

Some New Perspectives on the Method of Control Variates

Peter W. Glynn and Roberto Szechtman
Email: glynn@stanford.edu, rober@stanford.edu

Department of Management Science and Engineering, Stanford University,
Stanford CA 94305, USA

Abstract. The method of control variates is one of the most widely used variance reduction techniques associated with Monte Carlo simulation. This paper studies the method of control variates from several different viewpoints, and establishes new connections between the method of control variates and: conditional Monte Carlo, antithetics, rotation sampling, stratification, and nonparametric maximum likelihood. We also develop limit theory for the method of control variates under weak assumptions on the estimator of the optimal control coefficient.

1 Introduction

Suppose that a simulation analyst wishes to compute a quantity α that can be expressed as the expectation of a real-valued random variable (rv) X , so that $\alpha = EX$. The conventional sampling-based algorithm for computing α involves simulating n independent and identically distributed (iid) copies of the rv X , denoted X_1, \dots, X_n . The corresponding estimator for α is then just the sample mean $\bar{X}_n = n^{-1} \sum_{i=1}^n X_i$.

However, in many situations, the analyst can take advantage of existing problem structure so as to create a more efficient means of computing α . One powerful approach to exploiting problem structure is the method of control variates. Specifically, suppose that there exists a random variable Y , jointly distributed with X , for which EY is known. Then, the control variate $C = Y - EY$ is guaranteed to be a “mean zero” random variable, so that $X(\lambda) = X - \lambda C$ is an estimator for α . Consequently, this creates the possibility of computing α by generating iid copies $(X_1, C_1), \dots, (X_n, C_n)$ of the pair (X, C) and estimating α via $\bar{X}_n(\lambda) = \bar{X}_n - \lambda \bar{C}_n$, where $\bar{C}_n = n^{-1} \sum_{i=1}^n C_i$ and $\lambda \in \mathbb{R}$ is an arbitrary scalar. By choosing the “control coefficient” λ judiciously, one can therefore obtain a variance reduction relative to the conventional estimator \bar{X}_n .

In particular, it is natural to choose λ so as to minimize the variance of $X(\lambda)$. The variance-minimizing choice is just

$$\lambda^* = \text{cov}(X, C) / \text{var } C.$$

Since λ^* involves moment quantities that are generally unknown to the analyst, it must be estimated somehow. Fortunately, one of the great strengths of Monte Carlo sampling-based methodology is its ability to internally estimate such problem-dependent parameters by using the sample moments of the (X_i, C_i) 's to compute the population moments of (X, C) . In particular, λ^* can be estimated via

$$\lambda_1(n) = \frac{\sum_{i=1}^n (X_i - \bar{X}_n)(C_i - \bar{C}_n)}{\sum_{i=1}^n (C_i - \bar{C}_n)^2},$$

thereby suggesting the estimator $\alpha_1(n) = \bar{X}_n - \lambda_1(n)\bar{C}_n$. This estimator for α can easily be implemented in many practical problem settings. In particular, relative to the conventional estimator \bar{X}_n , the estimator $\alpha_1(n)$ requires only that the simulation code collect the control variate outcomes C_1, \dots, C_n , compute $\lambda_1(n)$, and form $\alpha_1(n)$.

Note that this variance-reduction method is “non-invasive”, in the sense that it requires only that additional statistics be collected during the simulation run, and it does not require that the simulation analyst modify the code that is used to generate the X_i 's themselves. Thus, control variates can be implemented (for example) in the background, while some visualization involving the X_i 's proceeds in the foreground.

Because of the relative ease with which control variates can be implemented, it is perhaps the most widely applied of all variance reduction techniques. It is therefore of some interest and importance to understand this method in greater depth. This paper is intended to provide new insights into the foundations of the method of control variates, and to describe certain connections between the method of control variates and some other widely used variance reduction techniques. The following results are among the major contributions of this paper:

1. a complete discussion, in Section 2, of the method of control variates from the Hilbert space viewpoint (associated with the L^2 space of square-integrable random variables);
2. limit theory (see Theorem 1) that establishes that any consistent estimator of λ^* yields “first order” asymptotic optimality;
3. a discussion, in Section 3, of how to construct control variates in the setting of terminating simulations that involve a random number of input random variables;
4. an argument, in Section 5, that proves that control variates reduces variance “cooperatively” with the method of conditional Monte Carlo;
5. a discussion, in Sections 4, 6, and 7, of the relationships between control variates and numerical integration, antithetics, stratification, and rotation sampling;

6. an introduction to “constrained Monte Carlo”, as well as a proof that constrained Monte Carlo coincides asymptotically with control variates in the equality-constrained context.

For additional discussion of the basic ideas underlying the method of control variates (as well as a range of different applications in which the idea is relevant), see, for example, [12], [4], [11], or [6].

2 Multiple Control Variates

The method of control variates easily generalizes to the setting in which there is a collection of random variables Y_1, \dots, Y_d (jointly distributed with X) for which EY_i is known for $1 \leq i \leq d$. Set

$$\mathbf{C} = (Y_1 - EY_1, \dots, Y_d - EY_d)^T.$$

Let $(X, \mathbf{C}_1), \dots, (X_n, \mathbf{C}_n)$ be n iid replicates of the random variable (X, \mathbf{C}) . For $\boldsymbol{\lambda} \in \mathbb{R}^d$ (with $\boldsymbol{\lambda}$ encoded as a column vector), let $X(\boldsymbol{\lambda}) = X - \boldsymbol{\lambda}^T \mathbf{C}$ and put

$$\bar{X}_n(\boldsymbol{\lambda}) = \bar{X}_n - \boldsymbol{\lambda}^T \bar{\mathbf{C}}_n.$$

As in the scalar setting, the goal is to now choose $\boldsymbol{\lambda}$ so as to minimize the variance of $X(\boldsymbol{\lambda})$. Assume that $E(X^2 + \mathbf{C}^T \mathbf{C}) < \infty$. Then,

$$\text{var } X(\boldsymbol{\lambda}) = \text{var } X - 2\boldsymbol{\lambda}^T E X \mathbf{C} + \boldsymbol{\lambda}^T E \mathbf{C} \mathbf{C}^T \boldsymbol{\lambda}.$$

To compute the minimizing $\boldsymbol{\lambda}^*$, there are several equivalent approaches one can follow, each of which offers insight into the general principles underlying the method of control variates.

Approach 1. Use calculus to minimize $\text{var } X(\boldsymbol{\lambda})$ over $\boldsymbol{\lambda} \in \mathbb{R}^d$. Assuming the covariance matrix $E \mathbf{C} \mathbf{C}^T$ is non-singular, the “first-order optimality” conditions yield

$$\boldsymbol{\lambda}^* = (E \mathbf{C} \mathbf{C}^T)^{-1} E X \mathbf{C}. \quad (1)$$

It is easily shown, via the “second-order optimality conditions”, that $\boldsymbol{\lambda}^*$ is indeed the variance-minimizing choice of $\boldsymbol{\lambda}$.

Approach 2. Let L^2 be the space of all square-integrable random variables that are defined on the same probability space as X , equipped with inner product $\langle Z_1, Z_2 \rangle = E Z_1 Z_2$ and norm $\|Z\| = (E Z^2)^{1/2}$ for $Z, Z_1, Z_2 \in L^2$. Also, let \mathcal{G} be the linear subspace of all random variables $W \in L^2$ such that $W = \boldsymbol{\lambda}^T \mathbf{C}$ for some $\boldsymbol{\lambda} \in \mathbb{R}^d$. It is easily seen that \mathcal{G} is a closed linear subspace of L^2 . Observe that

$$\min_{\boldsymbol{\lambda} \in \mathbb{R}^d} \text{var } X(\boldsymbol{\lambda}) = \min_{W \in \mathcal{G}} \|X - \alpha - W\|,$$

so that minimizing variance is equivalent to finding the closest point $W \in \mathcal{G}$ to $X - \alpha$. The Hilbert space projection theorem ([9]) asserts that the minimizing $W^* = \boldsymbol{\lambda}^{*T} \mathbf{C} \in \mathcal{G}$ is characterized via the relation

$$\langle X - W^*, W \rangle = 0$$

for $W \in \mathcal{G}$. In other words,

$$EWW^* = E(X - \alpha)W = EXW$$

for $W \in \mathcal{G}$. Hence,

$$\boldsymbol{\lambda}^T ECC^T \boldsymbol{\lambda}^* = \boldsymbol{\lambda}^T EXC$$

for all $\boldsymbol{\lambda} \in \mathbb{R}^d$. This implies that $\boldsymbol{\lambda}^*$ must satisfy the equations

$$ECC^T \boldsymbol{\lambda}^* = EXC.$$

We will exploit this Hilbert space viewpoint in Section 5's discussion of the relationship of control variates with conditional Monte Carlo.

