Proceedings of the 2015 Winter Simulation Conference
L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, eds.

UNBIASED MONTE CARLO FOR OPTIMIZATION AND FUNCTIONS OF EXPECTATIONS VIA MULTI-LEVEL RANDOMIZATION

Jose H. Blanchet

Department of IEOR and Department of Statistics
Columbia University
500 W 120th St, 3rd Floor
New York, NY 10027, USA

Peter W. Glynn

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

ABSTRACT

We present general principles for the design and analysis of unbiased Monte Carlo estimators for quantities such as $\alpha = g(E(X))$, where $E(X)$ denotes the expectation of a (possibly multidimensional) random variable X, and $g(\cdot)$ is a given deterministic function. Our estimators possess finite work-normalized variance under mild regularity conditions such as local twice differentiability of $g(\cdot)$ and suitable growth and finite-moment assumptions. We apply our estimator to various settings of interest, such as optimal value estimation in the context of Sample Average Approximations, and unbiased steady-state simulation of regenerative processes. Other applications include unbiased estimators for particle filters and conditional expectations.

1 INTRODUCTION

In many applications of the Monte Carlo method it is of interest to compute $\alpha = g(E(X))$, where $E(X)$ denotes the expectation of a (possibly multidimensional) random variable X, and $g(\cdot)$ is a given deterministic function. In the setting of steady-state simulation, for example, $X = (X_1, X_2)$, where X_1 and $X_2 > 0$ are one dimensional random variables which can be easily simulated, and

$$\alpha = g(E(X_1), E(X_2)) = \frac{E(X_1)}{E(X_2)}. \tag{1}$$

Other examples of interest which take advantage of ratio estimators such as (1), include unbiased estimators of conditional probabilities and unbiased estimators in the context of particle filters. The paper Blanchet, Chen, and Glynn (2015) discusses many other examples which can be treated with the methodology that we develop here, such as unbiased estimators for quantities such as covariances and correlations, etc. (We shall comment on the differences between the approach that we develop here and the one studied in Blanchet, Chen, and Glynn (2015).)

Stochastic optimization is another area of great importance which provides a wealth of situations of interest in which it is desirable to study unbiased Monte Carlo estimators of functions of expectations. In particular, note that if $EX = (E(X_1), ..., E(X_d))$, then the best performance among these d systems is given by

$$\alpha = g(E(X_1), ..., E(X_d)) = \max(E(X_1), ..., E(X_d)). \tag{2}$$

The natural Monte Carlo estimator for α, namely, $\max_{i=1}^{d} \bar{X}_i(n)$, where $\bar{X}_i(n)$ is the empirical mean of n i.i.d. copies of X_i, is known to exhibit a positive bias which is quantified using various bounding techniques, see (Mak, Morton, and Wood 1999, Kleywegt, Shapiro, and Homem-de Mello 2001). Our methodology here applies directly to unbiased estimation of the optimal value in the context of discrete stochastic optimization.

More generally, as we shall see, our ideas here also allow us to obtain unbiased estimators for the natural continuous extension of (2), namely,

$$\alpha = \max_{\theta \in R^d} \{h(\theta) = E(H(\theta, X))\}, \tag{3}$$

for a suitably smooth and convex function $h(\cdot)$. In this setting, the natural estimator for α, based on Sample Average Approximation, also exhibits a positive bias which can actually be quite significant (see, for instance, the discussion following Theorem 11 in Kim, Pasupathy, and Henderson (2015)). Moreover, we believe that having an unbiased estimator is specially convenient in the setting of stochastic optimization given that i.i.d. replications can be easily implemented in parallel computations.

Our estimator builds on the multilevel Monte Carlo method introduced in Giles (2008), and the de-biasing idea studied recently in Rhee and Glynn (2011) and McLeish (2011); related ideas are discussed in Bujok, Hambly, and Reisinger (2013) and Section 9 of Giles (2015). In particular, we note that if $\bar{X}(n) = (\bar{X}_1(n), \ldots, \bar{X}_d(n))$ is the vector of empirical means of i.i.d. replications of X then, under mild assumptions, we have that

$$\sum_{n=1}^{\infty} \{Eg(\bar{X}(n+1)) - Eg(\bar{X}(n))\} = Eg(E(X)) - Eg(X),$$

the previous telescopic representation lies at the core of the multilevel Monte Carlo method in Giles (2008). The approach in Rhee and Glynn (2011) and McLeish (2011) basically involves introducing independent randomization in order to represent the infinite sum in the previous display as an expectation with respect to an integer valued random variable, N, after dividing the n-th term by $P(N = n) = p(n)$. Then, the effort in the construction of the estimator is focused on finding a suitable random variable Δ_n such that $E(\Delta_n) = Eg(\bar{X}(n+1)) - Eg(\bar{X}(n))$ and also satisfying that $\sum E\left(|\Delta_n|^2\right)/p(n) < \infty$, (implying that the tails of N should be sufficiently heavy). If the random variable Δ_n is well constructed (i.e. is suitably small), then the tails of N can be well controlled (i.e. made not too heavy) and therefore the computational effort (which is directly a function of N) will be suitably controlled as well.

