

UNBIASED MONTE CARLO COMPUTATION OF SMOOTH FUNCTIONS OF EXPECTATIONS VIA TAYLOR EXPANSIONS

Jose H. Blanchet

Nan Chen

Department of Ind. Eng. & Operations Research and
Department of Statistics
Columbia University
500 W 120th St.,
New York, NY, 07670, USA

Department of Syst. Eng. and Eng. Management
The Chinese University of Hong Kong
Shatin, N.T., Hong Kong.

Peter W. Glynn

Department of Management Science and Eng.
Stanford University
475 Via Ortega,
Stanford, CA.

ABSTRACT

Many Monte Carlo computations involve computing quantities that can be expressed as $g(EX)$, where g is nonlinear and smooth, and X is an easily simulatable random variable. The nonlinearity of g makes the conventional Monte Carlo estimator for such quantities biased. In this paper, we show how such quantities can be estimated without bias. However, our approach typically increases the variance. Thus, our approach is primarily of theoretical interest in the above setting. However, our method can also be applied to the computation of the inner expectation associated with $Eg(EX|Z)$, and in this setting, the application of this method can have a significant positive effect on improving the rate of convergence relative to conventional “nested schemes” for carrying out such calculations.

1 INTRODUCTION

Many applications of simulation involve the computation of “smooth functions” of expectations, specifically quantities of the form

$$\alpha = g(EX), \quad (1.1)$$

where X is an \mathbb{R}^d -valued simulatable random variable (rv) and $g: \mathbb{R}^d \rightarrow \mathbb{R}$ is “smooth” in some neighborhood \mathcal{N} of κ containing $\mu \triangleq EX$. By smooth in \mathcal{N} , we mean that g is *analytic* within \mathcal{N} (in the sense that g is infinitely differentiable at κ , and g 's Taylor expansion converges absolutely in \mathcal{N}). The reader should think of κ as an auxiliary parameter in the design of the estimator, often one can select $\kappa = 0$.

One important problem of this form is the *ratio estimation problem*, in which $d = 2$ and $g(x_1, x_2) = x_1/x_2$. This, in turn, arises naturally in a number of different application settings.

Setting 1: Suppose that we wish to compute the average productivity per hour of labor expended. Suppose that Y_i is the total number of items produced on day i , and τ_i is the total number of labor hours expended on day i . Assuming the (Y_i, τ_i) 's are independent and identically distributed (iid), the average productivity per labor hour expended over the first n days is $(Y_1 + \dots + Y_n)/(\tau_1 + \dots + \tau_n)$, which converges almost

surely (a.s.) to $EY_1/E\tau_1$ as $n \rightarrow \infty$, assuming that $E|Y_1| < \infty$. So, the long-run productivity per labor hour is given by $\alpha = EY_1/E\tau_1$.

Setting 2: Consider a nonnegative positive recurrent real-valued regenerative stochastic process, with regeneration times $0 \leq T(0) < T(1) < \dots$. Then, $t^{-1} \int_0^t X(s)ds \rightarrow \alpha = EY_1/E\tau_1$ a.s. as $t \rightarrow \infty$, where

$$Y_1 = \int_{T(0)}^{T(1)} X(s)ds,$$

$$\tau_1 = T(1) - T(0).$$

In this case, $\alpha = EY_1/E\tau_1$ is the steady-state mean of X (see, for example, p.170 of (Asmussen 2000)).

Setting 3: Suppose here that X is a nonnegative non-delayed real-valued regenerative stochastic process. Observe that the expected infinite-horizon discounted “reward”

$$\alpha = E \int_0^\infty e^{-\lambda t} X(t) dt$$

can be expressed as $\alpha = EY_1/E\tau_1$, where

$$Y_1 = \int_0^{T(1)} e^{-\lambda t} X(t) dt,$$

$$\tau_1 = 1 - \exp(-\lambda T(1));$$

see (Fox and Glynn 1989). Thus, α can be expressed as a ratio of two expectations.