Approach 3. By applying Gram-Schmidt orthogonalization ([9]), we may assume that the d components of the control variate vector are orthogonal in our L^2 inner product. Equivalently, since ECC^T is symmetric and non-negative definite, one can express ECC^T as RDR^T , where D is a diagonal matrix with non-negative diagonal entries and $RR^T = I$. Set $\mathbf{B} = \mathbf{R}\mathbf{C}$. Then $E\mathbf{B}\mathbf{B}^T = D$, so that the components of \mathbf{B} are orthogonal in our inner product, and the variance of the i 'th component of \mathbf{B} is D_{ii} (for $1 \leq i \leq n$).

We may now choose to view \mathbf{B} as our vector of controls (rather than \mathbf{C}). This suggests finding $\boldsymbol{\beta}^* \in \mathbb{R}^d$ so as to minimize $\text{var}(X - \boldsymbol{\beta}^T \mathbf{B})$ over $\boldsymbol{\beta} \in \mathbb{R}^d$. Our discussion above establishes that

$$\beta_i^* = \frac{\langle X, B_i \rangle}{\langle B_i, B_i \rangle}$$

(provided $\langle B_i, B_i \rangle \neq 0$), so that

$$\boldsymbol{\beta}^* \mathbf{B} = \sum_{i=1}^d \frac{\langle X, B_i \rangle}{\langle B_i, B_i \rangle} B_i.$$

The right hand side of this expression is precisely what is obtained when one attempts to express X via an ‘‘orthogonal expansion’’ in terms of the rv's B_1, \dots, B_d . Thus, the optimal linear combination of control variates can be viewed in terms of an orthogonal expansion. If X is in the ‘‘span’’ of the ‘‘basis vectors’’ B_1, \dots, B_d , the variance of $X(\boldsymbol{\lambda}^*)$ is reduced to zero. In general, $X(\boldsymbol{\lambda}^*)$ has variability in proportion to that component of X that can not be explained in terms of the basis B_1, \dots, B_d (i.e. lies in the ‘‘orthogonal

complement" to B_1, \dots, B_d).

The above analysis provides several different theoretical justifications for the formula (1) for λ^* . As in the case of a scalar control, λ^* must typically be estimated from the sample data $(X_1, \mathbf{C}_1), \dots, (X_n, \mathbf{C}_n)$. The vector analog of $\lambda_1(n)$ is, of course,

$$\lambda_1(n) = \left(\sum_{i=1}^n (\mathbf{C}_i - \bar{\mathbf{C}}_n)(\mathbf{C}_i - \bar{\mathbf{C}}_n)^T \right)^{-1} \sum_{i=1}^n (X_i - \bar{X}_n)(\mathbf{C}_i - \bar{\mathbf{C}}_n).$$

It should be noted, however, that because it is known that $E\mathbf{C} = 0$, we can alternatively estimate λ^* via any of the following estimators:

$$\begin{aligned} \lambda_2(n) &= \left(\sum_{i=1}^n \mathbf{C}_i \mathbf{C}_i^T \right)^{-1} \sum_{i=1}^n (X_i - \bar{X}_n)(\mathbf{C}_i - \bar{\mathbf{C}}_n), \\ \lambda_3(n) &= \left(\sum_{i=1}^n (\mathbf{C}_i - \bar{\mathbf{C}}_n)(\mathbf{C}_i - \bar{\mathbf{C}}_n)^T \right)^{-1} \sum_{i=1}^n X_i \mathbf{C}_i, \\ \lambda_4(n) &= \left(\sum_{i=1}^n \mathbf{C}_i \mathbf{C}_i^T \right)^{-1} \sum_{i=1}^n X_i \mathbf{C}_i. \end{aligned}$$

In addition, one can apply the above idea component-wise, choosing to center some components of \mathbf{C}_i by their corresponding sample means and leaving the other components uncentered.

But other estimators for λ^* may also be available. In many applications settings, the entire joint distribution of \mathbf{C} may be known in closed form, so that the simulation analyst may know $E\mathbf{C}\mathbf{C}^T$. This suggests the estimators

$$\begin{aligned} \lambda_5(n) &= (E\mathbf{C}\mathbf{C}^T)^{-1} \sum_{i=1}^n (X_i - \bar{X}_n)(\mathbf{C}_i - \bar{\mathbf{C}}_n), \\ \lambda_6(n) &= (E\mathbf{C}\mathbf{C}^T)^{-1} \sum_{i=1}^n X_i \mathbf{C}_i \end{aligned}$$

for λ^* . Furthermore, it may be that only part of the covariance structure of \mathbf{C} is known to the analyst. Among the possibilities that exist are:

1. the joint distribution of only a subset of the control vector components is known, so that the elements of a principal submatrix of $E\mathbf{C}\mathbf{C}^T$ are known, with the other elements unknown. (By re-labeling the components of \mathbf{C} if necessary, this is equivalent to assume that some "northwest" square sub-block of $E\mathbf{C}\mathbf{C}^T$ is known, with the rest of $E\mathbf{C}\mathbf{C}^T$ unknown.)
2. Some of the control variate components are known to be orthogonal with respect to one another, so that the off-diagonal entries corresponding to

pairs of such orthogonal random variables are guaranteed to be zero. This can arise if some subset of controls forms a collection of mutually independent random variables. It may also occur in some settings in which X arises from the output of a stochastic process which is correlated with an associated martingale; the martingale differences can then be used as individual control variates, and such martingale differences are guaranteed to be orthogonal with respect to one another.

In both of the above circumstances, one may choose to substitute the known covariances for the corresponding covariance estimator, with or without “centering by the sample mean”. Thus, in the context of a vector-valued control \mathbf{C} , an enormous number of different estimators for $\boldsymbol{\lambda}^*$ suggest themselves. The following result shows that regardless of the choice of estimator for $\boldsymbol{\lambda}^*$, the resulting estimator for α achieves the maximal degree of variance reduction possible (subject only to the stipulation that the estimator for $\boldsymbol{\lambda}^*$ be consistent).

Let \Rightarrow denote weak convergence and suppose that $N(\mu, \sigma^2)$ is a normally distributed random variable with mean μ and variance σ^2 .

Theorem 1. Assume that $E\mathbf{C}\mathbf{C}^T$ is non-singular and that $E(X^2 + \mathbf{C}^T\mathbf{C}) < \infty$. Suppose that $E\mathbf{C} = \mathbf{0}$ and that $\boldsymbol{\lambda}(n) \Rightarrow \boldsymbol{\lambda}^*$ as $n \rightarrow \infty$. Set $\alpha_c(n) = \bar{X}_n - \boldsymbol{\lambda}(n)^T \bar{\mathbf{C}}_n$. Then,

$$n^{1/2}(\bar{X}_n(\boldsymbol{\lambda}^*) - \alpha_c(n)) \Rightarrow 0$$

as $n \rightarrow \infty$, so that

$$n^{1/2}(\alpha_c(n) - \alpha) \Rightarrow \sqrt{\text{var } X(\boldsymbol{\lambda}^*)}N(0, 1)$$

as $n \rightarrow \infty$. In addition,

$$\frac{1}{n} \sum_{i=1}^n (X_i - \boldsymbol{\lambda}(n)^T \mathbf{C}_i)^2 \Rightarrow EX(\boldsymbol{\lambda}^*)^2$$

as $n \rightarrow \infty$.