A suitable Δ_n with the properties described in the previous paragraph may be constructed using the $\bar{X}(n)$'s directly, but such estimator, it turns out, typically will require information about the gradient of $g(\cdot)$. So, instead, here we construct another estimator which can be implemented without the need of knowing the gradient of $g(\cdot)$. The description of this estimator is given in Section 2.

After describing the precise form of our estimator and discussing the conditions required to verify that our estimator indeed is unbiased, we proceed, in Section 3, to discuss its variance properties and the associated computational effort to generate it. We try to keep our discussion intuitive, avoiding technicalities as much as possible. However, in our main result, which is Theorem 1, we provide precise sufficient conditions for the validity (in terms of finite variance and finite expected cost) of our estimator. A more in-depth discussion and significant extensions of our results here are given in Blanchet and Glynn (2015).

We shall return to the motivating examples in the last two sections of the paper. In particular, Section 4 contains the application of our development to ratio estimators such as (1), and unbiased estimation of optimal values such as (2) and (3) are discussed in Section 5

Now, before we move on to the main thrust of our paper we would like to discuss the differences between the estimator that we develop here and the one discussed in Blanchet, Chen, and Glynn (2015). There are several crucial differences. First, the estimator in Blanchet, Chen, and Glynn (2015) requires the function $g(\cdot)$ to be analytic in a neighborhood of $E(X)$ (that is $g(\cdot)$ must be infinitely differentiable and its Taylor expansion must converge absolutely and uniformly in a neighborhood of $E(X)$). Second, the implementation of the estimator in Blanchet, Chen, and Glynn (2015) requires being able to compute the k-th derivative of $g(\cdot)$ for any given $k \geq 1$. Third, the variance of the estimator is likely to deteriorate as the dimension increases (because the number of terms in the Taylor expansion grows dramatically as the dimension increases).

In contrast, as we shall see, the estimator that we propose here requires that $g(\cdot)$ be twice differentiable in a neighborhood of $E(X)$. Our estimator does not require gradient information and the variance of the estimator can be controlled in terms of reasonable growth conditions in $g(\cdot)$ and suitable control on the Hessian of $g(\cdot)$ in a neighborhood of $E(X)$. Finally, we point out that in Blanchet and Glynn (2015), we discuss a slight variation of our estimator here which, in addition of being unbiased, achieves the same asymptotic variance as the canonical Monte Carlo estimator for α, namely $g(\bar{X}(n))$.

2 OUR ESTIMATOR: DESCRIPTION AND INTUITIVE BIAS ANALYSIS

In order to define our estimator we need to introduce some notation. We shall use $(X(k) : k \geq 1)$ to denote an i.i.d. sequence of copies of the random variable X which is assumed to take values in R^d. We define

$$X_O(k) = X(2k-1) \quad \text{and} \quad X_E(k) = X(2k),$$

for $k \geq 1$. Note that the X_O's correspond to $X(k)$'s indexed by odd values of k, while the X_E's correspond to the $X(k)$'s for which k is even.

Now, define

$$S(m) = X(1) + ... + X(m)$$

and similarly write

$$S_O(m) = X_O(1) + ... + X_O(m),$$
$$S_E(m) = X_E(1) + ... + X_E(m).$$

The estimator, as we explain next, is straightforward to construct. We will proceed informally for the moment, but Theorem 1 below will contain precise conditions which can be verified in order to ensure the validity of the bias and variance properties of our estimator.

First, define

$$\Delta_n = g\left(\frac{S(2^{n+1})}{2^{n+1}}\right) - \frac{1}{2}\left\{g\left(\frac{S_O(2^n)}{2^n}\right) + g\left(\frac{S_E(2^n)}{2^n}\right)\right\}. \tag{4}$$

Under natural integrability conditions we have that

$$E\left(g\left(\frac{S(2^n)}{2^n}\right)\right) = E\left(g\left(\frac{S_O(2^n)}{2^n}\right)\right) = E\left(g\left(\frac{S_E(2^n)}{2^n}\right)\right). \tag{5}$$

Therefore,

$$E\Delta_n = Eg\left(\frac{S(2^{n+1})}{2^{n+1}}\right) - E\left(g\left(\frac{S(2^n)}{2^n}\right)\right).$$

Assuming mild integrability conditions implying that

$$\lim_{m\to\infty} Eg\left(\frac{S(m)}{m}\right) = g(E(X)), \tag{6}$$

we can conclude that

$$\sum_{n=0}^{\infty} E\Delta_n = g(E(X)) - Eg(X). \tag{7}$$

As mentioned in the Introduction, the construction of our estimator exploits an idea that was been developed previously in Rhee and Glynn (2011) and McLeish (2011). The idea is to introduce an integer-valued random variable N, independent of the $X(k)$'s and let $p(k) = P(N = k)$ for $k \geq 0$, where $p(k) > 0$ for each $k \geq 0$. Define the estimator

$$\bar{Z} = \frac{\Delta_N}{p(N)} + g(X), \tag{8}$$

where X is taken to be independent of the $X(k)$'s and of N. One could actually take $X = X(1)$, but assuming that X is independent of everything else is convenient for the purpose of doing variance analysis. In fact, as we shall see, a reasonable selection is $N = N_*$ geometric with success parameter $r_* = 1 - 2^{-3/2}$ (i.e. $p(n) = r_*(1 - r_*)^n$). Such a selection, as we shall see in the next subsection, controls the work-normalized variance of the estimator (i.e. the product of the variance times the computational effort is finite).