Of course, many other important quantities take the form of (1.1):

variance of W:

$$\alpha = g(EW^2, EW), \text{ where } g(x_1, x_2) = x_1 - x_2^2.$$

standard deviation of W:

$$\alpha = g(EW^2, EW), \text{ where } g(x_1, x_2) = (x_1 - x_2^2)^{\frac{1}{2}}.$$

squared coefficient of variation of W:

$$\alpha = g(EW^2, EW), \text{ where } g(x_1, x_2) = (x_1 - x_2^2)/x_2^2.$$

coefficient of correlation between W and Z:

$$\alpha = g(EW^2, EZ^2, EW, EZ, EWZ), \text{ where } g(x_1, \dots, x_5) = (x_5 - x_3x_4)/((x_1 - x_3^2)(x_2 - x_4^2))^{\frac{1}{2}}.$$

Such problem structures can also arise, for example, in the context of optimization in the presence of uncertainty. Suppose, for example, that we wish to do portfolio optimization using a Markowitz formulation in which the asset returns are characterized through their expectations, variances, and covariances; see p.67-72 of (Lai and Xing 2010) for details. In this case, one can view the maximum expected return of the variance-constrained portfolio as a smooth (deterministic) function of the underlying asset means, variances, and covariances. Therefore, if one has a complicated asset return model in which simulations are used to generate potential asset return realizations so as to numerically estimate the asset means/variances/covariances, the

computation of the maximum expected portfolio return is then of the type discussed in this paper. (Of course, computing the derivatives of g is more cumbersome for this problem, but nevertheless is possible.)

Suppose that we simulate X_1, \dots, X_n from the distribution of X and form the sample mean $\bar{X}_n \triangleq n^{-1} \sum_{i=1}^n X_i$. It is well-known that the conventional Monte Carlo estimator $g(\bar{X}_n)$ for α is biased. In particular, under modest additional conditions on X and g ,

$$Eg(\bar{X}_n) = \alpha + \frac{1}{n} \sum_{i=1}^d \sum_{j=1}^d H_{ij} C_{ij} + \delta_n \tag{1.2}$$

as $n \rightarrow \infty$, where $H = (H_{ij} : 1 \leq i, j \leq d)$ is the Hessian of g evaluated at μ , $C = (C_{ij} : 1 \leq i, j \leq d)$ is the covariance matrix of X , and $\delta_n = o(1/n)$ as $n \rightarrow \infty$; see (Cramér 1999) p.346-349 for a rigorous discussion of moment expansions closely related to (1.2). (Recall that a deterministic sequence $(\delta_n : n \geq 0)$ is said to be $o(1/n)$ as $n \rightarrow \infty$ if $n\delta_n \rightarrow 0$ as $n \rightarrow \infty$.)

Because the magnitude of the bias term in (1.2) is typically unknown a priori, it is of interest to explore the degree to which unbiased estimation is possible. This paper is therefore concerned with the theoretical question of whether such smooth functions of expectations can be estimated without bias, in such a way that the canonical “square root convergence rate” of typical Monte Carlo computations is presented. We caution that the unbiased estimators proposed here will typically exhibit central limit theorems (CLT’s) displaying slower rates of convergence than those associated with the conventional estimator $g(\bar{X}_n)$. Thus, the methods discussed here are primarily of theoretical interest, although perhaps relevant in applications settings where concern about bias is of paramount interest.

In (Blanchet and Glynn 2015), we discuss an alternative approach that builds on a randomization idea explored in (Rhee and Glynn 2012). Both approaches can be used in the setting of nested simulation methods for computing quantities of the form $\alpha = Eg(E(X|Z))$ (with g smooth) to obtain “square root convergent” algorithms, thereby beating the best currently available Monte Carlo convergence rates; see (Andradóttir and Glynn 2015).

This paper is organized as follows. Section 2 develops the basic idea used here to construct unbiased estimators for $g(EX)$, based on Taylor expansions, while Section 3 provides an alternative estimator. Finally, Section 4 shows how either of the estimators constructed in Sections 2 and 3 can be used to obtain “square root convergent” estimators for $Eg(E(X|Z))$.

