What You Should Know About Simulation and Derivatives

Michael C. Fu

Robert H. Smith School of Business & Institute for Systems Research, University of Maryland, College Park, Maryland 20742

Received 18 June 2008; accepted 18 June 2008
DOI 10.1002/nav.20313
Published online 23 September 2008 in Wiley InterScience (www.interscience.wiley.com).

Abstract: Derivatives (or gradients) are important for both sensitivity analysis and optimization, and in simulation models, these can often be estimated efficiently using various methods other than brute-force finite differences. This article briefly summarizes the main approaches and discusses areas in which the approaches can most fruitfully be applied: queueing, inventory, and finance. In finance, the focus is on derivatives of another sort. Naval Research Logistics 55: 723–736, 2008

Keywords: stochastic simulation; Monte Carlo simulation; derivatives; gradient estimation; perturbation analysis; likelihood ratio method; score function method; weak derivatives; financial Greeks; Malliavin calculus

1. INTRODUCTION

Math programmers are of course primarily interested in optimization and also in sensitivity analysis. The availability of derivatives for sensitivity analysis is taken for granted in the traditional linear programming community. Similarly, for convex optimization and more general nonlinear programming approaches that incorporate gradient-based search, derivatives are key. Although researchers on the stochastic side of operations research are aware of the central role that derivatives play in traditional optimization, deterministic optimizers may not be aware of derivative (or gradient) estimation techniques in simulation. However, they often employ simulation in their research, so in conducting sensitivity analysis for a simulation model, they would simply use "brute-force" resimulation of the model by varying the values of the parameters of interest. If there are a lot of parameters, such an approach can clearly become highly inefficient and prohibitively expensive in terms of computational cost. Furthermore, in optimization applications, numerical inaccuracies (bias) in finite difference estimates can lead to slower convergence rates. The purpose of this article is to raise awareness of the vast array of tools available for estimating derivatives in a stochastic simulation that can be orders of magnitude more efficient than the "brute-force" finite difference approach.

The word "derivatives" in the title of this article has two meanings. The primary usage is the usual calculus definition, but the pricing and hedging of *financial derivatives* (a financial asset whose value is "derived" from some other underlying assets or entities) is one of the most successful applications for the techniques discussed in this article, so examples from that domain will be covered.

Here is something you should know about financial derivatives: They were designed to be used for hedging to reduce risk, but they can also be used for speculation (leveraging your bets) that can increase risk. Thus, the benefits and risks of financial derivatives are controversial, as evidenced by the following opinions featured on the front page of the Business section of the March 6, 2003, Washington Post [3], from two of the most prominent figures in the business world.

> Warren Buffet: "Derivatives are financial weapons of mass destruction, carrying dangers that, while now latent, are potentially lethal. . . . We view them as time bombs, both for the parties that deal in them and the economic system."
>
> Alan Greenspan: "These increasingly complex financial instruments have especially contributed, particularly over the past couple of stressful years, to the development of a far more flexible, efficient and resilient financial system than existed just a quarter-century ago."

The pricing and hedging of financial derivatives is the primary goal of financial engineering, and the two most widely utilized numerical techniques are Monte Carlo simulation and partial differential equations methods. To hedge a financial derivative, one takes positions (buying and selling/shorting) in other assets (which could include other financial derivatives) to offset the risk in the currently held or sold financial

Correspondence to: M.C. Fu (mfu@rhsmith.umd.edu)

instrument. For example, for a mortgage-backed security (a type of financial derivative based on holding a basket of mortgages), the primary risks are associated with interest rates and default (which are usually correlated). To hedge when using simulation, price sensitivity (calculus) derivatives—which we will call *price sensitivities* for short when used in the context of financial derivatives—play the central role in quantifying and managing risk. Therefore, the ability to estimate them efficiently in simulation is critical, as stated emphatically in the Lancaster Prize-winning book, *Monte Carlo Methods in Financial Engineering* [29, p. 377]:

> "Whereas the prices themselves can often be observed in the market, their sensitivities cannot, so accurate calculation of sensitivities is arguably even more important than calculation of prices."

However, the state of widespread knowledge in 2000 regarding simulation and both types of derivatives is exemplified by the following prescription in the 4th edition of a very popular textbook on derivatives pricing [41, p. 410]:

> "Suppose that we are interested in the partial derivative of f with q, where f is the value of the derivative and q is the value of an underlying variable or a parameter. First, Monte Carlo simulation is used in the usual way to calculate an estimate, f, for the value of the derivative. A small increase, Δq, is then made in the value of q, and a new value for the derivative, f^*, is calculated. An estimate for the hedge parameter is given by
>
> $$\frac{f^* - f}{\Delta q}$$
>
> In order to minimize to [sic] standard error of the estimate of the Greek letter, the number of time intervals, N, the random number streams, and the number of trials, M, should be the same for estimating both f and f^*."

In contrast to this brief recipe, all of Chapter 7 of Glasserman [29] is devoted to the topic.

Another important application of simulation and (calculus) derivatives is supply chain management. In the October 30, 2000 issue of Fortune Magazine, an article entitled, "New Victories in the Supply-Chain Revolution" [54] describes "a classic distribution challenge: how to avoid lost sales without incurring the cost of carrying extra inventory" when Caterpillar, the "world's largest builder of construction equipment... posed daunting supply chain questions" regarding the distribution of a new line of "compact" construction machines. A team from Carnegie-Mellon led by Sridhar Tayur and Alan Scheller-Wolf came up with a solution to "determine the appropriate inventory levels for the U.S., the most important market... Among the techniques the Carnegie-Mellon group used to attack this complex problem was so-called infinitesimal perturbation analysis..." Infinitesimal perturbation analysis (IPA) is one of the derivative estimation techniques discussed in detail in this article.

These two examples provide a glimpse at the opportunity provided by the availability of derivatives in the context of simulation. The rest of this article presents a high-level overview of three main approaches for estimating derivatives in simulation (see Fu [13], upon which most of the material in the next section is based, for more technical details): perturbation analysis (PA), the likelihood ratio/score function (LR/SF) method, and weak derivatives (WD); along with some further descriptions of three main application areas: queueing, inventory, and finance. For more details regarding simulation optimization applications and issues, see Fu [14].

2. OVERVIEW OF MAIN APPROACHES

Let $J(\theta)$ denote the real-valued performance measure of interest, where θ is the parameter with which we would like to estimate the derivative $\frac{dJ}{d\theta}$. In the optimization context, θ would generally be referred to as a decision variable and comprise a vector, and a constraint set would also be specified. Here, for expositional ease, we will assume that θ is scalar. Examples of performance measures are expected waiting time in a queueing system, expected costs in an inventory control system, and a stock option price in financial engineering. Possible parameters could be mean service time in a queueing system, a reorder point in an inventory control system, and the current stock price in the option pricing problem.

In simulation, it is assumed that the performance measure must be estimated. Again, for expositional ease, we will assume that the performance measure is an expectation, which covers most quantities of interest including probabilities, but does not include quantiles or the mode. Thus, we write $J = E[L]$, where L will be referred to as the *sample* performance, because the idea of Monte Carlo simulation (or stochastic simulation, or in this article, just simulation) is that by taking a large number of usually independent simulation samples, one can get a good estimate of the performance measure of interest (by the law of large numbers). Thus, the problem is to estimate

$$\frac{dE[L]}{d\theta} \tag{1}$$

using simulation.

A "brute-force" finite difference estimate is obtained by doing additional simulations at parameter value $(\theta + \Delta\theta)$, subtracting from these samples of L the simulation estimates of L at the original parameter value θ, and then dividing by the perturbation $\Delta\theta$. However, if the number of parameters is large (e.g., large number of decision variables in an optimization problem), then this could entail a large number of additional simulations. Furthermore, one has to select a value of $\Delta\theta$, which usually entails a trade-off between variance and bias, where larger $\Delta\theta$ generally means lower variance but higher bias when the goal is to estimate (1). The order of the bias can be decreased by using symmetric differences, i.e., simulate at both $(\theta + \Delta\theta)$ and $(\theta - \Delta\theta)$, but the downside is that this would require basically double the number of simulations.

