence of discontinuities, conditioning can be used to integrate (or "smooth out") such discontinuities, see [5].

Two approaches can be identified. The first seeks an estimator of the form

$$\frac{d}{d\theta}\mathbb{E}[g(X_\theta(k+1))|X_\theta(k)=x] \;=\; \mathbb{E}\left[\frac{d}{d\theta}\mathbb{E}[g(X_\theta(k+1))|X_\theta(k)=x,\mathcal{G}]\,\bigg|\, X_\theta(k)=x\right],$$

where $\mathcal{G}$ is a smoothing σ-field, and $H(\theta) = \mathbb{E}[g(X_\theta(k+1))|X_\theta(k)=x,\mathcal{G}]$ is a.s. Lipschitz continuous. It is often very difficult to identify such conditioning fields in practice.

The second approach [7] prescribes an analysis of a perturbed path using the same trajectory ω for θ and for $\theta \pm \Delta$ and conditioning on the (rare) events where discontinuities may occur. This formulation implies that the *nominal* and *perturbed* processes $X_\theta, X_{\theta+\Delta\theta}$ share a common filtration. Let g be a real valued integrable function and call $\{X_\theta(k,x), k \in \mathbb{N}\}$ the process started at $X_\theta(0) = x$ and consider evaluating the sensitivities of the one-step expectation. By the Markov property, this sensitivity is:

$$\frac{d}{d\theta}\bigg|_{\theta=\theta_0}\mathbb{E}[g(X_\theta(k+1))|X_{\theta_0}(k)=x] = \lim_{\Delta\theta\to 0}\mathbb{E}\left[\frac{g(X_{\theta_0+\Delta\theta}(1,x)) - g(X_{\theta_0}(1,x))}{\Delta\theta}\right],$$

The random variable $H_x(\theta) = g(X_\theta(1,x))$ may fail to be a.s. Lipschitz continuous, therefore it is possible to divide the state space into a set $\mathcal{A}_x^*(\Delta\theta,\theta)$ containing only trajectories where $H_x(\theta)$ is Lipschitz continuous, and the so-called *critical set*

$$\mathcal{A}_x(\Delta\theta,\theta) = \{\omega : |g(X_{\theta+\Delta\theta}(1,x;\omega)) - g(X_\theta(1,x;\omega))| > \alpha\,\Delta\theta\}\,,$$

for some $\alpha > 0$. It is assumed here that for each state x the limit $\mathcal{A}_x(\Delta\theta, \theta) \to \mathcal{A}_x(\theta)$ exists, for some measurable set $\mathcal{A}_x(\theta)$ and that the following limit (called *critical rate*)

$$\lim_{\Delta\theta \to 0} \frac{\mathbb{P}(\mathcal{A}_x(\Delta\theta, \theta))}{\Delta\theta} = p'_\theta(x) > 0$$

exists and is finite. This implies that for each $x \in S$ $\lim_{\Delta\theta \to 0} \mathbb{P}(\mathcal{A}^*(\Delta\theta, \theta)) = 1$ In addition, if the discontinuity itself is absolutely integrable, that is: $\mathbb{E}[|g(X_{\theta+\Delta\theta}(1,x)) - g(X_\theta(1,x))| \,|\, \mathcal{A}_x(\Delta\theta, \theta)] < \infty$ then Dominated Convergence yields:

$$\lim_{\Delta\theta \to 0} \mathbb{E}\left[\frac{g(X_{\theta+\Delta\theta}(1,x)) - g(X_\theta(1,x))}{\Delta\theta}\right]$$
$$= \mathbb{E}\left[\frac{d}{d\theta}H_x(\theta)\right] + \mathbb{E}[g(X_{\theta+}(1,x)) - g(X_\theta(1,x))|\mathcal{A}_x(\theta)]\, p'_\theta(x)\,,$$

where $X_{\theta+}(1,x)$ denotes the limit of $X_{\theta+\Delta\theta}(1,x)$ as $\Delta\theta \downarrow 0$. The term inside the first expectation is known as the *IPA term* and the second, as the *SPA term* of the derivative estimator. The effect of conditioning on the so-called critical events is to partially integrate the discontinuities via the critical rate p'_θ. For an example with zero IPA contribution, see [8].

In the foregoing, only one transition was affected by the perturbation of θ. When studying the process $\{X_\theta(n)\}$, the perturbations affect the entire trajectories. As done in [5], the expectation is re-written in terms of Filtered Monte Carlo, conditioning on each step. Under the assumed integrability conditions, the overall effect of the SPA term is obtained as if only one-step transitions were perturbed at a time, that is (assuming no IPA contribution):

$$\frac{d}{d\theta}\mathbb{E}[g(X_\theta(k+1))]$$
$$= \mathbb{E}\left[\sum_{i=0}^{k} \mathbb{E}\left[g(X_\theta^{(i)}(k+1)) - g(X_\theta(k+1))\Big| \mathcal{A}_{X_\theta(i)}(\theta)\right] p'_\theta(X_\theta(i))\right]\,, \quad (2)$$

where the process $\{X_\theta^{(i)}(n)\}$ is the limiting process from a perturbation $\theta + \Delta\theta$ *at the i-th transition only*. To show the validity of the expression above and to obtain the sample path estimators, the crucial step when using the pathwise analysis is to show that for small changes in θ, the discontinuous effect of the perturbation of the whole trajectory is only *local*: discontinuities initiate at each transition and then propagate. This can be done as in [1] by showing that the contribution from paths with discontinuities in multiple transitions has a vanishing effect in the limit as $\Delta \to 0$. To summarize, SPA involves a careful pathwise analysis of the propagation of delays and their effect on a given performance measure.

While SPA offers great flexibility, proofs of unbiasedness are often very cumbersome, which stems from the fact the effect of a perturbation on the *entire* sample path has to be studied. Furthermore, generally the results for SPA only hold for individual performance functions and changing the considered performance function leads to a entirely new proof.

Under our interpretation in terms of MVD, the term p'_θ in (2) represents the derivative of a probability distribution, and calculating the overall gradient in (1)

4

corresponds to applying a "product rule" of differentiation to a product measure. In this paper we derive such a product rule of measure-valued differentiation, which can be applied to SPA. More precisely, our product rule for measure-valued differentiation provides the sensitivity of the entire sample path out of a local analysis. Put another way, the analysis of the propagation of delays is not necessary anymore.

2.2 Weak Derivatives

In this section, we briefly review the concept of weak differentiation of probability measures as introduced by Pflug, see [15]. Let $(S, \mathcal{S})$ denote a Polish measurable space[1]. For most applications, $S \subset \mathbb{R}^d$ and $\mathcal{S}$ represents the σ-field of events that are Borel subsets of $\mathbb{R}^d$. Let $\mathcal{M} = \mathcal{M}(S, \mathcal{S})$ denote the set of finite signed measures on $(S, \mathcal{S})$, and $\mathcal{M}_1 = \mathcal{M}_1(S, \mathcal{S}) \subset \mathcal{M}$ the set of probability measures. Denote by $C_b := C_b(S)$ the set of bounded continuous mappings $g : S \to \mathbb{R}$. For any signed measure ν on $(S, \mathcal{S})$ there exists a set $G \in \mathcal{S}$, such that $[\nu]^+(A) := \nu(A \cap G) \geq 0$ and $[\nu]^-(A) := -\nu(A \cap G^c) \geq 0$ for any $A \in \mathcal{S}$, see, for example, [12]. In particular, the set G is implicitly defined via

$$\nu(G) = \sup\{A \in \mathcal{S} : \nu(A)\}. \tag{3}$$

The measures $[\nu^+]$ and $[\nu^-]$ are positive measures on $(S, \mathcal{S})$ and the pair $([\nu]^+, [\nu]^-)$ is called the *Hahn-Jordan decomposition* of ν. The Hahn-Jordan decomposition is unique in the sense that if $\hat{G}$ is another set, such that $\nu(A \cap \hat{G}) \geq 0$ and $\nu(A \cap \hat{G}^c) \leq 0$ for any $A \in \mathcal{S}$, then $\nu(A \cap G) = \nu(A \cap \hat{G})$ for any $A \in \mathcal{S}$. A signed measure is called *finite* if $[\nu]^+$ and $[\nu]^-$ are finite measures. Integration with respect to a signed measure is defined through

$$\int_S g(s)\, \nu(ds) = \int_S g(s)\, [\nu]^+(ds) - \int_S g(s)\, [\nu]^-(ds),$$

provided that the terms on the right-hand side of the above formula are finite. See [12] for more details.

Definition 1 *A measure $\mu_\theta \in \mathcal{M}_1$ is called* weakly differentiable at θ *if a signed finite measure $\mu'_\theta \in \mathcal{M}$ exists, such that for all $g \in C_b$ it holds that*

$$\lim_{\Delta \to 0} \frac{1}{\Delta} \left(\int_S g(s)\mu_{\theta+\Delta}(ds) - \int_S g(s)\mu_\theta(ds) \right) = \int_S g(s)\, \mu'_\theta(ds).$$

Note that $\mu'_\theta(S) = \int_S \mu'_\theta(ds) = 0$ (take $g = 1$), so that μ'_θ can be written as difference between two probability measures (apply, for example, the Hahn-Jordan decomposition).

Definition 2 *A triple $(c_\theta, \mu_\theta^+, \mu_\theta^-)$ is called a* weak derivative *of μ_θ, where $\mu_\theta^\pm \in \mathcal{M}_1$, if for all continuous bounded functions $g \in C_b$ it holds that*

$$\int_S g(s)\, \mu'_\theta(ds) = \lim_{\Delta \to 0} \frac{1}{\Delta} \left(\int_S g(s)\mu_{\theta+\Delta}(ds) - \int_S g(s)\mu_\theta(ds) \right)$$

$$= c_\theta \left(\int_S g(s)\mu_\theta^+(ds) - \int_S g(s)\mu_\theta^-(ds) \right). \tag{4}$$

---

[1]A topological space is called *separable* if it contains a countable dense set. It is called *Polish* if there exists a metric compatible with the topology under which the space is complete and separable; see e.g. [4].

The probability measure μ_θ^+ is called the (normalized) positive part of μ_θ' and μ_θ^- is called the (normalized) negative part of μ_θ', respectively. Note that the weak derivative is not unique. We illustrate this with the following example.

EXAMPLE 1. Let $S = [0, \infty)$ and η_θ the exponential distribution with mean θ. Let $f_\theta(x) = \theta \exp(-\theta x)$ denote the Lebesgue density of η_θ, then it holds for any $g \in C_b$

$$
\begin{aligned}
\frac{d}{d\theta} \int g(x)\, \eta_\theta(dx) &= \frac{d}{d\theta} \int g(x)\, f_\theta(x)\, dx \\
&= \int g(x)\, \frac{d}{d\theta} f_\theta(x)\, dx \\
&= \int g(x)(1 - \theta x)e^{-\theta x}\, dx \\
&= \frac{1}{\theta} \left(\int g(x) f_\theta(x)\, dx \;-\; \int g(x) h_\theta(x)\, dx \right),
\end{aligned}
$$

where h_θ is the density of the Gamma $(2,\theta)$ distribution, denoted by $\Gamma(2,\theta)$. Hence, η_θ is weakly differentiable and an instance of a weak derivative of μ_θ is given by

$$
\Big(1/\theta \,,\, \eta_\theta \,,\, \Gamma(2,\theta) \Big). \tag{5}
$$

On the other hand, the Hahn-Jordan decomposition leads to the representation $((\theta\, e)^{-1}, \mu_\theta^+, \mu_\theta^-)$ with

$$
\mu_\theta^+(A) \;=\; \int_0^{1/\theta} 1_A(x)\, (\theta - \theta^2 x)\, e^{1 - \theta x}\, dx
$$

and

$$
\mu_\theta^-(A) \;=\; \int_{1/\theta}^{\infty} 1_A(x)\, (\theta^2 x - \theta)\, e^{1 - \theta x}\, dx \,,
$$

for any measurable set A.

<div align="right">❈❈❈</div>

From the definition, it is clear that weak derivatives yield results which hold on C_b. Furthermore, a product rule of weak differentiation for products of independent measures exist, see [15]. However, whether there also exists a product rule of weak differentiation for conditional measures, like Markov kernels, is still an open question.

The MVD approach that we introduce in Section 3 extends the results of weak derivatives in two aspects: the product form is now established for the product of Markov kernels, and also admissible performance functions are more general, no longer requiring (piece-wise) continuity and boundedness.