Proof. Note that $\bar{X}_n(\boldsymbol{\lambda}^*) - \alpha_c(n) = (\boldsymbol{\lambda}(n) - \boldsymbol{\lambda}^*)^T \bar{\mathbf{C}}_n$. Furthermore, the multivariate central limit theorem (CLT) ensures that

$$n^{-1/2} \bar{\mathbf{C}}_n \Rightarrow N(\mathbf{0}, E\mathbf{C}\mathbf{C}^T)$$

as $n \rightarrow \infty$, where $N(\mathbf{0}, E\mathbf{C}\mathbf{C}^T)$ is a d -dimensional multivariate normal random vector with mean vector $\mathbf{0}$ and covariance matrix $E\mathbf{C}\mathbf{C}^T$. Since $\boldsymbol{\lambda}(n) \Rightarrow \boldsymbol{\lambda}^*$, it follows from a “converging together” argument (see [2], p. 27) that

$$n^{1/2}(\bar{X}_n(\boldsymbol{\lambda}^*) - \alpha_c(n)) = (\boldsymbol{\lambda}(n) - \boldsymbol{\lambda}^*)^T \cdot n^{1/2} \bar{\mathbf{C}}_n \Rightarrow \mathbf{0} \cdot N(\mathbf{0}, E\mathbf{C}\mathbf{C}^T) = 0$$

as $n \rightarrow \infty$, proving the first assertion. For the second result, observe that

$$(X_i - \boldsymbol{\lambda}^{*T} \mathbf{C}_i)^2 - (X_i - \boldsymbol{\lambda}(n)^T \mathbf{C}_i)^2 = (2X_i - (\boldsymbol{\lambda}(n) + \boldsymbol{\lambda}^*)^T \mathbf{C}_i)(\mathbf{C}_i^T (\boldsymbol{\lambda}(n) - \boldsymbol{\lambda}^*)),$$

so

$$\begin{aligned} & \frac{1}{n-1} \sum_{i=1}^n X_i (\boldsymbol{\lambda}^*)^2 - \frac{1}{n-1} \sum_{i=1}^n (X_i - \boldsymbol{\lambda}(n)^T \mathbf{C}_i)^2 \\ &= \frac{2}{n-1} \sum_{i=1}^n X_i \mathbf{C}_i^T (\boldsymbol{\lambda}(n) - \boldsymbol{\lambda}^*) - (\boldsymbol{\lambda}(n) + \boldsymbol{\lambda}^*)^T \frac{1}{n-1} \sum_{i=1}^n \mathbf{C}_i \mathbf{C}_i^T (\boldsymbol{\lambda}(n) - \boldsymbol{\lambda}^*) \\ & \qquad \qquad \qquad \Rightarrow 2EX\mathbf{C}^T \cdot \mathbf{0} - 2\boldsymbol{\lambda}^{*T} \cdot E\mathbf{C}\mathbf{C}^T \cdot \mathbf{0} = \mathbf{0} \end{aligned}$$

as $n \rightarrow \infty$. Since $n^{-1} \sum_{i=1}^n X_i (\boldsymbol{\lambda}^*)^2 \rightarrow EX(\boldsymbol{\lambda}^*)^2$ a.s. as $n \rightarrow \infty$, this establishes the second assertion ■

Theorem 1 asserts that all estimators of α in which $\boldsymbol{\lambda}(n)$ is chosen to be consistent for $\boldsymbol{\lambda}^*$ are within $o_p(n^{-1/2})$ of one another (where $o_p(a_n)$ denotes a sequence of rv's $(\xi_n : n \geq 1)$ such that $a_n^{-1} \xi_n \Rightarrow 0$ as $n \rightarrow \infty$). Thus, at the level of “first-order” central limit theory, the choice of $\boldsymbol{\lambda}(n)$ is irrelevant. The specific choice of the estimator $\boldsymbol{\lambda}(n)$ can therefore only make a difference at the “second-order” level.

Of course, it should be noted that knowledge of all or part of the covariance structure of \mathbf{C} permits one to potentially use the corresponding sample covariances (centered by their known population values) as control variates. By adding such additional control variates, one can obtain further (asymptotic) variance reduction that would reflect itself at a “first order” level by reducing the magnitude of the variance associated with $X(\boldsymbol{\lambda}^*)$.

The second assertion of Theorem 1 is important in that it offers theoretical support for the variance estimation needed to construct confidence intervals for α . Specifically, let $v(n) = (n^{-1} \sum_{i=1}^n (X_i - \boldsymbol{\lambda}(n)^T \mathbf{C}_i))^2 - \alpha_c(n)^2$. Under the conditions of Theorem 1 (and assuming that $\text{var } X(\boldsymbol{\lambda}^*) > 0$), Theorem 1 permits us to assert that

$$P(\alpha \in [L_n, R_n]) \rightarrow 1 - \delta$$

as $n \rightarrow \infty$, where

$$\begin{aligned} L_n &= \alpha_c(n) - z\sqrt{v(n)/n}, \\ R_n &= \alpha_c(n) + z\sqrt{v(n)/n}, \end{aligned}$$

and z is selected so that $P(-z \leq N(0, 1) \leq z) = 1 - \delta$. Hence the interval $[L_n, R_n]$ is guaranteed to be an asymptotic $100(1 - \delta)\%$ confidence interval for α (with asymptotic half-width $n^{-1/2} \cdot z\sqrt{\text{var } X(\boldsymbol{\lambda}^*)}$).

3 Construction of Control Variates for Terminating Simulations

Suppose that generation of the rv X involves simulating an associated stochastic process over some finite time horizon. In many applications, the stochastic process is itself driven by an associated iid “noise” sequence of random input vectors. The number τ of such input vectors needed to simulate the process over the given time horizon is generally random. For example, in simulating a continuous-time Markov chain to a deterministic time t , the number of exponential random variables required is random.

To make our idea more concrete, suppose that $\mathbf{V}_1, \mathbf{V}_2, \dots, \mathbf{V}_\tau$ is the sequence of iid random (column) vectors required to generate X . Typically $E\mathbf{V}_j$ is known in closed form. It is therefore tempting to use

$$\frac{1}{\tau} \sum_{i=1}^{\tau} \mathbf{V}_i - E\mathbf{V}_i$$

as a control variate. While this centered sample mean is indeed a control (with mean zero) when τ is deterministic, it generally is not a control when τ is random, for in that case

$$E(\tau^{-1} \sum_{i=1}^{\tau} \mathbf{V}_i) \neq E\mathbf{V}_i$$

usually holds. (However, if τ is large, the difference between the left-hand side and right-hand side is generally small.)

Fortunately, there is typically an easy way to remedy the situation. In general, the rv τ will be a stopping time that is adapted to a filtration to which the \mathbf{V}_j 's are also adapted. If $E\tau < \infty$ and $E(\sum_{i=1}^{\tau} |\mathbf{V}_j|) < \infty$, Wald's first moment identity asserts ([3]) that

$$E\left(\sum_{i=1}^{\tau} \mathbf{V}_i\right) = E\mathbf{V}_i \cdot E\tau,$$

so that $\mathbf{C} = \sum_{i=1}^{\tau} (\mathbf{V}_i - E\mathbf{V}_i)$ is a control variate satisfying $E\mathbf{C} = \mathbf{0}$. Note that \mathbf{C} is a legitimate mean-zero control even if the \mathbf{V}_i 's are non-identically distributed in i (provided that they are independent in i); this follows from the generalized Wald's first moment identity.

Furthermore, if the \mathbf{V}_i 's are iid with finite second moments and $E\tau^2 < \infty$, then Wald's second moment identity ensures that

$$E\mathbf{C}\mathbf{C}^T = E\mathbf{V}_1\mathbf{V}_1^T \cdot E\tau.$$

Consequently, if the covariance structure of \mathbf{V}_1 is known, the covariance structure of \mathbf{C} is known up to the scalar $E\tau$. This offers the opportunity, as discussed in Section 2, to develop estimators of the optimal control coefficient λ^* that can take advantage of the known covariance structure of the \mathbf{V}_i 's.

For additional discussion of the subtleties that can arise in constructing control variates in the terminating simulating setting, see [10].

4 Connections to Numerical Integration

The method of control variates has a counterpart that arises in the theory of numerical integration. Consider, for example, the problem of numerically computing the integral

$$\alpha = \int_0^1 f(x)dx,$$

where $f(x) \sim x^{-1/2}$ as $x \downarrow 0$. A standard “trick” in numerical integration is to “subtract out the singularity” (namely, the singularity in f at $x = 0$), so that we write

$$\alpha = \int_0^1 [f(x) - x^{-1/2}]dx + \int_0^1 x^{-1/2} = \int_0^1 g(x)dx + 2,$$

where $g(x) = f(x) - x^{-1/2}$ is then free of the singularity at zero (and consequently easier to numerically integrate via, for example, a quadratic rule). The key to the idea is to take advantage of the fact that a closely related integrand can be analytically integrated; see [5], p. 470, for further details on this technique. Clearly, the philosophy underlying this approach is essentially identical to that associated with the method of control variates.