The central message of this article is that research over the last three decades has resulted in many more accurate and efficient ways to estimate derivatives from simulation. The methods specifically discussed in this article fall into two main approaches, which depend on where the θ dependence appears. To be specific, we express the expectation as follows:

$$E[L(X)] = \int y dF_L(y) = \int L(x) dF_X(x), \qquad (2)$$

where X is a *vector* of all the input random variables used in the simulation, F_L is the distribution of L, F_X is the (joint) distribution of X, and the second integral is multidimensional (corresponding to the dimension of the vector X). In most cases, F_L is not known explicitly; else, simulation would not be needed. However, in the simulation setting F_X is known, because it must be used to generate the input process to the simulation model. In a queueing system, F_X would include all the interarrival times, service times, routing indicators, and other randomness in the system. Note that we have not put the parameter θ anywhere in the expression. The dependence on θ is the crux of the derivative estimation problem.

We consider two cases on where the θ dependence occurs: though the input random variables X themselves or in the distribution (measure) of the input random variables F_X, corresponding to the left-most or right-most side of Eq. (2), respectively:

$$E[L(X)] = \int_0^1 [L(X(\theta; u))du, \qquad (3)$$

and

$$E[L(X)] = \int_{-\infty}^{\infty} L(x) f(x;\theta) dx, \qquad (4)$$

where in the first case the underlying randomness is a vector of random numbers that generate the input random variables X depending on the parameter θ, and the second case assumes the existence of a (joint) density f of all of the input random variables corresponding to the distribution F_X (in the discrete case, the integral would be replaced by a sum, and the probability density function by a probability mass function). Note that the integrals should be multidimensional in both cases.

To illustrate the ideas more concretely, we begin with a specific example involving just two input random variables, assumed to be generated independently, with $X_1 \sim \exp(\theta)$ and $X_2 \sim U(0,1)$, i.e., the first random variable is exponentially distributed with mean θ and the second is uniformly distributed over [0,1]. Then Eqs. (3) and (4) are written, respectively, as

$$E[L(X)] = \int_0^1 \int_0^1 L(X_1(u_1;\theta), u_2) du_1 du_2, \qquad (5)$$

where $X_1(u;\theta) = -\theta \ln u$ is the usual way to generate an $\exp(\theta)$ random variate, and

$$E[L(X)] = \int_0^1 \int_0^{\infty} L(x_1, x_2) \frac{1}{\theta} e^{-x_1/\theta} dx_1 dx_2. \qquad (6)$$

Differentiating Eqs. (5) and (6) assuming that the differentiation operator can be brought inside the integral—which can be verified directly in the second case, but depends on the specific form of L in the first case—we get for the first case:

$$\frac{dE[L(X)]}{d\theta} = \int_0^1 \int_0^1 \frac{\partial L(X_1(u_1;\theta), u_2)}{\partial X_1} \frac{dX_1(\theta; u_1)}{d\theta} du_1 du_2, \qquad (7)$$

and for the second case, we further split into two different expressions:

$$\frac{dE[L(X)]}{d\theta} = \int_0^1 \int_0^{\infty} L(x_1, x_2) \left[\frac{1}{\theta}\left(\frac{x_1}{\theta} - 1\right)\right] \times \frac{1}{\theta} e^{-x_1/\theta} dx_1 dx_2 \qquad (8)$$

$$= \frac{1}{\theta}\left\{\int_0^1 \int_0^{\infty} L(x_1, x_2) \frac{x_1}{\theta^2} e^{-x_1/\theta} dx_1 dx_2 - \int_0^1 \int_0^{\infty} L(x_1, x_2) \frac{1}{\theta} e^{-x_1/\theta} dx_1 dx_2\right\}, \qquad (9)$$

with respective estimators corresponding to Eqs. (7)–(9):

$$\frac{\partial L(X_1, X_2)}{\partial X_1} \frac{dX_1}{d\theta}, \quad L(X_1, X_2)\left[\frac{1}{\theta}\left(\frac{X_1}{\theta} - 1\right)\right], \quad \frac{1}{\theta}\{L(X_1^*, X_2) - L(X_1, X_2)\},$$

where in the third estimator, X_1^* has an Erlang distribution with shape parameter 2, denoted by Erl(2, θ). The key points to note for each of the three estimators are as follows. To finalize the first estimator requires some knowledge of the *specific form* of L (to calculate the first term) and also defining what is the *derivative of a random variable* with respect to a parameter (to calculate the second term); intuitively, though, differentiating $X_1(u;\theta) = -\theta \ln u$ gives

$\frac{dX_1(u;\theta)}{d\theta} = -\ln u = \frac{X_1(u;\theta)}{\theta}$. The second estimator can be implemented *without any further knowledge or simulation* over what is required to estimate the original performance measure by $L(X_1, X_2)$, i.e., the original performance measure estimator is simply multiplied by a "weight" function. The third estimator is the *difference* of two quantities, one of which requires an *additional* simulation at a *different* setting of the first input random variable. This simple example contains the main ingredients of the general case for the three derivative estimation approaches that we consider.

Returning to the more general settings of Eqs. (3) and (4), again differentiating each, assuming an interchange of integration and differentiation is permissible,

$$\frac{dE[L(X)]}{d\theta} = \int_0^1 \frac{dL(X(\theta;u))}{d\theta} du, \tag{10}$$

and

$$\frac{dE[L(X)]}{d\theta} = \int_{-\infty}^{\infty} L(x) \frac{\partial f(x;\theta)}{d\theta} dx. \tag{11}$$

For expositional ease in introducing the approaches, we begin by assuming that the parameter only appears in X_1, which is generated *independently* of the other input random variables. So for the case of (10), we use the chain rule to write

$$\begin{aligned}\frac{dE[L(X)]}{d\theta} &= \int_0^1 \frac{dL(X_1(\theta;u_1), X_2, \ldots)}{d\theta} du \\ &= \int_0^1 \frac{\partial L}{\partial X_1} \frac{dX_1(\theta;u_1)}{d\theta} du \end{aligned} \tag{12}$$

In other words, the estimator takes the same general form as before:

$$\frac{\partial L(X)}{\partial X_1} \frac{dX_1}{d\theta}, \tag{13}$$

where the parameter appears in the transformation from random number to random variate, and the derivative is expressed as the product of a sample path derivative and derivative of a random variable. The issue of what constitutes the latter will be taken up shortly, but this approach is called infinitesimal perturbation analysis (IPA). For a first-come, first-served single-server queue, the sample path derivative could be derived using Lindley's equation, relating the time in system of a customer to the service times (and interarrival times, which are not a function of the parameter). It is just as simple for multi-server queues and the natural extension to queueing networks; however, depending on other characteristics of the network, the resulting IPA estimator may or may not be biased.

Assume that X_1 has marginal p.d.f. $f_1(\cdot;\theta)$ and that the joint p.d.f. for the remaining input random variables $(X_2, \ldots)$ is given by f_{-1}, which has no (functional) dependence on θ. Then the assumed independence gives $f = f_1 f_{-1}$, and the expression (11) involving differentiation of a density (measure) can be further manipulated using the product rule of differentiation to yield the following:

$$\frac{dE[L(X)]}{d\theta} = \int_{-\infty}^{\infty} L(x) \frac{\partial f_1(x_1;\theta)}{\partial\theta} f_{-1}(x_2, \ldots) dx \tag{14}$$

$$= \int_{-\infty}^{\infty} L(x) \frac{\partial \ln f_1(x_1;\theta)}{\partial\theta} f(x) dx. \tag{15}$$

In other words, the likelihood ratio or score function (LR/SF) estimator takes the form

$$L(X) \frac{\partial \ln f_1(X_1;\theta)}{\partial\theta}, \tag{16}$$

where the second term in (16) is called the score function in statistics, hence one of the monikers for this method. The other name is due to differentiating the likelihood ratio function given by

$$\frac{f_1(\cdot;\theta)}{f_1(\cdot;\theta_0)},$$

with respect to θ to obtain

$$\frac{\partial f_1(\cdot;\theta)/\partial\theta}{f_1(\cdot;\theta_0)},$$

which also equals the second term in (16) for $\theta_0 = \theta$.