2.3 Score Function

In this section, we briefly review some basic facts on the Score Function method. For details we refer to [17, 18]. Assume that $\nu \in \mathcal{M}$ exists, such that μ_θ is absolutely continuous with respect to ν for all $\theta \in \Theta$ and denote the ν-density of μ_θ by f_θ. If f_θ is ν almost surely differentiable with respect to θ and

$$
\int_S \sup_{\theta \in \Theta} \left| \frac{d}{d\theta} f_\theta(u) \right| \nu(du) \;<\; \infty \,,
$$

then for any $g \in C_b$

$$\frac{d}{d\theta} \int g(u)\,\mu_\theta(du) = \int g(u) \frac{d}{d\theta} f_\theta(u)\,\nu(du)$$

$$= \int g(u) \frac{d}{d\theta} \ln(f_\theta(u))\,\mu_\theta(du) . \qquad (6)$$

The mapping $d\ln(f_\theta(u))/d\theta$ is called *score function*.

The score function approach works on C_b. Furthermore, standard calculus implies a product rule for the score function. The key condition for the above approach is that $f_\theta(u) = 0$ implies $df_\theta(u)/d\theta = 0$. In other words, the measure μ'_θ given by $\mu'_\theta(B) = \int_B \frac{d}{d\theta} f_\theta(u)\,\nu(du)$ for $B \in \mathcal{S}$, is absolutely continuous w.r.t. ν and μ_θ. As we will illustrate in section 2.4, this restricts the applicability of the Score Function approach. In addition, the Score Function estimates suffer typically from variance problems.

2.4 A Note on Domination

Note that, for any $g \in C_b$, if (6) holds then μ_θ is weakly differentiable. However, the converse is not true. To see this, the key observation is that the above Score Function approach requires that μ_θ as well as μ'_θ are absolutely continuous with respect to the *same* measure ν, which is not required for the measure-valued concept of differentiation. The following simple example illustrates this issue.

EXAMPLE 2. For $\theta > 0$, let $\mathcal{U}_{[0,\theta]}$ denote the uniform distribution on the interval $[0, \theta]$. For any continuous mapping g,

$$\frac{d}{d\theta} \int g(x)\,\mathcal{U}_{[0,\theta]}(dx) = \frac{d}{d\theta} \left(\frac{1}{\theta} \int_0^\theta g(x)\,dx \right)$$

$$= \frac{1}{\theta} g(\theta) - \frac{1}{\theta^2} \int_0^\theta g(x)\,dx$$

$$= \frac{1}{\theta} \left(\int g(x)\,\delta_\theta(dx) - \int g(x)\,\mathcal{U}_{[0,\theta]}(dx) \right) ,$$

where δ_θ denotes the Dirac measure in θ. Hence, $(1/\theta, \delta_\theta, \mathcal{U}_{[0,\theta]})$ is a weak derivative of μ_θ. Moreover, the positive part of the weak derivative is a discrete measure whereas the negative part is absolutely continuous with respect to the Lebesgue measure. In other words, μ'_θ fails to be absolutely continuous with respect to either μ_θ or the Lebesgue measure. The Score function method is therefore not applicable.

<div align="right">✷✷✷</div>

Problems with domination can be easily dealt with by an *a posteriori* analysis. Indeed, the reason why the Score Function fails in the above example is that the Dirac measure in point θ is not absolutely continuous with respect to the uniform distribution on $[0, \theta]$. Switching from $\mathcal{U}_{[0,\theta]}$ as driving probability measure to, say, $\nu_\theta = \frac{1}{2}\mathcal{U}_{[0,\theta]} + \frac{1}{2}\delta_\theta$ solves this problem, since δ_θ and $\mathcal{U}_{[0,\theta]}$ are now absolutely continuous with respect to the new probability measure ν_θ; see [9] for more details.

3 Measure-Valued Differentiation

In this section, we formally present the concept of measure-valued differentiation (MVD), inspired by the concept of weak differentiation, but as we will soon establish, our methodology does not rely on weak topology only. The main result of this section is the proof of the product rule for MVD. Furthermore, using a conditioning approach, we show how the measure-valued differentiability of a Markov kernel can be deduced from that of more elementary distributions.

3.1 Basic Concepts and Definitions

Let $(S_2, \mathcal{S}_2)$ and $(S_1, \mathcal{S}_1)$ be Polish measurable spaces. Recall that $\mathcal{M}(S_2, \mathcal{S}_2)$ denotes the set of finite (signed) measures on $(S_2, \mathcal{S}_2)$ and $\mathcal{M}_1(S_2, \mathcal{S}_2)$ that of probability measures on $(S_2, \mathcal{S}_2)$.

Definition 3 *The mapping $P : \mathcal{S}_2 \times S_1 \to [0, 1]$ is called a* transition kernel *on $(\mathcal{S}_2, S_1)$ if*

(a) *$P(\cdot\,; s) \in \mathcal{M}(S_2, \mathcal{S}_2)$ for all $s \in S_1$; and*

(b) *$P(B; \cdot)$ is $\mathcal{S}_1$ measurable for all $B \in \mathcal{S}_2$.*

If, in condition (a), $\mathcal{M}(S_2, \mathcal{S}_2)$ can be replaced by $\mathcal{M}_1(S_2, \mathcal{S}_2)$, then P is called a Markov kernel *on $(\mathcal{S}_2, S_1)$.*

Denote the set of transition kernels on $(\mathcal{S}_2, S_1)$ by $\mathcal{K}(\mathcal{S}_2, S_1)$ and the set of Markov kernels on $(\mathcal{S}_2, S_1)$ by $\mathcal{K}_1(\mathcal{S}_2, S_1)$. If $(S_2, \mathcal{S}_2) \neq (S_1, \mathcal{S}_1)$, then the transition (respectively, Markov) kernel is called *inhomogeneous*, whereas for $(S, \mathcal{S}) := (S_2, \mathcal{S}_2) = (S_1, \mathcal{S}_1)$ it is called *homogenous* and P is then called a transition (respectively, Markov) kernel on $(\mathcal{S}, S)$.

Consider a family of Markov kernels $(P_\theta : \theta \in \Theta)$ on $(\mathcal{S}_2, S_1)$, where $\Theta \subset \mathbb{R}$ is a compact set, and let $L^1(P_\theta; \Theta) \subset \mathbb{R}^{S_2}$ denote the set of measurable mappings $g : S_2 \to \mathbb{R}$, such that $\int_{S_2} |g(u)| P_\theta(du; s)$ is finite for all $\theta \in \Theta$ and $s \in S_1$.

Definition 4 *Let $P_\theta \in \mathcal{K}(\mathcal{S}_2, S_1)$, for $\theta \in \Theta$, and let $\mathcal{D} \subset L^1(P_\theta; \Theta)$. We call P_θ differentiable at θ with respect to $\mathcal{D}$, or $\mathcal{D}$-differentiable for short, if for any $s \in S_1$ a $P'_\theta(\cdot\,; s) \in \mathcal{M}(S_2, \mathcal{S}_2)$ exists, such that, for any $s \in S_1$ and for all $g \in \mathcal{D}$*

$$\frac{d}{d\theta} \int_{S_2} g(u) \, P_\theta(du; s) \, = \, \int_{S_2} g(u) \, P'_\theta(du; s) \,. \tag{7}$$

If the left-hand side of equation (7) equals zero for all $g \in \mathcal{D}$, then we say that P'_θ is not significant.

Recall that $C_b(S)$ denotes the set of continuous bounded mappings from S to $\mathbb{R}$.

Lemma 1 *Let $P_\theta \in \mathcal{K}_1(\mathcal{S}_2, S_1)$ be $\mathcal{D}$-differentiable. If $C_b(S_2) \subset \mathcal{D}$, then $P'_\theta(\cdot\,; s) \in \mathcal{M}(S_2, \mathcal{S}_2)$ is uniquely defined for any $s \in S_1$.*

Proof: Fix $s \in S_1$, then $P'_\theta(\cdot; s)$ is a signed measure on $(S_2, \mathcal{S}_2)$ and its Hahn-Jordan decomposition is uniquely defined. Let $\nu_\theta \neq P'_\theta(\cdot; s)$ be another signed measure on $(S_2, \mathcal{S}_2)$, such that

$$\forall g \in \mathcal{D} : \qquad \frac{d}{d\theta} \int_{S_2} g(u)\, \nu_\theta(du) \;=\; \int_{S_2} g(u)\, P'_\theta(du; s)\,. \tag{8}$$

Taking the Hahn-Jordan decomposition of ν_θ, it follows from (8) that

$$\forall g \in \mathcal{D} : \qquad \frac{d}{d\theta} \int_{S_2} g(u)\, [\nu_\theta]^\pm(du) \;=\; \int_{S_2} g(u)\, [P'_\theta]^\pm(du; s)\,, \tag{9}$$

where we write $[\cdot]^\pm$ to indicate that the above equation holds for the positive and the negative part. Form $C_b(S_2) \subset \mathcal{D}$ it follows by (9) that for any bounded continuous mapping from S_2 to $\mathbb{R}$ integration with respect to $[\nu_\theta]^\pm$ and $[P'_\theta]^\pm(\cdot; s)$ yields the same results, hence, the measures are equal, that is, $[\nu_\theta]^\pm(A) = [P'_\theta]^\pm(A; s)$ for any $A \in \mathcal{S}_2$. Thus $\nu_\theta(\cdot) = P'_\theta(\cdot; s)$, which concludes the proof of the lemma.

$\square$

If P'_θ exists, then the fact that $P'_\theta(\cdot; s)$ fails to be a probability measure poses the problem of sampling from $P'_\theta(\cdot; s)$. For $s \in S_1$ fixed, we can represent $P'_\theta(\cdot; s)$ by its Hahn-Jordan decomposition as a difference between two probability measures. More precisely, this Hahn-Jordan decomposition is obtained as follows. Let

$$c_{P_\theta}(s) = [P'_\theta]^+(S; s) = [P'_\theta]^-(S; s) \tag{10}$$

and

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

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

$$\int_{S_2} g(u)\, P'_\theta(du; s) = c_{P_\theta}(s) \left(\int_{S_2} g(u)\, P_\theta^+(du; s) - \int_{S_2} g(u)\, P_\theta^-(du; s) \right)\,. \tag{11}$$

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 $P_\theta^+(A; \cdot)$ and $P_\theta^-(A; \cdot)$ for any $A \in \mathcal{S}_2$. This problem is equivalent to showing that $c_{P_\theta}(\cdot)$ in (10) is measurable as a mapping from S_1 to $\mathbb{R}$. Unfortunately, only sufficient conditions are known. If, for example, σ-field $\mathcal{S}_2$ has at most countably many elements (that is, if S_2 is finite), then the following lemma establishes measurability.

Lemma 2 *If $\mathcal{S}_2$ is a countable set, then c_{P_θ} defined in (10) is measurable as mapping from S_1 to $\mathbb{R}$.*

Proof: For $s \in S_1$, let G_s denote a set, such that

$$P'_\theta(G_s; s) \;=\; \sup\{A \in \mathcal{S}_2 \,:\, P'_\theta(A; s)\}\,.$$

The limit of a sequence of measurable mappings is again measurable. Hence, for every $A \subset \mathcal{S}_2$, $P'_\theta(A; \cdot)$ is a measurable mapping from S_1 to $\mathbb{R}$. Since the

supremum is taken over countably many sets, it follows that $P'_\theta(G_{(\cdot)}; \cdot)$ is measurable as a mapping from S_1 to $\mathbb{R}$. Equation (3) implies $c_{P_\theta}(s) = P'_\theta(G_s; s)$, which concludes the proof of the lemma. $\square$

In applications c_{P_θ} is calculated explicitly and its measurability is therefore established case by case. Specifically, in most of the examples presented in this paper, c_{P_θ} turns out to be a constant and measurability is thus guaranteed.

As explained in [10], a general sufficient condition for P'_θ to be a transition kernel is the following: for all $s \in S_1$ it holds that

$$\sup_{g \in C_b(S_2),\, |g| \leq 1} \left| \int_{S_2} P'_\theta(du; s) g(u) \right| < \infty \,.$$

In Section 4.3 we will show that measurability of c_{P_θ} defined in (10) holds for general state-space S_2 whenever P'_θ is absolutely continuous with respect to another kernel.

To conclude this section, we now introduce the notion of $\mathcal{D}$-derivative, which extends the concept of a weak derivative.

Definition 5 *Let P_θ be $\mathcal{D}$-differentiable. Any triple $(c_{P_\theta}(\cdot), P_\theta^+, P_\theta^-)$, with $P_\theta^\pm \in \mathcal{K}_1(\mathcal{S}_2, S_1)$ and c_{P_θ} a measurable mapping from S_2 to $\mathbb{R}$, that satisfies (11) 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 $c_{P_\theta}(\cdot)$ is called the* normalizing factor.