2 THE BASIC IDEA

We describe here the basic idea underlying our Taylor expansion-based approach to computing unbiased estimators for $\alpha = g(EX)$. Since the conventional estimator $g(\bar{X}_n)$ is unbiased when g is affine, the simplest such example where bias appears occurs in the quadratic case. However, if $d = 1$ and $g(x) = x^2$, $V \triangleq X_1 X_2$ is an unbiased estimator for $g(EX)$, provided that we take X_1 and X_2 to be two independent copies of X . Hence, if we generate iid replicates V_1, \dots, V_n of V and form $\bar{V}_n = n^{-1}(V_1 + \dots + V_n)$, \bar{V}_n will be an unbiased estimator of α that will satisfy a central limit theorem (CLT) with associated square-root convergence rate, provided that $\text{var} X < \infty$.

To carry this idea out more generally, assume (for notational convenience) that $d = 1$. Expanding g about κ , we get

$$g(EX) = \sum_{k=0}^{\infty} \frac{g^{(k)}(\kappa)}{k!} (\mu - \kappa)^k.$$

Suppose now that N is a \mathbb{Z}_+ -valued rv with associated probability mass function $(p_n : n \geq 0)$. Then,

$$g(EX) = E \frac{g^{(N)}(\kappa)}{N!} \frac{(\mu - \kappa)^N}{p_N}.$$

To construct an unbiased estimator for $g(EX)$, we start by generating N . Conditional on N , we simulate N iid replicates X_1, \dots, X_N of the rv X , and set

$$V = \frac{g^{(N)}(\kappa) \prod_i^N (X_i - \kappa)}{N! p_N}.$$

The unbiased estimator for α is then \bar{V}_n , where \bar{V}_n is the sample mean constructed from n iid replicates of V .

In order that \bar{V}_n exhibit a “square-root convergence rate” to $g(EX)$, we require that $\text{var} V < \infty$. But

$$\begin{aligned} EV^2 &= \sum_{k=0}^{\infty} \frac{g^{(k)}(\kappa)^2 (E(X_1 - \kappa)^2)^k}{k!^2 p_k} \\ &\triangleq \sum_{k=0}^{\infty} \frac{a_k}{p_k}. \end{aligned}$$

Note that the Cauchy-Schwarz inequality implies that

$$\begin{aligned} \sum_{k=0}^{\infty} \frac{a_k}{p_k} &= \sum_{k=0}^{\infty} p_k \left(\frac{\sqrt{a_k}}{p_k} \right)^2 \\ &\geq \left(\sum_{k=0}^{\infty} p_k \frac{\sqrt{a_k}}{p_k} \right)^2 \\ &= \left(\sum_{k=0}^{\infty} \sqrt{a_k} \right)^2. \end{aligned}$$

Hence, if we set

$$\tilde{p}_k = \frac{\sqrt{a_k}}{\sum_{j=0}^{\infty} \sqrt{a_j}},$$

then

$$\begin{aligned} EV^2 &\geq \left(\sum_{k=0}^{\infty} \sqrt{a_k} \right)^2 \\ &= \sum_{k=0}^{\infty} \frac{a_k}{\tilde{p}_k} \\ &\triangleq \tilde{E}V^2, \end{aligned}$$

where $\tilde{E}(\cdot)$ is the expectation operator corresponding to the probability under which N has probability mass function $(\tilde{p}_k : k \geq 0)$. In other words, the variance-minimizing choice for the distribution of N is

$$\tilde{p}_k = \frac{|g^{(k)}(\kappa)|}{k!} (E(X_1 - \kappa)^2)^{\frac{k}{2}},$$

in which case the variance of V is given by

$$\text{var} V = \left(\sum_{k=0}^{\infty} \frac{|g^{(k)}(\kappa)|}{k!} (E(X_1 - \kappa)^2)^{\frac{k}{2}} \right)^2 - \alpha^2.$$

Of course, this assumes that $(E(X_1 - \kappa)^2)^{\frac{1}{2}} \in \mathcal{N}$.