Note, in addition to (6), that if $\bar{Z}$ is integrable then,

$$E\bar{Z} = E\left(\frac{\Delta_N}{p(N)}\right) + Eg(X) = \sum_{n=0}^{\infty} E(\Delta_n) + Eg(X) = g(E(X)). \tag{9}$$

So, $\bar{Z}$ will typically yield an unbiased estimator for $\alpha = g(E(X))$.

We remark that there is a variation of the estimator given in (8) which is also unbiased and which only requires $P(N \geq n) > 0$ for $n \geq 0$ (as opposed to $p(n) > 0$ for $n \geq 0$). So such estimator is slightly more desirable than the one we consider here and in fact Rhee and Glynn (2015) use this flexibility on N in order to optimize its selection. However, such estimator requires the introduction of an auxiliary sequence in order to facility variance analysis. Here we shall consider the estimator $\bar{Z}$ only because it is slightly easier to manipulate and it does not require the introduction of the independence sequence, but we discuss the other alternative in Blanchet and Glynn (2015) and ultimately here, as we mentioned earlier, also discuss a reasonable selection N_*.

3 WORK-NORMALIZED VARIANCE AND STATEMENT OF MAIN RESULT

The fact that Z is an unbiased estimator for $g(E(X))$ is an interesting property. However, it is important to study the variance of the estimator and make sure that ultimately we still can estimate $g(E(X))$ with a rate of convergence of order $O\left(b^{-1/2}\right)$, where $b > 0$ is a given computational budget. We let T the computational cost (as measured by the number of independent copies of the $X(k)$'s) required to generate a single replication of $\bar{Z}$. In order to account for the number of i.i.d. replications of $\bar{Z}$ which can be generated with a budget b, we introduce the renewal process generated by i.i.d. copies of T. More precisely, let $A(0) = 0$ and $A(n) = T(1) ... + T(n)$, where the $T(k)$'s are i.i.d. copies of T, then set $N(b) = \max\{n \geq 0 : A(n) \leq b\}$. Our overall Monte Carlo estimator for $\alpha = g(E(X))$ is given by

$$\alpha(b) = \frac{1}{N(b)} \sum_{i=1}^{N(b)} \bar{Z}(i), \tag{10}$$

where the $\bar{Z}(i)$'s are i.i.d. copies of $\bar{Z}$. Assuming that $E(T) < \infty$ and $Var(\bar{Z}) < \infty$ we have that

$$\sqrt{b} \cdot (\alpha(b) - \alpha) \Rightarrow \sqrt{E(T) \cdot Var(\bar{Z})} \cdot N(0,1);$$

see, for instance, Glynn and Whitt (1992). This rate of convergence result emphasizes the importance of the so-called "work-normalized variance" measured by the product of the variance and the cost per replication.

Suppose that $(p(n) : n \geq 0)$ satisfies

$$E\left(\left(\frac{\Delta_N}{p(N)}\right)^2\right) = \sum_{n=0}^{\infty} \frac{E\left(\Delta_n^2\right)}{p(n)} < \infty, \tag{11}$$

and also assume that

$$E\left(g(X)^2\right) < \infty. \tag{12}$$

Then, clearly we have that $\bar{Z} = \Delta_N / p(N) + g(X)$ has finite variance.

Now, each replication of $\bar{Z}$ involves simulating $T = 2^{N+1}$ independent copies of X. So, enforcing $E(T) < \infty$ requires that we choose N so that $E\left(2^N\right) < \infty$, or, equivalently, that

$$\sum_{n=0}^{\infty} p(n)\, 2^n < \infty. \tag{13}$$

Equation (13) imposes a constrain which indicates that the $p(n)$'s must decrease to zero sufficiently fast to ensure integrability of T. On the other hand, such speed of convergence to zero must be slow enough to make sure that the series in (11) converges (because $p(n)$ appears in the denominator). Therefore, in order to understand the selection of $p(n)$ we must study the behavior of Δ_n.

Now, in order to understand the size of Δ_n, for large n, consider for simplicity the one dimensional case (i.e. $X \in R$) and let us write $\mu = E(X)$. We shall proceed informally for the moment to highlight the main ideas, but a rigorous treatment is given in the proof of Theorem 1, provided in the Appendix.

