



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

A New Unbiased Stochastic Derivative Estimator for Discontinuous Sample Performances with Structural Parameters

Yijie Peng, Michael C. Fu, Jian-Qiang Hu, Bernd Heidergott

To cite this article:

Yijie Peng, Michael C. Fu, Jian-Qiang Hu, Bernd Heidergott (2018) A New Unbiased Stochastic Derivative Estimator for Discontinuous Sample Performances with Structural Parameters. *Operations Research*

Published online in Articles in Advance 02 Feb 2018

<https://doi.org/10.1287/opre.2017.1674>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2018, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

A New Unbiased Stochastic Derivative Estimator for Discontinuous Sample Performances with Structural Parameters

Yijie Peng,^a Michael C. Fu,^b Jian-Qiang Hu,^c Bernd Heidergott^d

^a Department of Industrial Engineering and Management, College of Engineering, Peking University, Beijing 100871, China;

^b The Robert H. Smith School of Business, Institute for Systems Research, University of Maryland, College Park, Maryland 20742;

^c Department of Management Science, School of Management, Fudan University, Shanghai 200433, China; ^d Department of Econometrics and Operations Research, Vrije Universiteit Amsterdam, 1081 HV Amsterdam, Netherlands

Contact: pengyijie@pku.edu.cn,  <http://orcid.org/0000-0003-2584-8131> (YP); mfu@rhsmith.umd.edu (MCF); hujq@fudan.edu.cn (J-QH); b.f.heidergott@vu.nl (BH)

Received: January 12, 2016

Revised: October 12, 2016; June 12, 2017; June 27, 2017

Accepted: July 27, 2017

Published Online in Articles in Advance: February 2, 2018

Subject Classifications: simulation; inventory/production: sensitivity analysis; inventory/production: uncertainty: stochastic

Area of Review: Simulation

<https://doi.org/10.1287/opre.2017.1674>

Copyright: © 2018 INFORMS

Abstract. In this paper, we propose a new unbiased stochastic derivative estimator in a framework that can handle discontinuous sample performances with structural parameters. This work extends the three most popular unbiased stochastic derivative estimators: (1) infinitesimal perturbation analysis (IPA), (2) the likelihood ratio (LR) method, and (3) the weak derivative method, to a setting where they did not previously apply. Examples in probability constraints, control charts, and financial derivatives demonstrate the broad applicability of the proposed framework. The new estimator preserves the single-run efficiency of the classic IPA-LR estimators in applications, which is substantiated by numerical experiments.

Funding: This work was supported in part by the National Natural Science Foundation of China [Grants 71571048, 71720107003, 71690232, 71371015], by the National Science Foundation [Grants CMMI-0856256, CMMI-1362303, CMMI-1434419], by the Air Force Office of Scientific Research [Grant FA9550-15-10050], and by the Science and Technology Agency of Sichuan Province [Grant 2014GZX0002].

Supplemental Material: The online appendix is available at <https://doi.org/10.1287/opre.2017.1674>.

Keywords: simulation • stochastic derivative estimation • discontinuous sample performance • likelihood ratio • perturbation analysis • weak derivative

1. Introduction

Stochastic derivative estimation is an active research area in simulation optimization, because it plays a central role in both sensitivity analysis and gradient-based optimization (see Asmussen and Glynn 2007). Three of the most popular unbiased stochastic derivative estimators are infinitesimal perturbation analysis (IPA), the likelihood ratio (LR) method (also known as the score function method), and the weak derivative (WD) method; see Ho and Cao (1991), Glasserman (1991), Rubinstein and Shapiro (1993), Pflug (1996), and Fu (2006, 2008, 2015). In this work, we propose a new unbiased stochastic derivative estimator, which generalizes three methods in a given framework.

We consider stochastic models with discontinuous sample performances in the presence of structural parameters (parameters appearing directly in the sample performance). IPA requires continuity of the sample performance, whereas LR can easily handle discontinuous sample performances but only deals with distributional parameters (parameters appearing in the distribution of input random variables) and not structural parameters. The proposed generalized likelihood ratio (GLR) method allows sample performances that may

be discontinuous with respect to structural parameters, where the “generalized” reflects the property that the GLR estimator reduces to the classic LR estimator when there are no structural parameters.

A unified IPA-LR estimator was given by L’Ecuyer (1990) defined on a general probability space requiring—expressed in our framework—continuity of the sample performance when applied to a structural parameter. We provide a representation for the bias of this IPA-LR estimator because of discontinuities using a surface integration, and the GLR estimator is an unbiased modification of IPA-LR. More specifically, GLR is basically a summation of the classic LR estimator and an additional term because of discontinuities. The technique used in the development of GLR can also define a WD estimator for a distribution whose classical derivative cannot be directly defined, and thereby extends WD estimators to a setting that covers a broad range of applications with discontinuities.

Examples in probability constraints (Andrieu et al. 2010), control charts (Fu and Hu 1999), and financial derivatives (Wang et al. 2012), including new problems that cannot be easily addressed by existing methods, for example, compound options (see Online Appendix C), are treated by the GLR framework. The GLR estimator

has an analytical form, thus is computationally efficient to implement and preserves the single-run property of the classic IPA-LR estimator. GLR does not always give the most statistically efficient estimator, for example, for settings in which IPA applies, but is widely applicable and easily implementable.

We briefly review existing methods that address discontinuities in various settings. Smoothed perturbation analysis (SPA), an extension of IPA, is designed to deal with discontinuous sample performances, but finding what to condition on is generally problem dependent, and the estimation of the resulting conditional expectation may require function inversion and additional simulation (Gong and Ho 1987, Fu and Hu 1997). Push-out LR treats discontinuous sample performances with structural parameters, but in general requires an explicit function inversion to push the structural parameters out of the sample performance and into the density (Rubinstein and Shapiro 1993). In Online Appendix B, we show the GLR estimator is an extension of the push-out estimator when the structural parameters cannot be explicitly pushed out.

Early work on addressing discontinuous sample paths in derivative estimation for discrete event systems (DES), for example, queueing, inventory, and maintenance systems, includes Fu and Hu (1993), Fu (1994), and Heidergott (1999). Work focused on treating discontinuous payoffs in Greeks estimation for financial derivatives can be found in Fu and Hu (1995), Broadie and Glasserman (1996), Fournié et al. (1999), Heidergott (2001), Chen and Glasserman (2007), Liu and Hong (2011), and Wang et al. (2012). More recent work dealing with discontinuities in quantile, conditional value-at-risk, and distortion probabilities includes Hong and Liu (2009), Fu et al. (2009), Hong and Liu (2009), Cao and Wan (2014), Jiang and Fu (2015), and Heidergott and Volk-Makarewicz (2016).

We summarize the contributions of our work as follows:

- We derive a new stochastic derivative estimator in a general framework that handles discontinuous sample performances with structural parameters.
- The new estimator is unbiased and single run, involving no explicit functional inversions.
- The new method uses a single framework to treat a broad class of applications, many of which have been treated separately in the literature.

This work focuses on the theoretical foundations of the GLR estimator in a general framework. Concurrently, we have applied GLR to quantile sensitivity, distribution sensitivity, and additional problems in Peng et al. (2017, 2016a, b).

The rest of the paper is organized as follows. General theory for the proposed GLR estimator is established in Section 2. Section 3 provides some applications and numerical performances of the new method. The last section offers conclusions.

2. General Theory

In this section, we formulate the problem in Section 2.1 and provide a general theory for the GLR estimator of problem (2) beginning in Section 2.2 with an overview of the key ideas for its derivation without getting into the technical details, followed by a representation for the bias of IPA-LR in Section 2.3; and concluding in Section 2.4 with the conditions justifying unbiasedness and corresponding theoretical results.

2.1. Problem Formulation

We consider the derivative of an expectation with respect to a scalar parameter θ . Taking θ scalar is without loss of generality, since a gradient could be obtained by taking a vector of the derivatives. Suppose the sample performance (output random variable) is of the following form:

$$\varphi(g(X;\theta)) = \varphi(g_1(X_1, \dots, X_n; \theta), \dots, g_m(X_1, \dots, X_n; \theta)), \quad m \leq n, \quad (1)$$

where

- X denotes an n -dimensional real-valued random vector, in symbols, $X \doteq (X_1, \dots, X_n) \in \mathbb{R}^n$, X_i , $i = 1, \dots, n$, where X_i are input random variables with joint density $f(\cdot; \theta)$,
- $g(\cdot; \theta) \doteq (g_1(\cdot; \theta), \dots, g_m(\cdot; \theta))$, $g_i(\cdot; \theta)$, $i = 1, \dots, m$, are functions that have sufficient differentiability with respect to both the argument and parameter θ , and
- $\varphi: \mathbb{R}^m \rightarrow \mathbb{R}$ is a measurable function that is *continuous almost everywhere* (a.e.).

In practice, a.e. continuity is often satisfied. For example, it covers the case $\varphi(y) = \prod_{i=1}^m h_i(y_i)$, where $h_i(y_i)$ has countably many discontinuity points, which covers all examples in Section 3 and Online Appendix C. In Online Appendix A, the theory for general measurable function φ can be found. In Section 2.4, we will provide sufficient conditions for our GLR estimator specifying the setup in the above bulleted list.

Our theoretical results will focus on the problem of estimating the derivative of the expectation of sample performance (1), i.e.,

$$\frac{\partial}{\partial \theta} \mathbb{E}[\varphi(g(X;\theta))] = \frac{\partial}{\partial \theta} \int_{\mathbb{R}^n} \varphi(g(x;\theta)) f(x;\theta) dx, \quad (2)$$

where $x \doteq (x_1, \dots, x_n)$.

2.2. Overview of GLR

In this subsection, we provide a concise overview for deriving the GLR estimator for problem (2). Three ingredients for the derivation, i.e., function smoothing, integration by parts, and taking limits, will be presented successively.