$\mathcal{D}$-derivatives are not unique. To see this, consider $P_\theta \in \mathcal{K}_1(\mathcal{S}_2, S_1)$ with $\mathcal{D}$-derivative $(c_{P_\theta}, P_\theta^+, P_\theta^-)$ and take $Q \in \mathcal{K}(\mathcal{S}_2, S_1)$ so that $\int_{S_2} g(u)\, Q(du; s)$ is finite for any $g \in \mathcal{D}$ and $s \in S_1$. Set

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

Equation (11) implies for all $g \in \mathcal{D}$ and all $s \in S_1$ that

$$\frac{d}{d\theta} \int_{S_2} g(u)\, P_\theta(du; s) = 2 c_{P_\theta}(s) \left(\int_{S_2} g(u)\, \tilde{P}_\theta^+(du; s) - \int_{S_2} g(u)\, \tilde{P}_\theta^-(du; s) \right) \,.$$

3.2 The Product Rule of Measure-Valued Differentiation

For the finite horizon problem, as stated in (1), the transition kernel P_θ in Definition 4 is the n step transition probability of the Markov chain $\{X_\theta(m), m = 0, 1, \ldots\}$. In general, it is often very hard to write down the n step transition probability, and studying its differentiability properties is practically impossible. However, the n step transition probability is composed out of one step transition probabilities, that is, transition kernels, which are comparably easier to analyze.

This section establishes the main property of $\mathcal{D}$-differentiable transition kernels, namely, that the product of $\mathcal{D}$-differentiable Markov kernels is again $\mathcal{D}$-differentiable and that the $\mathcal{D}$-derivative can be expressed in terms of the $\mathcal{D}$-derivatives of the transition kernels.

Let P be a Markov kernel on $(\mathcal{S}_2, S_1)$ and Q a Markov kernel on $(\mathcal{S}_1, S_0)$, where $(S_0, \mathcal{S}_0)$ is a measurable Polish space. The product of transition kernels Q, P on $(\mathcal{S}_2, S_0)$ is defined as follows. For $s \in S_0$ and $B \in \mathcal{S}_2$ set $P\,Q(B; s) = (P \circ Q)(B, s) = \int_{S_1} P(B; z)\, Q(dz; s)$. Let $\mathcal{D}_2 \subset L^1(P)$ and $\mathcal{D}_1 \subset L^1(Q)$.

Definition 6 *Let $\mathcal{D}_2$ be a set of measurable mappings $g : S_2 \to \mathbb{R}$ and let $\mathcal{D}_1 \subset \mathbb{R}^{S_1}$. Transition kernel P_θ is called $(\mathcal{D}_2, \mathcal{D}_1)$-mapping if*

$$\forall g \in \mathcal{D}_2 : \quad \int_{S_2} g(u) P_\theta(du; \cdot) \in \mathcal{D}_1 .$$

If $\mathcal{D} = \mathcal{D}_1 = \mathcal{D}_2$, then P_θ satisfying the above condition is called $\mathcal{D}$-preserving.

A sufficient condition for $\int g(u)(PQ)(du; s)$ to exist for any $g \in \mathcal{D}_2$ and any $s \in S$, is that P is a $(\mathcal{D}_2, \mathcal{D}_1)$-mapping.

Definition 7 *Let $P_\theta \in \mathcal{K}(\mathcal{S}_2, S_1)$ and $\mathcal{D}^2 \subset L^1(P_\theta; \Theta)$, $\mathcal{D}^1 \subset \mathbb{R}^{S_1}$. We call P_θ $(\mathcal{D}^2, \mathcal{D}^1)$-Lipschitz continuous if for any $g \in \mathcal{D}^2$ a $K_g \in \mathcal{D}^1$ exists, such that for any $\Delta > 0$ with $\theta + \Delta \in \Theta$*

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

If $\mathcal{D} = \mathcal{D}^2 = \mathcal{D}^1$, then we call P_θ simply $\mathcal{D}$-Lipschitz continuous.

The following theorem presents the key technical result of this section. The proof is given in the Appendix.

Theorem 1 *Let $((S_i, \mathcal{S}_i) : 0 \leq i \leq n)$ be a sequence of Polish measurable spaces. For $1 \leq i \leq n$, let $P_{\theta,i}$ be a transition kernel on $(\mathcal{S}_i, S_{i-1})$, such that $P_{\theta,i}$ is $\mathcal{D}_i$-differentiable. Furthermore, set $\mathcal{D}_0 = \mathbb{R}^{S_0}$.*

We introduce the following assumptions: for each $i, 1 \leq i \leq n$,

(A0) *if $g, f \in \mathcal{D}_i$ then it holds that $f + g \in \mathcal{D}_i$,*

(A1) *$P_{\theta,i}$ is a $(\mathcal{D}_i, \mathcal{D}_{i-1})$-mapping,*

(A2) *$P_{\theta,i}$ is $(\mathcal{D}_i, \mathcal{D}_{i-1})$-Lipschitz continuous,*

(A3) *$P_{\theta,i}$ is $\mathcal{D}_i$-differentiable such that $P'_{\theta,i} \in \mathcal{K}(S_i, S_{i-1})$ and $P'_{\theta,i}$ is a $(\mathcal{D}_i, \mathcal{D}_{i-1})$-mapping.*

The following statements hold true:

*(i) Under Assumptions **(A0)**, **(A1)** and **(A2)**, the product $\prod_{i=1}^{n} P_{\theta,i}$ is $(\mathcal{D}_n, \mathcal{D}_0)$-Lipschitz continuous.*

*(ii) Under Assumptions **(A0)**, **(A1)**, **(A2)** and **(A3)** the product rule holds:*

$$\left(\prod_{i=1}^{n} P_{\theta,i} \right)' = \sum_{j=1}^{n} \prod_{i=j+1}^{n} P_{\theta,i} \, P'_{\theta,j} \prod_{i=1}^{j-1} P_{\theta,i} .$$

Following the line of proof of the above theorem, one obtains the following chain rule of differentiation.

Corollary 1 *Consider a $\mathcal{D}$-differentiable Markov kernel P_θ such that $P'_\theta \in \mathcal{K}$. Let $g_\theta \in \mathcal{D}$ and assume that $K_g \in \mathcal{D}$ exists, such that for any $\Delta \in \mathbb{R}$ with $\theta + \Delta \in \Theta$*

$$|g_{\theta+\Delta}(s) - g_\theta(s)| \leq \Delta K_g(s) .$$

If g_θ is differentiable at θ, then

$$\frac{d}{d\theta} \int g_\theta(u) \, P_\theta(du; s) = \int \left(\frac{d}{d\theta} g_\theta(u) \right) P_\theta(du; s) + \int g_\theta(u) \, P'_\theta(du; s) ,$$

for any $s \in S$.

Remark: When the performance function depends explicitly on θ the first term is recognizably the so-called *IPA term*, and the corresponding integrability assumption is given as a "weak $\mathcal{D}$" Lipschitz continuity assumption. In particular, when the kernel is independent of θ the corollary recovers the usual IPA formulation. It is worthwhile to notice that the pathwise analysis common to SPA/IPA formulations requires explicit construction of the trajectories to evaluate the propagation of the perturbations: our formalism implicitly deals with this propagation through a simple chain rule of differentiation.

Often, one is interested in evaluating expectations of an entire trajectory rather than a particular n step transition. In terms of (1) one is interested in

$$\frac{d}{d\theta} \mathbb{E}[g(X_\theta(n), \ldots, X_\theta(1))] .$$

As the following corollary shows, the product rule of measure-valued differentiation can be applied to the above problem too. However, $\mathcal{D}$-differentiability of, say, $P_{\theta,i}$ only controls the differentiability of functions on S_i. In order to ensure that a function of a sample path is differentiable, it is necessary to ensure that the sample path functions are "locally" in the corresponding sets $\mathcal{D}_i$. To make this precise, consider the following notation. Let $P_{\theta,i}$ be a Markov kernel on $(\mathcal{S}_i, S_{i-1})$, for $1 \leq i \leq n$. For $s \in S_0$, we denote the product measure on $\times_{i=1}^n S_i$ by $\left(\prod_{i=1}^n P_{\theta,i} \right) (\cdot; s)$, that is, for $B_i \in \mathcal{S}_i$, with $1 \leq i \leq n$,

$$\left(\prod_{i=1}^n P_{\theta,i} \right) (B_n, \ldots, B_1; s) = \int_{B_n} \ldots \int_{B_1} P_{\theta,n}(ds_n; s_{n-1}) \ldots P_{\theta,1}(ds_1; s) .$$

Corollary 2 *Let $P_{\theta,i}$ be a Markov kernel on $(S_{i-1}, \mathcal{S}_i)$, for $1 \leq i \leq n$, such that $P_{\theta,i}$ is $\mathcal{D}_i$-differentiable for each i. Assume* **(A0)**, **(A1)**, **(A2)** *and* **(A3)**. *Let $g : \times_{i=1}^n S_i \to \mathbb{R}$. If, for each $i \in \{1, \ldots, n\}$ and any $s_j \in S_j$ $(1 \leq j \leq n, \, j \neq i)$*

$$g(s_n \ldots, s_{i+1}, \cdot, s_{i-1}, \ldots, s_1) \in \mathcal{D}_i , \tag{12}$$

then it holds true that

$$\frac{d}{d\theta} \left(\int g(s_n, \ldots, s_1) \left(\prod_{i=1}^n P_{\theta,i} \right) (ds_n, \ldots, s_1; s) \right)$$

$$= \sum_{j=1}^n \int g(s_n, \ldots, s_1) P_{\theta,n}(ds_n; s_{n-1}) \ldots P'_{\theta,j}(ds_j; s_{j-1}) \ldots P_{\theta,1}(ds_1; s) ,$$

for $s \in S$.

Proof: See the Appendix. □

Remark: In applications, one is often interested in the sample mean as overall performance measure, that is, one considers

$$g(s_n, \ldots, s_1) = \frac{1}{n} \sum_{i=1}^{n} h(s_i) \, .$$

for some performance function h. In this case condition (12) reduces to: $h \in \cap_{i=1}^{n} \mathcal{D}_i$. Moreover, if the Markov chain is homogeneous, we recover the natural condition $h \in \mathcal{D}$.

Theorem 1 establishes a product rule of measure-valued differentiation for Markov kernels. In order to check whether a given Markov kernel satisfies the conditions in Theorem 1, it is often helpful to separate the parts of the transition kernel that depend on θ and those that are independent of θ. We illustrate this conditioning approach with the following example.

EXAMPLE 3. For $\Theta = [0, 1]$, let $\eta_\theta \in \{0, 1\}$ be Bernoulli-θ-distributed on $S_\eta = \{0, 1\}$, with

$$\mathbb{P}(\eta_\theta = 0) =: \mu_\theta(0) = \theta$$
$$\mathbb{P}(\eta_\theta = 1) =: \mu_\theta(1) = 1 - \theta.$$

For any $g = (g_0, g_1) \in \mathbb{R}^2$ it holds that

$$\frac{d}{d\theta} \int_{S_\eta} g_s \, \mu_\theta(ds) = g_0 - g_1 = \int_{S_\eta} g_s \, \delta_0(ds) - \int_{S_\eta} g_s \, \delta_1(ds) \, ,$$

where δ_y denotes the Dirac measure on y. Thus, μ_θ has $\mathbb{R}^2$-derivative $(1, \delta_0, \delta_1)$.

Let $\{X_\theta(n)\}$ denote the queue-length processes of a Markovian queueing network. Denote the transition kernel of $\{X_\theta(n)\}$ by P_θ, where θ is a routing parameter. The routing decision is made as follows. If, at the nth state transition, a customer leaves a particular sever of the network, a Bernoulli-(θ)-distributed random variable $\eta_\theta(n)$ is generated independent of everything else. For $\eta_\theta(n) = 0$ the customer is routed to a particular server, say j, and for $\eta_\theta(n) = 1$ he/she is routed to a server, say j', with $j \neq j'$. Using the fact that $\{\eta_\theta(n)\}$ is an i.i.d. sequence, we can draw a sample of $\eta_\theta(n)$ at each transition. Let $Q(\cdot; s, \cdot)$ denote the transition kernel of $X_\theta(n)$ given that $\eta_\theta(n) = s$ and let $\mathcal{D}$ be the set of all g, such that for any possible queue-length vector x and $s = 0, 1$:

$$\mathbb{E}\Big[|g(X_\theta(n + 1))| \,\Big|\, X_\theta(n) = x, \eta_\theta = s \Big] < \infty \, . \tag{13}$$

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

$$\frac{d}{d\theta} \int g(u) \, P_\theta(du; s) = \frac{d}{d\theta} \int_{S_\eta} \int g(u) \, Q(du; \eta, s) \, \mu_\theta(d\eta)$$

$$= \int g(u) \, Q(du; 0, s) - \int g(u) \, Q(du; 1, s)$$

$$= \int g(u) \, P_0(du; s) - \int g(u) \, P_1(du; s) \, .$$

Hence, a $\mathcal{D}$-derivative of P_θ can be obtained from $(1, P_0, P_1)$. Moreover, the $\mathcal{D}$-derivative is independent of θ and P_θ is thus $\mathcal{D}$-Lipschitz continuous (for a proof use the Mean Value Theorem). Linearity of the expected value in (13) implies that for $f, g \in \mathcal{D}$ it holds that $f + g \in \mathcal{D}$. Hence, provided that $\int g(u) \, P_\theta(du; \cdot) \in \mathcal{D}$ for any $g \in \mathcal{D}$ and $\theta \in [0, 1]$, the product rule applies to P_θ.