However, control variates and conventional numerical integration can interact in other ways, as well. For example, in certain contexts, it may be that the distribution of the rv Y (jointly distributed with X) is known, and yet no analytic closed-form for EY may be available. Typically, computing EY will then involve a one-dimensional integration with respect to the probability density of Y . This integration can be implemented numerically via an integration scheme that is not sampling-based, such as a quadrature scheme (or even an approach based on quasi-random sequences). Suppose that $\mu(n)$ is the resulting numerical approximation to EY obtained after expending a computational effort of n floating-point operations. Typically, such numerical integration schemes enjoy a rate of convergence of the form

$$\mu(n) = EY + O(n^{-1/2-p})$$

as $n \rightarrow \infty$, for some $p > 0$ (where $O(a_n)$ represents a deterministic sequence having an absolute value that can be bounded by some finite multiple of a_n).

Given a computer budget of c floating point operations, suppose that we then allocate $\lfloor c^\delta \rfloor$ operations to the numerical computation of EY and the remainder of the budget allocated to sampling of the pair (X, Y) , with $\delta = (1 + p)/(1 + 2p)$. We estimate $\alpha = EX$ via a sample mean of the rv $X - \lambda(Y - \mu(\lfloor c^\delta \rfloor))$. But $c^{1/2}(\mu(\lfloor c^\delta \rfloor) - EY) \rightarrow 0$, as $c \rightarrow \infty$ and the fraction of the budget allocated to computing $\mu(\lfloor c^\delta \rfloor)$ is converging to zero as $c \rightarrow \infty$. It follows easily that the sample mean has precisely the same “first-order” central limit behavior as that obtained by replicating the rv $X - \lambda(Y - EY)$. In other words, at the level of such a “first-order” analysis, our need to compute EY numerically via such an integration scheme has no impact whatsoever on the asymptotic convergence rate of our control variates methodology.

5 Connections to Conditional Monte Carlo

One powerful variance reduction technique that is often available in the presence of suitable problem structure is “conditional Monte Carlo”. Suppose that the probability space supporting the rv X is the triplet (Ω, \mathcal{F}, P) . It is often possible to explicitly compute $Z = E(X|\mathcal{H})$ for some sub- σ -field $\mathcal{H} \subseteq \mathcal{F}$, at the same time as one generates X . This provides us with two unbiased estimators of α , namely X and Z . As a consequence, $C = X - Z$ can be used as a control variate.

In this special setting, λ^* can be explicitly computed without a need to resort to estimation via sample covariances and sample variances.

Theorem 2. Suppose that $EX^2 < \infty$ and set $X(\lambda) = X - \lambda(X - E(X|\mathcal{H}))$. Then $\text{var } X(\lambda)$ is minimized over λ at $\lambda^* = 1$.

Proof. Suppose that $W \in L^2$ is \mathcal{H} -measurable. Then, the Cauchy-Schwarz inequality implies that XW and ZW are integrable rv’s. Furthermore, the conditional expectation Z satisfies

$$EXW = EZW$$

for all such W . In particular,

$$EXZ = EZ^2,$$

so that $EZ(X - Z) = 0$. It follows that $EXC = EX(X - Z) = E(X - Z + Z)(X - Z) = E(X - Z)^2 + EZ(X - Z) = EC^2$, so that $\lambda^* = EXC/EC^2 = 1$. ■

Thus, the best possible linear combination of X and Z is just $Z = E(X|\mathcal{H})$ itself. In other words, the conditional Monte Carlo estimator for α , based on a sample mean of iid copies of Z , is also optimal when viewed from a control

variates perspective.

Recall that for $X \in L^2$, the variance decomposition formula

$$\text{var } X = E \text{ var}(X|\mathcal{H}) + \text{var } E(X|\mathcal{H})$$

holds; see, for example, [4]. Use of conditional Monte Carlo reduces variance from $\text{var } X$ (the variance associated with the conventional Monte Carlo estimator for α , based on replicating X itself) to $\text{var } E(X|\mathcal{H})$. Thus, conditional Monte Carlo eliminates the contribution to $\text{var } X$ from $E \text{ var}(X|\mathcal{H})$.

We will now discuss how the method of control variates can frequently be used to reduce the contribution to $\text{var } X$ that is contributed by $\text{var } E(X|\mathcal{H})$. In particular, suppose that \mathcal{H} is the σ -algebra that is generated by a scalar rv Γ , and that the distribution of Γ is known. More precisely, suppose that the simulation analyst knows the entire moment sequence $(E\Gamma^k : k \geq 1)$ for the rv Γ . As a consequence, $Y_i = \Gamma^i$ has known mean for $i \geq 1$, so that for each $d \geq 1$,

$$\mathbf{C}(d) = (\Gamma - E\Gamma, \Gamma^2 - E\Gamma^2, \dots, \Gamma^d - E\Gamma^d)^T$$

is a d -dimensional control variate with mean $\mathbf{0}$. We will now analyze the behavior of $\mathbf{C}(d)$ as $d \rightarrow \infty$.

Let $W^*(d) = \boldsymbol{\lambda}^*(d)^T \mathbf{C}(d)$ be the associated linear combination of the components of $\mathbf{C}(d)$ that minimizes $\text{var}(X - \boldsymbol{\lambda}^T \mathbf{C}(d))$ over $\boldsymbol{\lambda} \in \mathbb{R}^d$.

Theorem 3. Suppose that $E|\Gamma|^i < \infty$ for $i \geq 1$, and that $EX^2 < \infty$. Then,

$$\|W^*(d) - E(X - \alpha|\Gamma)\| \rightarrow 0$$

as $d \rightarrow \infty$, and $\text{var}(X - W^*(d)) \rightarrow E \text{ var}(X|\Gamma)$ as $d \rightarrow \infty$.

Proof. Let \mathcal{G}_d be the linear subspace of all random variables $W \in L^2$ such that $W = \boldsymbol{\lambda}^T \mathbf{C}(d)$ for some $\boldsymbol{\lambda} \in \mathbb{R}^d$. Then, as discussed in Section 2,

$$\langle X - W^*(d), W \rangle = 0 \tag{2}$$

for $W \in \mathcal{G}_d$. Note that for $d_1 \leq d_2$, it follows that

$$EXW^*(d_1) = EW^*(d_1)W^*(d_2). \tag{3}$$

In addition,

$$\begin{aligned} EX^2 &= E(X - W^*(d) + W^*(d))^2 \\ &= E(X - W^*(d))^2 + 2EW^*(d)(X - W^*(d)) + EW^*(d)^2 \\ &= E(X - W^*(d))^2 + EW^*(d)^2. \end{aligned} \tag{4}$$

Hence

$$\begin{aligned}
& \|W^*(d_2) - W^*(d_1)\|^2 \\
&= EW^*(d_2)^2 + EW^*(d_1)^2 - 2EW^*(d_1)W^*(d_2) \\
&= EW^*(d_2)^2 + EW^*(d_1)^2 - 2EXW^*(d_1) \text{ (using (3))} \\
&= EW^*(d_2)^2 + EW^*(d_1)^2 - 2EXW^*(d_1) + EX^2 - EX^2 \\
&= E(W^*(d_1) - X)^2 + EW^*(d_2)^2 - EX^2 \\
&= EX^2 - EW^*(d_1)^2 + EW^*(d_2)^2 - EX^2 \text{ (using (4))} \\
&= \|W^*(d_2)\|^2 - \|W^*(d_1)\|^2,
\end{aligned}$$

so $(\|W^*(d)\|^2 : d \geq 1)$ is a non-decreasing sequence. Furthermore, it is bounded above by $\|X\|^2$.

Clearly, for $\epsilon > 0$, there exists $d = d(\epsilon)$ such that $\|W^*(d_2) - W^*(d_1)\|^2 < \epsilon$ for $d_1, d_2 \geq d$. In other words, $(W^*(d) : d \geq 1)$ is a Cauchy sequence (in L^2) of \mathcal{H} -measurable random variables. It follows that there exists a \mathcal{H} -measurable rv $W^*(\infty) \in L^2$ such that

$$\|W^*(d) - W^*(\infty)\| \rightarrow 0 \quad (5)$$

as $d \rightarrow \infty$.