1. *Function smoothing*. Notice that a major difference between our setting (1) and the classic IPA-LR framework is that φ is not necessarily continuous, so we want

to first “tame” it to be more tractable. For φ continuous a.e., we will show the existence of a smooth sequence of functions that approximate a nonsmooth function in a classic sense. To avoid imposing extra regularity conditions on the smoothed sequence for interchanging limit and integration later, we first truncate function φ to be bounded and then smooth the truncated function.

The indicator function of a set S is $\mathbf{1}_S(z) = 1$ if $z \in S$, 0 otherwise. We truncate φ to $\varphi_L(y) \doteq (\tau_L \circ \varphi)(y) \cdot \mathbf{1}_{S_L}(y)$, where L is a positive number and \circ denotes the composition of two functions,

$$\begin{aligned} \tau_L(z) &\doteq z\mathbf{1}_{[-L,L]}(z) + L\mathbf{1}_{(L,\infty)}(z) - L\mathbf{1}_{(-\infty,-L)}(z), \\ S_L &\doteq \{y \in \mathbb{R}^m: y_i \in [-L, L], i = 1, \dots, m\}. \end{aligned}$$

Note that τ_L endows φ_L with boundedness, and if φ is continuous at y , so is $\tau_L \circ \varphi$ because τ_L is continuous. φ_L has bounded support S_L , and the discontinuity points of φ_L are contained by the union of the discontinuous points of φ and the boundary of S_L that has zero Lebesgue measure. Thus, the truncated function has the following two properties: (1) if φ is continuous a.e., so is φ_L ; (2) for any $y \in \mathbb{R}^m$, $|\varphi_L(y)| \leq L$, $|\varphi_L(y)| \leq |\varphi(y)|$, and

$$\lim_{L \rightarrow \infty} \varphi_L(y) \rightarrow \varphi(y).$$

Next, we show the existence of a “mollifier” that approximates and smooths the truncated measurable function under a continuity condition for φ .

(A.0) $\varphi: \mathbb{R}^m \rightarrow \mathbb{R}$ is a measurable function that is continuous almost everywhere (a.e.).

Theorem 1. For any φ satisfying condition (A.0), there exists a sequence of smooth functions $\varphi_{\epsilon,L}$ such that $\|\varphi_{\epsilon,L}\|_\infty = \sup_{y \in \mathbb{R}^m} |\varphi_{\epsilon,L}(y)| \leq L$, and

$$\lim_{\epsilon \rightarrow 0} \varphi_{\epsilon,L}(y) = \varphi_L(y) \quad a.e.$$

Remark 1. Theorem 1 can be proved by using Lemma 1, which provides an approximation of φ_L by a continuous function, and Lemma 2, which provides an approximation of the continuous function by a smooth function in Online Appendix A, where the function approximation for any measurable function in an \mathbb{L}_p space can also be found.

While the mollifier mappings $\varphi_{\epsilon,L}(y)$ in Theorem 1 are only implicitly given, it is possible to provide an explicit construction in simple cases, as we show in the following example for the case of an indicator mapping.

Example 1. A continuous approximation function for an indicator function $\mathbf{1}_{(-\infty,0]}(\cdot)$ that has a discontinuity point at zero is given by

$$\chi_\epsilon(z) = \begin{cases} 1 & z < -\epsilon, \\ 1 - (z + \epsilon)/(2\epsilon) & -\epsilon \leq z \leq \epsilon, \\ 0 & z > \epsilon. \end{cases} \quad (3)$$

Note that as $\epsilon \rightarrow 0$, $\chi_\epsilon(\cdot) \rightarrow \mathbf{1}_{(-\infty,0]}(\cdot)$ a.e. Moreover, $\partial\chi_\epsilon(z)/\partial z = -1/2$ for $-\epsilon < z < \epsilon$, and 0 for $|z| > \epsilon$, while $\chi_\epsilon(z)$ fails to be differentiable at the boundary points $-\epsilon, \epsilon$.

Suppose the support $\Omega \subset \mathbb{R}^n$ of the density $f(\cdot)$ is independent of θ ; otherwise, a change of variable can be implemented to push parameter θ out of the support (see Wang et al. 2012). Then, we derive the derivative estimator for the expectation of the truncated and smoothed sample performance. Assuming sufficient smoothness on g and f , and that derivative and expectation can be interchanged,

$$\begin{aligned} \frac{\partial}{\partial \theta} \mathbb{E}[\varphi_{\epsilon,L}(g(X; \theta))] &= \frac{\partial}{\partial \theta} \int_{\Omega} \varphi_{\epsilon,L}(g(x; \theta)) f(x; \theta) dx \\ &= \int_{\Omega} s_{\epsilon,L}(x; \theta) f(x; \theta) dx, \end{aligned} \quad (4)$$

where

$$\begin{aligned} s_{\epsilon,L}(x; \theta) &\doteq \underbrace{\sum_{i=1}^m \frac{\partial \varphi_{\epsilon,L}(y)}{\partial y_i} \bigg|_{y=g(x;\theta)} \frac{\partial g_i(x; \theta)}{\partial \theta}}_{\text{IPA part}} \\ &\quad + \underbrace{\varphi_{\epsilon,L}(g(x; \theta)) \frac{\partial \ln f(x; \theta)}{\partial \theta}}_{\text{LR part}}. \end{aligned} \quad (5)$$

Notice that this is a straightforward application of the product rule of differentiation form analysis; see also the discussion in L’Ecuyer (1990), where the above way of organizing the derivatives is referred to as IPA-LR.

The above estimator will be biased because of the truncation and smoothing. In general, we should not expect that directly letting ϵ go to zero could eliminate the bias, because the limit of the smoothed function after differentiation would be ill-behaved at the discontinuity points being smoothed. Taking the explicit approximation function χ_ϵ given by (3), for example, the smoothed function after differentiating $\chi'_\epsilon(\cdot)$ would go to infinity at the discontinuity point being smoothed as ϵ goes to zero, thus the integrand in (4) cannot be dominated by an integrable function. Actually, we know the limit of $\chi'_\epsilon(\cdot)$ is zero a.e., so letting ϵ go to zero in (4) cannot lead to an unbiased estimator in this specific case. A sample performance with indicator functions is a typical setting where the IPA-LR estimator fails, and how to derive an unbiased estimator for this problem has been studied widely in the literature (e.g., Glasserman and Gong 1990, Fu and Hu 1997, Hong and Liu 2010, Heidergott and Volk-Makarewicz 2016).

2. *Integration by parts.* The following treatment uses integration by parts (Evans 1998), derived from the

Gauss-Green Theorem, to deal with the bias introduced by truncation and smoothing in (4). Denote the Jacobian matrix of the vector of functions g by

$$J_g(x; \theta) \doteq \begin{pmatrix} \frac{\partial g_1(x; \theta)}{\partial x_1} & \frac{\partial g_2(x; \theta)}{\partial x_1} & \cdots & \frac{\partial g_m(x; \theta)}{\partial x_1} \\ \frac{\partial g_1(x; \theta)}{\partial x_2} & \frac{\partial g_2(x; \theta)}{\partial x_2} & \cdots & \frac{\partial g_m(x; \theta)}{\partial x_2} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial g_1(x; \theta)}{\partial x_n} & \frac{\partial g_2(x; \theta)}{\partial x_n} & \cdots & \frac{\partial g_m(x; \theta)}{\partial x_n} \end{pmatrix},$$

assuming g is continuously differentiable. By the chain rule,

$$\frac{\partial \varphi_{\epsilon, L}(g(x; \theta))}{\partial x_i} = \sum_{j=1}^m \frac{\partial \varphi_{\epsilon, L}(y)}{\partial y_j} \Big|_{y=g(x; \theta)} \frac{\partial g_j(x; \theta)}{\partial x_i},$$

$$i = 1, \dots, n,$$

or in a matrix form:

$$\nabla_x(\varphi_{\epsilon, L} \circ g)(x; \theta) = J_g(x; \theta) \cdot \nabla_y \varphi_{\epsilon, L}(y) \Big|_{y=g(x; \theta)},$$

where

$$\nabla_x(\varphi_{\epsilon, L} \circ g)(x; \theta) = \left(\frac{\partial \varphi_{\epsilon, L}(g(x; \theta))}{\partial x_1}, \dots, \frac{\partial \varphi_{\epsilon, L}(g(x; \theta))}{\partial x_n} \right)^T,$$

$$\nabla_y \varphi_{\epsilon, L}(y) \Big|_{y=g(x; \theta)} = \left(\frac{\partial \varphi_{\epsilon, L}(y)}{\partial y_1}, \dots, \frac{\partial \varphi_{\epsilon, L}(y)}{\partial y_m} \right)^T \Big|_{y=g(x; \theta)},$$

the superscript T denoting transpose. We assume the following regularity condition.