To emphasize the potential benefits of the MVD approach, we stress that from the simple formulas for the weak derivative of a Bernoulli random variable it is now possible to reconstruct the MVD formulas for the routing sensitivities in the whole network, via the product rule.

The above conditioning approach can be interpreted as a particular kind of conditioning within the SPA setting (see Section 2.1), although MVD does not yield a pathwise estimator, but a closed formula for the distributions. In Section 5 we specifically deal with the construction of various estimators from MVD formulas.

<div align="center">✲✲✲</div>

4 Setting The Product Rule To Work

In this section, we discuss various meaningful ways of interpreting the conditions in Theorem 1. Simple examples will be given to illustrate the situations we have in mind. For the sake of simplicity, consider homogeneous Markov chains and denote the state-space by $(S, \mathcal{S})$. To simplify the notation, drop the explicit dependence on the state-space whenever this causes no confusion. For example, we will write C_b instead $C_b(S)$ for the set of bounded continuous functions.

4.1 Bounded Performance Functions

As a first choice for $\mathcal{D}$ take the set of bounded measurable mappings, denoted by $\mathcal{D}^0$. Note that $\mathcal{D}^0$ satisfies **(A0)**.

Lemma 3 *Let P_θ be a Markov kernel that is $\mathcal{D}^0$-differentiable on Θ with $\mathcal{D}^0$-derivative $((c_{P_\theta}(s), P_\theta^+(\cdot; s), P_\theta^-(\cdot; s)) : s \in S)$. If*

$$\sup_{\theta \in \Theta} c_{P_\theta}(\cdot) \in \mathcal{D}^0 \,,$$

then

$$(P_\theta^n)' = \sum_{j=1}^n P_\theta^{n-j} \, P_\theta' \, P_\theta^{j-1} \,.$$

Proof: Conditions **(A1)** and **(A3)** in Theorem 1 are trivially satisfied and we turn to condition **(A2)**: Lipschitz continuity. For any $g \in \mathcal{D}^0$,

$$
\begin{aligned}
\left| \int g(u) P_\theta'(du; s) \right| &= c_{P_\theta}(s) \left| \int g(u) P_\theta^+(du; s) \, - \, \int g(u) P_\theta^-(du; s) \right| \\
&\leq c_{P_\theta}(s) \left(\int |g(u)| \, P_\theta^+(du; s) \, + \, \int |g(u)| \, P_\theta^-(du; s) \right) \\
&\leq 2 \, c_{P_\theta}(s) \, \|g\|_\infty \,,
\end{aligned}
$$

where $||g||_\infty = \sup(|g(u)| : u \in S)$ denotes the sup-norm of g. Hence, applying the Mean Value Theorem, it follows that for all Δ with $\theta + \Delta \in \Theta$

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

$$\leq \Delta \sup_{\theta \in \Theta} \left| \int g(s) P'_\theta(ds; \cdot) \right|$$

$$\leq 2 ||g||_\infty \sup_{\theta \in \Theta} c_{P_\theta}(s),$$

which is in $\mathcal{D}^0$ by assumption. Hence, **(A2)** in Theorem 1 is satisfied and proof of the lemma follows directly from Theorem 1 (ii). $\square$

The above lemma extends the product rule of weak differentiation of independent measures, as established in [15], to (a) products of conditional measures, and (b) bounded measurable performance functions (in contrast to bounded piece-wise continuous).

EXAMPLE 4. Let $X_\theta(n)$ be the discrete-time queue length process of an M/M/1 queue with arrival rate λ and service rate θ, with $\theta \geq a > 0$. The transition kernel is then given in matrix form by

$$P_\theta = \begin{bmatrix} 0 & 1 & 0 \\ \frac{\lambda}{\lambda+\theta} & 0 & \frac{\theta}{\lambda+\theta} & 0 \\ 0 & \frac{\lambda}{\lambda+\theta} & 0 & \frac{\theta}{\lambda+\theta} & 0 & \cdots \\ & & & \ddots \\ & & & & 1 & 0 \end{bmatrix}$$

Let

$$c_{P_\theta}(k) = \begin{cases} 0, & k = 0, \\ \frac{\lambda}{(\lambda+\theta)^2}, & 0 < k < m, \\ 0, & k = m. \end{cases}$$

then a weak derivative of $P_\theta(\cdot; k)$ can be obtained as follows. For $k = 0$ and $k = m$,

$$(c_{P_\theta}(k), P_\theta^+(\cdot; k), P_\theta^-(\cdot; k)) = (0, P_\theta(\cdot; k), P_\theta(\cdot; k))$$

and for $0 < k < m$

$$(c_{P_\theta}(k), P_\theta^+(\cdot; k), P_\theta^-(\cdot; k)) = \left(\frac{\lambda}{(\lambda + \theta)^2}, \delta_{k+1}(\cdot), \delta_{k-1}(\cdot) \right).$$

Since $\theta \geq a > 0$, $\sup_{\theta \in \Theta} c_{P_\theta}(\cdot) \in C_b$ and Lemma 3 yields, for example, a closed form expression for the derivative of any moment of the queue length at the nth state with respect to the service rate.

For this setup the $\mathcal{D}$-derivative can be represented in a concise form through matrix notation. To see this, define the matrix C_{P_θ} by

$$C_{P_\theta} = \begin{bmatrix} 0 \\ & \frac{1}{(\lambda+\theta)^2} \\ & & \ddots \\ & & & \frac{1}{(\lambda+\theta)^2} \\ & & & & 0 \end{bmatrix}$$

and matrices P^+, P^- by

$$P^+ = \begin{bmatrix} 0 & 0 & & & \\ & 0 & 1 & & \\ & & \ddots & & \\ & & & 0 & 1 \\ & & & 0 & 0 \end{bmatrix} \qquad P^- = \begin{bmatrix} 0 & 0 & & & \\ & 1 & 0 & & \\ & & \ddots & & \\ & & & 1 & 0 \\ & & & 0 & 0 \end{bmatrix} ,$$

then

$$\frac{d}{d\theta} P_\theta = C_{P_\theta} \left(P^+ - P^- \right),$$

and (with slight abuses of notation) the triple (C_{P_θ}, P^+, P^-) may serve as matrix-valued $\mathcal{D}$-derivative of P_θ. The statement of Lemma 3 then reads

$$\frac{d}{d\theta} P_\theta^n = \sum_{j=1}^n P_\theta^{n-j} C_{P_\theta} P^+ P_\theta^{j-1} - \sum_{j=1}^n P_\theta^{n-j} C_{P_\theta} P^- P_\theta^{j-1} .$$

<div align="right">✻✻✻</div>

4.2   Performance Functions Bounded By A Polynomial

In applications, to assume that the sample performance is bounded ($g \in \mathcal{D}^0$) is often too restrictive. A convenient set of performance functions is the set $\mathcal{D}^p$ of polynomially bounded performance functions defined by

$$\mathcal{D}^p = \left\{ g : S \to \mathbb{R} \; \middle| \; g(x) \leq \sum_{i=0}^p \kappa_i \, ||x||^i \, , \; \kappa_i \in \mathbb{R}, 0 \leq i \leq p \right\} , \qquad (14)$$

for some $p \in \mathbb{N}$, where $|| \cdot ||$ denotes a norm on S (assuming that S is indeed equipped with a norm). Most cases of interest in applications fall within this setting. Note that $\mathcal{D}^p$ satisfies **(A0)** and that $\mathcal{D}^p \subset L^1(P_\theta : \Theta)$ if $P_\theta(\cdot; s)$ has finite pth moment for any $s \in S$ and $\theta \in \Theta$. The above definition recovers $\mathcal{D}^0$ as the set of bounded functions.

Lemma 4 *Let $p \in \mathbb{N}$. Consider a (homogeneous) Markov kernel P_θ, with finite pth moment for any $s \in S$ and $\theta \in \Theta$. Assume that P_θ is $\mathcal{D}^p$-differentiable on Θ with $\mathcal{D}^p$-derivative $((c_{P_\theta}(s), P_\theta^+(\cdot; s), P_\theta^-(\cdot; s)) : s \in S)$.*
If P_θ' is $\mathcal{D}^p$-preserving and a $K(\cdot) \in \mathcal{D}^p$ exists, such that

$$\sup_{\theta \in \Theta} \left(c_{P_\theta}(\cdot) \int (1 + ||s||^p) \, P_\theta^\pm(ds; \cdot) \right) \leq K(\cdot),$$

then

$$(P_\theta^n)' = \sum_{j=1}^n P_\theta^{n-j} \, P_\theta' \, P_\theta^{j-1} .$$

Proof: Assumption **(A1)** is satisfied and condition **(A3)** holds by assumption. To apply Theorem 1 it remains to be shown that under the conditions in the

lemma condition (**A2**) holds as well. Using the $\mathcal{D}^p$-derivative of P_θ, it is easily checked that for any $g \in \mathcal{D}^p$ and $s \in S$

$$\sup_{\theta \in \Theta} \left| \int g(u) P'_\theta(du; s) \right|$$

$$\leq \sum_{i=0}^p \kappa_i \sup_{\theta \in \Theta} \left\{ c_{P_\theta}(s) \left(\int ||u||^i P_\theta^+(du; s) + \int ||u||^i P_\theta^-(du; s) \right) \right\} .$$

Define $H : S \to [0, \infty)$ by

$$H(s) = \sum_{i=0}^p \kappa_i K(s), \quad s \in S .$$

where $K(\cdot) \in \mathcal{D}^p$ by assumption. Hence,

$$\sup_{\theta \in \Theta} \left| \int g(u) P'_\theta(du; s) \right| \leq H(s), \quad s \in S ,$$

and $H(\cdot) \in \mathcal{D}^p$, because the set $\mathcal{D}^p$ is closed with respect to finite summation. The proof of the lemma follows from the Mean Value Theorem.□

EXAMPLE 5. Let $X_\theta(n)$ denote the nth waiting time at a GI/F_θ/1 queue with generally distributed inter-arrival times. Let η_θ denote the exponential distribution with mean $1/\theta$ and $\Gamma(2, \theta)$ the Gamma$(2, \theta)$ distribution. The service times are governed by the distribution

$$F_\theta = \theta \, \eta_{\theta_0} + (1 - \theta) \, \Gamma(2, \theta_0), \quad \theta \in \Theta = [0, 1] ,$$

that is, with probability θ the service time is exponentially distributed with mean θ_0 and with probability $1 - \theta$ is is distributed like the sum of two independent exponentials with mean θ_0 each. We have is $S = \mathbb{R}$ and we take the usual norm on $\mathbb{R}$ for $|| \cdot ||_S$. Observe that F_θ is $\mathcal{D}^p$-differentiable for any p and a $\mathcal{D}^p$-derivative is given by

$$\left(1, \, \eta_{\theta_0}, \, \Gamma(2, \theta_0) \right), \tag{15}$$

which is independent of θ. Let $\{A(n)\}$ be the i.i.d. sequence of inter-arrival times and $\{S_\theta(n)\}$ the i.i.d. sequence of service times, respectively. Lindley's recursion yields:

$$X_\theta(n + 1) = \max(X_\theta(n) + S_\theta(n) - A(n + 1), 0), \quad n > 1 ,$$

and $X_\theta(1) = 0$. As performance function, take the pth moment of the waiting times (which is not in $\mathcal{D}^0$). Let $G(\cdot)$ denote the distribution of the inter-arrival time and assume that the first p moments of G are finite. For $w > 0$, the transition kernel for the waiting times is given by

$$P_\theta((0, w]; v) = \int_0^\infty \int_{s+v-w}^{s+v} G(da) \, F_\theta(ds)$$

$$=: \int Q((0, w]; s, v) \, F_\theta(ds)$$

17

and

$$P_\theta(\{0\}; v) = \int_0^\infty \int_{s+v}^\infty G(da)\, F_\theta(ds)$$

$$=: \int Q(\{0\}; s, v)\, F_\theta(ds)$$

For any $g \in \mathcal{D}^p$, it then holds that $\int g(u) P_\theta(du; v) \in \mathcal{D}^p$ and P_θ is thus $\mathcal{D}^p$-preserving.