If $g(\cdot)$ is twice continuously differentiable then

$$g\left(\frac{S\left(2^{n+1}\right)}{2^{n+1}}\right) = g(\mu) + g'(\mu)\left(\frac{S\left(2^{n+1}\right)}{2^{n+1}} - \mu\right) + \frac{g''(\xi(n+1))}{2}\left(\frac{S\left(2^{n+1}\right)}{2^{n+1}} - \mu\right)^2, \tag{14}$$

for some $\xi(n+1)$ laying somewhere between μ and $S\left(2^{n+1}\right)/2^{n+1}$. Similarly, we have that for $\xi_O(n)$ and $\xi_E(n)$ laying between μ and $S_O(2^n)/2^n$ and, respectively, $S_E(2^n)/2^n$, we have that

$$g\left(\frac{S_O(2^n)}{2^n}\right) = g(\mu) + g'(\mu)\left(\frac{S_O(2^n)}{2^n} - \mu\right) + \frac{g''(\xi_O(n))}{2}\left(\frac{S_O(2^n)}{2^n} - \mu\right)^2, \tag{15}$$

$$g\left(\frac{S_E(2^n)}{2^n}\right) = g(\mu) + g'(\mu)\left(\frac{S_E(2^n)}{2^n} - \mu\right) + \frac{g''(\xi_E(n))}{2}\left(\frac{S_E(2^n)}{2^n} - \mu\right)^2. \tag{16}$$

Therefore, since

$$\frac{S\left(2^{n+1}\right)}{2^{n+1}} = \frac{1}{2}\left\{\frac{S_O(2^n)}{2^n} + \frac{S_E(2^n)}{2^n}\right\}, \tag{17}$$

we expect that under reasonable technical conditions, plugging in (14) to (17) into the definition of Δ_n in (4), we obtain that

$$\Delta_n = O_p\left(\frac{S\left(2^{n+1}\right)}{2^{n+1}} - \mu\right)^2 + O_p\left(\frac{S_O(2^n)}{2^n} - \mu\right)^2 + O_p\left(\frac{S_E(2^n)}{2^n} - \mu\right)^2.$$

Consequently, by the Central Limit Theorem, under suitable integrability conditions, we arrive at the estimate

$$2^{2n} E\Delta_n^2 = O\left(2^{2n} E\left(\left(\frac{S(2^n)}{2^n} - \mu\right)^4\right)\right) = O(1) \tag{18}$$

as $n \to \infty$.

In summary, if $g(\cdot)$ is sufficiently smooth, basically twice differentiable, and suitable integrability conditions are enforced, we typically will have that

$$E\left(\Delta_n^2\right) \le c 2^{-2n}, \tag{19}$$

for some $c > 0$. Estimate (19) is, in the end, the key estimate behind the validity of our estimator. In particular, owing to (19) we obtain that in order to enforce (11) we must have that

$$\sum_{n=0}^{\infty} \frac{2^{-2n}}{p(n)} < \infty. \tag{20}$$

We conclude, then, that it suffices to pick N geometrically distributed with success parameter $r \in (1/2, 3/4)$, so that $p(n) = (1-r)^n r$, in order to ensure both that $Var(\bar{Z}) < \infty$ and that $E(T) < \infty$.

In order to find a reasonable selection for $p(n)$ it is sensible to combine (13) and (20) with the idea of controlling the work-normalized variance. For instance, one might take (20) as a proxy for the variance of $\bar{Z}$ and then minimize

$$\sum_{n=0}^{\infty} p(n) 2^n \times \sum_{n=0}^{\infty} \frac{2^{-2n}}{p(n)} = \sum_{n=0}^{\infty} p(n) 2^n \times \sum_{n=0}^{\infty} \frac{2^{-n}}{p(n) 2^n}.$$

Applying logarithms and differentiating with respect to $p(k)$ we can easily solve for the optimal selection, $p_*(k)$, which yields N_* distributed geometric with parameter $r_* = 1 - 2^{-3/2}$; that is,

$$p^*(k) = \left(\frac{1}{2^{3/2}}\right)^k \left(1 - \frac{1}{2^{3/2}}\right).$$

We close this section with the statement of our main result, Theorem 1, and a discussion of the precise assumptions that we shall impose and which can be used to justify our previous development. The emphasis here is not on generality but on exposing the key types of regularity conditions behind our analysis. The proof of Theorem 1 is given in the Appendix of this paper (for the one dimensional case only, the multidimensional case, and extensions are given in Blanchet and Glynn (2015)).

Assumption 1 (growth): Suppose that $g : R^d \to R$ satisfies a growth condition of the form $|g(x)| \leq c(1 + \|x\|_1)$ for some $c \geq 0$, where $\|\cdot\|_1$ denotes the l_1 norm in Euclidian space.

Assumption 2 (local twice differentiability): Suppose that $g(\cdot)$ is twice continuously differentiable in a neighborhood of $\mu = E(X)$.

Assumption 3 (moment conditions): Assume that $E\|X\|_1^6 < \infty$.

Theorem 1 If N is geometrically distributed with success parameter $r \in (1/4, 1/2)$, and Assumptions 1 to 3 are in force, then $\bar{Z}$ is an unbiased estimator for $\alpha = g(E(X))$ and the overall estimator $\alpha(b)$ defined in (10) for a computational budget of size $b > 0$, is asymptotically normal with finite variance asymptotically equivalent to $b^{-1} Var(T) \cdot Var(\bar{Z})$ as $b \to \infty$.