On the other hand, if instead of expressing the right-hand side of (14) as (15), the density derivative is written as

$$\frac{\partial f_1(x_1;\theta)}{\partial\theta} = c(\theta)\big(f_1^{(2)}(x_1;\theta) - f_1^{(1)}(x_1;\theta)\big),$$

where $f_1^{(2)}$ and $f_1^{(1)}$ are probability density functions, then the following relationship is obtained:

$$\begin{aligned}\frac{dE[L(X)]}{d\theta} = c(\theta) \bigg(&\int_{-\infty}^{\infty} L(x) f_1^{(2)}(x_1;\theta) f_{-1}(x_2, \ldots) dx \\ &- \int_{-\infty}^{\infty} L(x) f_1^{(1)}(x_1;\theta) f_{-1}(x_2, \ldots) dx \bigg). \end{aligned} \tag{17}$$

Any such (nonunique) triple $(c(\theta), f_1^{(2)}, f_1^{(1)})$ constitutes a weak derivative (WD) for f_1, giving the corresponding WD estimator

$$c(\theta)\big(L\big(X_1^{(2)}, X_2, \ldots\big) - L\big(X_1^{(1)}, X_2, \ldots\big)\big), \tag{18}$$

where $X_1^{(1)} \sim f_1^{(1)}$, $X_1^{(2)} \sim f_1^{(2)}$, and it is assumed that common random numbers is used for the other input random variables, in which a single realization of $X_2, X_3, \ldots$, is used in both simulation estimates of L (note that Eq. (17) only implies that they have the same joint distribution f_{-1}). The term "weak" derivative comes about from the possibility that $\frac{\partial f_1(\cdot;\theta)}{\partial\theta}$ in (14) may not be proper, but its *integral* may be well-defined, e.g., it might involve delta-functions (impulses), corresponding to mass functions of discrete distributions. In other words, it is Eq. (17) that defines the WD approach and not Eq. (14).

The IPA estimator (13) requires the notion of a derivatives of a random variable, which intuitively coincides with the definition in the usual sense by defining the limiting difference on a common probability space and fixing the sample outcome; see [13, 28, 29] for technical details. Here, we will consider a couple of useful cases.

The first case assumes the existence of a density, for which we have

$$\frac{dX(\theta)}{d\theta} = -\left.\frac{\partial F(x;\theta)/\partial\theta}{\partial F(x;\theta)/\partial x}\right|_{x=X}, \tag{19}$$

where F is the cumulative distribution function of X, so that the denominator is simply the probability density function.

Another useful case in practice is when the parameter is a location, scale, or generalized scale parameter, defined as follows in terms of the random variable (rather than the distribution, as is often done, cf. [13]):

- θ is a *scale* parameter for X if there exists a random variable Z that has no dependence on θ such that $X = Z\theta$ with probability 1.
- θ is a *location* parameter for X if there exists a random variable Z that has no dependence on θ such that $X = Z + \theta$ with probability 1.
- θ is a *generalized* scale parameter for X if there exist a constant $\bar\theta$ and a random variable Z that has no dependence on θ such that $X = \bar\theta + Z\theta$ with probability 1.

Then one can use the following sample derivatives for the three respective cases (location, scale, generalized scale):

$$\frac{dX}{d\theta} = 1, \quad \frac{dX}{d\theta} = \frac{X}{\theta}, \quad \frac{dX}{d\theta} = \frac{X-\bar\theta}{\theta}.$$

An example is the exponential distribution, for which the mean is a scale parameter, verifying the result $\frac{dX}{d\theta} = \frac{X}{\theta}$ that was intuitively derived earlier through direct differentiation of the expression $X(u,\theta) = -\theta \ln u$, and which could also be derived using Eq. (19). Nearly every distribution has at least one parameter that is location or (generalized) scale, and many have both (e.g., normal, Cauchy, Gumel, logistic, general uniform).

To illustrate the ideas concretely, we consider two sample performance functions of two random variables, but this time we do not specify the actual distributions.

EXAMPLE 1: $L(X) = X_1 + X_2$.

- IPA estimator:

$$\frac{dX_1}{d\theta} + \frac{dX_2}{d\theta},$$

regardless of any dependence between X_1 and X_2.
- LR/SF estimators:

$$(X_1 + X_2)\frac{\partial \ln f(X_1, X_2;\theta)}{\partial\theta},$$

for the general bivariate case, and

$$(X_1 + X_2)\left(\frac{\partial \ln f_1(X_1,\theta)}{\partial\theta} + \frac{\partial \ln f_2(X_2,\theta)}{\partial\theta}\right),$$

for the independent case.
- WD estimators:

$$c(\theta)\big(\big(X_1^{(2)} + X_2^{(2)}\big) - \big(X_1^{(1)} + X_2^{(1)}\big)\big),$$

where $(X_1^{(1)}, X_2^{(1)}) \sim f^{(1)}$, $(X_1^{(2)}, X_2^{(2)}) \sim f^{(2)}$ for the general bivariate case with WD $(c(\theta), f^{(2)}, f^{(1)})$, and

$$c_1(\theta)\big(\big(X_1^{(2)} + X_2\big) - \big(X_1^{(1)} + X_1\big)\big) + c_2(\theta)\big(\big(X_1 + X_2^{(2)}\big) - \big(X_1 + X_2^{(1)}\big)\big),$$

where $X_1^{(1)} \sim f_1^{(1)}$, $X_1^{(2)} \sim f_1^{(2)}$, $X_2^{(1)} \sim f_2^{(1)}$, $X_2^{(2)} \sim f_2^{(2)}$ for the independent case with WDs $(c_1(\theta), f_1^{(2)}, f_1^{(1)})$ and $(c_2(\theta), f_2^{(2)}, f_2^{(1)})$.

EXAMPLE 2: $L(X) = \max(X_1, X_2)$.

- IPA estimator:

$$\frac{dX_1}{d\theta}\mathbf{1}\{X_1 > X_2\} + \frac{dX_2}{d\theta}\mathbf{1}\{X_1 \le X_2\},$$

regardless of any dependence between X_1 and X_2.
- LR/SF estimators:

$$\max(X_1, X_2)\frac{\partial \ln f(X_1, X_2;\theta)}{\partial\theta},$$

for the general bivariate case, and

$$\max(X_1, X_2)\left(\frac{\partial \ln f_1(X_1,;\theta)}{\partial\theta} + \frac{\partial \ln f_2(X_2,;\theta)}{\partial\theta}\right),$$

for the independent case.
- WD estimators:

$$c(\theta)\big(\max\big(X_1^{(2)}, X_2^{(2)}\big) - \max\big(X_1^{(1)}, X_2^{(1)}\big)\big),$$

where $(X_1^{(1)}, X_2^{(1)}) \sim f^{(1)}$, $(X_1^{(2)}, X_2^{(2)}) \sim f^{(2)}$ for the general bivariate case with WD $(c(\theta), f^{(2)}, f^{(1)})$, and

$$c_1(\theta)\big(\max\big(X_1^{(2)}, X_2\big) - \max\big(X_1^{(1)}, X_1\big)\big) + c_2(\theta)\big(\max\big(X_1, X_2^{(2)}\big) - \max\big(X_1, X_2^{(1)}\big)\big),$$

where $X_1^{(1)} \sim f_1^{(1)}$, $X_1^{(2)} \sim f_1^{(2)}$, $X_2^{(1)} \sim f_2^{(1)}$, $X_2^{(2)} \sim f_2^{(2)}$ for the independent case with WDs $(c_1(\theta), f_1^{(2)}, f_1^{(1)})$ and $(c_2(\theta), f_2^{(2)}, f_2^{(1)})$.

In Example 1, IPA also has linearly increasing variance; however, most sample performances are not just simple sums, but involve other operations, such as the maximum operation, as in Example 2. In both examples, notice that the LR/SF and WD estimators do not depend on the form of the sample performance, but on the number of random variables involved. Also, the simple examples illustrate that where the parameter appears is in some sense under the control of the modeler/analyst. So a service-time parameter can be viewed as a parameter of the underlying distribution generating the service time random variate, or it can be viewed as a parameter directly influencing the value of the service time random variable itself, where the underlying uncertainty is simply a stream of random numbers. In some cases, the parameter could even appear in both places at the same time [44].

Generalizing these two examples leads to the following general forms of estimators, where the parameter can possibly occur in every input random variable (otherwise, that term is simply zero):

- IPA estimator:

$$\sum_i \frac{\partial L(X)}{\partial X_i}\frac{dX_i}{d\theta},$$

- LR/SF estimators: (multivariate, independent)

$$L(X)\frac{\partial \ln f(X;\theta)}{\partial\theta}, L(X)\sum_i \frac{\partial \ln f_i(X_i;\theta)}{\partial\theta},$$

- WD estimators: (multivariate, independent)

$$c(\theta)(L(X^{(2)}) - L(X^{(1)})),$$
$$\sum_i c_i(\theta)\big(L\big(X_1,\ldots,X_i^{(2)},\ldots\big) - L\big(X_1,\ldots,X_i^{(1)},\ldots\big)\big),$$

where $X^{(1)} \sim f^{(1)}$, $X^{(2)} \sim f^{(2)}$ for the general multivariate case with WD $(c(\theta), f^{(2)}, f^{(1)})$, and $\{X_i^{(1)} \sim f_i^{(1)}, X_i^{(2)} \sim f_i^{(2)}\}$, for the independent case with WDs $\{(c_i(\theta), f_i^{(2)}, f_i^{(1)})\}$.