Note that (2) implies that

$$EX\Gamma^i = EW^*(d)\Gamma^i$$

for $d \geq i$. But the Cauchy-Schwarz inequality and (5) together establish that

$$E(W^*(\infty) - W^*(d))\Gamma^i \rightarrow 0$$

as $d \rightarrow \infty$. We conclude that

$$EX\Gamma^i = EW^*(\infty)\Gamma^i$$

for $i \geq 1$. Consequently, the collection of all \mathcal{H} -measurable rv's $W \in L^2$ for which

$$E(\alpha + W^*(\infty))W = EXW \quad (6)$$

contains the collection $\{\Gamma^i : i \geq 0\}$. Evidently, (6) therefore holds for all \mathcal{H} -measurable $W \in L^2$ ([3]). Since $W^*(\infty)$ is itself \mathcal{H} -measurable, we conclude that $W^*(\infty) = E(X|\Gamma) - \alpha$, proving the first assertion of the theorem.

For the second assertion,

$$\begin{aligned}
& \|X - W^*(\infty)\| - \|W^*(\infty) - W^*(d)\| \\
& \leq \|X - W^*(d)\| \leq \|X - W^*(\infty)\| + \|W^*(\infty) - W^*(d)\|,
\end{aligned}$$

so $\|X - W^*(d)\|^2 \rightarrow \|X - W^*(\infty)\|^2$ as $d \rightarrow \infty$. Hence, $\text{var}(X - W^*(d)) \rightarrow \text{var}(X - W^*(\infty))$ as $d \rightarrow \infty$. But $\text{var}(X - W^*(\infty)) = E(X - Z)^2$, where $Z = E(X|I)$. Since $\langle X - Z, Z \rangle = 0$, evidently

$$\begin{aligned} EX^2 &= E(X - Z + Z)^2 \\ &= E(X - Z)^2 + EZ^2, \end{aligned}$$

so $E(X - Z)^2 = EX^2 - EZ^2 = E(E(X^2|I)) - (E(X|I))^2 = E \text{var}(X|I)$, proving the second assertion. ■

Theorem 3 proves that control variates eliminates that variability in X that is “orthogonal” to the variability eliminated by conditional Monte Carlo, in that control variates (based on powers of I) and conditional Monte Carlo (based on conditioning on I) are guaranteed to work cooperatively with one another in reducing variance. We note, in passing, that a weaker version of Theorem 3, with stronger hypotheses, appears in [10].

The principles that Theorem 3 illustrates hold much more generally. If conditional Monte Carlo is based on conditioning on a sub- σ -field \mathcal{H} , then \mathcal{H} -measurable control variates always reduce variance cooperatively. Furthermore, if a sequence of control variates is selected that ultimately spans (in L^2) the space of \mathcal{H} -measurable random variables, then the combination of conditional Monte Carlo and control variates can asymptotically reduce the variance to zero.

6 Connections to Antithetics

The method of antithetic variates is a sampling-based means of computing $\alpha = EX$ that attempts to induce an advantageous correlation between the copies of the rv X that are simulated. In particular, the idea is to generate a pair of copies of the rv X , say $X(0)$ and $X(1)$, that are correlated more advantageously than would be the case if the $X(i)$ ’s had been generated via iid sampling. Given that

$$\text{var}\left(\frac{1}{2}(X(0) + X(1))\right) = \frac{1}{2}(\text{var } X + \text{cov}(X(0), X(1))),$$

it follows that a necessary and sufficient condition for such a variance reduction (relative to iid sampling) is that $\text{cov}(X(0), X(1)) \leq 0$. See, for example, [12] for one means of inducing such a correlation. In the presence of such negative correlation, an estimator based on n iid samples of $(X(0) + X(1))/2$ has lower variance than one based on $2n$ iid samples of X .

However, a “control variates” possibility also presents itself. In particular, note that $C = X(0) - X(1)$ is a control variate having mean zero. Consequently,

$$X(\lambda) = X(0) - \lambda(X(0) - X(1))$$

is a rv, having mean α , that can be replicated to produce an estimator for $\alpha = EX$. The method of antithetics is recovered when $\lambda = 1/2$. We now prove that $\lambda = 1/2$ is indeed the universally optimal choice of the control coefficient.

Theorem 4. Suppose that $X(0) \stackrel{\mathcal{D}}{=} X(1) \stackrel{\mathcal{D}}{=} X$ (where $\stackrel{\mathcal{D}}{=}$ denotes “equality in distribution”). If $EX^2 < \infty$, then $\text{var } X(\lambda)$ is minimized over λ at $\lambda^* = 1/2$.

Proof. Note that

$$\text{var } X(\lambda) = [(1 - \lambda)^2 + \lambda^2 + 2\lambda(1 - \lambda)\rho] \text{var } X,$$

where ρ is the coefficient of correlation between $X(0)$ and $X(1)$. At a minimizer λ^* , the derivative of $\text{var } X(\lambda)$ must vanish, so that

$$2(\lambda^* - 1) + 2\lambda^* + (2 - 4\lambda^*)\rho = 0,$$

which has solution $\lambda^* = 1/2$. Since the second derivative is positive at λ^* if $\rho < 1$ (and non-negative if $\rho = 1$), it follows that $\lambda^* = 1/2$ is the minimizer. ■

Thus, regardless of the covariance (or the sign of the covariance), $\lambda^* = 1/2$ is the optimal control coefficient (so that “equal weighting” of $X(0)$ and $X(1)$ minimizes variance).

A related concept is that of “rotation sampling”; see [7] and [8]. Here, the idea is to generate $d + 1$ correlated copies of X at a time, call them $X(0), \dots, X(d)$, and to then replicate iid copies of the random vector $(X(0), \dots, X(d))^T$, from which an estimator of α can then be constructed. The term “rotation sampling” arises from the fact that the $X(i)$ ’s are generated, via inversion, by adding $i(d + 1)^{-1}$ to each of the uniform rv’s associated with $X(0)$, where addition is performed in “modulo one” arithmetic. Thus, the uniforms for the $X(i)$ ’s can be viewed as the uniforms for $X(0)$, suitably “rotated” around a circle of unit circumference. Note that we can embed $(X(i) : 0 \leq i \leq d)$ in a doubly infinite stationary sequence $\mathbf{X} = (X(i) : -\infty < i < \infty)$, in which $X(i) = X(i + d + 1)$ for all i .

As in the antithetic setting, a “control variates” possibility presents itself. In particular, let $C_i = X(0) - X(i)$ for $1 \leq i \leq d$, and put $\mathbf{C} = (C_1, \dots, C_d)$. Then

$$X(\boldsymbol{\lambda}) = X(0) - \boldsymbol{\lambda}^T \mathbf{C}$$

is an unbiased estimator for $\alpha = EX$. An important issue is the choice of $\boldsymbol{\lambda}$ that minimizes the variance of $X(\boldsymbol{\lambda})$.

Theorem 5. Suppose that $\mathbf{X} = (X(i) : -\infty < i < \infty)$ is a stationary process for which $X(i) = X(i + d + 1)$ for all i . If $EX^2(0) < \infty$, then

$\text{var } X(\boldsymbol{\lambda})$ is minimized over $\boldsymbol{\lambda}$ by $\lambda_i = 1/(d+1)$ for $1 \leq i \leq d$.

Proof. Let U be a rv, independent of \mathbf{X} , that is uniformly distributed on the integers $0, 1, \dots, d$. Note that by stationarity of \mathbf{X} ,

$$X(0) - \boldsymbol{\lambda}^T \mathbf{C} = (1 - \sum_{i=1}^d \lambda_i)X(0) + \sum_{i=1}^d \lambda_i X(i) \stackrel{\mathcal{D}}{=} (1 - \sum_{i=1}^d \lambda_i)X(U) + \sum_{i=1}^d \lambda_i X(U+i).$$

Hence, by the Cauchy-Schwarz inequality,

$$\begin{aligned} EX(\boldsymbol{\lambda})^2 &= E \left(E \left[\left((1 - \sum_{i=1}^d \lambda_i)X(U) + \sum_{i=1}^d \lambda_i X(U+i) \right)^2 \middle| \mathbf{X} \right] \right) \\ &\geq E \left(E \left[\left(1 - \sum_{i=1}^d \lambda_i \right)^2 X(U)^2 + \sum_{i=1}^d \lambda_i^2 X(U+i)^2 \middle| \mathbf{X} \right] \right). \end{aligned}$$

But

$$E[X(U+i) | \mathbf{X}] = \frac{1}{d+1} \sum_{j=0}^d X(j).$$

So,

$$E \left[\left(1 - \sum_{i=1}^d \lambda_i \right)^2 X(U)^2 + \sum_{i=1}^d \lambda_i^2 X(U+i)^2 \middle| \mathbf{X} \right] = \frac{1}{d+1} \sum_{j=0}^d X(j)^2,$$

and consequently

$$\begin{aligned} EX(\boldsymbol{\lambda})^2 &\geq E \left(\frac{1}{d+1} \sum_{j=0}^d X(j)^2 \right)^2 \\ &= EX(\boldsymbol{\lambda}^*)^2, \end{aligned}$$

proving the result. ■

This theorem proves that the choice of “uniform weighting” on the $X(i)$'s is the variance-minimizing choice for the control coefficients.