The first step is to calculate the $\mathcal{D}^p$-derivative of P_θ. For any $v \geq 0$ and $g \in \mathcal{D}^p$, $\int g(s) Q(dr; \cdot, v)$ is again in $\mathcal{D}^p$ and since F_θ is $\mathcal{D}^p$-differentiable it easily follows that $P_\theta(\cdot; v)$ is $\mathcal{D}^p$-differentiable. A $\mathcal{D}^p$-derivative of F_θ is given in (15) and a $\mathcal{D}^p$-derivative of P_θ can therefore be obtained from

$$P^+((0, w]; v) = \int_0^\infty Q((0, w]; s, v)\, \eta_{\theta_0}(ds), \quad w > 0,$$

$$P^+(\{0\}; v) = \int_0^\infty Q(\{0\}; s, v) \eta_{\theta_0}(ds),$$

$$P^-((0, w]; v) = \int_0^\infty Q((0, w]; s, v)\, \Gamma(2, \theta_0)(ds), \quad w > 0,$$

$$P^-(\{0\}; v) = \int_0^\infty Q(\{0\}; s, v) \Gamma(2, \theta_0)(ds),$$

and

$$c_{P_\theta} = 1.$$

Note that this simple calculation implies that $P'_\theta = P^+_\theta - P^-_\theta$ is a transition kernel.

Longer service times lead to longer waiting times, which implies the following chain of inequalities, for any $\theta \in [0, 1]$,

$$\int (1 + u)^p P^+(du; v) \leq \int (1 + u)^p P_\theta(du; v)$$

$$\leq \int (1 + u)^p P^-(du; v)$$

$$= \int (1 + u)^p P_1(du; v),$$

for $v \geq 0$. Note that $\int (1+u)^p P_1(du; \cdot) =: K(\cdot) \in \mathcal{D}^p$. Hence, P'_θ is $\mathcal{D}^p$-preserving. Moreover, elaborating on the fact that $c_{P_\theta} = 1$ and that $P^\pm$ are independent of θ, it readily follows that

$$\sup_{\theta \in \Theta} \int (1 + u)^p P^\pm(du; v) = \int (1 + u)^p P^\pm(du; v) \leq K(v).$$

Hence, Lemma 4 yields, for example, a closed form expression for the derivative of the pth moment of the nth waiting time at a $GI/F_\theta/1$ queue.

❋❋❋

4.3   Markov Kernels With Differentiable Densities

As already illustrated in Section 2.3, 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 product rule of measure valued differentiation simplify under the presence of differentiable densities.

For $P, Q \in \mathcal{K}_1$, let P be absolutely continuous with respect to Q, in symbols: $P \ll Q$. 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 is given through integrating the positive and negative parts of the derivative of $[dP'_\theta/dQ](r; s)$, and the corresponding the normalizing factor is measurable. The precise statement is given in the following lemma.

Lemma 5 *Let $P_\theta, Q \in \mathcal{K}_1$, for $\theta \in \Theta$. Assume that*

- *P_θ is $\mathcal{D}$-differentiable at θ, and*

- *$P'_\theta \ll Q$.*

Then $P'_\theta \in \mathcal{K}$ and $(c_{P_\theta}, P_\theta^+, P_\theta^-)$, with

$$c_{P_\theta}(s) \;=\; \int_S \max\left(0, \left[\frac{dP'_\theta}{dQ}\right](r; s)\right) Q(dr; s), \quad s \in S,$$

for any $A \in \mathcal{S}$ and $s \in S$,

$$P_\theta^+(A; s) \;=\; \frac{1}{c_{P_\theta}(s)} \int_A \max\left(0, \left[\frac{dP'_\theta}{dQ}\right](r; s)\right) Q(dr; s)$$

and

$$P_\theta^-(A; s) \;=\; \frac{1}{c_{P_\theta}(s)} \int_A \max\left(0, -\left[\frac{dP'_\theta}{dQ}\right](r; s)\right) Q(dr; s),$$

is a $\mathcal{D}$-derivative of P_θ.

Proof: $\mathcal{D}$-differentiability of P_θ implies that, for any $g \in \mathcal{D}$ and any $s \in S$,

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

$$= \int_S g(u) \left[\frac{dP'_\theta}{dQ}\right](r; s) Q(dr; s),$$

where the last equation sign follows from $P'_\theta \ll Q$. We not turn to the $\mathcal{D}$-derivative. For $s \in S$, set

$$c_{P_\theta}(s) = \int_S \max\left(0, \left[\frac{dP'_\theta}{dQ}\right](r; s)\right) Q(dr; s).$$

Let $\{X(n)\}$ denote the Markov chain associated with Q, then the above equation can be written as follows

$$c_{P_\theta}(s) = \mathbb{E}\left[\max\left(0, \left[\frac{dP'_\theta}{dQ}\right](X(1); X(0))\right) \,\Big|\, X(0) = s\right], \quad s \in S.$$

The Radon-Nikodym derivative of P'_θ with respect to Q is measurable in both components and all operators in the expression on the righthand side of the above equation preserve measurability. Therefore, c_{P_θ} is measurable as mapping in s and furthermore $P'_\theta \in \mathcal{K}$.

$\mathcal{D}$-differentiability implies that $P'_\theta(\cdot; s)$ is a finite (signed) measure for any s, which implies that $c_{P_\theta}(s) < \infty$ for any s. Hence, for any $A \in \mathcal{S}$ and $s \in S$,

$$P_\theta^+(A; s) \;=\; \frac{1}{c_{P_\theta}(s)} \int_A \max\left(0, \left[\frac{dP'_\theta}{dQ}\right](r; s)\right) Q(dr; s)$$

and

$$P_\theta^-(A; s) \;=\; \frac{1}{c_{P_\theta}(s)} \int_A \max\left(0, -\left[\frac{dP'_\theta}{dQ}\right](r; s)\right) Q(dr; s)$$

are Markov kernels. It is easily checked that $(c_{P_\theta}, P_\theta^+, P_\theta^-)$ is a $\mathcal{D}$-derivative of P_θ.
$\square$

The following lemma establishes sufficient conditions for the product rule to hold in the presence of domination.

Lemma 6 *Let $p \in \mathbb{N}$. Consider a (homogeneous) Markov kernel P_θ, with finite pth moment for any $s \in S$ and $\theta \in \Theta$. Assume that P_θ is $\mathcal{D}^p$-preserving and $\mathcal{D}^p$-differentiable on Θ, and that $P'_\theta \ll P_\theta$. If P'_θ is $\mathcal{D}^p$-preserving and a $K(\cdot) \in \mathcal{D}^p$ exists, such that*

$$\sup_{\theta \in \Theta} \int (1 + \|s\|^p) \left|\left[\frac{dP'_\theta}{dP_\theta}\right](ds; \cdot)\right| P_\theta(ds; \cdot) \;\leq\; K(\cdot),$$

then

$$(P_\theta^n)' \;=\; \sum_{j=1}^n P_\theta^{n-j} P'_\theta P_\theta^{j-1}.$$

Proof: The proof follows the same line of argument as the proof of Lemma 4 except for the way in which the $\mathcal{D}^p$-Lipschitz constant is constructed. For any $g \in \mathcal{D}^p$ and $s \in S$

$$\sup_{\theta \in \Theta} \left|\int g(u) P'_\theta(du; s)\right|$$

$$\leq \sum_{i=0}^p \kappa_i \sup_{\theta \in \Theta} \int \|u\|^i \left|\left[\frac{dP'_\theta}{dP_\theta}\right](u, s)\right| P_\theta(du; s).$$

By assumption $K(\cdot) \in \mathcal{D}^p$ exists, such that for any $s \in S$

$$\sup_{\theta \in \Theta} \int (1 + \|u\|^p) \left|\left[\frac{dP'_\theta}{dP_\theta}\right](u, s)\right| P_\theta(du; s) \;\leq\; K(s).$$

Choose $H(\cdot) \in \mathcal{D}^p$ such that $H(\cdot) \geq p \max(\kappa_i, 1) K(\cdot)$, which is possible because $K \in \mathcal{D}^p$ implies $nK \in \mathcal{D}^p$ for any $n \in \mathbb{N}$. Hence,

$$\sup_{\theta \in \Theta} \left|\int g(u) P'_\theta(du; \cdot)\right| \;\leq\; H(\cdot) \in \mathcal{D}^p$$

and, by the Mean Value Theorem, $H(\cdot)$ serves as Lipschitz constant for $\int g(u) P'_\theta(du; \cdot)$. Using the fact that $H(\cdot) \in \mathcal{D}^p$ and noting that by Lemma 5 P'_θ is indeed a transition kernel, completes the proof. $\square$

The key to applying Lemma 6 is to compute the Radon-Nikodym derivative of P'_θ with respect to P_θ. In applications, P'_θ is typically of rather complex structure and computing the Radon-Nikodym derivative of P'_θ with respect to P_θ leads to cumbersome calculations, as will be illustrated in Example 6. However, using a conditioning argument, the assumption in Lemma 6 can be restated in terms of verifiable conditions. We omit the details for the sake of brevity.

EXAMPLE 6. Let $X_\theta(n)$ denote the n-th waiting time at a GI/M/1 queue. Let $\{A(n)\}$ be the sequence of inter-arrival times and $\{S_\theta(n)\}$ the sequence of exponentially distributed service times with mean $1/\theta$, respectively. Let $\Theta = [a, b] \subset (0, \infty)$. Denote the distribution of $S_\theta(n)$ by η_θ and the corresponding Lebesgue density by $f^S_\theta(x) = \theta e^{-\theta x}$. Let $A(n)$ have a finite pth moment and let f^A denote the Lebesgue density of the inter-arrival times. As performance measures of interest, consider the pth moment of the waiting time. Let P_θ denote the transition kernel of $\{X_\theta(n)\}$. Following the line of thought in Example 5, for any $w > 0, v \geq 0$, the transition kernel for the waiting times is given by

$$P_\theta((0, w]; v) = \int_0^\infty \int_{v+x-w}^{v+x} f^A(a) f^S_\theta(x) \, da \, dx =: \int_0^\infty R((0, w]; x, v) f^S_\theta(x) \, dx,$$

$$P_\theta(\{0\}; v) = \int_0^\infty \int_{v+x}^\infty f^A(a) f^S_\theta(x) \, da \, dx =: \int_0^\infty R(\{0\}; x, v) f^S_\theta(x) \, dx.$$

The exponential distribution is $\mathcal{D}^p$-differentiable for any p (see Example 1). Differentiating P_θ with respect to θ yields

$$P'_\theta((0, w]; v) = \int_0^\infty R((0, w]; x, v) (1 - \theta x) f^S_\theta(x) \, dx, \quad w > 0,$$

and

$$P'_\theta(\{0\}; v) = \int_0^\infty R(\{0\}; x, v) (1 - \theta x) f^S_\theta(x) \, dx.$$

A $\mathcal{D}^p$-derivative of P_θ can be obtained from setting $c_{P_\theta} = 1/\theta$ and, for any measurable set A

$$P^+_\theta(A; v) = \int_0^\infty R(A; x, v) f^S_\theta(x) \, dx = P_\theta(A; v)$$

and for $w > 0$

$$P^-_\theta((0, w]; v) = \int_0^\infty R((0, w]; x, v) h^S_\theta(x) \, dx = \int_0^\infty \int_{v+x}^{v+x-w} f^A(a) h^S_\theta(x) \, da \, dx,$$

$$P^-_\theta(\{0\}; v) = \int_0^\infty R(\{0\}; x, v) h^S_\theta(x) \, ds = \int_0^\infty \int_{v+x}^\infty f^A(a) h^S_\theta(x) \, da \, dx,$$

where h_θ denotes the density of the Gamma($2,\theta$) distribution. Let $\Theta = [a, b]$, with $a > 0$.

We now show that the product rule of measure-valued differentiation applies to P_θ. For any $g \in \mathcal{D}^p$ it holds that

$$H_g(x,v) \;=\; \int_0^\infty g(u)\, R(du; x, v) \;=\; \int_{v+x}^\infty g(v + x - a)\, f^A(a)\, da \,,$$

assuming for the sake of simplicity that $g(0) = 0$. Because the inter-arrival times have finite pth moment (by assumption) it is easily verified that for $g \in \mathcal{D}^p$ the mapping $\int H_g(x,\cdot) f_\theta^S(x) dx$ is in $\mathcal{D}^p$ and P_θ is hence $\mathcal{D}$-preserving. Following the same line of argument, it is easily seen that $\int H_g(x,\cdot) h_\theta^S(x) dx$ is in $\mathcal{D}^p$ for any $g \in \mathcal{D}^p$ and thus $P'_\theta = \frac{1}{\theta}(P_\theta^+ - P_\theta^-)$ is $\mathcal{D}^p$-preserving.

Note that $\eta'_\theta \ll \eta_\theta$ for any $\theta \in \Theta = [a,b]$. Then, replacing η_θ by the corresponding density f_θ^S, it follows

$$\sup_{\theta \in [a,b]} \int_0^\infty (1 + u^p) \left| \frac{\frac{d}{d\theta} f_\theta^S(u)}{f_\theta^S(u)} \right| f_\theta^S(u)\, du \;\leq\; \frac{1}{a} \int_0^\infty (1 + u^p)(1 + bu) f_a^S(u)\, du \,,$$

which is finite for any $p \in \mathbb{N}$. In accordance with Lemma 4, for any $g \in \mathcal{D}^p$, with $p \in \mathbb{N}$, the product rule applies to P_θ and yields a closed-form expression for the derivative of the pth moment of the waiting time.