One might perhaps hope that setting $\kappa = \mu$ is an optimal variance-minimizing choice for the algorithmic parameter κ , from a theoretical perspective. (Of course, this could never be implemented at a practical level, because μ is unknown. But one could hope to choose κ close to μ , or to adaptively estimate the choice of κ based on prior samples.) However, this fails to be true, in general. To see this, suppose that $g(x) = x^2 - x$. In this case,

$$\begin{aligned} \sqrt{EV^2} &= |g(\kappa)| + |g'(\kappa)|\sqrt{E(X - \kappa)^2} + \frac{|g''(\kappa)|}{2}(E(X - \kappa)^2) \\ &= |\kappa^2 - \kappa| + |2\kappa - 1|\sqrt{\text{var}X + (\kappa - \mu)^2} + (\text{var}X + (\kappa - \mu)^2) \\ &\triangleq h(\kappa). \end{aligned}$$

So, if $\mu > 1$, $h'(\mu) = 2\mu - 1 + \sqrt{\text{var}X} > 0$, so $\kappa = \mu$ is not a minimizer of $h(\cdot)$.

The estimator described above is variance-minimizing. In most computational settings, the most sensible measure to optimize is the computational efficiency of the estimator. In this formulation, we seek the distribution for N that minimizes the asymptotic variance of the estimator available after expending c units of computer time. (In contrast to variance minimization, this requires taking into account the average time required to generate the V_i 's.) As is well-known (see (Glynn and Whitt 1992)), we should therefore choose the distribution of N to minimize

$$E(\text{time to generate } V) \cdot \text{var}V.$$

In this context, we shall assume that the expected time to generate V is proportional to N (since this corresponds to the number of copies of X that must be generated). Thus, we wish to find the distribution of N that minimizes

$$\left(\sum_{k=0}^{\infty} kp_k \right) \left(\sum_{k=0}^{\infty} \frac{a_k}{p_k} - \alpha^2 \right)$$

over all probability mass functions $(p_k : k \geq 0)$. According to Theorem 4 of (Rhee and Glynn 2015), the minimizer $(p_k^* : k \geq 0)$ is given by

$$p_k^* = \sqrt{\frac{a_k}{\alpha^2 + d^*k}},$$

where d^* is the unique root (in d) of the equation

$$\sum_{k=0}^{\infty} \sqrt{\frac{a_k}{\alpha^2 + dk}} = 1.$$

We note that guaranteeing unbiasedness may come at a significant computational cost. In particular, the asymptotic variance of even the optimal unbiased estimator associated with $(p_n^* : n \geq 0)$ will typically be larger than that associated with the conventional estimator $g(\bar{X}_n)$ in the Introduction. In order to see the types of issues that affect the variance of this estimator, just think of the case in which $g(x) = x^k$ for $k \in \mathbb{Z}_+$, and X is exponentially distributed with unit mean. Observe that the variance of the corresponding estimator $V = X_1X_2\dots X_k$ grows like 2^k , whereas the asymptotic variance in the conventional biased estimator will behave as k^2 , which arises in the corresponding Taylor expansion around $x = \mu$ in the delta method. Due to this issue we therefore take the view that this approach to construct unbiased estimators for $g(\mu)$ is primarily of theoretical interest

We close this section by briefly indicating how one could generate this to the higher-dimensional setting where $d > 1$. In this context,

$$\begin{aligned} g(\mu) &= \sum_{k=0}^{\infty} \frac{1}{k!} \sum_{i_1=1}^d \cdots \sum_{i_k=1}^d \frac{\partial^k g(\kappa)}{\partial \kappa_{i_1} \cdots \partial \kappa_{i_k}} \prod_{j=1}^k (\mu_{i_j} - \kappa_{i_j}) \\ &= E \frac{d^N}{p_N} \frac{\partial^N g(\kappa)}{\partial \kappa_{M_1} \cdots \partial \kappa_{M_N}} \prod_{j=1}^N (\mu_{M_j} - \kappa_{M_j}) \\ &= E \frac{d^N}{p_N} \frac{\partial^N g(\kappa)}{\partial \kappa_{M_1} \cdots \partial \kappa_{M_N}} \prod_{j=1}^N (X_{M_j}^j - \kappa_{M_j}) \\ &\triangleq EV, \end{aligned}$$

where N is a \mathbb{Z}_+ -valued rv with probability mass function $(p_k : k \geq 0)$, M_1, \dots, M_N are iid uniform rv's in $\{1, \dots, d\}$, and $X^1 = (X_1^1, \dots, X_d^1), \dots, X^N = (X_1^N, \dots, X_d^N)$ are N iid copies of X .