7 Connections to Stratification

Suppose that $\alpha = EX$, where X is jointly distributed with some rv Γ . Assume that the range of Γ can be partitioned into a collection of sets A_1, A_2, \dots, A_d for which $p_i = P(\Gamma \in A_i)$ is known for $1 \leq i \leq d$ and for which it is possible to generate variates from each of the d conditional distributions

$$P_i(\cdot) = P(X \in \cdot | \Gamma \in A_i).$$

For $1 \leq i \leq d$, let $\bar{X}_n(i)$ be an independently generated sample mean based on n iid replications of the rv $X(i)$. For a given total sample size n , we can then estimate α via

$$\sum_{i=1}^d p_i \bar{X}_{\lfloor c_i n \rfloor}(i),$$

where $(c_i : 1 \leq i \leq d)$ is a probability mass function that distributes the total sampling budget among the d “strata”. Proportional stratification is the special case where $c_i = p_i$ for $1 \leq i \leq d$. If $EX^2 < \infty$, then it is well known that

$$n^{1/2} \left(\sum_{i=1}^d p_i \bar{X}_{\lfloor p_i n \rfloor}(i) - \alpha \right) \Rightarrow \eta N(0, 1)$$

as $n \rightarrow \infty$, where $\eta^2 \leq \text{var } X$; see, for example, [12]. Thus, proportional stratification is always more efficient than iid sampling. Of course, one (major) disadvantage of stratification is its need to explicitly generate random variates from each of the distributions $p_i(\cdot)$, $1 \leq i \leq d$.

A more easily implemented alternative is post-stratification. Let $(X_1, \Gamma_1), \dots, (X_n, \Gamma_n)$ be n iid replications of the random vector (X, Γ) and put

$$\bar{X}(n, i) = \sum_{j=1}^n X_j I(\Gamma_j \in A_i) / \sum_{j=1}^n I(\Gamma_j \in A_i)$$

for $1 \leq i \leq d$. Note that conditional on $\sum_{j=1}^n I(\Gamma_j \in A_i) = k$, $\bar{X}(n, i) \stackrel{D}{=} \bar{X}_k(i)$. This suggests the estimator

$$\sum_{i=1}^d p_i \bar{X}(n, i).$$

If $EX^2 < \infty$, the post-stratified estimator satisfies the CLT

$$n^{1/2} \left(\sum_{i=1}^d p_i \bar{X}(n, i) - \alpha \right) \Rightarrow \eta N(0, 1)$$

as $n \rightarrow \infty$, so that the “first order” central limit theory allows us to conclude that the post-stratified estimator is as efficient as a stratified estimator with proportional stratification.

However, an obvious control variates possibility also presents itself. In particular, the assumption that the p_i 's are known opens the possibility of using

$$\mathbf{C} = (I(\Gamma \in A_1) - p_1, \dots, I(\Gamma \in A_d) - p_d)$$

as a d -dimensional control variate (having mean $\mathbf{0}$). Consequently,

$$X(\boldsymbol{\lambda}) = X - \boldsymbol{\lambda}^T \mathbf{C}$$

is a control variates estimand from which α can be estimated. A natural question here is the relationship between the variance minimizing $X(\boldsymbol{\lambda}^*)$ and the estimator obtained through post-stratification.

Theorem 6. Suppose that $EX^2 < \infty$. Then, $\text{var}(\boldsymbol{\lambda})$ is minimized over $\boldsymbol{\lambda}$ by the choice $\boldsymbol{\lambda}^*$, where $\lambda_i^* = E(X|\Gamma \in A_i)$ for $1 \leq i \leq d$. Furthermore,

$$n^{1/2} \left(\sum_{i=1}^d p_i \bar{X}(n, i) - \bar{X}_n(\boldsymbol{\lambda}^*) \right) \Rightarrow 0$$

as $n \rightarrow \infty$, where $\bar{X}_n(\boldsymbol{\lambda}^*)$ is the sample mean based on $X_1 - \boldsymbol{\lambda}^{*T} \mathbf{C}_1, \dots, X_n - \boldsymbol{\lambda}^{*T} \mathbf{C}_n$.

Proof. Put $\mu_i = E(X|\Gamma \in A_i)$ for $1 \leq i \leq d$. It is easily seen that the i 'th diagonal entry of $E\mathbf{C}\mathbf{C}^T$ is $p_i(1-p_i)$, whereas the (i, j) 'th entry (for $i \neq j$) is $-p_i p_j$. Furthermore, the i 'th entry of $E\mathbf{X}\mathbf{C}$ is $(\mu_i - \alpha)p_i$. A simple computation establishes that the vector $\boldsymbol{\mu} = (\mu_1, \dots, \mu_d)$ satisfies $E\mathbf{C}\mathbf{C}^T \boldsymbol{\mu} = E\mathbf{X}\mathbf{C}$. Consequently, our discussion of Section 2 proves that $\boldsymbol{\lambda}^* = \boldsymbol{\mu}$.

For the second assertion, note that if $\hat{p}_i(n) = n^{-1} \sum_{j=1}^n I(\Gamma_j \in A_i)$, then

$$\begin{aligned} \sum_{i=1}^d p_i \bar{X}(n, i) &= \sum_{i=1}^d \hat{p}_i(n) \bar{X}(n, i) - \sum_{i=1}^d (\hat{p}_i(n) - p_i) \bar{X}(n, i) \\ &= \bar{X}_n - \sum_{i=1}^d (\hat{p}_i(n) - p_i) \bar{X}(n, i) \\ &= \bar{X}_n - \sum_{i=1}^d (\hat{p}_i(n) - p_i) \mu_i + o_p(n^{-1/2}) \\ &= \bar{X}_n(\boldsymbol{\lambda}^*) + o_p(n^{-1/2}), \end{aligned}$$

as required. ■

Thus, we may conclude that the post-stratified estimator is actually optimal from a control variates standpoint.

It is interesting to note that the control variates ‘‘correction’’ to \bar{X}_n is just

$$\sum_{i=1}^d (\hat{p}_i(n) - p_i) \mu_i = \frac{1}{n} \sum_{j=1}^n (g^*(R_j) - \alpha),$$

where $R_j = \sum_{i=1}^d iI(\Gamma_j \in A_i)$ and $g^*(R_j) = E(X_j|R_j)$. So, the optimal control variates estimator is a sample mean of iid replicates of $X_1 - (E(X_1|R_1) - \alpha)$. Hence, post-stratification effectively involves using the control variate $E(X_1|R_1) - \alpha$, leading to a variance (per replication) of $\text{var}(X_1 - E(X_1|R_1))$. Thus, post-stratification eliminates exactly the same variance component as does the control variate scheme described in Section 5.

8 Connection to Nonparametric Maximum Likelihood Estimation

We conclude this paper with a discussion of the connection between the method of control variates and the method of nonparametric maximum likelihood estimation (MLE). Nonparametric maximum likelihood arises naturally in the setting of “constrained Monte Carlo”.

Monte Carlo computation, in the presence of constraints, has recently been investigated by [13]; see also [14]. The general statement of the problem involves the computation of $\alpha = EX$, where X is the first component of a simulatable random vector $\mathbf{X} = (X(0), X(1), \dots, X(d))$. In other words, we wish to compute $\alpha = EX$, where $X \stackrel{D}{=} X(0)$. Constrained Monte Carlo deals with the situation in which there exists a given set $B \subseteq \mathbb{R}^{d+1}$ for which it is known that $E\mathbf{X} \in B$. Computation, in the presence of such constraints, arises in many different problem settings; see [13] for details.