Assumption 2 is not surprising given the formal development leading to equation (18), which, in turn, highlights the role played by the existence of the fourth moment of X. The appearance of the fourth moment gives a clue for the presence high order moments in Assumption 3. The additional two moments (from four to six) imposed in Assumption 3 arise because (18) neglects the potential contribution of the second derivative evaluated at points such as $\xi(n+1)$. One can see that some sort of growth condition must be present for $g(\cdot)$ in order to even make sense of $Eg(X)$. Such growth condition could also be imposed, for example, in terms of the second derivative itself, coupled with a suitable control on the moments, as given in Assumption 3.

4 APPLICATIONS TO STEADY-STATE SIMULATION AND OTHER SETTINGS

4.1 Unbiased Steady-state Regenerative Simulation

The context of steady-state simulation provides an important instance in which developing unbiased estimators is desirable. Recall that if $(W(n) : n \geq 0)$ is a positive recurrent regenerative process taking values on some space $\mathscr{Y}$, then for all (measurable) A, we have that the following limit holds with probability one

$$\pi(A) := \lim_{m \to \infty} \frac{1}{m} \sum_{n=0}^{m} I(W(n) \in A) = \frac{E_0\left(\sum_{n=0}^{\tau-1} I(W(n) \in A)\right)}{E_0(\tau)},$$

where the notation $E_0(\cdot)$ indicates that $W(\cdot)$ is zero-delayed under the associated probability measure $P_0(\cdot)$, and τ denotes the first regeneration time of the process $W(\cdot)$. The limiting measure $\pi(\cdot)$ is the unique stationary distribution of the process $W(\cdot)$; for additional discussion on regenerative processes see the appendix on regenerative processes in Asmussen and Glynn (2008), and also Asmussen (2000).

Most ergodic Markov chains that arise in practice are regenerative; certainly all irreducible and positive recurrent countable state-space Markov chains are regenerative.

A canonical example which is useful to keep in mind to conceptualize a regenerative process is the waiting time sequence of the single server queue. In which case, it is well known that the waiting time of the n-th customer, $W(n)$, satisfies the recursion $W(n+1) = \max(W(n)+Y(n+1), 0)$, where the $Y(n)$'s form an i.i.d. sequence of random variables with negative mean. The waiting time sequence regenerates at zero, so, if $W(0) = 0$, the waiting time sequence forms a zero-delayed regenerative process.

Let $f(\cdot)$ be a bounded measurable function and write

$$X_1 = \sum_{n=0}^{\tau-1} f(W(n)) \quad \text{and } X_2 = \tau,$$

then we can estimate the stationary expectation $E_\pi f(W)$ via the ratio

$$E_\pi f(W) = \frac{E_0(X_1)}{E_0(X_2)}. \tag{21}$$

Since $\tau \geq 1$ it follows that for $g(x_1,x_2) = x_1/x_2$ Assumptions 1 and 2 can be easily verified. Moreover, if $E\left(\tau^6\right) < \infty$; we have that Assumption 3 is satisfied and therefore Theorem 1 applies directly. For instance, in the setting of the waiting time sequence of the single server queue, Assumption 3 is satisfied if $E\left(Y(1)^6\right) < \infty$, see Asmussen (2000).

4.2 Additional Applications

In addition to steady-state simulation, ratio estimators such as (21) arise in the context of particle filters and state-dependent importance sampling for Bayesian computation, see (Del Moral 2004, Liu 2008).

In the context of Bayesian inference, one is interested in estimating expectations from some density $(\pi(y) : y \in \mathscr{S})$ of the form $\pi(y) = h(y)/\kappa$, where $h(\cdot)$ is a non-negative function with a given (computable) functional form and $\kappa > 0$ is a normalizing constant which is not computable, but is well defined (i.e. finite) and ensures that $\pi(\cdot)$ is indeed a well defined density on $\mathscr{S}$. Since $\kappa > 0$ is unknown one must resort to techniques such as Markov chain Monte Carlo or sequential importance sampling to estimate $E_\pi f(Y)$ (for any integrable function $f(\cdot)$), see for instance (Liu 2008).

Ultimately, the use of sequential importance samplers or particle filters relies on the identity

$$E_\pi f(Y) = E_q\left(\frac{h(Y)}{q(Y)} f(Y)\right) / E_q\left(\frac{h(Y)}{q(Y)}\right), \tag{22}$$

where $(q(y) : y \in \mathscr{S})$ is a density on $\mathscr{S}$ and $E_q(\cdot)$ denotes the expectation operator associated to $q(\cdot)$ (and we use $P_q(\cdot)$ for the associated probability). Of course, we must have that the likelihood ratio $\pi(Y)/q(Y)$ is well defined almost surely with respect to $P_q(\cdot)$ and

$$E_q\left(\frac{\pi(Y)}{q(Y)}\right) = 1.$$

And, thus, by using sequential importance sampling or particle filters one produces a ratio estimator (22) and therefore the application of our result in this setting is very similar to the one described in the previous subsection. The verification of Theorem 1 requires additional assumption on the selection of $q(\cdot)$, which should have heavier tails that $\pi(\cdot)$ in order to satisfy Assumption 3.