<div align="right">✻✻✻</div>

4.4 The Influence Of The Normalizing Factor

In general, c_{P_θ} is a function on the state-space and introducing it may lead to problems with respect to integrability. For example, if only the first p moments of $P_\theta(\cdot; s)$ are finite for any s but c_{P_θ} is a polynomial of degree larger than p, then $\int c_{P_\theta}(u)\, P_\theta(du; s)$ fails to be bounded. In the previous sections, $c_{P_\theta}(s)$ turned out to be bounded in all examples. We conclude this section with an example of the case where $c_{P_\theta}(s)$ is only polynomial bounded.

Unbounded normalizing constants typically stem from the fact that a random variable, say, $X(n)$ that is independent of θ is taken as input of a function h_θ in order to form a new random variable $h_\theta(X(n))$. If the derivative of h_θ is unbounded as a function of $X(n)$, one obtains an unbounded normalizing constant. We explain this concept with the following academic example.

EXAMPLE 7. Let $\{X(n)\}$ be an i.i.d. sequence with $X(n) \in [0, \infty)$ defined on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let $\Theta = [a, b]$, with $a < b$, and define $\{Y_\theta(n)\}$, with $Y_\theta(n) \in \{0, 1\}$, as follows

$$\begin{aligned}
\mathbb{P}(\, Y_\theta(n) = 0 \,|\, X(n-1)) &= \frac{1}{2} + \frac{1}{4} \sin(\theta^2 (X(n-1))^m) \\
&= 1 - \mathbb{P}(\, Y_\theta(n) = 1 \,|\, X(n-1)) \,,
\end{aligned}$$

for $n \geq 1$. Consider the Markov process $Z_\theta(n) = (Y_\theta(n), X(n))$ with transition kernel

$$P_\theta(\, \eta \, \times [0, w); (x, y)\,) \;=\; \mathbb{P}(\, Y_\theta(n) = \eta \,|\, X(n-1) = x)\, \mathbb{P}(\, X(n) \leq w\,) \,,$$

for $\eta \in \{0, 1\}$ and $x \geq 0$. From

$$\frac{d}{d\theta} \mathbb{P}(\, Y_\theta(n) = 0 \,|\, X(n-1) = x\,) \;=\; \frac{1}{2} \theta x^m \cos(\theta^2 x^m)$$

$$\frac{d}{d\theta}\mathbb{P}(\,Y_\theta(n) = 1 \,|\, X(n-1) = x\,) \;=\; -\frac{1}{2}\theta x^m \cos(\theta^2 x^m),$$

the derivative of $P_\theta(\,\eta \times [0, w); Z_\theta(n-1))$ can be written as

$$\left(\frac{1}{2}\theta(X(n-1))^m \cos(\theta^2 X(n-1)^m)\,,\, \delta_0(\eta)\,\mathbb{P}(\,X(n) \le w\,)\,,\, \delta_1(\eta)\,\mathbb{P}(\,X(n) \le w\,)\right),$$

$$(16)$$

with normalizing constant

$$c_{P_\theta}(x) \;=\; \frac{1}{2}\theta x^m \cos(\theta^2 x^m)\,,$$

which is not uniformly bounded in x. Fix $p \ge 0$ and consider the set $\mathcal{D}^p$, see (14) for a definition. Assumptions **(A0)** and **(A1)** are trivially satisfied. If the first p moments of $X(n)$ are finite, then assumption **(A2)** in Theorem 1 is satisfied and we turn to the $\mathcal{D}^p$-differentiability condition **(A3)**. For any $g \in \mathcal{D}^p$

$$\left| \int g(z)\, P'_\theta(dz;(y,x)) \right| \le \theta x^m \cos(\theta^2 x^m) \sum_{i=0}^{p} \kappa_i \int \eta^i\, \mu(d\eta)\,,$$

where μ denotes the distribution of $X(n)$. Since the first p moments of $X(n)$ are finite, the sum of integrals on the right-hand side of the above expression equals some number $d \in [0, \infty)$. This yields, for any $g \in \mathcal{D}^p$

$$\sup_{\theta \in [a,b]} \left| \int g(z)\, P'_\theta(dz; y, x) \right| \le bx^m\, d\,,$$

for $b \in \mathbb{R}$. The expression on the right-hand side of the above inequality lies in $\mathcal{D}^p$ if $m \le p$. In case $m > p$ we extend the polynomial in the definition of $\mathcal{D}^p$ up to the mth power, that is, we take $p = k$. Hence, if the first $\max(m, p)$ moments of $X(n)$ are finite, then Theorem 1 applies.

<div align="center">❅❅❅</div>

In Example 7, the first $\max(p, m)$ moments of $X(n)$ must be bounded in order to obtain the sensitivities of performance functions that are bounded by a polynomial of degree p. This is in contrast to the situation in Example 5, where only the existence of the first p moments is required. The obvious reason for this difference is that the normalizing factor in Example 7 is a polynomial in $X(n-1)$. One obtains the following rule of thumb: If g is bounded by a polynomial of degree p in $X(n)$ and the normalizing constant is bounded by a polynomial of degree m in $X(n-1)$, then the first $\max(p, m)$ moments of $X(n)$ have to exist in order to satisfy the conditions in Theorem 1. In particular, if one is interested in measurable bounded performance functions, then $p = 0$ is sufficient, but because the normalizing factor is bounded by a polynomial of degree m, it is now necessary that the first m moments of $X(n)$ be finite in order to apply the product rule of measure-valued differentiation. This exemplifies the tradeoff between generality with respect to the performance functions and that with respect to the transition probabilities.

5   Gradient Estimation

While MVD offers a methodology that helps to establish a closed formula for (1), in practice one wishes to construct an estimator based on observations (or simulations) of the underlying Markov process.

Let $(P_{\theta,i} : 1 \leq i \leq n)$ be a family of $\mathcal{D}$-differentiable Markov kernels on $(S, \mathcal{S})$. The product rule of measure-valued differentiability yields

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

$$= \sum_{j=1}^n \iiint g(s_n,\ldots,s_1)\prod_{i=j+1}^n P_{\theta,i}(ds_i; s_{i-1})\, P'_{\theta,j}(ds_j; s_{j-1})\prod_{i=1}^{j-1} P_{\theta,i}(ds_i; s_{i-1}),$$

with $s_0 \in S$. How to transform the above into an unbiased estimator?

5.1 Phantom Estimators

In this section, we establish sufficient conditions for phantom type estimators to be unbiased. Let the conditions in Theorem 1 be in force. Then one can express $P'_\theta(ds_j; s_{j-1})$ in terms of the normalized difference between two expectations, or $(c_{P_\theta}(s_{j-1}), P^+_{\theta,j}(ds_j; s_{j-1}), P^-_{\theta,j}(ds_j; s_{j-1}))$. The product form can be rewritten in terms of the processes as follows:

$$\frac{d}{d\theta}\mathbb{E}[g(X_\theta(n),\ldots,X_\theta(1))] = \sum_{j=1}^n \Bigg(\mathbb{E}\Big[c_{P_\theta}(X_\theta(j-1))\,g(X^+_{\theta,j}(n),\ldots,X^+_{\theta,j}(1))\Big]$$

$$-\mathbb{E}\Big[c_{P_\theta}(X_\theta(j-1))\,g(X^-_{\theta,j}(n),\ldots,X^-_{\theta,j}(1))\Big]\Bigg), \tag{17}$$

where the processes $\{X^\pm_{\theta,j}(i)\}$ are the Markov chains that follow the transition kernels $P_{\theta,i}(ds_i; s_{i-1}), i \neq j$ and where $\mathbb{P}(X^\pm_{\theta,j}(j) \in \cdot | X^\pm_{\theta,j}(j-1) = s_{j-1}) = P^\pm_{\theta,j}(\cdot; s_{j-1})$. The chains $\{X^\pm_{\theta,j}(i)\}$ are called *phantoms* in the literature.

Equation (17) has the following interpretation: the processes $\{X^\pm_{\theta,j}(n), j = 1,\ldots,N\}$ follow the transition kernel of the process $\{X_\theta(n)\}$ up to $n = j-1$. Next, $\mathbb{P}(X^\pm_{\theta,j}(j) \in \cdot | X^\pm_{\theta,j}(j-1) = x)$ follows the kernel $P^\pm_\theta(\cdot; x)$. After this transition, again the one step transition kernel of the processes $X^\pm_{\theta,j}$ follow $P_\theta(\cdot; x)$.

EXAMPLE 8. Consider a standard periodic review inventory model with backlog. Consecutive demands $\{D(n)\}$ are assumed continuous with Lebesgue density $f(\cdot)$, so that the inventory level $\{X_\theta(n)\}$ at the review epochs is Markovian, $X_\theta(0) = 0$ and for $n \geq 0$:

$$X_\theta(n+1) = \begin{cases} X_\theta(n) - D(n+1) & \text{if } X_\theta(n) - D(n+1) \geq \theta \\ S & \text{otherwise} \end{cases} \qquad \{D(n)\} \text{ i.i.d.} \sim F$$

and θ represents the threshold for the ordering policy. Clearly this is a Markov chain in $\mathbb{R}^+$. Call P_θ the corresponding kernel. The cost per period is:

$$\tilde{g}(X_\theta(n), D(n)) = h(\tilde{X}_\theta(n) - D(n))\mathbf{1}_{\left\{D(n) < \tilde{X}_\theta(n)\right\}} + p(D(n) - \tilde{X}_\theta(n))\mathbf{1}_{\left\{D(n) > \tilde{X}_\theta(n)\right\}}$$

$$+ K\mathbf{1}_{\left\{D(n) > \tilde{X}_\theta(n) - \theta\right\}},$$

where h is unit holding cost, K is ordering cost and p is a backlog penalty. Define the integrated cost per period at state $X_\theta(n) = x$ by $g(X_\theta(n)) = \mathbb{E}[\tilde{g}(X_\theta(n), D(n)) | X_\theta(n)]$. The finite horizon cost is:

$$J(\theta) = \sum_{n=1}^N \mathbb{E}[g(X_\theta(n))],$$

This cost function is not a.s. Lipschitz continuous. Moreover, because θ is a threshold parameter then actually

$$\frac{d}{d\theta} g(X_\theta(n)) = 0 \quad a.s.,$$

so that

$$\mathbb{E}\left[\sum_{n=1}^{N} \frac{d}{d\theta} g(X_\theta(n))\right] \neq \frac{d}{d\theta} \sum_{n=1}^{N} \mathbb{E}\left[g(X_\theta(n))\right].$$

The problem has been studied using the SPA pathwise methodology, see [5, 20]. It is clear that for any bounded and continuous function $g \in \mathcal{D}^0$, and all $x \in (\theta, S]$

$$\mathbb{E}[g(X_\theta(n+1))|X_\theta(n) = x] = \int_0^{x-\theta} g(x-y) f(y) \, dy + g(S) (1 - F(x-\theta))$$

where f is the Lebesgue density of the demand D_n. The derivative is calculated directly, yielding:

$$\frac{d}{d\theta}\mathbb{E}[g(X_\theta(n+1))|X_\theta(n)] = f(X_\theta(n) - \theta)(g(S) - g(\theta))$$

$$= c_\theta(X_\theta(n)) \, \mathbb{E}[g(X_\theta^+(n+1)) - g(X_\theta^-(n+1))],$$

where $c_\theta(x) = f(x - \theta)$ for $x \in \mathbb{R}$ and the random variables $X_\theta^{\pm}(n+1)$ are concentrated at the mass points S and θ respectively (note that $c_\theta(\cdot)$ is measurable). Because

$$\sup_{\theta \in [0,S]} c_\theta(\cdot) = \sup_{\theta \in [0,S]} f((\cdot) - \theta) \in \mathcal{D}^0$$

the product rule of measure-valued differentiation applies to P_θ, see Lemma 3. Specifically, the product rule for $\mathcal{D}$-differentiability prescribes defining the processes $\{X_{\theta,j}^{\pm}(i)\}$ as follows: the first $j - 1$ transitions are governed by the kernel P_θ as the inventory process itself. Next $X_{\theta,j}^+(j) = S, X_{\theta,j}^-(j) = \theta$ and the rest of the transitions are again governed by P_θ. The product rule of $\mathcal{D}$-differentiation, see Theorem 1, yields

$$\frac{d}{d\theta} J(\theta) =$$

$$= \sum_{j=1}^{N} \int \cdots \int \sum_{n=1}^{N} g(x_n) \prod_{i=j+1}^{N} P_\theta(dx_i; x_{i-1}) P'_\theta(dx_j,; x_{j-1}) \prod_{i=1}^{j-1} P_\theta(dx_i; x_{i-1})$$

$$= \mathbb{E}\left[\sum_{j=1}^{N} c_\theta(X_\theta(j-1)) \sum_{n=1}^{N} \left(g(X_{\theta,j}^+(n)) - g(X_{\theta,j}^-(n))\right)\right]$$

$$= \mathbb{E}\left[\sum_{j=1}^{N} c_\theta(X_\theta(j-1)) \sum_{n=j}^{N} \left(g(X_{\theta,j}^+(n)) - g(X_{\theta,j}^-(n))\right)\right].$$

The SPA formulation yields an estimator which is an instance of the above processes, using common random variables for the past history up to transition j. In this example, decoupling occurs because $f(X_\theta(j) - \theta)$ is independent of

$X_{\theta,j}^{\pm}(i), i > j$, which allows for several implementations. Two implementations have been proposed: one where the future of the processes $X_{\theta,j}^{\pm}(i), i > j$ is simulated "off line", and one where the same underlying demands are used to drive all parallel processes (called "phantom estimation"). Let

$$g_j^N(\theta, x) = \mathbb{E}\left[\sum_{n=j}^{N} g(X_\theta(n)) \,\middle|\, X_\theta(j) = x, D(j+1), \dots, D(N)\right]$$

the phantom method of [20] uses the estimator:

$$\frac{d}{d\theta} J(\theta) = \sum_{j=1}^{N} \mathbb{E}\left[c_\theta(X_\theta(j-1))\left(g_j^N(\theta, S) - g_j^N(\theta, \theta)\right)\right] \ .$$

Derivation of this formula via SPA takes about three to four pages of careful analysis of critical events and evaluation of the critical rates. We believe that the product rule provides here a method that is easier to implement. As well, coupling via the use of common random variables has been shown in [20] to decrease both variance as well as computational effort.