Pure equality constraints arise when it is known that $\mathbf{Y} \triangleq (X(1), \dots, X(2))$ satisfies the constraint $E\mathbf{Y} = \mathbf{y}$ for some known \mathbf{y} . Thus, if we set $\mathbf{C} = \mathbf{Y} - \mathbf{y}$, we can view \mathbf{C} as a d -dimensional vector of control variates having mean $\mathbf{0}$.

One approach to the problem of constrained Monte Carlo is nonparametric MLE. Specifically, suppose that we have generated n iid copies $\mathbf{X}_1, \dots, \mathbf{X}_n$ of the random vector \mathbf{X} . Nonparametric MLE involves finding a probability distribution of the form

$$\sum_{i=1}^n p_i \delta_{\mathbf{X}_i}(\cdot)$$

that explains the simulated data set best, subject to the constraint that the probability distribution must satisfy the known constraints on $E\mathbf{X}$. (Here, $\delta_{\mathbf{X}_i}(\cdot)$ is a unit point mass on \mathbf{X}_i and the collection $(p_i : 1 \leq i \leq n)$ is a probability mass function on $\{1, \dots, n\}$.) Set $\mathbf{p} = (p_1, \dots, p_n)$. The “highest likelihood” solution for the weights \mathbf{p} involves maximizing the likelihood

$$L(\mathbf{p}) = \prod_{i=1}^n p_i \tag{7}$$

subject to $p_i \geq 0$, $1 \leq i \leq n$, and

$$\sum_{i=1}^n p_i = 1 \quad (8)$$

$$\sum_{i=1}^n p_i \mathbf{X}_i \in B. \quad (9)$$

Assuming that $\mathbf{p}_n^* = (p_{1n}^*, \dots, p_{nn}^*)$ is the maximizer of the above optimization problem, the estimator for α is then

$$\begin{aligned} \tilde{\alpha}(n) &= \int_{\mathbb{R}^{d+1}} x_0 \sum_{i=1}^n p_{in}^* \delta_{\mathbf{X}_i}(d\mathbf{x}) \\ &= \sum_{i=1}^n p_{in}^* X_i(0), \end{aligned} \quad (10)$$

where $X_i(0)$ is the zero'th component of \mathbf{X}_i . As indicated above, there is a clear connection between constrained Monte Carlo and the method of control variates when the constraint takes the form $E\mathbf{Y} = \mathbf{y}$. The remainder of this section is devoted to exploring this connection in greater detail.

In the setting of such an equality constraint, the simulated data can be represented as n iid copies $(X_1, \mathbf{C}_1), \dots, (X_n, \mathbf{C}_n)$ of a random vector (X, \mathbf{C}) , in which it is known that $E\mathbf{C} = \mathbf{0}$. The optimization problem (7)-(10) then becomes:

$$\text{Maximize } L(\mathbf{p}) = \prod_{i=1}^n p_i \quad (11)$$

subject to $p_i \geq 0$, $1 \leq i \leq n$, and

$$\sum_{i=1}^n p_i = 1 \quad (12)$$

$$\sum_{i=1}^n p_i \mathbf{C}_i = \mathbf{0}. \quad (13)$$

If \mathbf{p}^* is the maximizer of (11)-(13), the ‘‘constrained Monte Carlo’’ estimator for α then takes the form

$$\tilde{\alpha}(n) = \sum_{i=1}^n p_{in}^* X_i. \quad (14)$$

Of course, maximizing (11) is equivalent to maximizing the log-likelihood, namely

$$\mathcal{L}(\mathbf{p}) = \sum_{i=1}^n \log p_i.$$

Note that the feasible region described by (12), (13), and the non-negativity constraints is non-empty for n sufficiently large, since $\mathbf{0}$ must lie in the convex hull of the \mathbf{C}_i 's for n large enough (on account of the fact that $E\mathbf{C} = \mathbf{0}$). Furthermore, the feasible region is clearly convex. Because \mathcal{L} is strictly concave, it follows that the above optimization problem can have at most one global maximizer, namely \mathbf{p}_n^* .

The standard means of solving such an optimization problem is to temporarily ignore the non-negativity constraints and to introduce Kuhn-Tucker multipliers $\gamma_n \in \mathbb{R}$ for (12) and $\boldsymbol{\nu}_n \in \mathbb{R}^d$ for (13). As usual, we choose to encode \mathbf{p}_n^* and $\boldsymbol{\nu}_n$ as column vectors. At the maximizer \mathbf{p}_n^* , the multipliers should satisfy

$$(1/p_{in}^*)\mathbf{C}_i - \boldsymbol{\nu}_n^T \mathbf{C}_i - \gamma_n = 0 \quad (15)$$

for $1 \leq i \leq n$. Multiplying through (15) by p_{in}^* and summing over i , we get

$$n - \boldsymbol{\nu}_n^T \sum_{i=1}^n p_{in}^* \mathbf{C}_i - \gamma_n \sum_{i=1}^n p_{in}^* = 0.$$

Hence, $\gamma_n = n$ and (15) implies that

$$p_{in}^* = (n + \boldsymbol{\nu}_n^T \mathbf{C}_i)^{-1} \quad (16)$$

for $1 \leq i \leq n$. Multiplying through (16) by \mathbf{C}_i and summing over i , we arrive at the equation

$$\sum_{i=1}^n \mathbf{C}_i (n + \boldsymbol{\nu}_n^T \mathbf{C}_i)^{-1} = \mathbf{0}.$$

In other words, the multiplier $\boldsymbol{\nu}_n$ that characterizes \mathbf{p}_n^* via relation (16) is a root of the equation $f_n(\boldsymbol{\nu}_n) = \mathbf{0}$, where $f_n : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is given by

$$f_n(\boldsymbol{\nu}) = \sum_{i=1}^n \mathbf{C}_i (n + \boldsymbol{\nu}_n^T \mathbf{C}_i)^{-1}.$$

Given the root $\boldsymbol{\nu}_n \in \mathbb{R}^d$, $\hat{\alpha}_n$ is then defined by

$$\hat{\alpha}(n) = \sum_{i=1}^n X_i (n + \boldsymbol{\nu}_n^T \mathbf{C}_i)^{-1}.$$

If the root $\boldsymbol{\nu}_n$ yields p_{in}^* 's that are all positive, then the relaxation obtained by ignoring the non-negativity constraints yields a solution to the fully constrained problem. The proof of Theorem 7 will confirm that the above approach is asymptotically valid (so that the p_{in}^* 's defining $\hat{\alpha}(n)$ can be obtained from the root $\boldsymbol{\nu}_n$ of $f_n(\boldsymbol{\nu}) = \mathbf{0}$). More importantly, Theorem 7 shows that the constrained Monte Carlo estimator $\hat{\alpha}(n)$ is asymptotically identical to the optimal control variates estimator $\bar{X}_n(\boldsymbol{\lambda}^*) = \bar{X}_n - \boldsymbol{\lambda}^{*T} \bar{\mathbf{C}}_n$. This result

extends, to $d \geq 2$, an identical conclusion reached by [13] in the setting of scalar control variates.

For a vector $\mathbf{x} \in \mathbb{R}^d$, let $\|\mathbf{x}\|_\infty \triangleq \max\{|x_i| : 1 \leq i \leq d\}$, where the x_i 's are the components of \mathbf{x} .

Theorem 7. Suppose that $ECCT^T$ is non-singular and that $E|X|^2 < \infty$ and $E\|\mathbf{C}\|_\infty^4 < \infty$. Then,

$$\hat{\alpha}(n) = \bar{X}_n(\boldsymbol{\lambda}^*) + o_p(n^{-1/2})$$

as $n \rightarrow \infty$.

Proof. We will start by establishing that for n sufficiently large, the equation $f_n(\boldsymbol{\nu}) = \mathbf{0}$ has a root $\boldsymbol{\nu}_n$ for which all the quantities p_{in}^* , $1 \leq i \leq n$, defined by (16) are positive. It will then follow by convexity and from Theorem 4.38 of [1] that $\boldsymbol{p}_n^* = (p_{1n}^*, \dots, p_{nn}^*)^T$ is the global solution of the fully constrained problem.

