

CONSTRAINED MONTE CARLO AND THE METHOD OF CONTROL VARIATES

Roberto Szechtman
Peter W. Glynn

Department of Management
Science and Engineering
Stanford University
Stanford, CA 94305 U.S.A.

ABSTRACT

A constrained Monte Carlo problem arises when one computes an expectation in the presence of a priori computable constraints on the expectations of quantities that are correlated with the estimand. This paper discusses different applications settings in which such constrained Monte Carlo computations arise, and establishes a close connection with the method of control variates when the constraints are of equality form.

1 INTRODUCTION

This paper is concerned with “constrained Monte Carlo” computation. In particular, suppose that we wish to compute an expectation of the form $\alpha = EX$ via a sampling-based method. In many applications settings, there exists a random vector Y , correlated with X , for which one has a priori knowledge that EY lies in some known set B . For example, if B is a singleton μ , we are dealing with an equality constraint of the form $EY = \mu$. This, of course, is precisely the setting of the traditional method of control variates. However, as we shall see in Section 2, the set B can also describe inequality constraints, in which case a general methodology capable of dealing with both equality and inequality constraints on EY is needed.

In this paper, we will show how inequality constraints can arise quite naturally in many applications. This is the subject of Section 2. In the remainder of the paper, we discuss one approach to the problem of constrained Monte Carlo, and show how the proposed methodology coincides asymptotically with the traditional method of control variates when B exclusively describes equality constraints. In particular, in Section 3, we describe a nonparametric maximum likelihood approach and relate its large-sample behavior to control variates. In Section 4, we illustrate our theory with an Asian option pricing example.

We are unaware of any previous literature on the problem of constrained Monte Carlo, at the level of generality

described here. However, recent work by Avellaneda and his collaborators on a closely related “calibration” problem that arises in the finance context contains some ideas that parallel our proposals; see Avellaneda et al. (2000); and Avellaneda and Gamba (2000). Nevertheless, the theorem that we present in Section 3 appears to be new.

As mentioned above, this paper is largely concerned with motivating the need for constrained Monte Carlo methods (Section 2), providing an analysis of the proposed methodology in the context of equality constraints (Section 3), and giving numerical evidence of our results (Section 4). A complete discussion of constrained Monte Carlo in the inequality setting can be found in Szechtman and Glynn (2001).

2 CONSTRAINED MONTE CARLO COMPUTATION

Suppose that we wish to compute $\alpha = EX$, where X is a real-valued random variable. We assume that we can simulate independent and identically distributed (iid) replicates $\mathbf{X}_1, \mathbf{X}_2, \dots$ of the \mathbb{R}^d -valued random vector \mathbf{X} , where $\mathbf{X} = (X_1, \dots, X_d)$ and $X_1 \stackrel{D}{=} X$ (where $\stackrel{D}{=}$ denotes “equality in distribution”). In other words, the estimand X is just the first component of the simulatable random vector \mathbf{X} .

Constrained Monte Carlo deals with the situation in which there exists a given set $B \subseteq \mathbb{R}^d$ for which it is known that $E\mathbf{X} \in B$. As mentioned in the Introduction, one important such context is the setting of equality constraints. Pure equality constraints arise when it is known that $Y \triangleq (X_2, \dots, X_d)$ satisfies the constraint $EY = \mu \in \mathbb{R}^{d-1}$ with μ given. This is the problem environment within which the traditional method of control variates is relevant.

However, there are many applications contexts within which more complex constraints arise.

Example 1 Suppose that we wish to compute the probability $\alpha = P(A)$ of the event A via importance sampling. In this context, the estimator X for α takes

the form $X = I(A)L$, where $I(A)$ is the indicator of the event A (i.e. $I(A)$ is 1 or 0, depending in whether or not the A occurs), and L is a suitable likelihood ratio; see, for example, Glynn and Iglehart (1989). If we set $\mathbf{X} = (I(A)L, L) = (X_1, X_2)$, note that $E\mathbf{X} \in B$, where

$$B = \{(x_1, x_2) : 0 \leq x_1 \leq 1, x_2 = 1\}.$$

Example 2 There are many classical areas of applied probability in which various bounds on expectations and probabilities have been deduced over the years. A good example of such a bound is the Cramér-Lundberg inequality.