3 AN ALTERNATIVE ESTIMATOR

There is an alternative randomization that we can use to develop an unbiased estimator for $\alpha = g(EX)$, based on the Taylor expansion ideas of Section 2. Focusing, for simplicity, on $d = 1$, we write

$$\begin{aligned} g(EX) &= \sum_{k=0}^{\infty} \frac{g^{(k)}(\kappa)}{k!} (\mu - \kappa)^k \\ &= E \sum_{k=0}^N \frac{g^{(k)}(\kappa)}{k!} \frac{(\mu - \kappa)^k}{P(N \geq k)} \\ &= E \sum_{k=0}^N \frac{g^{(k)}(\kappa)}{k!} \frac{\prod_{j=1}^k (X_j - \kappa)}{P(N \geq k)} \\ &\triangleq EV, \end{aligned}$$

where N is a \mathbb{Z}_+ -valued rv, and X_1, X_2, \dots are iid copies of X independent of N . The choice $(P^*(N \geq n) : n \geq 0)$ that maximizes the computational efficiency of V is the minimizer of

$$\left(\sum_{n=0}^{\infty} \frac{\bar{\beta}_n}{\bar{F}_n} \right) \cdot \left(\sum_{n=0}^{\infty} \bar{F}_n \right), \tag{3.1}$$

subject to the constraints $\bar{F}_0 = 1$, $\bar{F}_i > 0$, and $\bar{F}_i \geq \bar{F}_{i+1}$ for all $i \geq 0$, where

$$\bar{\beta}_0 = g(\kappa) - \alpha^2$$

and

$$\bar{\beta}_k = \frac{g^{(k)}(\kappa)^2}{(k!)^2} (E(X - \kappa)^2)^k + \frac{2g^{(k)}(\kappa)}{k!} \sum_{j=k+1}^{\infty} \frac{g^{(j)}(\kappa)}{j!} (E(X - \kappa)^2)^k (\mu - \kappa)^{j-k}.$$

In deriving (3.1), we assume that the computational time required to generate V is proportional to the number N of X_i 's that must be simulated per V . We refer the reader to Section 3 of (Rhee and Glynn 2015) for a discussion of the minimizer of (3.1).

4 IMPLICATIONS FOR COMPUTING $Eg(E(X|Z))$

Problems in which computations of $\alpha \triangleq Eg(E(X|Z))$ arise (with g smooth and nonlinear) occur in several different applications settings; see (Andradóttir and Glynn 2015). We assume that almost surely conditional on $Z = z$, we can simulate iid copies $X_1(z), X_2(z), \dots$ from the distribution $P(\cdot, z) \triangleq P(X \in \cdot | Z = z)$. The natural estimator for α here involves first generating an iid sample Z_1, \dots, Z_m from the distribution of Z . Then, for each value Z_i , we can generate $X_1(Z_i), \dots, X_n(Z_i)$ from $P(\cdot, Z_i)$ and form the sample mean $\bar{X}_n(Z_i) = (1/n) \sum_{j=1}^n X_j(Z_i)$. Finally, we estimate α via $(1/m) \sum_{i=1}^m g(\bar{X}_n(Z_i))$. This class of estimators has been considered in (Andradóttir and Glynn 2015). They show that the best possible convergence rate, in the computational budget c , is achieved when m is order $c^{2/3}$ and n is of order $c^{1/3}$, in which the associated rate of convergence is of order $c^{-1/3}$. Thus, this estimator does not achieve the best possible “square root convergence rate” associated with Monte Carlo.

On the other hand, suppose that we apply one of the unbiased estimation methods discussed earlier in this paper. In this algorithm, we first draw Z_1 from the distribution of Z . Conditional on $Z = z$, we now generate an unbiased estimator $V_1(z)$ for $g(E(X|Z = z))$ using one of the two “V estimators” described in Sections 2 and 3. Hence, $W_1 \triangleq V_1(Z_1)$ is unbiased as an estimator for α . Hence, by repeating this process $n - 1$ additional iid times, thereby producing W_1, \dots, W_n , we can obtain $\bar{W}_n = (1/n) \sum_{i=1}^n W_i$. The CLT then implies that the estimator \bar{W}_n enjoys a rate of convergence of order $c^{-1/2}$ in the computational budget c (i.e. a “square root convergence rate”) whenever $\text{var}W_1 < \infty$.