Table 1, taken from Fu [13], provides specific expressions needed to implement the IPA, WD, or LR/SF estimators for many common input distributions, both continuous and discrete.

2.1. Guidelines

In this section, we provide brief guidelines on the challenges and implementation issues that face each of the approaches, and then conclude with some general remarks on the research literature.

2.1.1. Challenges

- Both the LR/SF and WD approaches are likely to encounter computational challenges in the case where the same parameter recurs, e.g., if the simulation model involves i.i.d. input random variables where θ is a parameter in the common distribution. Examples include service times or interarrival times in queueing and customer demands in inventory control. More details are provided below when implementation is discussed.
- The main requirement for IPA to succeed is that the sample performance be continuous with respect to the parameter of interest. Although technically this is only a sufficient and not necessary condition (i.e., it is not difficult to construct examples where IPA works for discontinuous sample performances), for all practical purposes this is also a necessary condition. In practice, a discontinuity manifests itself in two principal forms: inherently discontinuous sample performances, e.g., probability performance measures (where the sample performances are indicator functions) or any performance measure whose sample performance is discrete-valued (except in the trivial case where it is constant as a function of the parameter of interest); and simulation models where small changes in the parameter of interest could cause an abrupt change in system behavior such as a different event to occur, e.g., a customer in a queueing system balking rather than being served; one way of checking the latter is the "commuting condition" [28].
- LR/SF and WD may encounter difficulties in handling *nondistributional* parameters, such as inventory control parameters, although a change of variables sometimes can be used to move the parameter into a distribution. The ability to move the parameter into or out of the underlying input distributions are referred to as "push-out" and "push-in" methods in the SF literature [45].
- The LR/SF method cannot handle parameters that change the underlying support of the distribution. The

Table 1. Derivatives for some common/simple input distributions.

Input dist $X \sim F$	WD $(c, F^{(2)}, F^{(1)})$	IPA $\frac{dX}{d\theta}$	LR/SF $\frac{\partial \ln f(X;\theta)}{\partial\theta}$
$\mathrm{Ber}(\theta; a, b)$	$(1, a, b)$	NA	$\frac{1}{\theta}\mathbf{1}\{X=a\}$ $-\frac{1}{1-\theta}\mathbf{1}\{X=b\}$
$\mathrm{Ber}(p; \theta, b)$	NA	$\mathbf{1}\{X=\theta\}$	NA
$\mathrm{geo}(\theta)$	$\left(\frac{1}{\theta}, \mathrm{geo}(\theta), \mathrm{negbin}(2,\theta)\right)$	NA	$\frac{1}{\theta} + \frac{1-X}{1-\theta}$
$\mathrm{bin}(n,\theta)$	$(n, 1+\mathrm{bin}(n-1,\theta), \mathrm{bin}(n-1,\theta))$	NA	$\frac{X}{\theta} - \frac{n-X}{1-\theta}$
$\mathrm{Poi}(\theta)$	$(1, 1+\mathrm{Poi}(\theta), \mathrm{Poi}(\theta))$	NA	$\frac{X}{\theta} - 1$
$N(\theta, \sigma^2)$	$\left(\frac{1}{\sqrt{2\pi}\sigma}, \theta + \mathrm{Wei}\left(2, \frac{1}{2\sigma^2}\right), \theta - \mathrm{Wei}\left(2, \frac{1}{2\sigma^2}\right)\right)$	1	$\frac{X-\theta}{\sigma^2}$
$N(\mu, \theta^2)$	$\left(\frac{1}{\theta}, \mathrm{Mxw}(\mu,\theta^2), N(\mu,\theta^2)\right)$	$\frac{X-\mu}{\theta}$	$\frac{1}{\theta}\left[\left(\frac{x-\mu}{\theta}\right)^2 - 1\right]$
$U(0,\theta)$	$\left(\frac{1}{\theta}, \theta, U(0,\theta)\right)$	$\frac{X}{\theta}$	NA
$U(\theta-\gamma, \theta+\gamma)$	$\left(\frac{1}{2\gamma}, \theta+\gamma, \theta-\gamma\right)$	1	NA
$U(\mu-\theta, \mu+\theta)$	$\left(\frac{1}{\theta}, \mathrm{Ber}(0.5; \mu-\theta, \mu+\theta), U(\mu-\theta, \mu+\theta)\right)$	$\frac{X-\mu}{\theta}$	NA
$\exp(\theta)$	$\left(\frac{1}{\theta}, \mathrm{Erl}(2,\theta), \exp(\theta)\right)$	$\frac{X}{\theta}$	$\frac{1}{\theta}\left(\frac{X}{\theta} - 1\right)$
$\mathrm{Wei}(\alpha,\theta)$	$\left(\frac{\alpha}{\theta}, F^*(\alpha,\theta), \mathrm{Wei}(\alpha,\theta)\right)$	$\frac{X}{\theta}$	$\frac{1}{\theta}\left[\left(\frac{X}{\theta}\right)^\alpha - \alpha\right]$
$\mathrm{gam}(\alpha,\theta)$	$\left(\frac{\alpha}{\theta}, \mathrm{gam}(\alpha+1,\theta), \mathrm{gam}(\alpha,\theta)\right)$	$\frac{X}{\theta}$	$\frac{1}{\theta}\left(\frac{X}{\theta} - \alpha\right)$
$\mathrm{Par}(\alpha,\theta)$	$\left(\frac{\alpha}{\theta}, \mathrm{Par}(\alpha,\theta), \theta\right)$	$\frac{X}{\theta}$	NA

"NA" = not applicable, indicating that the estimator cannot be implemented for that particular distribution and corresponding parameter; "Ber" = Bernoulli; "geo" = geometric; "bin" = binomial; "negbin" = negative binomial; "Poi" = Poisson; "exp" = exponential; "Wei" = Weibull; "gam" = gamma; "Par" = Pareto; "Mxw" = Maxwell; "Erl" = Erlang; please refer to Fu [13] for the specifics of the actual parameterization of each distribution and also for the definition of the two-parameter distribution indicated by F^*.

three uniform distribution cases in Table 1 illustrate this.

- For *discrete* distributions, IPA works if the parameter occurs in the support *values*, whereas LR/SF and WD work if the parameter occurs in the support *probabilities*. Examples of this can be seen in the first five lines of Table 1.
- Higher *derivative* estimates are usually easiest to derive using LR/SF, but even then the variance of the resulting estimators may be problematic. A combination of two approaches might lead to the estimator with the best variance properties, e.g., IPA for the first derivative and then LR/SF for the derivative of that, or vice versa (see [29] for specific examples and numerical comparisons of performance in the finance setting).

2.1.2. Implementation Issues

- Implementation of the LR/SF estimator is usually the most straightforward. How ever, if the input process involves an oft-repeated (e.g., i.i.d.) random variable whose common distribution depends on the parameter of interest, then the straightforward estimator should not be used, because the *variance of the estimator will grow linearly* with the number of times the input random variable that contains the parameter is used in the simulaton, so the estimator will quickly become practically useless. This can be mitigated by batching. For example instead of averaging over 1000 customers, take an average of 100 samples of 10 customers each. Using regenerative cycles is another way around the problem (see Fu [13] and the queueing example in the next section).

- Implementation of a WD estimator is also relatively straightforward, but will generally require additional simulations, and again if the input process involves a repeated random variable, then the number of *additional simulations will grow linearly* with the number of times the random variable is used; however, this does not necessarily result in any degradation in the variance of the estimator, but it will increase the computational burden. Because a weak derivative is not unique, a WD estimator involves the choice of the WD representation for the probability density/mass function derivative; Table 1 provides one possibility for many commonly encountered distributions, but see Fu [13] and the references therein for further guidance.
- Implementation of the IPA estimator generally requires *further knowledge of the dynamics of the system*, essentially to be able to implement the $\partial L/\partial X_i$ portion of the chain rule, where X_i is the input random variable on which the parameter dependence enters. In many cases of practical interest, this is not difficult to implement, and the additional computational burden is often relatively minimal.