(A.1) *Matrix invertibility*: There exists an $m \times m$ sub-matrix $\bar{J}_g(x; \theta)$ of the Jacobian $J_g(x; \theta)$ such that $\bar{J}_g(x; \theta)$ is invertible a.e., i.e., $\nu(\mathcal{N}) = 0$, where ν is the Lebesgue measure, and

$$\mathcal{N} \doteq \{x \in \mathbb{R}^n: \det(\bar{J}_g(x; \theta)) = 0\}.$$

To simplify the notation, we assume $m = n$ throughout Section 2.2, so $\bar{J}_g(x; \theta) = J_g(x; \theta)$. Let J_g^{-1} be the inverse of J_g . For $m < n$, \bar{J}_g^{-1} is a generalized inverse of J_g , which will be discussed in Section 2.4. With the notation and assumption (A.1), we have (a.e.)

$$\nabla_y \varphi_{\epsilon, L}(y) \Big|_{y=g(x; \theta)} = J_g^{-1}(x; \theta) \nabla_x(\varphi_{\epsilon, L} \circ g)(x; \theta). \quad (6)$$

Equation (6) transforms the gradient operation with respect to y to a gradient operation with respect to x , which is pivotal for applying integration by parts, so we call (6) the *transformational equation*. With the transformational equation,

$$\int_{\Omega} s_{\epsilon, L}(x; \theta) f(x; \theta) dx = \int_{\Omega} \varphi_{\epsilon, L}(g(x; \theta)) \frac{\partial}{\partial \theta} f(x; \theta) dx$$

$$+ \int_{\Omega} (\partial_{\theta} g(x; \theta))^T J_g^{-1}(x; \theta) \nabla_x(\varphi_{\epsilon, L} \circ g)(x; \theta) f(x; \theta) dx, \quad (7)$$

where $\partial_{\theta} g(x; \theta) \doteq (\partial g_1(x; \theta)/\partial \theta, \dots, \partial g_m(x; \theta)/\partial \theta)^T$. By integration by parts, the second term on the right-hand side of (7) is equal to

$$\int_{\partial \Omega} (\varphi_{\epsilon, L} \circ g) f(\partial_{\theta} g)^T J_g^{-1} \hat{v} ds$$

$$- \int_{\Omega} \varphi_{\epsilon, L}(g(x; \theta)) \operatorname{div}((\partial_{\theta} g(x; \theta))^T J_g^{-1}(x; \theta) f(x; \theta)) dx, \quad (8)$$

where the first term in (8) is a surface integration along the boundary $\partial \Omega$ with the unit norm vector \hat{v} pointing outward, ds is the surface measure on $\partial \Omega$, and for a vector of functions $h(x) = (h_1(x), \dots, h_n(x))$, its divergence is

$$\operatorname{div}(h(x)) = \sum_{i=1}^n \frac{\partial h_i(x)}{\partial x_n}.$$

We will provide sufficient condition for the surface integral to be zero, so that by introducing appropriate likelihood ratios, we obtain

$$\int_{\Omega} s_{\epsilon, L}(x; \theta) f(x; \theta) dx$$

$$= \int_{\Omega} \varphi_{\epsilon, L}(g(x; \theta)) \frac{\partial \ln f(x; \theta)}{\partial \theta} f(x; \theta) dx$$

$$- \int_{\Omega} \varphi_{\epsilon, L}(g(x; \theta)) \operatorname{div}((\partial_{\theta} g(x; \theta))^T J_g^{-1}(x; \theta) f(x; \theta)) dx.$$

In obtaining the estimator, we could alternatively split $\partial f(x; \theta)/\partial \theta$ into a difference of two densities obtained from normalizing the positive and negative part of $\partial f(x; \theta)/\partial \theta$; see Pflug (1996) or Heidergott and Leahu (2010) for details. In this paper, we keep the classic LR term.

3. *Taking limits*. Notice that there is no differentiation operation on the truncated smoothed function $\varphi_{\epsilon, L}$ in (8) after applying integration by parts to move the differentiation to the other terms. Under appropriate regularity conditions that will be formalized in Section 2.4, we can take limits in (7) to obtain an unbiased estimator. Specifically, by letting $\epsilon \rightarrow 0$ and $L \rightarrow \infty$ in (4), (7), and (8), assuming suitable smoothness, and that limit, derivative, integration can be interchanged,

$$\frac{\partial}{\partial \theta} \mathbb{E}[\varphi(g(X; \theta))]$$

$$= \int_{\Omega} \varphi(g(x; \theta)) \left(\frac{\partial \ln f(x; \theta)}{\partial \theta} + d(x; \theta) \right) f(x; \theta) dx, \quad (9)$$

where

$$d(x; \theta) \doteq -\operatorname{div}((\partial_{\theta} g(x; \theta))^T J_g^{-1}(x; \theta) f(x; \theta)) / f(x; \theta).$$

Notice that the smoothing function sequence is not in (9) anymore after taking limits.

The additional term $d(\cdot)$ to the classic LR score function in the GLR estimator is because of the presence of structural parameters. An analytical expression for $d(\cdot)$ given by (11) is derived in Section 2.4, yielding a single-run unbiased estimator.

If we assume the density f goes to zero as x approaches the boundary $\partial\Omega$, the surface integration in (9) could be zero under an appropriate integrability condition (see Online Appendix A). We also provide here a high-level view to understand taking the limit to achieve unbiasedness when the surface integration part is zero. For two measurable functions ψ_1 and ψ_2 under an appropriate integrability condition, we denote the linear operation in the \mathbb{L}_p (inner product in \mathbb{L}_2) space by Rudin (1987)

$$\langle \psi_1, \psi_2 \rangle = \int_{\mathbb{R}^n} \psi_1(y)\psi_2(y) dy.$$

For simplicity of illustration, we assume that g is invertible in this section. By the change of variables $y = g(x; \theta)$, the right-hand side of (9) can be written as

$$\langle \varphi, \psi \cdot \mathbf{1}_{\mathcal{S}} \rangle,$$

where $\psi(y; \theta) \doteq \omega(g^{-1}(y; \theta); \theta) |\det(J_g(g^{-1}(y; \theta); \theta))|^{-1} f(g^{-1}(y; \theta); \theta)$,

$$\omega(x; \theta) \doteq \frac{\partial \ln f(x; \theta)}{\partial \theta} + d(x; \theta),$$

and

$$\mathcal{S}(\theta) \doteq \{y \in \mathbb{R}^m: y = g(x; \theta), x \in \mathbb{R}^n\}, \quad (10)$$

with dependence on θ suppressed in $\mathbf{1}_{\mathcal{S}}$. Similarly, $\langle \varphi_{\epsilon, L}, \psi \cdot \mathbf{1}_{\mathcal{S}} \rangle$ is the right-hand side of (7) subtracting the surface integration term in (8). If φ is continuous a.e., the limit can be taken in a classic sense and unbiasedness of GLR is obtained by interchanging limit and integration, justified by the dominated convergence theorem (Rudin 1987), i.e.,

$$\lim_{\epsilon \rightarrow 0} \langle \varphi_{\epsilon, L}, \psi \cdot \mathbf{1}_{\mathcal{S}} \rangle = \langle \lim_{\epsilon \rightarrow 0} \varphi_{\epsilon, L}, \psi \cdot \mathbf{1}_{\mathcal{S}} \rangle = \langle \varphi, \psi \cdot \mathbf{1}_{\mathcal{S}} \rangle.$$

In Online Appendix A, we show that for any measurable function φ , the limits can be taken in an \mathbb{L}^p space with integrability of an order higher than that of the function $\psi \cdot \mathbf{1}_{\mathcal{S}}$, which is stronger than the first-order integrability condition required when φ is continuous a.e. Therefore, an interesting insight on how to deal with discontinuous sample performances is that less smoothness in the sample performance can be compensated by stronger integrability conditions.

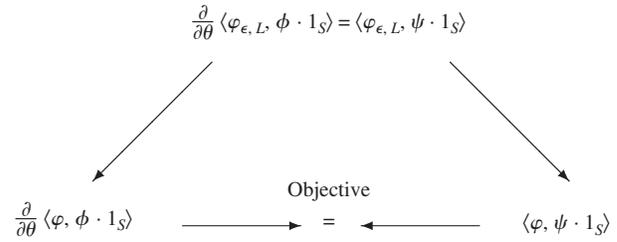
By a change of variables,

$$\mathbb{E}[\varphi(g(X; \theta))] = \langle \varphi, \phi \cdot \mathbf{1}_{\mathcal{S}} \rangle,$$

where

$$\phi(y; \theta) \doteq |\det(J_g(g^{-1}(y; \theta); \theta))|^{-1} f(g^{-1}(y; \theta); \theta),$$

Figure 1. Expanding the Unbiasedness of the GLR Estimator from Smoothed and Truncated $\varphi_{\epsilon, L}$ to the General φ That Is Not Necessarily Smooth and Bounded



and $\phi \cdot \mathbf{1}_{\mathcal{S}}$ is the density of a distribution supported on the image space $\mathcal{S}(\theta)$. Figure 1 illustrates expanding the unbiasedness of the GLR estimator from smoothed and truncated $\varphi_{\epsilon, L}$ to the general φ that is not necessarily smooth and bounded through the three ingredients of the derivation. Notice that through the three ingredients, a WD $\psi \cdot \mathbf{1}_{\mathcal{S}}$ is defined for a distribution $\phi \cdot \mathbf{1}_{\mathcal{S}}$ whose classical derivative cannot be directly defined in general, because the image $\mathcal{S}(\theta)$ may be dependent on the parameter. GLR extends the classical setting of WD (see Pflug 1996 and Heidergott and Leahu 2010) by allowing the existence of structural parameters.

2.3. Bias Representation of IPA-LR

In this subsection, we offer a representation of the bias for IPA-LR because of discontinuities in the sample performance. To illustrate in a simplified setting, if θ is the parameter of interest and $\vartheta(X; \theta)$ is the output sample performance written as a function of the vector of input random variables X following joint density f , then IPA treats integrals of the form

$$\int_{\Omega} \vartheta(x; \theta) f(x) dx,$$

for ϑ continuous, whereas LR treats integral of the form

$$\int_{\Omega} \vartheta(x) f(x; \theta) dx,$$

where ϑ can be discontinuous. A natural generalization (L'Ecuyer 1990) is

$$\int_{\Omega} \vartheta(x; \theta) f(x; \theta) dx,$$

but again ϑ must be continuous with respect to θ , so that IPA can be applied. Differentiating and interchanging with integration then yields an IPA-LR estimator.