<div align="center">※※※</div>

Section 4.1 and Section 4.2 studied different choices for the space functions $\mathcal{D}$. Obviously, the minimal condition on $\mathcal{D}$ is absolute integrability with respect to $P^n \mu$ for any n, where μ is the initial distribution of the Markov chain.

Consider a Markov chain $\{X_\theta(n)\}$ with transition kernel P_θ, and use the notation P_θ^n to indicate the n-step transition probability. In what follows, assume that the initial distribution μ of the Markov chain is fixed. Let

$$\mathcal{D}_\mu = \left\{g : S \to \mathbb{R} \,\middle|\, \forall n \in \mathbb{N} \, \forall \theta \in \Theta : \quad \int |g(u)|(P_\theta^n \mu)(du) < \infty\right\},$$

or, in terms of random variables

$$\mathcal{D}_\mu = \left\{g : S \to \mathbb{R} \,\middle|\, \forall n \in \mathbb{N} \, \forall \theta \in \Theta : \quad \int \mathbb{E}_s[|g(X_\theta(n))|]\, \mu(ds) < \infty\right\},$$

where $\mathbb{E}_s$ denotes the expected value conditioned on the event $X(0) = s$, for $s \in S$. In the presence of domination it is possible to state a sufficient condition for the product rule to hold on $\mathcal{D}_\mu$.

Lemma 7 *Let $\{X_\theta(n)\}$ be a homogeneous Markov chain with transition kernel P_θ, where Θ is a neighborhood of θ^* and assume that $P_\theta \ll P_{\theta^*}$. Let P_θ be $\mathcal{D}_\mu$-differentiable, such that $P_\theta^+, P_\theta^- \ll P_{\theta^*}$.*

If, for $s \in S$,

$$(a) \qquad\qquad \mathbb{E}_s[\,|g(X_\theta(n))|\,] < \infty\,,$$

for any n, and

$$(b) \qquad\qquad \mathbb{E}_s\left[\sup_{\theta \in \Theta} |g(X_\theta(n))| \left|\frac{dP_\theta'}{dP_\theta}(X_\theta(j), X_\theta(j-1))\right|\right] < \infty\,,$$

for any j and any n, then

$$\frac{d}{d\theta}\mathbb{E}_s[g(X_\theta(n))] = \sum_{j=1}^{n} \mathbb{E}_s\left[c_{P_\theta}(X_\theta(j-1))\left(g(X_{\theta,j}^+(n)) - g(X_{\theta,j}^-(n))\right)\right] \qquad (18)$$

Proof: Notice that $\mathcal{D}_\mu$ satisfies **(A0)** and that condition **(A3)** is satisfied by assumption. In order to apply Theorem 1 it then suffices to verify the conditions **(A1)** and **(A2)**.

We now turn to condition **(A1)**. For any $k, m \in \mathbb{N}$, assumption (a) implies that

$$\int |g(u)|\,(P^{k+m}\mu)(du) = \int\int |g(u)|\,P^k(du; s)(P^m\mu)(ds) < \infty\,.$$

Taking $k = 1 = m$ this implies that $\int |g(u)|\,P(du; \cdot)$ is absolutely integrable with respect to $P\mu$ and thus in $\mathcal{D}_\mu$.

Consider now the Lipschitz continuity condition **(A2)**. To apply the Mean Value Theorem it suffices to show that under the conditions of the lemma

$$\sup_{\theta \in \Theta}\left| \int g(u)\,P'_\theta(du; \cdot) \right| \leq K(\cdot) \in \mathcal{D}_\mu\,.$$

Assumption (b) implies that for any $n \in \mathbb{N}$ it holds that

$$\mathbb{E}_s\left[\sup_{\theta \in \Theta} |g(X_\theta(n+1))| \left| \left[\frac{dP'_\theta}{dP_\theta}\right](X_\theta(n+1), X_\theta(n)) \right| \right]$$

$$\leq \sup_{\theta \in \Theta} \int \mathbb{E}_s\left[\sup_{\theta \in \Theta} |g(X_\theta(n+1))| \left| \left[\frac{dP'_\theta}{dP_\theta}\right](X_\theta(n+1), x) \right| \,\Big|\, X_\theta(n) = x \right] P^n_\theta\mu(dx)\,,$$

which implies for any $\theta_0 \in \Theta$ that

$$\int \sup_{\theta \in \Theta} \int g(u)\,P'_\theta(du; s)\,(P^n_{\theta_0}\mu)(ds) < \infty$$

and thus

$$\sup_{\theta \in \Theta} \int g(u)\,P'_\theta(du; \cdot) \in \mathcal{D}_\mu\,,$$

which concludes the proof of the lemma. $\square$

5.2 Single Run Estimation

Consider a family $(P_{\theta,i} : 1 \leq i \leq n)$ of $\mathcal{D}$-differentiable Markov kernels on $(S, \mathcal{S})$, such that $P'_{\theta,i}$ is absolutely continuous with respect to $P_{\theta,i}$, in symbols, $P'_{\theta,i} \ll P_{\theta,i}$ for any $\theta \in \Theta$ and $1 \leq i \leq n$. Under uniform integrability conditions, using the same arguments as in Lemma 6 it can be shown that

$$\frac{d}{d\theta}\mathbb{E}[g(X_\theta(n), \dots, X_\theta(1)) \mid X_\theta(0) = s_0] = \frac{d}{d\theta} \int g(s_n, \dots, s_1) \prod_{i=1}^n P_{\theta,i}(ds_i; s_{i-1})$$

$$= \sum_{j=1}^n \iiint g(s_n, \dots, s_1) \prod_{i=j+1}^n P_{\theta,i}(ds_i; s_{i-1})\, P'_{\theta,j}(ds_j; s_{j-1}) \prod_{i=1}^{j-1} P_{\theta,i}(ds_i; s_{i-1})$$

$$= \sum_{j=1}^n \iiint g(s_n, \dots, s_1) \frac{dP'_{\theta,j}}{dP_{\theta,j}}(s_j, s_{j-1}) \prod_{i=1}^n P_{\theta,i}(ds_i; s_{i-1})$$

$$= \mathbb{E}\left[g(X_\theta(n), \dots, X_\theta(1)) \sum_{j=1}^n \frac{dP'_{\theta,j}}{dP_{\theta,j}}(X_\theta(j); X_\theta(j-1)) \,\Big|\, X_\theta(0) = s_0 \right]\,,$$

with $s_0 \in S$. Note that

$$\sum_{j=1}^{n} \frac{dP'_{\theta,j}}{dP_{\theta,j}}(\cdot;\cdot) = \sum_{j=1}^{n} \frac{d}{d\theta} \ln\left(P_{\theta,j}(\cdot;\cdot)\right),$$

which recovers the estimator called the Score Function, see Section 2.3.

It is worth noting that a single-run estimator can also be constructed if $P'_{\theta,i}$ fails to be dominated by P_θ. In such a case on can take $Q_\theta = \frac{1}{3}P^+_{\theta,i} + \frac{1}{3}P^-_{\theta,i} + \frac{1}{3}P_{\theta,i}$ as Markov kernel. Then, $P^\pm_{\theta,i}, P_{\theta,i} \ll Q_\theta$ and single-run estimator of the above type can be found, see [9] for details. However, manipulating the underlying Markov kernel in the above way is not always feasible and increases the variance of the estimator.

6   Discussion and Further Research

Building an estimator from the measure-valued differentiation formulas can be performed in a number of ways, depending on the implementation chosen. The estimator should

1. be easy to implement,

2. have low variance, and

3. have a low computational effort.

Item (1) is often a matter of taste, while the two remaining criteria determine the *efficiency* of an estimator in simulation, and are often problem dependent.

Two measures μ, ν on $(S, \mathcal{S})$ are *orthogonal* if $A \in \mathcal{S}$ exists, such that $\mu(A) = 0$ and $\nu(A^c) = 0$; in symbols $\mu \perp \nu$. Applying the Hahn-Jordan decomposition for the $\mathcal{D}$-derivative of each of the one-step transition kernels P'_θ, the resulting measures are orthogonal: $P^+(\cdot;s) \perp P^-(\cdot;s)$ for all $s \in S$. For an example of such an orthogonal derivative, see Example 4. There is no guarantee that an orthogonal representation is always the one with the smallest variance: see Example 4.12 on page 238 in [15]. However, as numerical examples show, see e.g. [15] Example 4.28 on page 250, it is safe to say that the orthogonal representation usually has a small variance (and in many cases indeed minimizes the variance).

Apart from the particular decomposition, $\mathcal{D}$-derivatives offer a further "degree of freedom". A $\mathcal{D}$-derivative only describes the marginal distribution of $(X^+_{\theta,j}(i), X^-_{\theta,j}(i))$ but not the *joint* distribution. A particular implementation of the estimation is to use the same underlying random variables to drive the evolution of each of the pairs $\{X^\pm_{\theta,j}(i), i = 1, 2, \ldots\}$, thus making these adapted to the natural filtration. Use of common random numbers for these processes further simplifies the estimation into:

$$\frac{d}{d\theta}\mathbb{E}[g(X_\theta(n), \ldots, X_\theta(1))] =$$

$$= \sum_{j=1}^{n} \mathbb{E}\left[c_{P_\theta}(X_\theta(j-1))\left(g(X^+_{\theta,j}(n), \ldots, X^+_{\theta,j}(1)) - g(X^-_{\theta,j}(n), \ldots, X^-_{\theta,j}(1))\right)\right].$$

Coupling via common random numbers is not necessarily optimal (in terms of variance reduction) for every performance function g. Nonetheless, examples abound where the "difference processes" $g(X_{\theta,j}^{+}) - g(X_{\theta,j}^{-})$ can be calculated recursively via the so-called modified Lindley equations ([2], [19]). The resulting estimators can have extremely low computational overhead, thus rendering very efficient estimation.

Lastly, it is not obvious when to choose a single-run estimator and when to implement a phantom estimator. For a given gradient estimation problem it generally depends on the particular problem which type of estimator is more efficient in terms of computation time. A thorough analysis of the trade-off between the two types of estimators in topic of further research.

Acknowledgements

References

[1] P. Brémaud and F. Vázquez-Abad. On the Pathwise Computation of Derivatives with Respect to the Rate of a Point Process: The Phantom RPA Method. *Queueing Systems*, **10**: 249-270, 1992.

[2] C. Cassandras and V. Julka. Scheduling policies using marked/phantom slot algorithms. *Queueing Systems*, **20**: 207-254, 1995.

[3] X.-R. Cao. *Realization Probabilities: The Dynamics of Queueing Networks.* Springer, New-York,1994.

[4] C. Dellacherie and P. Meyer. *Probabilities and Potential.* North-Holland, Amsterdam, 1978.

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

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

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

[8] B. Heidergott. Optimization of a single-component maintenance system: A smoothed perturbation analysis approach. *EJOR*, **119**: 181-190, 1999.

[9] B. Heidergott and A. Hordijk. Single-run gradient estimation via measure-valued differentiation. to appear in *IEEE Trans. Automatic Control*, 2004.

[10] B. Heidergott, A. Hordijk and H. Weisshaupt. Derivatives of Markov kernels and their Jordan decomposition. *submitted*, 2002.