2.1.3. Research Notes and Further Reading

- What is often referred to as "perturbation analysis" (PA) is just one of an array of methods (see Fu [13, 15] for more details and references), the main idea being to study the effect on a performance measure of interest when a parameter value is perturbed slightly. In most of the PA methods, the parameter is never actually changed in either the analysis or in the simulation. The analysis usually entails the limit as the size of the perturbation vanishes. This is also called the "pathwise" method when applied in finance [29]. Two classic books that treat IPA are Ho and Cao [36] and Glasserman [28]; however, since they are both nearly two decades old now, the examples are queueing systems (see also Cao [7]). The book by Fu and Hu [24] focuses on IPA and smoothed perturbation analysis (SPA), which is described in Section 2.3, and includes examples in inventory control and finance. Perhaps the only *textbook* that includes an entire chapter on PA is Cassandras and Lafortune [8], covering IPA, SPA, and finite PA.
- The LR/SF method was introduced by Reiman and Weiss [51], Rubinstein [52], Glynn [31] and is treated in depth by Rubinstein and Shapiro [53], which also includes some discussion of IPA (Chapter 5 on the "push in" method); see also Melamed and Rubinstein [45]. In the finance setting, Chaper 7 in Glasserman [29] includes a discussion of both IPA (known as the pathwise method) and the LR/SF method.
- The WD method was introduced by Pflug [46, 47]. More recently, it has been put into a more general framework of "measure-valued differentiation" [35].

2.2. Validity of Interchange

Validating the exchange of limit and expectation operators required in Eqs. (10) or (11) can be equated with the concept of uniform integrability, which is a necessary and sufficient condition (see Glasserman [28, 29]); however, it is difficult to verify in practical applications and will not be discussed here. In practice, the key result used in the theoretical proofs of unbiasedness is the (Lebesgue) dominated convergence theorem:

(DOMINATED CONVERGENCE THEOREM): If a sequence of random variables $\{Y_n\}$ converges to a random variable Y with probability 1, where $|Y_n| \leq M\ \forall n$ with probability 1 and $E[M] < \infty$, then $\lim_{n\to\infty} E[Y_n] = E[Y]$.

The dominated convergence theorem is used to justify (10) and (11), by taking $\Delta\theta \to 0$ for $n \to \infty$, and taking

$$Y_{\Delta\theta} = \frac{L(\theta + \Delta\theta) - L(\theta)}{\Delta\theta},$$
$$Y_{\Delta\theta} = L(x)\frac{f(x;\theta + \Delta\theta) - f(x;\theta)}{\Delta\theta},$$

for IPA and LR/SF, respectively, with Y the corresponding derivative estimator, and then finding a bound corresponding to the random variable M. The required technical conditions for the WD approach involve the existence of a weak derivative satisfying the defining relationship of the form of Eq. (17).

For IPA, a sufficient condition is that the sample performance be continuous with respect to the parameter, which translates into requirements on the form of the performance measure and on the dynamics of the underlying stochastic system. A nice set of sufficient conditions that is often used requires Lipschitz continuity (see Glasserman [29]). In the framework of generalized semi-Markov processes (GMSPs) as a model for stochastic discrete-event simulation, Glasserman [28] provides a useful set of structural conditions for checking the underlying system.

For the LR/SF method, the bound is applied to the (joint) density (or mass) function, where the bound is for $f(x;\theta)$ with respect to the parameter θ and not its usual argument x. This is more intuitively viewed as a requirement of absolute continuity of $f(\cdot;\theta + \Delta\theta)$ with respect to $f(\cdot;\theta)$, whereby for all x, if $f(x;\theta) = 0$, then $f(x;\theta + \Delta\theta) = 0$ (Glasserman [29]).

The previous examples can be used to show in very simple cases where difficulties arise. Consider the $U(0,\theta)$ probability density function

$$f(x;\theta) = \frac{1}{\theta}\mathbf{1}\{0 < x < \theta\},$$

where the LR/SF method does not apply. In this case, f viewed as a function of θ for fixed x has a discontinuity at $\theta = x$; in terms of absolute continuity, for $\Delta\theta > 0$, at the point $x = \theta + \Delta\theta/2$, the function fails to be absolutely continuous, since $f(\theta+\Delta\theta/2;\theta) = 0$, but $f(\theta+\Delta\theta/2;\theta+\Delta\theta) = 1/(\theta + \Delta\theta)$. Similarly, for probability performance measures, the sample performance is an indicator function, which generally implies a discontinuity that leads to a biased IPA estimator.

2.3. Another Related Approach

When IPA fails, many extensions have been proposed, so "perturbation analysis" actually refers to all of these different approaches and not just IPA; see Fu [15] for more details and references. Here, we will just briefly discuss perhaps the most developed extension called smoothed perturbation analysis (SPA), which is based on conditional Monte Carlo, a well-known variance reduction technique in simulation. Specifically, the idea is to choose a random variable (or set of random variables), represented here by Z, such that

(a) $E[L|Z]$ is easily computable;
(b) $\mathrm{Var}(E[L|Z])$ is small.

In variance reduction applications, it can be shown that $\mathrm{Var}(E[L|Z]) \le \mathrm{Var}L$, so conditional Monte Carlo estimators are guaranteed to do no worse in terms of variance (this does not consider the additional computation that might be required).

Analogously, the main idea of SPA is to choose Z such that

(a) $\frac{d}{d\theta}E[L|Z]$ is easily computable;
(b) $E[\frac{d}{d\theta}E[L|Z]] = \frac{d}{d\theta}E[E[L|Z]] = \frac{dE[L]}{d\theta}$.

Thus, the main challenges in applying SPA are (i) the choice of what to condition on and (ii) how to compute (estimate) the resulting conditional expectation. Challenge (i) is somewhat analogous to the WD representation choice. Challenge (ii) may lead to many additional simulations. The book of Fu and Hu [24] treats SPA (introduced by Gong and Ho [32] and Suri and Zazanis [55]) in detail.

We can demonstrate how this works for the two simple examples by taking as the performance measure a probability, i.e., $P(L(X) \le x)$ for some fixed x. In this case, IPA fails because the sample performance is an indicator function. For simplicity of exposition, take the case where the two random variables are independent and assume θ enters into the distribution of the first random variables X_1. For $L(X) = X_1 + X_2$, noting that $P(X_1 + X_2 \le x) = E[\mathbf{1}\{X_1 + X_2 \le x\}]$, we simply condition on the other random variable:

$$E[1\{X_1 + X_2 \le x\}|X_2] = P(X_1 \le x - X_2|X_2) = F_1(x - X_2;\theta),$$

where F_1 is the cumulative distribution function for X_1. Differentiating gives the density as the SPA estimator:

$$\frac{d}{d\theta}E[\mathbf{1}\{X_1 + X_2 \le x\}|X_2] = f_1(x - X_2;\theta).$$

Similarly, for the other sample performance $L(X) = \max(X_1, X_2)$, we have

$$E[\mathbf{1}\{\max(X_1, X_2) \le x\}|X_2] = P(X_1 \le x|X_2 \le x) \mathbf{1}\{X_2 \le x\} + 0 \cdot \mathbf{1}\{X_2 > x\} = F_1(x;\theta)\mathbf{1}\{X_2 < x\},$$

and differentiating gives the SPA estimator:

$$f_1(x;\theta)\mathbf{1}\{X_2 \le x\}.$$

3. APPLICATIONS

We discuss the three main application areas of derivative estimation in simulation and provide some simple illustrative examples for each. The first application area, starting in the 1970s, was queueing systems, which dominated the research focus for almost two decades; the most important subsequent application areas over the past decade or so have been inventory control and financial engineering.

3.1. Queueing

Because queueing networks can be used to model so many different classes of large-scale systems, including communications networks, manufacturing systems, supply chains, or transportation network/systems, queueing network simulation models cover a huge range of applications. We start with estimating the sensitivity to a single random variable in the system, e.g., a service time, an interarrival time, or a routing probability. In this setting, it is easy to apply an LR/SF estimator, as long as the parameter of interest does not violate the prescriptions in the previous section, e.g., $U(0,\theta)$ would not work. The estimator would require no additional simulations, and the variance properties should be reasonably good, especially if there are a lot of other input random variables in the simulation. It is also easy to apply a WD estimator in this setting, but it will generally require two additional simulations, although for some distributions only one additional simulation is required, e.g., for $U(0,\theta)$ and $\exp(\theta)$. Again, the variance properties should be good for even a large system. Lastly, IPA could be easily implemented, requiring no additional simulation and having perhaps the best variance properties of all, but following the caveats of the previous section concerning the sample performance and the structure of the system, e.g., if the parameter is in a routing probability, then a small change could cause a customer to switch from one destination queue to another, leading to an IPA estimator that is biased. Multiclass queueing networks can present such problems for IPA. In these cases, SPA can be applied, which requires more work, in the sense that the estimators may require some additional simulation. The books [28, 36, 53] all have many examples of queueing systems for IPA and LR/SF, whereas developing corresponding WD estimators is straightforward for the distributions given in Table 1, as described in Fu [13].