Thus, the use of the unbiased estimation methods of this paper can have a profound impact on the quality of the estimators available in this setting, serving to improve the convergence rate from $c^{-1/3}$ to $c^{-1/2}$.

5 ACKNOWLEDGEMENTS

The first and third authors gratefully acknowledge the support of the National Science Foundation under Award Number DMS-1320158, as well as Rob J. Wang for his assistance with L^AT_EX.

REFERENCES

- Andradóttir, S., and P. W. Glynn. 2015. “Computing Bayesian Means Using Simulation”. *To appear in ACM Transactions on Computer Modeling and Simulation*.
- Asmussen, S. 2000. *Applied Probability and Queues (2nd. Ed.)*. Springer-Verlag, New York.
- Blanchet, J. H., and P. W. Glynn. 2015. “Unbiased Monte Carlo for Optimization and Functions of Expectations via Multi-level Randomization”. In *Proceedings of the 2015 Winter Simulation Conference*, edited by L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, To appear.
- Cramér, H. 1999. *Mathematical Methods of Statistics*. Princeton University Press, Princeton.
- Fox, B. L., and P. W. Glynn. 1989. “Simulating Discounted Costs”. *Management Science* 35 (11): 1297–1315.
- Glynn, P. W., and W. Whitt. 1992. “The Asymptotic Efficiency of Simulation Estimators”. *Operations Research* 40 (3): 505–520.
- Lai, T. L., and H. Xing. 2010. *Statistical Models and Methods for Financial Markets*. Springer Science+Business Media, New York.
- Rhee, C., and P. W. Glynn. 2012. “A New Approach to Unbiased Estimation for SDE’s”. In *Proceedings of the 2011 Winter Simulation Conference*, edited by C. Laroque, J. Himmelspach, R. Pasupathy, O. Rose, and A. M. Uhrmacher, 495–503.
- Rhee, C., and P. W. Glynn. 2015. “Unbiased Estimation with Square Root Convergence for SDE Models”. *To appear in Operations Research*.

AUTHOR BIOGRAPHIES

JOSE H. BLANCHET is an Associate Professor in the Departments of IEOR and Statistics at Columbia University. Jose holds a Ph.D. in Management Science and Engineering from Stanford University. Prior to joining Columbia, he was a faculty member in the Department of Statistics at Harvard University. Jose is a recipient of the 2009 Best Publication Award given by the INFORMS Applied Probability Society and of the 2010 Erlang Prize. He also received a PECASE award given by NSF in 2010. He worked as an analyst in Protego Financial Advisors, a leading investment bank in Mexico. He has research interests in applied probability and Monte Carlo methods. He serves in the editorial board of *Advances in Applied Probability*, *Journal of Applied Probability*, *Mathematics of Operations Research*, and *QUESTA*. His email is jose.blanchet@columbia.edu.

NAN CHEN is an Associate Professor in the Department of Systems Engineering and Engineering Management at The Chinese University of Hong Kong, which he joined in 2006. Nan received his M.Sc. degree in Probability and Statistics in 2001 at Peking University and his M.Phil. and Ph.D. degrees in 2006 from Columbia University. His research interests include financial Engineering, Monte Carlo simulation, and applied probability. His email address is nchen@se.cuhk.edu.hk.

PETER W. GLYNN is currently the Chair of the Department of Management Science and Engineering and Thomas Ford Professor of Engineering at Stanford University. He is a Fellow of INFORMS and of the Institute of Mathematical Statistics, has been co-winner of Best Publication Awards from the INFORMS Simulation Society in 1993 and 2008, and was the co-winner of the John von Neumann Theory Prize from INFORMS in 2010. In 2012, he was elected to the National Academy of Engineering. His research interests lie in stochastic simulation, queueing theory, and statistical inference for stochastic processes. His email address is glynn@stanford.edu.