Other applications of ratio estimators and additional functions of expectations are discussed in Blanchet, Chen, and Glynn (2015). In the settings discussed there, once again Theorem 1 can be used to justify unbiased randomized multilevel Monte Carlo estimators.

5 APPLICATIONS TO STOCHASTIC OPTIMIZATION

We shall discuss some important applications of Theorem 1 in the context of stochastic optimization algorithms. We first study the case in which the feasible region is discrete, and then the case in which the feasible region is continuous. One of the standard tools in these settings is the method of Sample Average Approximation (SAA), which consists in replacing the expectations to optimize by their empirical means; for an excellent exposition of SAA see Shapiro, Dentcheva, and Ruszczyński (2009).

5.1 Unbiased Optimization via Simulation of Discrete Systems

An important problem in the area of simulation optimization is that of computing the best performance among d systems, that is,

$$\alpha = \max\left(E\left(X_1\right), ...E\left(X_d\right)\right),$$

where X_i denotes a typical random outcome obtained by running the i-th system. A standard approach in this setting consists in using SAA, which in this case consists in producing the estimator

$$\widehat{\alpha}_n = \max\left(\frac{1}{n}\sum_{i=1}^{n} X_1\left(i\right), ..., \frac{1}{n}\sum_{i=1}^{n} X_d\left(i\right)\right).$$

Nevertheless, it follows immediately that

$$E\left(\widehat{\alpha}_n\right) \geq \frac{1}{n} E \sum_{i=1}^{n} X_k\left(i\right) = EX_k, \tag{23}$$

for any k, therefore $\widehat{\alpha}_n$ has a positive bias which has been studied, for instance, in Mak, Morton, and Wood (1999).

Our goal here is to produce, using Theorem 1, an unbiased estimator for α with guaranteed $O\left(b^{-1/2}\right)$ convergence rate, where b is the available computational budget. We shall assume for the moment that there is a unique maximum among $E\left(X_i\right)$'s; the more general case is studied in Blanchet and Glynn (2015). In order to verify Assumptions 1 to 3 define

$$g\left(x_1, ..., x_d\right) = \max\left(x_1, ..., x_d\right),$$

and note that clearly $\left|g\left(x_1, ..., x_d\right)\right| \leq 1 + \left\|x\right\|_1$, so Assumption 1 is satisfied. Since we are assuming that there is a unique i such that $E\left(X_i\right) = \max\left(E\left(X_j\right) : 1 \leq j \leq d\right)$, we clearly have that Assumption 2 is satisfied; in fact, $D^2 g\left(x\right) = 0$, for x in a neighborhood of μ.

Finally, we shall impose the condition that $E\left\|X\right\|_1^6 < \infty$ and thus Assumption 3 follows immediately. We then conclude that $\bar{Z}$ provides an unbiased estimator for $\max\left(E\left(X_1\right), ..., E\left(X_d\right)\right)$. The general case, in which more than one system can achieve maximum performance is studied in Blanchet and Glynn (2015).

5.2 Unbiased Optimization via Randomized Sample Average Approximations

We shall illustrate the power of the ideas behind our estimator in continuous settings. As indicated earlier, SAA is a classical tool in the current setting, and just as we obtained via inequality (23), we also have in this context that the associated SAA objective value estimator has a positive bias (see Theorem 5.7 in Shapiro, Dentcheva, and Ruszczyński (2009) and also the discussion after Theorem 11 in Kim, Pasupathy, and Henderson (2015)).

Our goal is to explain how the ideas behind our methodology apply in order to obtain unbiased estimation of the objective value

$$\min_{\theta \in R^d} h\left(\theta\right), \tag{24}$$

where $h(\theta) := E(H(\theta, X))$ is a strictly concave function. Our discussion is informal, with the objective of exposing the power of the ideas presented in Section 3. A more complete analysis is given in Blanchet and Glynn (2015).

In order to simplify our exposition, let us assume that $h(\cdot)$ is concave, twice continuously differentiable, and that a unique optimal solution θ_* exists. So, the Hessian $D^2 h(\theta_*)$ is positive definite and if we assume that $H(\cdot, X)$ is a twice differentiable and concave then we have that

$$0 = Dh(\theta_*) = E(DH(\theta_*, X)), \tag{25}$$

where DH is the gradient of $H(\cdot)$ with respect to θ.