Along the lines just discussed, we illustrate using the simplest queueing example—a first-come, first-served single-server queue, taking the performance measure as the expected waiting time (in queue) of the nth customer ($n > 1$), denoted by $E[W_n]$, with the parameter θ in the service time distribution, where X_i denotes the service time of the ith customer. To be concrete, we take θ as the mean of an exponential distribution, and thus also illustrate the use of the corresponding line for $\exp(\theta)$ in Table 1.

We first consider the case where the parameter θ affects only a single service time, specifically $X_{n-1} \sim \exp(\theta)$, for which we have the following estimators for $\frac{dE[W_n]}{d\theta}$:

- IPA

$$\frac{dX_{n-1}}{d\theta}\mathbf{1}\{W_n > 0\} = \frac{X_{n-1}}{\theta}\mathbf{1}\{W_n > 0\},$$

  which reflects the notion that the service time of a customer only affects the waiting times of subsequent customers in the same busy period, where the condition $W_n > 0$ implies that the nth and $(n-1)$st customer are in the same busy period.
- LR/SF

$$\frac{W_n}{\theta}\left(\frac{X_{n-1}}{\theta} - 1\right).$$

- WD

$$\frac{1}{\theta}\left[W_n\left(X_{n-1}^*\right) - W_n(X_{n-1})\right],$$

  where $X_{n-1}^* \sim \mathrm{Erl}(2,\theta)$ and only the dependence on the $(n-1)$st service time is displayed for W_n, with the other service times and interarrival times assumed the same in both expressions. Note that this would require one additional simulation.

Now consider the case where the parameter affects all of the service times, i.e., θ is the common mean for the i.i.d. exponentially distributed service times $\{X_i\}$, which leads to the following estimators for $\frac{dE[W_n]}{d\theta}$:

- IPA

$$\sum_{n^* \le i < n} \frac{dX_i}{d\theta} = \sum_{n^* \le i < n} \frac{X_i}{\theta},$$

  where n^* denotes the index of the customer that begins the busy period of customer n, so the sum is empty if customer n begins the busy period.
- LR/SF

$$\frac{W_n}{\theta}\sum_{i<n}\left(\frac{X_i}{\theta} - 1\right),$$

  which basically has variance proportional to n.
- WD

$$\frac{1}{\theta}\sum_{i<n}\left[W_n\left(X_i^*\right) - W_n(X_i)\right],$$

  where $X_i^* \sim \mathrm{Erl}(2,\theta)$ (i.i.d.) and only the dependence on the different (ith) service time is displayed for W_n, with the other service times and interarrival times assumed the same in both expressions. The number of additional simulations is basically proportional to n.

Although the straightforward LR/SF estimator is computationally prohibitive for large n, by utilizing the independence of W_n from service times before the start of customers n's busy period, which is done implicitly in the IPA estimator, the variance problem can be ameliorated. In other words, for $i < n^*$, $E[W_n X_i] = E[W_n]E[X_i] \Rightarrow E[W_n(X_i/\theta - 1)] = 0$, so a much improved LR/SF estimator with the same expected value is given by

$$\frac{W_n}{\theta}\sum_{n^* \le i < n}\left(\frac{X_i}{\theta} - 1\right).$$

This illustrates one way regeneration can be used to mitigate potential variance problems in LR/SF estimators.

3.2. Inventory

Inventory control is the second class of applications to which derivative estimation has been successfully applied. The first application to inventory systems was Fu [16] for (s, S) policies, where both IPA and SPA estimators are derived; see also Fu and Hu [22] and Bashyam and Fu [1]. The

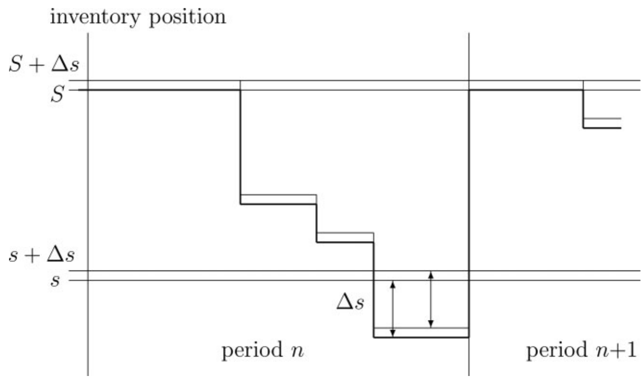

Figure 1. IPA is easy for (s, S) inventory system for the derivative with respect to s, assuming q held fixed, because the distance of the inventory position from the reorder point is the same in the perturbed (lighter) and the original (bolder) sample paths, as indicated by arrows.

Caterpillar supply chain success story described in the introduction used IPA estimators for base-stock policies along the lines of Glasserman and Tayur [30], as detailed in the *Operations Research* OR Chronicle article [50]; see also Kapuscinski and Tayur [42]. Other inventory optimization examples described in Fu et al. [19] include the following, which all use PA estimators: Fu and Healy [20, 21] consider the classical unconstrained cost minimization formulation, comparing a stochastic approximation algorithm with a "retrospective optimization" algorithm; Bashyam and Fu [2] minimize expected holding and ordering costs subject to a service level constraint that also requires simulation to estimate; Zhang and Fu [57] add a pricing decison variable to the unconstrained optimization problem. For these types of inventory systems, where the parameter of interest is one of the decision variables, the LR/SF and WD estimators cannot be applied in the usual manner, because the parameters do not occur naturally in any distribution; however, Pflug and Rubinstein [48] used a change of variables and conditioning argument to come up with unbiased estimators for the simplest periodic review (s, S) inventory system.

When simulating inventory control systems utilizing base-stock policies, IPA estimators can be easily implemented, which we will illustrate with a simple example. Because the sample performance is generally continuous with respect to the base-stock level, the resulting IPA estimators are unbiased. Difficulties arise when setup costs are included in the model, requiring a two-parameter order policy with a reorder point and order quantity or order-up-to level, because small changes in one of the parameters can cause discontinuities in the sample path, as we shall demonstrate graphically on an (s, S) inventory control example. Here, s is the reorder point and S is the order-up-to point. Instead of the parameters s and S, we will consider the parameters s and $q = S - s$, where the latter can be viewed as the (approximate) order quantity. The sample paths will show intuitively the continuity with respect to changes in s for q held constant, whereas there are discontinuities with respect to changes in q with s held constant. As a result, IPA estimators will be unbiased for performance measures with respect to s (with q held constant), whereas they would be biased for performance measures with respect to q, and SPA estimators would be required, as in [1, 16, 22].

Figure 1 shows two sample paths of the inventory position, where the vertical lines indicate the order decision points at the end of each period. Thus, if the inventory position is below the reorder point s, then an order up to S is placed. The bolder path is the original sample path—called the *nominal path*—in PA literature, whereas the lighter path—called the *perturbed path*—is the path after the perturbation by Δs with q held constant, which results in S also being increased by Δs. The key in the sample path analysis is that discontinuities in going from the nominal path to the perturbed path can only occur at the reorder points if an order decision becomes a no-order decision or vice versa. In Fig. 1, there is an order placed at the end of period n in the nominal path, and there is also an order placed at the end of period n in the perturbed path. Most importantly, the distance of the inventory position from the reorder point is unchanged before and after the perturbation, *regardless of the size (and sign) of the perturbation*, and thus the perturbed path is simply the nominal path changed by Δs.

Figure 2 shows two analogous sample paths when s is kept constant, but q is increased by $\Delta q > 0$. In the case shown, Δq is sufficiently large so that the perturbed path ends up not ordering at the end of period n, which results in a drastic divergence from the original (nominal) sample path, resulting in a discontinuity in a sample performance defined on the process.