In our framework, the bias of the IPA-LR estimator because of the discontinuity in the sample performance can be understood by using integration by parts. Let Ξ be the set of discontinuity points of φ , and the inverse image of g on Ξ is denoted by

$$\Gamma \doteq \{x \in \mathbb{R}^n: g(x; \theta) \in \Xi\}.$$

Suppose $\varphi \circ g$ is continuously differentiable outside of Γ , and Γ is a smooth surface on \mathbb{R}^n with zero Lebesgue measure. We can construct an open set $\Gamma_\epsilon \supset \Gamma$ with Lebesgue measure ϵ (see the proof of Lemma 1 in Online Appendix A). Assuming the surface integration is zero on the boundary $\partial\Omega$, similarly as the procedure to derive (9), we can use integration by parts to obtain

$$\int_{\Omega \setminus \Gamma_\epsilon} s(x; \theta) f(x; \theta) dx = \int_{\partial\Gamma_\epsilon} (\varphi \circ g) f(\partial_\theta g)^T J_g^{-1} \hat{\nu} ds + \int_{\Omega \setminus \Gamma_\epsilon} \varphi(g(x; \theta)) \omega(x; \theta) f(x; \theta) dx,$$

where

$$s(x; \theta) \doteq (\partial_\theta g(x; \theta))^T \nabla_y \varphi(y; \theta)|_{y=g(x; \theta)} + \varphi(g(x; \theta)) \frac{\partial \ln f(x; \theta)}{\partial \theta}.$$

With an appropriate integrability condition, we can shrink Γ_ϵ to Γ such that

$$\int_{\Omega \setminus \Gamma} s(x; \theta) f(x; \theta) dx = \mathcal{F}_\pm + \langle \varphi, \psi \cdot \mathbf{1}_\varphi \rangle,$$

where

$$\mathcal{F}_\pm \doteq \int_{\Gamma^+ \cup \Gamma^-} (\varphi \circ g) f(\partial_\theta g)^T J_g^{-1} \hat{\nu} ds,$$

Γ^+ and Γ^- are two directed surfaces with opposite orientations, corresponding to the undirected surface Γ , and \mathcal{F}_\pm is a surface integration along the two surfaces with opposite orientations. Fortunately, if $\varphi \circ g$ is continuous on Γ , then $\mathcal{F}_\pm = 0$. For the case where $\varphi \circ g$ is discontinuous on Γ , \mathcal{F}_\pm is a representation of the bias of the IPA-LR estimator $s(x; \theta)$. Introducing an appropriate measure on the surfaces, the difference \mathcal{F}_\pm could be estimated as a difference of two (usually rather complex) simulation experiments; see Pflug (1996). In our framework with discontinuities and structural parameters, the GLR estimator can be viewed as an unbiased modification of IPA-LR through the three ingredients presented in Section 2.2.

2.4. Technical Details for GLR

In this subsection, we provide rigorous conditions to justify the unbiasedness of the GLR estimator, and derive an analytical form for the GLR estimator that achieves single-run efficiency in simulation.

The boundary of Ω is given by $\partial\Omega = \bar{\Omega} \setminus \Omega^0$, where $\bar{\Omega}$ is the closure of Ω and Ω^0 is the interior of Ω . We introduce a smoothness condition for the function vector $g(x; \theta)$ and density $f(x; \theta)$.

(A.2) *Smoothness conditions*: $g(x; \theta)$ is twice continuously differentiable with respect to $(x, \theta) \in \Omega \times \Theta$, where Θ is an open neighborhood for the parameter of interest, and $f(x; \theta)$ is continuously differentiable with respect to $(x, \theta) \in \mathbb{R}^n \times \Theta$ and goes to zero as x approaches infinity.

Remark 2. If the support of the density is not all of \mathbb{R}^n , the density f is zero $\forall x \in \partial\Omega$ under condition (A.2); otherwise, the density will be strictly positive at $x_0 \in \partial\Omega$, thus is discontinuous at x_0 , because it is a cluster of points where the density is zero, which violates condition (A.2). Many distributions not supported on the whole space are not continuous on the whole space, for example, the exponential distribution, which is discontinuous at zero. However, in practice, we can implement a change of variables first to transform the support of the density to the whole space to satisfy condition (A.2), for example, for the exponential distribution, we can use transformation $z = \ln x$. When X_i , $i = 1, \dots, n$, are independent, a systematic way to carry out the change of variables is provided in Online Appendix A. An example of a change of variables in the dependent case can be found in the maintenance systems example in Online Appendix C.

For $m < n$, suppose condition (A.1) holds. Without loss of generality, we can assume the submatrix comprising the first m rows of J_g has rank m a.e.; otherwise, we can reorder the indices of g_i , $i = 1, \dots, n$, appropriately. By the chain rule,

$$\nabla_{\bar{x}}(\varphi_{\epsilon, L} \circ g)(x; \theta) = \bar{J}_g(x; \theta) \cdot \nabla_y \varphi_{\epsilon, L}(y; \theta)|_{y=g(x; \theta)},$$

where $\bar{x} \doteq (x_1, \dots, x_m)$, \bar{J}_g is the submatrix comprising of the first m rows of J_g , and

$$\nabla_{\bar{x}}(\varphi_{\epsilon, L} \circ g)(x; \theta) = \left(\frac{\partial \varphi_{\epsilon, L}(g(x; \theta))}{\partial x_1}, \dots, \frac{\partial \varphi_{\epsilon, L}(g(x; \theta))}{\partial x_m} \right)^T, \\ \nabla_y \varphi_{\epsilon, L}(y; \theta)|_{y=g(x; \theta)} = \left(\frac{\partial \varphi_{\epsilon, L}(y)}{\partial y_1}, \dots, \frac{\partial \varphi_{\epsilon, L}(y)}{\partial y_m} \right)^T \Big|_{y=g(x; \theta)},$$

so transformational Equation (6) is adjusted to be

$$\nabla_y \varphi_{\epsilon, L}(y; \theta)|_{y=g(x; \theta)} = \bar{J}_g^{-1}(x; \theta) \nabla_{\bar{x}}(\varphi_{\epsilon, L} \circ g)(x; \theta).$$

We only need to replace $J_g^{-1}(x; \theta)$ and $\nabla_x(\varphi_{\epsilon, L} \circ g)(x; \theta)$ in (7) with $\bar{J}_g^{-1}(x; \theta)$ and $\nabla_{\bar{x}}(\varphi_{\epsilon, L} \circ g)(x; \theta)$, respectively, to make the corresponding equation hold for $m < n$. Let $\bar{y} \doteq y$ if $m = n$ and $\bar{y} \doteq (y, x_{m+1:n})$ if $m < n$, and $\bar{g}(x; \theta) \doteq g(x; \theta)$ if $m = n$ and $\bar{g}(x; \theta) \doteq (g(x; \theta), x_{m+1:n})$ if $m < n$. We introduce the following function invertibility condition.

(A.3) *Function invertibility*: There exist sets Ω_i , $i = 1, \dots, l$, such that $\Omega \setminus \mathcal{N} = \bigcup_{i=1, \dots, l} \Omega_i$ and $\bar{g}(\cdot; \theta): \Omega_i \rightarrow \bar{\mathcal{Y}}_i(\theta)$ is invertible.

Remark 3. Under conditions (A.1) and (A.2), for any $x \in \Omega \setminus \mathcal{N}$, there exists a neighborhood $\mathcal{O}_x \times \mathcal{O}_{\bar{g}(x)}$ of $(x, \bar{g}(x; \theta))$ and a unique differentiable inverse function $\bar{g}^{-1}(\cdot; \theta): \mathcal{O}_{\bar{g}(x)} \rightarrow \mathcal{O}_x$, by the implicit function theorem (Lang 2013). If Ω is compact and $\mathcal{N} = \emptyset$, condition (A.1) implies (A.3) by the Heine-Borel theorem (Lang 2013).

Next, we introduce the remaining regularity conditions required to establish the unbiasedness of the GLR estimator for problem (2).

(A.4) Integrability conditions:

(i) For any $\epsilon, L > 0$,

$$\int_{\Omega} \sup_{\theta \in \Theta} |s_{\epsilon, L}(x; \theta) f(x; \theta)| dx < \infty,$$

where $s_{\epsilon, L}$ is defined by (5).

(ii) $d(\cdot)$ is absolutely integrable:

$$\mathbb{E}[|d(X; \theta)|] = \int_{\Omega} |d(x; \theta)| f(x; \theta) dx < \infty.$$

(iii) If $\{\Omega_{\epsilon}\}$ is a sequence of bounded open sets such that $\Omega_{\epsilon'} \subset \Omega_{\epsilon}$ if $\epsilon' < \epsilon$ and $\Omega = \bigcup_{\epsilon} \Omega_{\epsilon}$, then there exists $0 < c < 1$ such that

$$\lim_{\epsilon \rightarrow 0} \int_{\partial \Omega_{\epsilon}} |(\partial_{\theta} g)^T \bar{J}_g^{-1} \hat{v}| |f|^c ds < \infty.$$

(iv) For $i = 1, \dots, l$,

$$\int_{\mathbb{R}^n} |\varphi(y) \vee 1| |\phi(\bar{y}; \theta) \cdot \mathbf{1}_{\bar{\mathcal{F}}_i(\theta)}(\bar{y})| d\bar{y} < \infty,$$

$$\int_{\mathbb{R}^n} |\varphi(y) \vee 1| \sup_{\theta \in \Theta} |\psi(\bar{y}; \theta) \cdot \mathbf{1}_{\bar{\mathcal{F}}_i(\theta)}(\bar{y})| d\bar{y} < \infty,$$

where $z_1 \vee z_2 \doteq \max\{z_1, z_2\}$, and ϕ and ψ equal zero outside of $\bar{\mathcal{F}}_i$.