Note that $f_n(\boldsymbol{\nu}) = \mathbf{0}$ if and only if $T_n(\boldsymbol{\nu}) = \boldsymbol{\nu}$, where $T_n(\boldsymbol{\nu}) = \boldsymbol{\nu} + n(ECCT^T)^{-1}f_n(\boldsymbol{\nu})$. We shall shortly see that $\tilde{\boldsymbol{\nu}}_n = (ECCT^T)^{-1} \sum_{i=1}^n \mathbf{C}_i$ is an approximate ‘‘fixed point’’ of T_n , so that $\boldsymbol{\nu}_n$ is close to $\tilde{\boldsymbol{\nu}}_n$. In fact, for $0 < \delta < 1/2$, we will show that $\boldsymbol{\nu}_n$ is, for n large, in the set $B_n \triangleq \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x} - \tilde{\boldsymbol{\nu}}_n\|_\infty \leq n^\delta\}$. (Given that the Law of the Iterated Logarithm ensures that $\tilde{\boldsymbol{\nu}}_n = O((n \log \log n)^{1/2})$ a.s. as $n \rightarrow \infty$ and $\delta < 1/2$, $\boldsymbol{\nu}_n \in B_n$ is therefore indeed a statement of ‘‘closeness’’.)

Observe that for $|x| < 1/2$, $|(1+x)^{-1} - (1-x)| \leq 2x^2$. Also, the moment condition $E\|\mathbf{C}\|_\infty^4 < \infty$, in conjunction with the Borel-Cantelli lemma, implies that $\|\mathbf{C}_n\|_\infty = O(n^{1/4})$ a.s., from which it follows that $\sup\{|\boldsymbol{\nu}^T \mathbf{C}_i|/n : \boldsymbol{\nu} \in B_n\} \rightarrow 0$ a.s. as $n \rightarrow \infty$. Thus, for n sufficiently large,

$$\sup_{\boldsymbol{\nu} \in B_n} |(1 + \boldsymbol{\nu} \mathbf{C}_i^T/n)^{-1} - (1 - \boldsymbol{\nu} \mathbf{C}_i^T/n)| \leq 3.(\|\tilde{\boldsymbol{\nu}}_n\|_\infty/n)^2 \|\mathbf{C}_i\|_\infty^2. \quad (17)$$

So,

$$\begin{aligned}
T_n(\boldsymbol{\nu}) &= \boldsymbol{\nu} + (ECC^T)^{-1} \sum_{i=1}^n \mathbf{C}_i (1 + \boldsymbol{\nu}^T \mathbf{C}_i / n)^{-1} \\
&= \boldsymbol{\nu} + (ECC^T)^{-1} \left[\sum_{i=1}^n \mathbf{C}_i - \frac{1}{n} \sum_{i=1}^n \mathbf{C}_i \mathbf{C}_i^T \boldsymbol{\nu} \right] \\
&\quad + O\left(\left(\frac{\|\tilde{\boldsymbol{\nu}}_n\|_\infty}{n} \right)^2 \sum_{i=1}^n \frac{\|\mathbf{C}_i\|_\infty^3}{n} \right) \tag{18} \\
&= \boldsymbol{\nu} - \tilde{\boldsymbol{\nu}}_n - (ECC^T)^{-1} \frac{1}{n} \sum_{i=1}^n \mathbf{C}_i \mathbf{C}_i^T \boldsymbol{\nu} + O\left(\frac{\log \log n}{n} \right) \text{ a.s.} \\
&= \tilde{\boldsymbol{\nu}}_n + O(\log \log n) \text{ a.s.}
\end{aligned}$$

uniformly in $\boldsymbol{\nu} \in B_n$. Consequently, T_n maps B_n into itself for n sufficiently large.

Furthermore, for $\mathbf{x}, \mathbf{y} \in B_n$,

$$\begin{aligned}
&\|T_n(\mathbf{x}) - T_n(\mathbf{y})\|_\infty \\
&= \|(\mathbf{x} - \mathbf{y}) + (ECC^T)^{-1} \frac{1}{n} \sum_{i=1}^n \mathbf{C}_i [(1 + \mathbf{x} \mathbf{C}_i / n)^{-1} - (1 + \mathbf{y} \mathbf{C}_i / n)^{-1}]\|_\infty \\
&\leq \|I - (ECC^T)^{-1} \frac{1}{n} \sum_{i=1}^n \mathbf{C}_i \mathbf{C}_i^T (1 + \mathbf{x} \mathbf{C}_i / n)^{-1} (1 + \mathbf{y} \mathbf{C}_i / n)\|_\infty \cdot \|\mathbf{x} - \mathbf{y}\|_\infty \\
&= \|I - (ECC^T)^{-1} \frac{1}{n} \sum_{i=1}^n \mathbf{C}_i \mathbf{C}_i^T (1 + O(n^{-1/8}))^2\|_\infty \cdot \|\mathbf{x} - \mathbf{y}\|_\infty \\
&= O\left(\left(\frac{\log \log n}{n} \right)^{1/2} \right) \|\mathbf{x} - \mathbf{y}\|_\infty
\end{aligned}$$

uniformly in $\mathbf{x}, \mathbf{y} \in B_n$. It follows that for n sufficiently large, T_n is a strict contraction on the closed ball B_n . The Banach Fixed Point Theorem (see, for example, [9]) therefore guarantees the existence of a unique fixed point $\boldsymbol{\nu}_n \in B_n$ for the mapping T_n . Furthermore, because $\boldsymbol{\nu}_n \in B_n$, it is easily seen that the resulting p_{in}^* 's are uniformly positive for n large. As argued earlier, we may therefore conclude that $\boldsymbol{\nu}_n$ is the multiplier associated with the desired global maximizer to the fully constrained problem.

Substituting $\boldsymbol{\nu} = \boldsymbol{\nu}_n$ in (18), we find that

$$\boldsymbol{\nu}_n = \tilde{\boldsymbol{\nu}}_n + O(\log \log n) \text{ a.s.}$$

So, (17) shows that

$$\begin{aligned}
\tilde{\nu}_n &= \frac{1}{n} \sum_{i=1}^n X_i (1 + \nu_n^T \mathbf{C}_i / n)^{-1} \\
&= \frac{1}{n} \sum_{i=1}^n X_i (1 - \nu_n^T \mathbf{C}_i / n) + O\left(\left(\frac{\|\tilde{\nu}_n\|_\infty}{n}\right)^2 \frac{1}{n} \sum_{i=1}^n |X_i| \|\mathbf{C}_i\|^2 \infty\right) \\
&= \frac{1}{n} \sum_{i=1}^n X_i (1 - \tilde{\nu}_n^T \mathbf{C}_i / n) + O\left(\frac{\log \log n}{n}\right) \text{ a.s.} \\
&= \frac{1}{n} \sum_{i=1}^n X_i(\lambda^*) + O\left(\frac{\log \log n}{n}\right) \text{ a.s.},
\end{aligned}$$

proving the theorem. ■

References

1. Avriel, M. (1976) *Nonlinear Programming: Analysis and Methods*. Prentice-Hall, Englewood Cliffs.
2. Billingsley, P. (1968) *Convergence of Probability Measures*. John Wiley, New York.
3. Billingsley, P. (1995) *Probability and Measure*, 3rd edn. John Wiley, New York.
4. Bratley, P., Fox, B.L., and L.E. Schrage (1987) *A Guide to Simulation*. Springer-Verlag, New York.
5. Davis, P.J., and P. Rabinowitz (1984) *Methods of Numerical Integration*. Academic Press, New York.
6. Evans, M., and T. Swartz (2000) *Approximating Integrals Via Monte Carlo and Deterministic Methods*. Oxford University Press, Oxford.
7. Fishman, G.S., and B.D. Huang (1983) Antithetic Variates Revisited. *Comms. of the ACM* 26:11:964-971.
8. Glynn, P.W. , and W. Whitt (1992) The Asymptotic Efficiency of Simulation Estimators . *Op. Res.* 40:505-520.
9. Kreyszig, E. (1978) *Introductory Functional Analysis with Applications*. John Wiley, New York.
10. Loh, W.W. (1995) *On the Method of Control Variates*. PhD Thesis, Stanford University, Stanford.
11. Ross, S.M. (1990) *A Course in Simulation*. Macmillan Publishing Company, New York.
12. Rubinstein, R.Y. (1981) *Simulation and the Monte Carlo Method*. John Wiley, New York.
13. Szechtman, R., and P.W. Glynn (2001) *Constrained Monte Carlo and the Method of Control Variates*. Submitted to Win. Simulation Conf.
14. Szechtman, R. (2001) *Efficient Monte Carlo Simulation in the Presence of Constraints*. PhD Thesis, Stanford University, Stanford.