Now, the idea is to use sample average approximations in the following way. Let $\theta(n+1)$ solve the problem

$$\min_{\theta} \frac{1}{2^{n+1}} \sum_{j=1}^{2^{n+1}} H(\theta, X(j)),$$

and suppose that $\theta_O(n)$ and $\theta_E(n)$ solve the problems

$$\min_{\theta} \frac{1}{2^n} \sum_{j=1}^{2^n} H(\theta, X_O(j)), \quad \text{and} \quad \min_{\theta} \frac{1}{2^n} \sum_{j=1}^{2^n} H(\theta, X_E(j)),$$

respectively. Then write

$$\Delta_n = \frac{1}{2^{n+1}} \sum_{j=1}^{2^{n+1}} H(\theta(n+1), X(j)) - \frac{1}{2}\{\frac{1}{2^n} \sum_{j=1}^{2^n} H(\theta_O(n), X_O(j)) + \frac{1}{2^n} \sum_{j=1}^{2^n} H(\theta_E(n), X_E(j))\}. \tag{26}$$

Under mild regularity conditions, similar to Assumptions 1 to 3, we have that the estimator

$$\bar{Z} = \frac{\Delta_N}{p(N)} + \min_{\theta} H(\theta, X)$$

is an unbiased estimator for the value function (24). The key estimates that allows to control the work normalized variance of $\bar{Z}$ are the following. Note that under the assumption that the there is a unique optimizer θ_*, one has, under suitable regularity conditions, that

$$2^n \|\theta(n) - \theta_*\|_1^2 \Rightarrow R, \tag{27}$$

for some random variable R (see, for example, Shapiro, Dentcheva, and Ruszczyński (2009), and also Theorem 12 in Kim, Pasupathy, and Henderson (2015)). Using this observation we can proceed as we did in Section 3, using a Taylor development,

$$H(\theta(n+1), X(j)) = H(\theta_*, X(j)) + DH(\theta_*, X(j)) \cdot (\theta(n+1) - \theta_*) + O\left(\|\theta(n+1) - \theta_*\|_1^2\right),$$
$$H(\theta_O(n), X_O(j)) = H(\theta_*, X_O(j)) + DH(\theta_*, X_O(j)) \cdot (\theta_O(n) - \theta_*) + O\left(\|\theta_O(n) - \theta_*\|_1^2\right),$$
$$H(\theta_E(n), X_E(j)) = H(\theta_*, X_E(j)) + DH(\theta_*, X_E(j)) \cdot (\theta_E(n) - \theta_*) + O\left(\|\theta_E(n) - \theta_*\|_1^2\right).$$

Plugging in these estimates into (26), after simple manipulations, we obtain,

$$\Delta_n = \frac{1}{2^n} \sum_{j=1}^{2^n} DH(\theta_*, X_O(j)) \cdot (\theta_O(n) - \theta(n+1)) + \frac{1}{2^n} \sum_{j=1}^{2^n} DH(\theta_*, X_E(j)) \cdot (\theta_E(n) - \theta(n+1)) + 2^{-n} O_p(1),$$

where the quantity $O_p(1)$ denotes a random quantity whose distribution remains tight as $n \to \infty$ (due to (27)). The key observation is that (25) together with (27) implies that

$$\frac{1}{2^n}\sum_{j=1}^{2^n} DH(\theta_*, X(j)) \cdot (\theta_O(n) - \theta(n+1)) = 2^{-n} O_p(1) = \frac{1}{2^n}\sum_{j=1}^{2^n} DH(\theta_*, X(j)) \cdot (\theta_E(n) - \theta(n+1)).$$

Hence, we obtain that under suitable regularity conditions,

$$\Delta_n^2 = 2^{-2n} O_p(1),$$

which is in the end (up to integrability conditions) the key estimate given in (19) which is behind the validity of Theorem 1.

ACKNOWLEDGMENTS

Support from the National Science Foundation under Award Number DMS-1320158 is gratefully acknowledged.

6 APPENDIX: PROOF OF THEOREM 1

We shall present the proof of Theorem 1 in the one dimensional setting only; the general version is given in Blanchet and Glynn (2015). The sketch of the analysis has been given in Sections 2 and 3; here we will verify the validity of certain integrability conditions and estimates in view of Assumptions 1 to 3. First, from Assumption 1 we have that

$$|g(S_n/n)|^2 \leq c\left(1 + |S_n/n|^2\right)$$

and therefore, due to Assumption 3,

$$E|g(S_n/n)|^2 \leq c\left(1 + \frac{Var(X)}{n} + E(X)^2\right) < \infty.$$

This implies in particular that $g(S_n/n)$ is not only integrable, but, in fact, uniformly integrable and, because Assumption 2 guarantees that $g(\cdot)$ in particular is continuous in a neighborhood of the origin, we have that

$$\lim_{n\to\infty} Eg(S_n/n) = g(E(X)).$$

We therefore have verified that the equality in (5) is valid and the estimate in (6) holds. Now, we will verify the validity of the key estimate (19), namely,

$$E\Delta_n^2 \leq c2^{-2n}$$

for all $n \geq 0$. To see this, pick $\delta > 0$ small enough so that $g(\cdot)$ is twice continuously differentiable in a neighborhood of size δ around μ. Then write

$$\begin{aligned}
|\Delta_n| &= |\Delta_n| I(\max(|S_O(2^n)/2^n - \mu|, |S_E(2^n)/2^n - \mu|) > \delta/2) \\
&\quad + |\Delta_n| I(|S_O(2^n)/2^n - \mu| \leq \delta/2, |S_E(2^n)/2^n - \mu| \leq \delta/2) \\
&\leq |\Delta_n| I(|S_O(2^n)/2^n - \mu| > \delta/2) + |\Delta_n| I(|S_E(2^n)/2^n - \mu| > \delta/2) \\
&\quad + |\Delta_n| I(|S_O(2^n)/2^n - \mu| \leq \delta/2, |S_E(2^n)/2^n - \mu| \leq \delta/2).
\end{aligned}$$