Remark 4. Condition (i) justifies the interchange of derivative and integration in (4). Conditions (i)–(iii) together with condition (A.2) justify the use of integration by parts in (8). Condition (iii) together with condition (A.2) also guarantee the surface integration $\int_{\partial \Omega}$ is zero. Condition (iv) together with condition (A.4) justifies the unbiasedness of the GLR estimator by interchanging limits with respect to ϵ and L and derivative with respect to θ , which requires a certain uniform convergence. Under (A.2), the inverse image of f at zero, i.e., $\{x: f(x) = 0\}$, is a closed set; thus the support Ω is an open set and the monotone expansion from bounded open sets $\{\Omega_{\epsilon}\}$ to Ω in (iii) is always feasible.

Remark 5. The generality of GLR comes at the price that only the existence of $\varphi_{\epsilon, L}$ and \bar{g}^{-1} is provided, thus the integrability condition (A.4) cannot generally be checked in practice. An exception is the case that $\varphi_{\epsilon, L}$ and \bar{g}^{-1} can be explicitly given.

Theorem 2. Under continuity condition (A.0), the matrix invertibility condition (A.1), smoothness condition (A.2), function invertibility condition (A.3), and integrability condition (A.4),

$$\frac{\partial}{\partial \theta} \mathbb{E}[\varphi(g(X; \theta))] = \mathbb{E}[\varphi(g(X; \theta)) \omega(X; \theta)],$$

where $\omega(x; \theta) = \partial \ln f(x; \theta) / \partial \theta + d(x; \theta)$, and

$$d(x; \theta) = \sum_{i=1}^m (\bar{J}_g^{-1}(x; \theta) \partial_{x_i} \bar{J}_g(x; \theta) \bar{J}_g^{-1}(x; \theta) e_i)^T \partial_{\theta} g(x; \theta) - \text{trace}(\bar{J}_g^{-1}(x; \theta) \partial_{\theta} \bar{J}_g(x; \theta)) - (\partial_{\theta} g(x; \theta))^T \bar{J}_g^{-1}(x; \theta) \nabla_{\bar{x}} \ln f(x; \theta), \quad (11)$$

where $\partial_z \bar{J}_g$ means differentiating every element in the matrix \bar{J}_g with respect to z , and

$$\nabla_{\bar{x}} \ln f(x; \theta) = (\partial \ln f(x; \theta) / \partial x_1, \dots, \partial \ln f(x; \theta) / \partial x_m)^T.$$

Remark 6. The proof of the Theorem 2 can be found in Online Appendix A, where an alternative justification of the unbiasedness for any measurable function φ can also be found. The GLR estimator for problem (2) is given by $\varphi(g(X; \theta)) \omega(X; \theta)$, and has an analytical form that can be calculated by differentiation, matrix inversion, and elementary operations. If the sample performance contains no structural parameters, the GLR estimator simplifies to the classic LR estimator, because all the ∂_{θ} terms would be zero, and hence $d(x; \theta)$ would be zero. For general measurable φ discussed in Online Appendix A, it requires a stronger condition, i.e., condition (A.5), which is more difficult to verify than condition (A.4), to justify the unbiasedness of GLR.

We offer a more tractable version of conditions (A.1)–(A.4) for the special case that (a) the derivatives of $\varphi_{\epsilon, L}(y)$ exist a.e. for each ϵ and are uniformly bounded in y and ϵ on all points of differentiability; (b) the input random variables are independent and supported on the whole space; (c) $g: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a linear (affine) transformation:

$$g(x; \theta) = A(\theta)x + B(\theta),$$

where $A(\theta)$ is a parameterized $n \times n$ matrix and $B(\theta)$ is a parameterized $n \times 1$ vector, with sufficient smoothness. Push-out LR applies to this linear transformation and would lead to the same formula as GLR in this case (see Online Appendix B), but GLR applies more generally to nonlinear g where the structural parameter is not easy to be explicitly pushed out (see Online Appendix C).

Condition (a) is satisfied by the indicator function in Example 1 and all applications presented in Section 3; condition (c) applies to most cases where SPA and push-out LR are easily implementable. Thus, this special set of conditions applies to a rich class of problems that are of importance in applications. It is easy to show that conditions (A.1) and (A.2) are satisfied if $A(\theta)$ is invertible, and f satisfies the smoothness in condition (A.2). We can easily calculate

$$J_g(x; \theta) = A(\theta), \quad g^{-1}(y; \theta) = A^{-1}(\theta)(y - B(\theta)),$$

$$d(x; \theta) = -\text{trace}(A^{-1}(\theta) \partial_{\theta} A(\theta)) - (\partial_{\theta} A(\theta)x + \partial_{\theta} B(\theta))^T A^{-1}(\theta) \nabla_x \ln f(x; \theta).$$

Since g is invertible, condition (A.3) is satisfied. Integrability conditions (i)–(iii) in (A.4) can be simplified as follows:

(i') Suppose

$$\int_{\mathbb{R}^n} \sup_{\theta \in \Theta} |\partial_{\theta} f(x; \theta)| dx < \infty,$$

$$\max_{i=1, \dots, n} \int_{\mathbb{R}^n} \sup_{\theta \in \Theta} |e_i^T (\partial_{\theta} A(\theta)x + \partial_{\theta} B(\theta)) f(x; \theta)| dx < \infty.$$

(ii') For $i = 1, \dots, n$,

$$\int_{\mathbb{R}} |\partial_{x_i} f_i(x_i; \theta)| dx_i < \infty, \quad \int_{\mathbb{R}} |x_i \partial_{x_i} f_i(x_i; \theta)| dx_i < \infty,$$

where f_i is the marginal density of X_i .

(iii') For $i = 1, \dots, n$,

$$\lim_{x_i \rightarrow \infty} |x_i| f_i(x_i; \theta) = 0, \quad \int_{\mathbb{R}} |x_i| f_i(x_i; \theta) dx_i < \infty.$$

Notice that for this special case, the integrability condition (A.4) involves no implicitly constructed functions. Uniform integrability is often used to justify the unbiasedness of classic IPA and LR by applying the mean value theorem and dominated convergence theorem (Glasserman 1991). Thus, checking the integrability condition required to justify the unbiasedness of the GLR estimator for this special case is no more difficult than that of IPA and LR in principle. For f and A independent of θ ,

$$\omega(x; \theta) = -(\partial_{\theta} B(\theta))^T A^{-1} \nabla_x \ln f(x),$$

integrability condition (i') and the second inequality in (ii') are automatically satisfied, and integrability conditions (iv) can be further simplified as follows:

(iv') For $i = 1, \dots, n$,

$$\int_{\mathbb{R}^n} |\varphi(y) \vee 1| |f(A^{-1}(y - B(\theta)))| dy < \infty,$$

$$\int_{\mathbb{R}^n} |\varphi(y) \vee 1| \sup_{\theta \in \Theta} |\partial_{\theta} b_i(\theta) \partial_{x_i} f(x)|_{x=A^{-1}(y-B(\theta))} dy < \infty,$$

where $B(\theta) = (b_1(\theta), \dots, b_n(\theta))^T$.

3. Applications

In this section, we consider three applications previously analyzed by three different methods in the literature. Applications in maintenance systems and compound options can be found in Online Appendix C.

3.1. Probability Constraints (Andrieu et al. 2010)

We consider an investment problem where the capital is borrowed at interest rate r . Let θ_1 be the proportion of capital invested in a bond with fixed rate b , and θ_2 be the proportion of capital invested in a risky asset with a random rate X , where $\mathbb{E}[X] > r$, which means that there

is a positive risk premium. The remaining proportion of capital $1 - \theta_1 - \theta_2$ is for consumption. Andrieu et al. (2010) aim to maximize the sum of consumption satisfaction and the expected final capital subject to probability constraint

$$P((1+b)\theta_1 + (1+X)\theta_2 \geq 1+r) \geq \pi,$$

which means that the probability of repayment is at least as high as π . Since the parameters of interest θ_1 and θ_2 appear in the probability, which is the expectation of an indicator function that is discontinuous, IPA and LR do not apply. Andrieu et al. (2010) construct an explicit mollifier with a tuning parameter to address the discontinuity, similar to the ϵ -approximation in Example 1, which leads to a biased derivative estimator. For the actual optimization algorithm, they let the tuning parameter go to zero as the number of iteration steps in stochastic approximation (SA) goes to infinity. For such a Kiefer-Wolfowitz type algorithm (Kushner and Yin 2003), the step size and the tuning parameter in the mollifier needs to be well tuned at proper rates as the step goes to infinity in SA, which is a nontrivial task even in simple examples. Fortunately, as we will show in the following, applying GLR we obtain an unbiased derivative estimator, which allows for a Robbins-Monro type algorithm (Kushner and Yin 2003) and thereby overcomes the difficulty of reducing the bias induced by the mollifiers at the correct rate.

We estimate

$$\frac{\partial P((1+b)\theta_1 + (1+X)\theta_2 \geq 1+r)}{\partial \theta},$$

for $\theta = \theta_1 = 0.4$, $\theta_2 = 0.4$, $r = 0.05$, $b = 0.1$, X a normal random variable with mean $\mu = 0.2$ and variance $\sigma^2 = 0.04$. Since this example falls into the special case at the end of Section 2, conditions (A.1)–(A.4) can be easily checked for this example. The GLR estimator is

$$\mathbf{1}\{(1+b)\theta_1 + (1+X)\theta_2 \geq 1+r\} \frac{(1+b)(X-\mu)}{\theta_2 \sigma^2}.$$

Andrieu et al. (2010) provide an approximation by convolution (AC) method and propose

$$\frac{1}{\delta} h\left(\frac{1+r - (1+b)\theta_1 - (1+X)\theta_2}{\delta}\right) (1+b),$$

where $h(x) = 3(1-x^2)I(x)/4$, $I(x) = 1$ if $-1 \leq x \leq 1$ and $I(x) = 0$ otherwise. We compare GLR with AC and finite difference method with common random numbers (FDC). $AC(\delta)$ and $FDC(\delta)$ denote AC with tuning parameter δ and FDC with perturbation size δ , respectively. The standard error is defined by the estimated standard deviation divided by the square root of the number of samples.