[11] B. Heidergott and F. Vázquez-Abad, Measure-valued differentiation for stochastic processes: the random horizon case. *GERAD G-2001-18*, 2001

[12] E. Hewitt and K. Stromberg. *Real and Abstract Analysis.* Springer, Berlin, 1969.

[13] Y.-C. Ho and X.-R. Cao. *Perturbation Analysis of Discrete Event Dynamic Systems.* Kluwer Academic, Boston, 1991.

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

[15] G. Pflug. *Optimisation of Stochastic Models.* Kluwer Academic, Boston, 1996.

[16] G. Pflug. Gradient estimates for the performance of Markov chains and discrete event processes. *Annals of Operations Research*, 39, 1992:173-194.

[17] M. Reiman and A. Weiss. Sensitivity Analysis for Simulations via Likelihood Ratios. *Operations Research*, **37**: 830-844, 1989.

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

[19] F. Vázquez-Abad. The Case of the Swapping and of the Disappearing Phantoms. invited paper Proceedings of WODES98, Cagliari, August 1998, Edited by The Institute of Electrical Engineers, Great Britain : 102-107.

[20] F. Vázquez-Abad and M. Cepeda-Jüneman. Phantom SPA method: an inventory method revisited. *Proceedings of WSC99*, 1999.

The Appendix

Proof of Theorem 1

We give a proof by induction. Let $n = 2$ and for the sake of simplicity denote $P_{\theta,2} = P_\theta$, $P_{\theta,1} = Q_\theta$ and $\mathcal{D}_2 = \mathcal{D}_P$, $\mathcal{D}_1 = \mathcal{D}_Q$. We first address the second part of the theorem: result (ii).

To establish the induction hypothesis, it suffices to show that

$$(P_\theta Q_\theta)' = P_\theta Q_\theta' + P_\theta' Q_\theta .$$

Let $g \in \mathcal{D}_P$. By calculation,

$$\frac{1}{\Delta} \left| \int \int g(u) \, P_{\theta+\Delta}(du; r) Q_{\theta+\Delta}(dr; s) - \int \int g(u) \, P_\theta(du; r) Q_\theta(dr; s) \right|$$

$$\leq \underbrace{\frac{1}{\Delta} \left| \int \int g(u) \, P_\theta(du; r) \, Q_{\theta+\Delta}(dr; s) - \int \int g(u) \, P_\theta(du; r) \, Q_\theta(dr; s) \right|}_{(a)}$$

$$+ \underbrace{\frac{1}{\Delta} \left| \int \int g(u) \left(P_{\theta+\Delta}(du; r) - P_{\theta}(du; r) \right) \left(Q_{\theta+\Delta}(dr; s) - Q_{\theta}(dr; s) \right) \right|}_{(b)}$$

$$+ \underbrace{\frac{1}{\Delta} \left| \int \int g(u)\, P_{\theta+\Delta}(du; r) Q_{\theta}(dr; s) \;-\; \int \int g(u) P_{\theta}(du; r) Q_{\theta}(dr; s) \right|}_{(c)} .$$

We first deal with expression (a) and show that

$$\lim_{\Delta \to 0} (a) \;=\; \int \int g(u)\, P_{\theta}(du; r)\, Q_{\theta}{}'(dr; s) .$$

By assumption **(A1)**, $\int g(u) P_{\theta}(du; \cdot) =: g^P \in \mathcal{D}_Q$, and $\mathcal{D}$-differentiability of Q_θ yields

$$\left| \int \int g(u)\, P_{\theta}(du; r)\, Q_{\theta+\Delta}(dr; s) \;-\; \int \int g(u)\, P_{\theta}(du; r)\, Q_{\theta}(dr; s) \right|$$

$$\leq \Delta \int g^P(r)\, |Q_{\theta+\Delta}(dr; s) \;-\; Q_{\theta}(dr; s)|$$

$$\leq \Delta\, K^Q_{g^P}(s) ,$$

with $K^Q_{g^P} \in \mathcal{D}_0$, see assumption **(A2)**. Moreover, $\mathcal{D}^Q$-differentiability of Q implies that the limit of (a) as Δ tends to zero is

$$\lim_{\Delta \to 0} \frac{1}{\Delta} \left| \int \int g(u)\, P_{\theta}(du; r)\, Q_{\theta+\Delta}(dr; s) \;-\; \int \int g(u)\, P_{\theta}(du; r)\, Q_{\theta}(dr; s) \right|$$

$$= \int \int g(u)\, P_{\theta}(du; r) Q'_{\theta}(dr; s) .$$

We now turn to (b). Assumption **(A2)** implies

$$\left| \int g(u) \left(P_{\theta+\Delta}(du; s) - P_{\theta}(du; s) \right) \right| \leq \Delta\, K^P_g(s) , \tag{19}$$

with $K^P_g(\cdot) \in \mathcal{D}_Q$. Hence,

$$\left| \int \int g(u) \left(P_{\theta+\Delta}(du; r) - P_{\theta}(du; r) \right) \left(Q_{\theta+\Delta}(dr; s) - Q_{\theta}(dr; s) \right) \right|$$

$$\leq \left| \int \int g(u) \left(P_{\theta+\Delta}(du; r) - P_{\theta}(du; r) \right) Q_{\theta+\Delta}(dr; s) \right|$$

$$+ \left| \int \int g(u) \left(P_{\theta+\Delta}(du; r) - P_{\theta}(du; r) \right) Q_{\theta}(dr; s) \right|$$

$$\leq \Delta \int K^P_g(r)\, Q_{\theta+\Delta}(dr; s) + \Delta \int K^P_g(r) Q_{\theta}(dr; s)$$

$$\leq 2\Delta \int K^P_g(r)\, Q_{\theta}(dr; s) \;+\; \Delta \left| \int K^P_g(r) \left(Q_{\theta+\Delta}(dr; s) - Q_{\theta}(dr; s) \right) \right|$$

$$\leq \Delta K(s) \tag{20}$$

for Δ sufficiently small, with $K \in \mathcal{D}_0$. Moreover, (20) implies that

$$\lim_{\Delta \to 0} \frac{1}{\Delta} \left| \int \int g(u) \left(P_{\theta+\Delta}(du; r) - P_\theta(du; r) \right) Q_{\theta+\Delta}(dr; s) - Q_\theta(dr; s) \right| = 0 \,.$$

Turn now to the expression (c). Recall (see (19)) that a $K_g^P(\cdot) \in \mathcal{D}_Q$ exists, such that

$$\left| \int g(u) \, P_{\theta+\Delta}(du; \cdot) - \int g(u) \, P_\theta(du; \cdot) \right| \leq \Delta \, K_g^P(\cdot) \in \mathcal{D}_Q \,.$$

Using assumption **(A1)** on the measure Q_θ, and the fact that $K_g^P(\cdot) \in \mathcal{D}_1$, it follows

$$\left| \int \left(\int g(u) P_{\theta+\Delta}(du; r) - \int \int g(u) P_\theta(du; r) \right) Q_\theta(dr; \cdot) \right|$$
$$\leq \Delta \int K_g^P(r) \, Q_\theta(dr; \cdot)$$
$$= \Delta K_{K_g^P}^Q(\cdot) \,,$$

with $K_{K_g^P}^Q(\cdot) \in \mathcal{D}_0$. Thus, $\int K_g(r) Q_\theta(dr; s)$ is finite for all s, and the Dominated Convergence Theorem gives

$$\lim_{\Delta \to 0} (c) = \int \frac{d}{d\theta} \left(\int g(u) \, P_\theta(du; r) \right) Q_\theta(dr; s) \,.$$

Differentiability of the transition kernels (assumption **(A3)**) yields $(PQ)' = PQ' + P'Q$, which concludes the proof of the induction hypothesis for the result (ii) of the theorem.

Under Assumption **(A0)**, $\mathcal{D}_0$ is closed with respect to finite summation, which implies that

$$\left| \int \int g(u) \, P_{\theta+\Delta}(du; r) Q_{\theta+\Delta}(dr; \cdot) - \int \int g(u) \, P_\theta(du; r) Q_\theta(dr; \cdot) \right|$$
$$\leq \Delta \left(K_{g^P}^Q(\cdot) + K_{K_g^P}^Q(\cdot) K(\cdot) \right)$$
$$\in \mathcal{D}_0 \,,$$

which concludes the proof of the induction hypothesis for the result (i) of the theorem. Now suppose that the statement (i) has already been shown for $n > 2$. We show that the product rule also applies to $n + 1$. Applying the induction hypothesis to $\prod_{i=1}^{n} P_{\theta,i}$ yields $(\mathcal{D}_n, \mathcal{D}_0)$-Lipschitz continuity of this transition kernel. Applying the induction hypothesis again to $P_{\theta,n+1}$ and $\prod_{i=1}^{n} P_{\theta,i}$ yields $(\mathcal{D}_{n+1}, \mathcal{D}_0)$-Lipschitz continuity of $\prod_{i=1}^{n+1} P_{\theta,i}$, which establishes the result (i).

To conclude the proof of the statement (ii), it suffices now to show $\mathcal{D}_{n+1}$-differentiability of $\prod_{i=1}^{n+1} P_{\theta,i}$. Apply the induction hypothesis to the transition measures $\prod_{i=1}^{n} P_{\theta,i}$ and $P_{\theta,n+1}$:

$$\left(\prod_{i=1}^{n+1} P_{\theta,i} \right)' = P_{\theta,n+1} \left(\prod_{i=1}^{n} P_{\theta,i} \right)' + P'_{\theta,n+1} \left(\prod_{i=1}^{n} P_{\theta,i} \right) \,.$$

Since the product rule is assumed to hold for the first n terms, then

$$\left(\prod_{i=1}^{n+1} P_{\theta,i}\right)' = P_{\theta,n+1}\left(\prod_{i=1}^{n} P_{\theta,i}\right)' + \sum_{i=1}^{n} P'_{\theta,n+1} \prod_{i=j+1}^{n} P_{\theta,i} \; P'_{\theta,j} \prod_{i=1}^{j-1} P_{\theta,i}\,,$$

which concludes the proof of the Theorem 1.

6.1   Proof of Corollary 2

Call $\{X_\theta(m), m = 1, \dots, n\}$ the Markov chain governed by the kernels $P_{\theta,i}$ and let $Y_\theta(i) = (X_\theta(i), \dots, X_\theta(1)) \in S^i$, then for any $g : S^i \to \mathbb{R}$

$$g(X_\theta(i), \dots, X_\theta(1)) = g(Y_\theta(i))\,.$$

Denote the transition kernels from $Y_\theta(i)$ to $Y_\theta(i+1)$ by $P^Y_{\theta,i+1}$, given by

$$\mathbb{P}\Big(Y_\theta(i+1) \in \times_{j=1}^{i+1} B_j \,|\, Y_\theta(i) = (y_i, \dots, y_1)\Big)$$
$$= P^Y_{\theta,i+1}\Big(\times_{j=1}^{i+1} B_j; (y_i, \dots, y_1)\Big)$$
$$= 1_{(y_i,\dots,y_1)\in\times_{j=1}^{n} B_j} P_{\theta,i+1}\Big(B_{i+1}; y_i\Big)\,.$$

From the assumptions on $P_{\theta,i}$ it is straightforward to show that **(A1)** and **(A2)** are satisfied by $P^Y_{\theta,i}$. Indeed, for $g : \times_{i=1}^{n} S_i \to \mathbb{R}$ we obtain

$$\int g(s_n, \dots, s_1)\left(\prod_{i=1}^{n} P_{\theta,i}\right)(ds_n, \dots, ds_1; s)$$
$$= \int g(y) P^Y_{\theta,n}(dy; s)$$
$$= \int g(y_n)\left(\prod_{i=1}^{n} P^Y_{\theta,i}\right)(dy_n, \dots, dy_1; s)\,. \tag{21}$$

Set

$$\tilde{\mathcal{D}}_i = \{g : S^i \to \mathbb{R} \,|\, \forall (s_{i-1}, \dots, s_1) \in S^{i-1} : g(\cdot, s_{i-1}, \dots, s_1) \in \mathcal{D}_i\}\,.$$

Then, for any $g \in \tilde{\mathcal{D}}_i$, it holds

$$\frac{d}{d\theta} \int g(y) P^Y_{\theta,i}(dy; \hat{y}) = \frac{d}{d\theta} \int g(u, \hat{y}_{i-1}, \dots, \hat{y}_1)\, P_{\theta,i}(du; \hat{y}_{i-1})$$
$$= \int g(u, \hat{y}_{i-1}, \dots, \hat{y}_1)\, (P_{\theta,i})'(du; \hat{y}_{i-1})\,.$$

$\mathcal{D}_i$-differentiability of $P_{\theta,i}$ (see **(A3)**) therefore implies that $P^Y_{\theta,i}$ is $\tilde{\mathcal{D}}_i$-differentiable. Applying Theorem 1 to the product in (21) and substituting $Y_\theta(i)$ by the vector $(X_\theta(i), \dots, X_\theta(1))$ concludes the proof of Corollary 2.