Observe using Assumption 1 (linear growth) that

$$\begin{aligned}
& E\Delta_n^2 I(|S_O(2^n)/2^n - \mu| > \delta/2) \\
& \leq cE\left(\left(1 + |S_O(2^n)/2^n|^2 + |S_E(2^n)/2^n|^2\right)|S_O(2^n)/2^n - \mu|^4\right).
\end{aligned}$$

Using well known results from Von Bahr (1965) we obtain that if $EX^6 < \infty$, then

$$2^{2n} E\left(|S_O(2^n)/2^n|^2 |S_O(2^n)/2^n - \mu|^4\right) \to 2^{2n} \mu^2 Var(X)^2 E\left(N(0,1)^4\right)$$

as $n \to \infty$. A similar analysis yields that

$$E\Delta_n^2 I(\max(|S_O(2^n)/2^n - \mu|, |S_E(2^n)/2^n - \mu|) > \delta/2) = O\left(2^{-2n}\right).$$

The line of reasoning used in Section 3 can be directly applied to

$$|\Delta_n| I(|S_O(2^n)/2^n - \mu| \leq \delta/2, |S_E(2^n)/2^n - \mu| \leq \delta/2),$$

where the terms such as $\left(D^2 g\right)(\xi(n+1))$ can now be bounded by some deterministic constant c uniformly in n due to the presence of the indicator function. Consequently, we obtain that the key estimate (19) holds true and therefore the reasoning behind the conclusion that $\bar{Z}$ is unbiased reached in Section 2 is applicable and the work normalized estimates from Section 3 hold as well.

7 ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of the National Science Foundation under Award Number DMS-1320158 as well as stimulating discussions with Mike Giles.

REFERENCES

Asmussen, S. 2000. *Applied Probability and Queues (2nd. Ed.)*. Springer-Verlag, New York.

Asmussen, S., and P. Glynn. 2008. *Stochastic Simulation: Algorithms and Analysis*. New York, NY, USA: Springer-Verlag.

Blanchet, J. H., N. Chen, and P. W. Glynn. 2015. "Unbiased Monte Carlo Computation of Smooth Functions of Expectations via Taylor Expansions". 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.

Blanchet, J. H., and P. W. Glynn. 2015. "Unbiased Monte Carlo Computations for Optimization and Functions of Expectations". *In progress*.

Bujok, K., B. Hambly, and C. Reisinger. 2013. "Multilevel Simulation of Functionals of Bernoulli Random Variables with Application to Basket Credit Derivatives". *Methodology and Computing in Applied Probability*.

Del Moral, P. 2004. *Feynman-Kac Formulae Genealogical and Interacting Particle Systems with Applications*. Springer-Verlag, New York.

Giles, M. 2008. "Multilevel Monte Carlo Path Simulation". *Operations Research* 56 (3): 607–677.

Giles, M. 2015. "Multilevel Monte Carlo Methods". *Acta Numerica* 24:pp 259–328.

Glynn, P. W., and W. Whitt. 1992. "The Asymptotic Efficiency of Simulation Estimators". *Operations Research* 40 (3): 505–520.

Kim, S., R. Pasupathy, and S. Henderson. 2015. "A Guide to Sample Average Approximation". In *Handbook of Simulation Optimization*, edited by M. Fu, Volume 216. New York: International Series in Operations Research and Management Science, Springer.

Kleywegt, A., A. Shapiro, and T. Homem-de Mello. 2001. "The Sample Average Approximation Method for Stochastic Discrete Optimization". *SIAM Journal of Optimization* 12 (2): 479–502.

Liu, J. S. 2008. *Monte Carlo Strategies in Scientific Computing*. Springer, New York.

Mak, W. K., D. P. Morton, and R. K. Wood. 1999. "Monte Carlo Bounding Techniques for Determining Solution Quality in Stochastic Programs". *Operations Research Letters* 24:47–56.

McLeish, D. 2011. "A General Method for Debiasing a Monte Carlo Estimator". *Monte Carlo Methods and Applications* 17 (4): 301–315.

Rhee, C., and P. W. Glynn. 2011. "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*.

Shapiro, A., D. Dentcheva, and A. Ruszczyński. 2009. *Lectures on Stochastic Programming: Modeling and Theory*. MPS-SIAM Series on Optimization. Philadelphia, PA,.

Von Bahr, B. 1965. "On the Convergence of Moments in the Central Limit Theorem". *Ann. Math. Statist.* 36 (3): 808–818.

AUTHOR BIOGRAPHIES

JOSE H. BLANCHET 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. His email is jose.blanchet@columbia.edu

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