From Table 1, we see that both AC and FDC suffer from large bias when $\delta = 0.1$ and large variance

Table 1. Derivatives Estimates of the Probability Constraint with Respect to θ , Based on 10^6 Independent Replications (Mean \pm Standard Error)

| GLR | AC(0.1) | AC(0.01) | FDC(0.1) | FDC(0.01) |
|------------------|------------------|------------------|------------------|------------------|
| 1.46 \pm 0.006 | 1.77 \pm 0.003 | 1.47 \pm 0.010 | 3.49 \pm 0.005 | 1.62 \pm 0.009 |

when $\delta = 0.01$. FDC(0.001) outputs 1.46 \pm 0.003 (mean \pm standard error) based on 10^8 independent replications, which substantiates the unbiasedness of GLR. In addition, the performance of AC heavily relies on the choice of the mollifier h , and Andrieu et al. (2010) show that some choices perform very poorly, for example, $h(x) = I(x)$. In contrast, GLR is an unbiased derivative estimator without any mollifier and tuning parameter appearing in the final estimator.

3.2. Control Charts (Fu and Hu 1999)

In statistical process control, control charts are used to determine if a manufacturing or business process is in a state of statistical control. A performance of interest is the average run length

$$ARL(\theta) = \mathbb{E}[N] = \sum_{n=1}^{\infty} nP(N = n),$$

where N is the time when the system declares out-of-control, i.e.,

$$N = \min\{i: Y_i \notin [\theta_{1,i}, \theta_{2,i}]\},$$

and

$$Y_i = \psi_i(X_i, Y_{i-1}; \theta), \quad i > 1, Y_1 = X_1,$$

where θ is a generic parameter that can be $\theta_{1,i}$ and $\theta_{2,i}$ in this example, X_i is the output of the i th sample, Y_i is the (observable) test statistic after the i th sample, ψ_i is the transition function of a Markov chain $\{Y_i\}$, $\theta_{1,i}$ and $\theta_{2,i}$ are the lower control limit and upper control limit for the i th test statistic, respectively.

We want to estimate the sensitivity of the average run length with respect to parameter θ . To treat this problem in our framework, we can rewrite the sample performance as

$$\begin{aligned} V_n(X_1, \dots, X_n; \theta) &= \mathbf{1}\{N = n\} \\ &= \left(\prod_{i=1}^{n-1} \mathbf{1}\{0 < g_i^{(n)}(X_1, \dots, X_i; \theta) < 1\} \right) \\ &\quad \times (1 - \mathbf{1}\{0 < g_n^{(n)}(X_1, \dots, X_n; \theta) < 1\}), \\ g_i^{(n)}(X_1, \dots, X_i; \theta) &= \frac{Y_i - \theta_{1,i}}{\theta_{2,i} - \theta_{1,i}}, \end{aligned} \quad (12)$$

where $\prod_{i=1}^0 \doteq 1$. The system will output samples of different statistical behaviors when it is in control and out of control. Conditional on $\{Z = z\}$, where $Z = R/\Delta$, R is

a (unobservable) random duration for the system to go out of control and follows a distribution with density $q_0(\cdot)$, and Δ is the sampling interval (duration between two monitoring epochs), the conditional density of X_i is given by

$$f_i(x_i | z) = \mathbf{1}\{i < z\}q_1(x_i) + \mathbf{1}\{i \geq z\}q_2(x_i),$$

where $q_1(\cdot)$ and $q_2(\cdot)$ are the densities of the distributions of the sampling process when in control and out of control, respectively. Conditional on Z , (X_1, \dots, X_n) are independent, and their conditional joint density is given by

$$f^{(n)}(x_{1:n}; z) = \prod_{i=1}^n f(x_i | z),$$

where $x_{1:n} \doteq (x_1, \dots, x_n)$.

The derivative estimators in Fu and Hu (1999) require extra simulation to estimate a conditional expectation term and an explicit inverse for ψ_i , for example, for Shewhart chart, $Y_i = X_i$, and for EWMA chart, $Y_i = \alpha X_i + (1 - \alpha)Y_{i-1}$, whereas our framework does not have these requirements.

In the numerical experiment, we let $\theta_{1,i} = \theta_1$ and $\theta_{2,i} = \theta_2$ and test the sensitivity with respect to $\theta = \theta_2$, where classical IPA and LR do not apply. Suppose the time to go out of control follows an exponential distribution, in-control sample output follows a normal distribution with mean μ_0 and variance σ^2 , and out-of-control sample output follows a normal distribution with mean μ_1 and variance σ^2 . The GLR estimator for the EWMA control chart is given by

$$\begin{aligned} \nabla_{x_{1:n}} f^{(n)}(x_{1:n}; z) &= - \left(\frac{x_i - (\mu_0 \mathbf{1}\{i < z\} + \mu_1 \mathbf{1}\{i \geq z\})}{\sigma^2} \right)_{i=1, \dots, n}^T, \\ \partial_{\theta} g^{(n)}(x; \theta) &\doteq - \frac{1}{\theta_2 - \theta_1} (g_1^{(n)}(x_1; \theta), \dots, g_n^{(n)}(x_1, \dots, x_n; \theta))^T, \end{aligned}$$

and

$$J_g^{(n)} \doteq \frac{1}{\theta_2 - \theta_1} \begin{pmatrix} 1 & 1-\alpha & (1-\alpha)^2 & \dots & (1-\alpha)^{n-2} & (1-\alpha)^{n-1} \\ 0 & \alpha & \alpha(1-\alpha) & \dots & \alpha(1-\alpha)^{n-3} & \alpha(1-\alpha)^{n-2} \\ 0 & 0 & \alpha & \dots & \alpha(1-\alpha)^{n-4} & \alpha(1-\alpha)^{n-3} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & \dots & \alpha & \alpha(1-\alpha) \\ 0 & 0 & \dots & \dots & 0 & \alpha \end{pmatrix}_{n \times n}.$$

It is easy to verify the conditions (A.1)–(A.4) for the unbiasedness of the GLR estimator for $\partial_{\theta} P(N = n) / \partial \theta$, which falls into the special case at the end of Section 2. The GLR estimator for $\partial ARL / \partial \theta$ is given by

$$\begin{aligned} N \left[\frac{N}{\theta_2 - \theta_1} - (\partial_{\theta} g^{(N)}(X_1, \dots, X_N; \theta))^T \right. \\ \left. \cdot (J_g^{(N)})^{-1} \nabla_{x_{1:n}} f^{(n)}(x_{1:n}; Z) \Big|_{n=N, x_{1:n}=X_{1:n}} \right], \end{aligned}$$

Table 2. Derivative of Shewhart Control Chart with Respect to Upper Control Limit θ_2 (Mean \pm Standard Error)

| | $\mu_1 = 1.0$ | $\mu_1 = 3.0$ |
|-----------------|----------------|-----------------|
| 10,000 reps | | |
| FDC(0.1) | 70.9 \pm 2.2 | 3.63 \pm 0.34 |
| FDC(0.01) | 67.5 \pm 7.0 | 3.96 \pm 1.3 |
| PA _L | 59.5 \pm 1.3 | 3.64 \pm 0.06 |
| PA _R | 61.8 \pm 0.6 | 3.91 \pm 0.18 |
| GLR | 61.0 \pm 4.0 | 3.32 \pm 1.1 |
| 1,000,000 reps | | |
| FDC(0.1) | 71.2 \pm 0.2 | 3.54 \pm 0.03 |
| FDC(0.01) | 63.1 \pm 0.6 | 3.80 \pm 0.1 |
| GLR | 62.8 \pm 0.4 | 3.77 \pm 0.1 |
| True value | 63.0 | 3.73 |

where $X_{1:n} \doteq (X_1, \dots, X_n)$. The unbiasedness for the estimator of $\partial \text{ARL} / \partial \theta$ can be justified by imposing additional integrability condition

$$\sum_{n=1}^{\infty} n \sup_{\theta \in \Theta} \left| \frac{\partial_{\theta} P(N = n)}{\partial \theta} \right| < \infty,$$

which justifies interchange of summation and differentiation by mean value theorem and dominated convergence theorem.

We test the performance for the Shewhart chart, because the analytical form of the derivative can be used to access the accuracy of the estimates. We compare GLR with the PA_L and PA_R estimators in Fu and Hu (1999), and FDC under the same setting as in Fu and Hu (1999): $\alpha = 1$, sampling frequency $\Delta = 1$, time to go out of control following an exponential distribution with mean 20, sample output following a normal distribution with in-control mean $\mu_0 = 0$ and out-of-control mean $\mu_1 = 1.0$ and 3.0 , and variance $\sigma = 1$, lower and upper control limits $\theta_1 = -2.81$ and $\theta_2 = 2.81$, respectively, for all $i \in \mathbb{N}$.

The results of PA_L and PA_R are taken from Fu and Hu (1999). From Table 2, we can see FDC suffers from large bias when $\delta = 0.1$ and large variance when $\delta = 0.01$. Although the variances of PA_L and PA_R are smaller than that of GLR, they achieve this by running extra simulation, so it is not clear which is superior after taking the extra simulation into account.