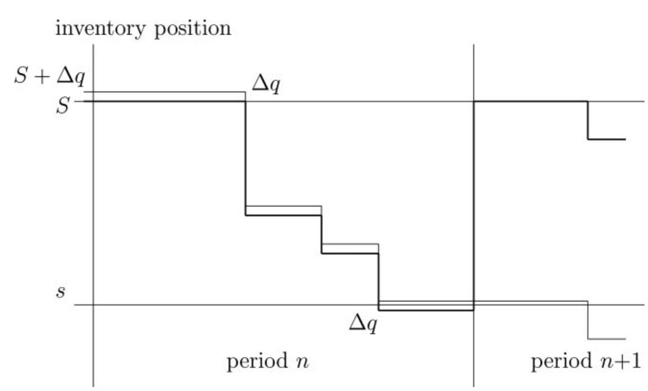

Figure 2. IPA fails for (s, S) inventory system for the derivative with respect to q (or S), assuming s held fixed, because a positive perturbation can cause an order decision in the original (bolder) path to become a no-order decision in the perturbed (lighter) path.

3.3. Finance

Probably the most flourishing research in derivative estimation these days is happening in financial engineering, where derivatives of derivatives are used for pricing and hedging, as mentioned in the introduction. Since simulation is widely applied, the techniques discussed in this article can be used to obtain estimates of the sensitivities with respect to various parameters without resimulation, i.e., estimates of the so-called "Greeks," which include for example the delta (Δ), which is the sensitivity with respect to the underlying stock price, and the vega, which is the sensitivity with respect to the volatility. This is described in detail in Chapter 7 of Glasserman [29]. In both pricing and hedging, the final problem usually comes down to an optimization problem. For hedging, one usually tries to minimize some deviation from a target payoff (or path of payoffs), whereas for pricing, the optimization problem is generally to maximize the expected payoff in a classical optimal stopping problem. A hedging example for mortgage-backed securities demonstrating orders of magnitude savings in simulation effort was reported in Chen and Fu [9]. In terms of pricing, the 1993 edition of one of the most popular textbooks for derivatives in finance proclaimed, "Monte Carlo simulation can only be used for European-style options" [40, p.363]. However, this was soundly refuted by a flurry of subsequent simulation research (see Fu et al. [27]), and one simulation approach for pricing American-style options uses gradient-based stochastic optimization, where the decision of whether or not to exercise an option is made according to a decision rule that is defined by a parameterized exercise boundary, and thus gradients with respect to the parameters are central to the method; see Fu and Hu [23], Heidergott [34], Wu and Fu [56] for examples of SPA and WD estimators.

We illustrate the ease with which these estimators can be implemented by considering a simple European call stock option, whose payoff (sample performance) is given by

$$(S_T - K)^+,$$

where S_t denotes the stock price at time t, T is the expiration time of the option, and K is the strike price. The option price is given by the expectation of the discounted payoff, where we assume discounting by a constant risk-free continuous interest rate r. The delta of this option is the sensitivity with respect to the current stock price S_0, i.e., we wish to estimate

$$\frac{d}{dS_0} E[e^{-rT}(S_T - K)^+].$$

The IPA estimator is straightforward to obtain, although the final form will depend on the form of underlying stock price model. To apply the LR/SF and WD approaches, the "parameter" of interest S_0 has to be incorporated into the underlying distribution of S_T. How this is done also depends on the underlying stock price model. Under the Black-Scholes (geometric Brownian motion process, lognormally distributed) stock price model with volatility σ, we would have the following estimators:

- IPA [23]

  $$e^{-rT}\frac{S_T}{S_0}\mathbf{1}\{S_T > K\},$$

  the form of which holds for many other stock price models in which the current stock price acts as a scale parameter for the future stock price (see Glasserman [29] for many other examples and Fu [18] for simulation of the variance-gamma asset pricing process).
- LR/SF [5]

  $$e^{-rT}(S_T - K)^+\frac{Z}{S_0\sigma\sqrt{T}},$$

  where $Z \sim N(0,1)$ is the standard normal random variate used to generate S_T from S_0 via $S_T = S_0 e^{(r-\sigma^2/2)T+\sigma\sqrt{T}Z}$.
- WD

  $$\frac{1}{S_0}[e^{-rT}(S_T - K)^+ - e^{-rT}(S_T^* - K)^+],$$

  where S_T^* is generated from a complicated density not written out explicitly here.

3.4. Other Areas

Although queueing, inventory, and finance are the most prominent application areas, there have been many others, of which the following is a small sample:

- stochastic activity networks (Rubinstein and Shapiro [53] for the LR/SF method, Bowman [4] for IPA, and Fu [17] for SPA and WD);
- preventive maintenance (Fu et al. [26]; Heidergott [33, 34]);
- statistical process control (Fu and Hu [25]);
- revenue management (Karaesmen and van Ryzin [43]), where LR/SF estimators were used in simulation optimization; and
- traffic light signal control (Howell and Fu [39]).

4. CONCLUSIONS AND FINAL THOUGHTS

The main message of this article is that sensitivity analysis and optimization in simulation can be done quite efficiently using the methods discussed here that have been developed over the past few decades. We have attempted to explain the intuition behind each of the three approaches and provide guidance as to which approach is most appropriate for a particular setting and what the actual implementation of the corresponding derivative estimator would likely entail. Simple examples and a range of application domains were used to illustrated the wide applicability of the approaches. We conclude with some final comments and thoughts on various topics of current interest.

Derivatives for quantiles, which do not fit in the problem framework of (1), were not discussed here, and have only been treated recently. Perturbation analysis estimators are discussed in Hong [37]; see also Hong and Liu [38]. These performance measures are used in service systems as an alternative or complement to probability performance measures, and are one of the risk measures used (often required by government regulation agencies) in financial risk management, under the name of Value at Risk (VaR).

In finance, the underlying models are generally continuous time, whereas the derivative estimation techniques discussed here originated from discrete-event simulation models such as those arising in queueing, where sample paths do not generally change continuously in time. The use of Malliavin calculus involves stochastic processes in continuous time (Wiener process, or Brownian motion), and application of results from this field was initiated mainly by physicists as another approach to estimating derivatives in stochastic simulation for finance (see Fournié et al. [11, 12]). Ultimately, however, implementation must be done using some sort of discretization of the continuous time processes, and Chen and Glasserman [10] explain the connection between that stream of research and some of the methods discussed here.

In many operations management models, much of the analysis goes into proving the structure of optimal policies, which often reduce to a few parameters to be optimized. In many cases, such as when Markov decision process (MDP) models are used, even after such structure is found, it is not easy to find the optimal policy using other numerical techniques, so simulation optimization is employed, and this is aided greatly by the availability of derivatives. Connections between MDPs, simulations, and derivatives are included in the recent books by Cao [6] and Powell [49]. Oftentimes, sample path analysis is used to establish such structural results, and similar analysis can be applied to derive derivative estimators using perturbation analysis. Thus, not only do we hope to reach deterministic optimizers who happen to use some simulation in their work, but also applied probabilists whose expertise in sample path analysis could be fruitfully applied to enhance derivative estimation research.

Finally, one question that naturally comes to mind is: Why haven't the commercial software vendors implemented any efficient gradient estimation into their offerings? We hope that this article has also conveyed the message that it would not be very difficult to implement many of the estimators described here in simulation software.

ACKNOWLEDGMENTS

This work has been supported in part by the National Science Foundation (NSF) under Grants DMI-0540312 and DMI-0323220, by the Air Force Office of Scientific Research (AFOSR) under Grant FA9550-07-1-0366, and by the Center for Intelligent and Networked Systems (CFINS) at Tsinghua University, China. The author thanks the Editor-in-Chief, Barry Nelson, for his guidance and for suggestions that have led to an improved exposition.

REFERENCES

[1] S. Bashyam and M.C. Fu, Application of perturbation analysis to a class of periodic review (s, S) inventory systems, Naval Res Logist 41 (1994), 47–80.

[2] S. Bashyam and M.C. Fu, Optimization of (s, S) inventory systems with random lead times and a service level constraint, Management Sci 44 (1998), S243–S256.

[3] J.M. Berry, Divided on derivatives, Washington Post, Business section (March 6), 2003.

[4] R.A. Bowman, Stochastic gradient-based time-cost tradeoffs in PERT network using simulation, Ann Oper Res 53 (1994), 533–551.

[5] M. Broadie and P. Glasserman, Estimating security price derivatives using simulation, Management Sci 42 (1996), 269–285.

[6] X.-R. Cao, Stochastic learning and optimization: A sensitivity-based approach, Springer, New York, NY, 2007.

[7] X.-R. Cao, Realization probabilities: The dynamics of queuing systems, Springer-Verlag, Boston, Massachusetts, 1994.

[8] C.G. Cassandras and S. Lafortune, Introduction to discrete event systems, Springer, New York, NY, 2008.