3.3. Barrier Options (Wang et al. 2012)

To cover more applications, we consider a sample performance more general than (1) as follows:

$$Q(X; \theta) = \vartheta(V(X; \theta); \theta), \quad (13)$$

where $\vartheta(v; \theta)$, $v \doteq (v_1, \dots, v_k)$, is differentiable with respect to θ and v , $V(x; \theta) \doteq (V_1(x; \theta), \dots, V_k(x; \theta))$,

$$V_i(x; \theta) \doteq \varphi_i(g_1(x; \theta), \dots, g_{m_i}(x; \theta)),$$

and $\varphi_i(\cdot)$ is a measurable function with $m_i \leq n$, $i = 1, \dots, k$. With additional mild regularity conditions, the GLR estimator for sample performance (13) is

$$\left. \frac{\partial \vartheta(v; \theta)}{\partial \theta} \right|_{v=V(X; \theta)} + Q(X; \theta) \omega(X; \theta),$$

where ω is defined by (11) with $g(x; \theta) = (g_1(x; \theta), \dots, g_m(x; \theta))$, where $m \doteq \max_{i=1, \dots, k} m_i$ and \bar{J}_g being an $m \times m$ submatrix of the Jacobian of g . The only major differences in the derivation are as follows: the IPA-LR $s_{\epsilon, L}(x; \theta)$ after truncation and smoothing in (4) needs to be changed to

$$(\partial_{\theta} g(x; \theta))^T \Sigma(x; \theta) \nabla_v \vartheta(v; \theta) \Big|_{v=V(x; \theta)},$$

where

$$\begin{aligned} \Sigma(x; \theta) &\doteq ((\nabla_y \hat{\varphi}_1(y; \theta))^T, \dots, (\nabla_y \hat{\varphi}_k(y; \theta))^T)^T \Big|_{y=g(x; \theta)}, \\ \nabla_y \hat{\varphi}_i(y; \theta) &= \left(\frac{\partial \hat{\varphi}_i(y)}{\partial y_1}, \dots, \frac{\partial \hat{\varphi}_i(y)}{\partial y_{m_i}}, 0, \dots, 0 \right)^T_{1 \times m}, \\ & \quad i = 1, \dots, k, \end{aligned}$$

$\hat{\varphi}_i$, $i = 1, \dots, k$, denote the truncated and smoothed functions, and by the chain rule,

$$\nabla_{\bar{x}}(\vartheta \circ V)(x; \theta) = \bar{J}_g(x; \theta) \Sigma(x; \theta) \nabla_v \vartheta(v; \theta) \Big|_{v=V(x; \theta)},$$

so the transformational Equation (6) is changed to

$$\begin{aligned} &(\partial_{\theta} g(x; \theta))^T \Sigma(x; \theta) \nabla_v \vartheta(v; \theta) \Big|_{v=V(x; \theta)} \\ &= (\partial_{\theta} g(x; \theta))^T \bar{J}_g^{-1}(x; \theta) \nabla_{\bar{x}}(\vartheta \circ V)(x; \theta). \end{aligned}$$

A similar analysis in Section 2 follows, and the unbiasedness of the GLR estimator for sample performance (1) can be straightforwardly extended to the more general case (13).

In a particular probability space where simulation can be naturally implemented, sample performance (13) covers the IPA-LR framework if we view $\vartheta(v; \theta)$ as a simulation model with structural parameter θ and a vector of input random variables $v = V(X; \theta)$, where X is the vector of uniform random numbers. Our framework also covers many discontinuous settings previously studied in literature, for example, the setting $\varphi(g_1(x; \theta), g_2(x; \theta))$, where $\varphi(y_1, y_2) = y_1 \mathbf{1}\{y_2 \leq 0\}$, in a kernel-based method (Liu and Hong 2011) and support independent unified likelihood ratio and infinitesimal perturbation analysis (SLRIPA) (Wang et al. 2012).

Framework (13) can easily handle sensitivity estimates of barrier options, which were first addressed in Wang et al. (2012) using SLRIPA. In this example, we assume the underlying asset process follows a geometric Brownian motion process, i.e., $S_t = S_0 e^{(r - \sigma^2/2)t + \sigma B_t}$, where S_t and B_t are the asset price and standard Brownian motion at time t , respectively.

The numerical results for the underlying asset following a jump-diffusion process can be found in Online Appendix C.

An up-and-out barrier option is a financial derivative that becomes worthless if the path of the underlying asset exceeds the barrier H . For a European up-and-out barrier option, the payoff in the discrete-time setting is given by

$$e^{-n\Delta}(S_{n\Delta} - K)\mathbf{1}\left\{\max_{i=1,\dots,n-1} S_{i\Delta} < H\right\}\mathbf{1}\{K < S_{n\Delta} < H\},$$

where $T = n\Delta$ is the time to maturity and K is the strike price ($K < H$). We have that for $i = 1, \dots, n - 1$,

$$S_{i\Delta}/H = \exp(g_i(X_1, \dots, X_i; \theta)),$$

$$S_{n\Delta}/K = \exp\{\log(H/K)g_n(X_1, \dots, X_n; \theta)\},$$

and by taking log and simple algebra,

$$\mathbf{1}\{S_{i\Delta} < H\} = \mathbf{1}\{g_i(X_1, \dots, X_i; \theta) < 0\},$$

$$\mathbf{1}\{K < S_{n\Delta} < H\} = \mathbf{1}\{1 < S_{n\Delta}/K < H/K\}$$

$$= \mathbf{1}\{0 < g_n(X_1, \dots, X_n; \theta) < 1\},$$

where for $i = 1, \dots, n - 1$,

$$g_i(X_1, \dots, X_i; \theta) = \log(S_0/H) + \sigma\sqrt{\Delta} \sum_{j=1}^i X_j + i\left(r - \frac{\sigma^2}{2}\right)\Delta,$$

$$g_n(X_1, \dots, X_n; \theta) = \frac{1}{\log(H/K)} \left(\log(S_0/K) + \sigma\sqrt{\Delta} \sum_{j=1}^n X_j + n\left(r - \frac{\sigma^2}{2}\right)\Delta \right).$$

Thus, the sample performance of the payoff for European up-and-out barrier option can be rewritten by

$$Q(X_1, \dots, X_n; \theta) = e^{-rn\Delta} K \{ \exp[\log(H/K)V_1(X_1, \dots, X_n; \theta)] - 1 \} \cdot V_2(X_1, \dots, X_n; \theta),$$

where θ is a generic parameter that can be r, Δ, K, H, S_0 , and σ in this example, Δ is the discrete step size, $X_i = (B_{i\Delta} - B_{(i-1)\Delta})/\sqrt{\Delta}$, $i = 1, \dots, n$, which is standard normally distributed if the underlying asset follows geometric Brownian motion,

$$V_1(X_1, \dots, X_n; \theta) = g_n(X_1, \dots, X_n; \theta),$$

$$V_2(X_1, \dots, X_n; \theta) = \varphi(g(X_1, \dots, X_n; \theta)),$$

where

$$\varphi(y_1, \dots, y_n) = \prod_{i=1}^{n-1} \mathbf{1}\{y_i < 0\} \mathbf{1}\{0 < y_n < 1\}.$$

Then we test the performance of GLR for the European barrier options, and compare it with the SLRIPA, SPA, and FDC, which are tested in Wang et al. (2012) using the following parameter values: $\sigma = 0.1, r = 0.05$,

$H = 110, S_0 = K = 100, n\Delta = 1$. For FDC, we perturb θ to $\theta + \delta$ with $\delta = 0.2$. As in Wang et al. (2012), we test sensitivity with respect to parameter $\theta = H$, where classical IPA and LR do not apply. GLR can also be applied straightforwardly to estimate the sensitivity with respect to other parameters.

In this example, we have $\nabla_x \ln f(x; \theta) = -(x_1, \dots, x_n)$, where f is the joint density of n i.i.d. standard normal random variables, $\partial_\theta g(x; \theta) = -(1/H)(1, \dots, 1, g_n(x_1, \dots, x_n; \theta)/\log(H/K))$,

$$J_g = \sigma\sqrt{\Delta} \begin{pmatrix} 1 & 1 & \dots & 1 & (\log(H/K))^{-1} \\ 0 & 1 & \dots & 1 & (\log(H/K))^{-1} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & (\log(H/K))^{-1} \\ 0 & 0 & \dots & 0 & (\log(H/K))^{-1} \end{pmatrix},$$

$$J_g^{-1} = \frac{1}{\sigma\sqrt{\Delta}} \begin{pmatrix} 1 & -1 & \dots & 0 & 0 \\ 0 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & -1 \\ 0 & 0 & \dots & 0 & \log(H/K) \end{pmatrix}.$$

We can construct a continuous and piecewise smooth function as

$$\varphi_\epsilon(y_1, \dots, y_n) = \prod_{i=1}^{n-1} \chi_\epsilon(y_i)(\chi_\epsilon(y_n - 1) - \chi_\epsilon(y_n)).$$

Notice that this example falls into the special case discussed at the end of Section 2. Conditions (A.1)–(A.3), and (i)–(iii) in (A.4) are obviously satisfied. Since the input random variables follow normal distributions, integrability condition (iv) in (A.4) can be justified similarly as in the proof of Lemma 2 in Online Appendix A. Thus, the GLR estimator is guaranteed to be unbiased and has the following form:

$$\frac{e^{-rT}K}{H} \mathbf{1}\left\{\max_{i=1,\dots,n-1} g_i(X_1, \dots, X_i; \theta) < 0\right\} \cdot \mathbf{1}\{0 < g_n(X_1, \dots, X_n; \theta) < 1\} \times \left\{ \exp\left[\log\left(\frac{H}{K}\right)g_n(X_1, \dots, X_n; \theta)\right] - 1 \right\} \cdot \left[\frac{(1 - g_n(X_1, \dots, X_n; \theta))X_n - X_1}{\sigma\sqrt{\Delta}} + \frac{1}{\log(H/K)} \right] + \exp\left[\log\left(\frac{H}{K}\right)g_n(X_1, \dots, X_n; \theta)\right] g_n(X_1, \dots, X_n; \theta) \}.$$

The SLRIPA and SPA estimators for this example can be found in Online Appendix C. The GLR and SLRIPA estimators have analytical forms and are single-run, while the SPA estimator includes a conditional expectation term that is not in analytical form and requires extra simulation to estimate.

The numerical results are shown in Table 3, where the results for SLRIPA, SPA, and FDC are taken from

Table 3. Sensitivity of European Barrier Option with Respect to Parameter H , Based on 2,000 Independent Replications (Mean \pm Standard Error)

| | $n = 10$ | $n = 20$ | $n = 30$ |
|--------|-------------------|-------------------|-------------------|
| FDC | 0.329 \pm 0.084 | 0.310 \pm 0.077 | 0.271 \pm 0.071 |
| SPA | 0.286 \pm 0.011 | 0.262 \pm 0.013 | 0.254 \pm 0.015 |
| SLRIPA | 0.276 \pm 0.018 | 0.251 \pm 0.024 | 0.266 \pm 0.028 |
| GLR | 0.278 \pm 0.020 | 0.263 \pm 0.025 | 0.255 \pm 0.029 |

Wang et al. (2012). Note that the SPA estimator requires additional simulation to estimate some conditional expectation terms (see Online Appendix C), and the variances of these estimator are not reflected in the reported standard errors in Table 3. All results are obtained from 2000 independent replications for each estimator. GLR and SLRIPA have comparable variances and superior computational efficiency. SPA has the lowest variance but requires a substantially higher amount of computation because of the additional simulations. FDC has the largest variances in this example, and is generally biased. The convergence rate of FDC in terms of mean-squared error can be found in L'Ecuyer and Perron (1994).

4. Conclusions

In this paper, we propose a systematic procedure using function smoothing and integration by parts for deriving an unbiased derivative estimator for discontinuous sample performances with structural parameters. For the particular model specified in Section 2.1, the GLR estimator extends IPA, LR, and WD to a setting where they did not previously apply. Applications in probability constraints, statistical process control, and financial derivatives can be treated in our general framework. Although GLR is not always the best alternative, numerical experiments substantiate its overall single-run efficiency. The proposed GLR estimator is compatible with variance reduction techniques such as splitting in the WD method, conditional Monte Carlo in SPA, and renewal theory in steady-state simulation. How to further reduce the variance of GLR is an important direction for future research.

Acknowledgments

The authors thank the referees, editors, and Cheng Jie for their comments and suggestions that have greatly improved the paper.

References

Andrieu L, Cohen G, Vázquez-Abad FJ (2010) Gradient-based simulation optimization under probability constraints. *Eur. J. Oper. Res.* 212(2):345–351.
 Asmussen S, Glynn PW (2007) *Stochastic Simulation: Algorithms and Analysis*, Vol. 57 (Springer, New York).
 Broadie M, Glasserman P (1996) Estimating security price derivatives using simulation. *Management Sci.* 42(2):269–285.

Cao X-R, Wan X (2014) Sensitivity analysis of nonlinear behavior with distorted probability. *Math. Finance* 27(1):115–150.
 Chen N, Glasserman P (2007) Malliavin Greeks without Malliavin calculus. *Stochastic Processes and Their Appl.* 117(11):1689–1723.
 Evans LC (1998) *Partial Differential Equations* (American Mathematical Society, Providence, RI).
 Fournié E, Lasry J-M, Lebuchoux J, Lions P-L, Touzi N (1999) Applications of Malliavin calculus to Monte Carlo methods in finance. *Finance and Stochastics* 3(4):391–412.
 Fu MC (1994) Sample path derivatives for (s, S) inventory systems. *Oper. Res.* 42(2):351–364.
 Fu MC (2006) Gradient estimation. Henderson SG, Nelson BL, eds. *Handbooks in Operations Research and Management Science*, Vol. 13 (Elsevier, Amsterdam), 575–616.
 Fu MC (2008) What you should know about simulation and derivatives. *Naval Res. Logist.* 55(8):723–736.
 Fu MC (2015) Stochastic gradient estimation. Fu MC, ed. *Handbook of Simulation Optimization*, Vol. 216 (Springer, New York), 105–147.
 Fu MC, Hu J-Q (1993) Second derivative sample path estimators for the GI/G/m queue. *Management Sci.* 39(3):359–383.
 Fu MC, Hu J-Q (1995) Sensitivity analysis for Monte Carlo simulation of option pricing. *Probab. Engrg. Information Sci.* 9(3):417–446.
 Fu MC, Hu J-Q (1997) *Conditional Monte Carlo: Gradient Estimation and Optimization Applications* (Kluwer Academic Publishers, Boston).
 Fu MC, Hu J-Q (1999) Efficient design and sensitivity analysis of control charts using Monte Carlo simulation. *Management Sci.* 45(3):395–413.
 Fu MC, Hong LJ, Hu J-Q (2009) Conditional Monte Carlo estimation of quantile sensitivities. *Management Sci.* 55(12):2019–2027.
 Glasserman P (1991) *Gradient Estimation via Perturbation Analysis* (Kluwer Academic Publishers, Boston).
 Glasserman P, Gong W-B (1990) Smoothed perturbation analysis for a class of discrete-event systems. *IEEE Trans. Automatic Control* 35(11):1218–1230.
 Gong W-B, Ho Y-C (1987) Smoothed perturbation analysis of discrete event dynamical systems. *IEEE Trans. Automatic Control* 32(10):858–866.
 Heidegott B (1999) Optimisation of a single-component maintenance system: A smoothed perturbation analysis approach. *Eur. J. Oper. Res.* 119(1):181–190.
 Heidegott B (2001) Option pricing via Monte Carlo simulation: A weak derivative approach. *Probab. Engrg. Information Sci.* 15(3):335–349.
 Heidegott B, Leahu H (2010) Weak differentiability of product measures. *Math. Oper. Res.* 35(1):27–51.
 Heidegott B, Volk-Makarewicz W (2016) A measure-valued differentiation approach to sensitivity analysis of quantiles. *Math. Oper. Res.* 41(1):293–317.
 Ho Y-C, Cao X-R (1991) *Discrete Event Dynamic Systems and Perturbation Analysis* (Kluwer Academic Publishers, Boston).
 Hong LJ, Liu G (2009) Simulating sensitivities of conditional value at risk. *Management Sci.* 55(2):281–293.
 Hong LJ, Liu G (2010) Pathwise estimation of probability sensitivities through terminating or steady-state simulations. *Oper. Res.* 58(2):357–370.
 Jiang G, Fu MC (2015) Technical note—On estimating quantile sensitivities via infinitesimal perturbation analysis. *Oper. Res.* 63(2):435–441.
 Kushner HJ, Yin GG (2003) *Stochastic Approximation and Recursive Algorithms and Applications* (Springer, New York).
 Lang S (2013) *Undergraduate Analysis* (Springer, New York).
 L'Ecuyer P (1990) A unified view of the IPA, SF, and LR gradient estimation techniques. *Management Sci.* 36(11):1364–1383.
 L'Ecuyer P, Perron G (1994) On the convergence rates of IPA and FDC derivative estimators. *Oper. Res.* 42(4):643–656.
 Liu G, Hong LJ (2011) Kernel estimation of the Greeks for options with discontinuous payoffs. *Oper. Res.* 59(1):96–108.
 Peng Y, Fu MC, Hu J-Q (2016a) Estimating distribution sensitivity using generalized likelihood ratio method. Cassandras CG, Giua A, Li Z, eds. *Proc. 13th Internat. Workshop on Discrete Event Systems, WODES '16* (IEEE, Piscataway, NJ), 123–128.
 Peng Y, Fu MC, Hu J-Q (2016b) On the regularity conditions and applications for generalized likelihood ratio method. *Proc. Winter Simulation Conf., WSC '16* (IEEE, Piscataway, NJ), 919–930.

- Peng Y, Fu MC, Glynn PW, Hu J-Q (2017) On the asymptotic analysis of quantile sensitivity estimation by Monte Carlo simulation. *Proc. Winter Simulation Conf., WSC '17* (IEEE, Piscataway, NJ).
- Pflug GC (1996) *Optimization of Stochastic Models* (Kluwer Academic, Boston).
- Rubinstein RY, Shapiro A (1993) *Discrete Event Systems: Sensitivity Analysis and Stochastic Optimization by the Score Function Method* (John Wiley & Sons, New York).
- Rudin W (1987) *Real and Complex Analysis* (McGraw-Hill Education, New York).
- Wang Y, Fu MC, Marcus SI (2012) A new stochastic derivative estimator for discontinuous payoff functions with application to financial derivatives. *Oper. Res.* 60(2):447–460.

Yijie Peng is an assistant professor in the Department of Industrial Engineering and Management at Peking University. His research interests include sensitivity analysis and ranking and selection in simulation optimization, with applications in manufacturing and financial engineering.

Michael C. Fu holds the Smith Chair of Management Science in the Robert H. Smith School of Business, with a joint

appointment in the Institute for Systems Research and affiliate faculty appointment in the Department of Electrical and Computer Engineering, all at the University of Maryland. He is a Fellow of INFORMS and IEEE.

Jian-Qiang Hu is a professor with the Department of Management Science, School of Management, Fudan University. His research interests include discrete-event stochastic systems, simulation, queueing network theory, stochastic control theory, with applications toward supply chain management, risk management in financial markets and derivatives, and healthcare.

Bernd Heidergott is a professor of stochastic optimization at the Department of Econometrics and Operations Research at the Vrije Universiteit Amsterdam, the Netherlands. He is programme director of the B.S.c and M.Sc. Econometrics and Operations Research, and research fellow of the Tinbergen Institute and of EURANDOM. His research interests are optimization and control of discrete event systems, perturbation analysis of Markov chains, max-plus algebra, and social networks.