[9] J. Chen and M.C. Fu, Hedging beyond duration and convexity, Proc 2002 Winter Simulation Conf, IEEE, Piscataway, NJ, 2002, pp. 1593–1599.

[10] N. Chen and P. Glasserman, Malliavin Greeks without Malliavin calculus, Stochastic Process Appl 117 (2007), 1689–1723.

[11] E. Fournié, J.-M. Lasry, J. Lebuchoux, P.-L. Lions, and N. Touzi, An application of Malliavin calculus to Monte Carlo methods in finance, Finance Stochastics 3 (1999), 391–412.

[12] E. Fournié, J.-M. Lasry, J. Lebuchoux, and P.-L. Lions, Applications of Malliavin calculus to Monte Carlo methods in finance II, Finance Stochastics 5 (2001), 201–236.

[13] M.C. Fu, "Gradient estimation," Handbooks in operations research and management science: Simulation, S.G. Henderson and B.L. Nelson (Editors), Elsevier, Boston, MA, 2006, pp. 575–616.

[14] M.C. Fu, Optimization for simulation: Theory vs. practice (feature article), INFORMS J Comput 14 (2002), 192–215.

[15] M.C. Fu, "Perturbation analysis," Encyclopedia of operations research and management science, 2nd ed., S.I. Gass and C.M. Harris (Editors), Kluwer Academic, Boston, MA, 2001, pp. 608–611.

[16] M.C. Fu, Sample path derivatives for (s, S) inventory systems, Oper Res 42 (1994), 351–364.

[17] M.C. Fu, "Sensitivity analysis in Monte Carlo simulation of stochastic activity networks," Perspectives in operations research: Papers in honor of Saul Gass' 80th birthday, F.B. Alt, M.C. Fu, and B.L. Golden (Editors), Springer, New York, NY, 2006, pp. 351–366.
[18] M.C. Fu, "Variance-gamma and Monte Carlo," Advances in mathematical finance, M.C. Fu, R.A. Jarrow, J.-Y. Yen, and R.J. Elliott (Editors), Birkhäuser, Boston, MA, 2007, pp. 21–35.
[19] M.C. Fu, C.H. Chen, and L. Shi, Some topics in simulation optimization, Proc 2008 Winter Simulation Conf, IEEE, Piscataway, NJ, 2008 [in press].
[20] M.C. Fu and K.J. Healy, Simulation optimization of (s, S) inventory systems, Proc 1992 Winter Simulation Conf, IEEE, Piscataway, NJ, 1992, pp. 506–514.
[21] M.C. Fu and K.J. Healy, Techniques for simulation optimization: An experimental study on an (s, S) inventory system, IIE Trans 29 (1997), 191–199.
[22] M.C. Fu and J.Q. Hu, (s, S) inventory systems with random lead times: Harris recurrence and its implications in sensitivity analysis, Probab Eng Info Sci 8 (1994), 355–376.
[23] M.C. Fu and J.Q. Hu, Sensitivity analysis for Monte Carlo simulation of option pricing, Probab Eng Info Sci 9 (1995), 417–446.
[24] M.C. Fu and J.Q. Hu, Conditional Monte Carlo: Gradient estimation and optimization applications, Kluwer Academic Publishers, 1997.
[25] M.C. Fu and J.Q. Hu, Efficient design and sensitivity analysis of control charts using Monte Carlo simulation, Management Sci 45 (1999), 385–413.
[26] M.C. Fu, J.Q. Hu, and L. Shi, An application of perturbation analysis to replacement problems in maintenance, Proc 1993 Winter Simulation Conf, IEEE, Piscataway, NJ, 1993, pp. 329–337.
[27] M.C. Fu, S.B. Laprise, D.B. Madan, Y. Su, and R. Wu, Pricing American options: A comparison of Monte Carlo simulation approaches, J Comput Finance 4 (2001), 39–88.
[28] P. Glasserman, Gradient estimation via perturbation analysis, Kluwer Academic Publishers, Boston, Massachusetts, 1991.
[29] P. Glasserman, Monte Carlo methods in financial engineering, Springer, New York, 2004.
[30] P. Glasserman and S. Tayur, Sensitivity analysis for base-stock levels in multi-echelon production-inventory systems, Management Sci 41 (1995), 263–281.
[31] P.W. Glynn, Likelihood ratio gradient estimation: An overview, Proc 1987 Winter Simulation Conf, IEEE, Piscataway, NJ, 1987, pp. 366–375.
[32] W.B. Gong and Y.C. Ho, Smoothed perturbation analysis of discrete-event dynamic systems, IEEE Trans Automatic Control AC-32 (1987), 858–867.
[33] B. Heidergott, Optimisation of a single-component maintenance system: A smoothed perturbation analysis approach, Eur J Oper Res 119 (1999), 181–190.
[34] B. Heidergott, Option pricing via Monte Carlo simulation—A weak derivative approach, Probab Eng Info Sci 15 (2001), 335–349.
[35] B. Heidergott and F. Vázquez-Abad, Measure-valued differentiation for stochastic processes: the finite horizon case, EURANDOM, 2000.
[36] Y.C. Ho and X.R. Cao, Perturbation analysis and discrete event dynamic systems, Kluwer Academic, Boston, MA, 1991.
[37] L.J. Hong, Estimating quantile sensitivities, Oper Res 56, (2008) [in press].
[38] L.J. Hong and G. Liu, Simulating sensitivities of conditional value-at-risk, Management Sci 54, (2008) [in press].
[39] W.C. Howell and M.C. Fu, Application of perturbation analysis to traffic light signal timing, Proc 42nd IEEE Conf Decision and Control, IEEE, Piscataway, NJ, 2003, pp. 4837–4840.
[40] J.C. Hull, Options, futures, and other derivative securities, 2nd ed., Prentice Hall, Englewood Cliffs, NJ, 1993.
[41] J.C. Hull, Options, futures, and other derivative securities, 4th ed., Prentice Hall, 2000.
[42] R. Kapuscinski and S.R. Tayur, "Optimal policies and simulation based optimization for capacitated production inventory systems," Quantitative models for supply chain management, S.R. Tayur, R. Ganeshan, and M.J. Magazine (Editors), Kluwer Academic, Boston, MA, 1999, pp. 7–40.
[43] I. Karaesmen and G. van Ryzin, Overbooking with substitutable inventory classes, Oper Res 52 (2004), 83–104.
[44] P. L'Ecuyer, A unified view of the IPA, SF, and LR gradient estimation techniques, Management Sci 36 (1990), 1364–1383.
[45] B. Melamed and R.Y. Rubinstein, Modern simulation and modeling, John Wiley & Sons, Upper Saddle River, NJ, 1998.
[46] G.C. Pflug, Sampling derivatives of probabilities, Computing 42 (1989), 315–328.
[47] G.C. Pflug, Optimization of stochastic models, Kluwer Academic, Boston, MA, 1996.
[48] G.C. Pflug and R.Y. Rubinstein, Inventory processes: Quasi-regenerative property, performance evaluation and sensitivity estimation via simulation, Stochastic Models 18 (2002), 469–496.
[49] W.B. Powell, Approximate dynamic programming: Solving the curses of dimensionality, Wiley, New York, NY, 2007.
[50] U. Rao, A. Scheller-Wolf, and S. Tayur, Development of a rapid-response supply chain at Caterpillar, Oper Res 48 (2000), 189–204.
[51] M.I. Reiman and A. Weiss, Sensitivity analysis for simulations via likelihood ratios, Oper Res 37 (1989), 830–844.
[52] R.Y. Rubinstein, Sensitivity analysis of computer simulation models via the score efficient, Oper Res 37 (1989), 72–81.
[53] R.Y. Rubinstein and A. Shapiro, Discrete event systems: Sensitivity analysis and stochastic optimization by the score function method, John Wiley & Sons, New York, NY, 1993.
[54] P. Siekman, New victories in the supply-chain revolution, Fortune, (October 30) (2000).
[55] R. Suri and M.A. Zazanis, Perturbation analysis gives strongly consistent sensitivity estimates for the $M/G/1$ queue, Management Sci 34 (1988), 39–64.
[56] R. Wu and M.C. Fu, Optimal exercise policies and simulation-based valuation for American-Asian options, Oper Res 51 (2003), 52–66.
[57] H. Zhang and M.C. Fu, "Sample path derivatives for (s, S) inventory systems with price determination," The next wave in computing, optimization, and decision technologies, E.A. Wasil, B.L. Golden, and S. Raghavan (Editors), Kluwer Academic, 2005, pp. 229–246.