Specifically, let $S = (S_n : n \geq 0)$ be a random walk that can be represented as $S_n = Z_1 + \dots + Z_n$, where the Z_i 's are iid real-valued random variables. Suppose that there exists $\theta^* > 0$ for which $E \exp(\theta^* Z_1) = 1$ and suppose that $E Z_1 < 0$. The negative drift of the random walk implies that $S_n \rightarrow -\infty$ a.s. as $n \rightarrow \infty$, so that $M = \max\{S_k : k \geq 0\} < \infty$ a.s. The random variable M is of central interest in several different contexts. In particular, $P(M > x)$ describes the probability of ruin for an insurance company that has initial reserve x . Also, M has precisely the same distribution as the steady-state waiting time (exclusive of service) in a single-server queue; see Asmussen (1987) for additional details. The Cramér-Lundberg inequality states that

$$P(M > x) \leq \exp(-\theta^* x) \tag{1}$$

for $x \geq 0$.

Note that if we let $M_m = \max\{S_k : 0 \leq k \leq m\}$ be our (simulatable) approximation to M , $P(M > x)$ can be approximately estimated by the estimator $X = I(M_m > x)$. If $p_m = P(M_m > x)$, it follows that a sample mean formed from iid copies of X will have a standard deviation $\sqrt{p_m(1 - p_m)/n}$, where n is the sample size. Given that $P(M > x) \sim c \exp(-\theta^* x)$ as $x \rightarrow \infty$ for some $c \in (0, 1)$ (here, $a(x) \sim b(x)$ means that $a(x)/b(x) \rightarrow 1$ as $x \rightarrow \infty$) is valid under modest additional regularity conditions (Asmussen 1987, p. 269), it follows that the bound (1) will be violated with substantial positive probability for values of n that are of equal or smaller order than $\exp(\theta^* x)$. Thus, this is a setting in which constrained Monte Carlo is potentially valuable, because the inequality constraint is likely to be active even for large sample sizes.

Example 3 There has been recent activity in exploiting the method of Lyapunov functions to obtain a priori bounds on steady-state expectations for Markov chains and processes; see, for example Bertsimas, Gamarnik, and Tsitsiklis (1998); Bertsimas, Gamarnik, and Tsitsiklis (2000); and Kumar and Kumar (1994).

We offer here a different bound, with an elementary proof. To set the stage, let $Z = (Z_n : n \geq 0)$ be an S -valued (time-homogeneous) Markov chain, and let P be the

(one-step) transition kernel given by

$$P(x, dy) = P(Z_1 \in dy | Z_0 = x)$$

for $x, y \in S$. Given a non-negative (measurable) function $h : S \rightarrow \mathbb{R}$, let Ph be the function defined by

$$(Ph)(x) = \int_S h(y)P(x, dy)$$

for $x \in S$. Furthermore, for a probability η defined on S , let ηP be the probability

$$(\eta P)(\cdot) = \int_S \eta(dx)P(x, \cdot)$$

and let $\eta h = \int_S h(y)\eta(dy)$.

Our goal is to provide a bound on πf , where π is a probability that is stationary for P in the sense that $\pi = \pi P$. The quantity πf can typically be viewed as the steady-state reward accrued per unit time, when $f(x)$ is interpreted as the reward received when Z spends one unit of time in $x \in S$.

Proposition 1. Suppose f and g are non-negative functions and π is a stationary distribution for which $\pi g < \infty$. If there exists a (measurable) set $K \subseteq S$ such that

$$(Pg)(x) \leq g(x) - f(x) \tag{2}$$

for $x \in K^c$, then

$$0 \leq \pi f \leq \sup\{f(x) + (Pg)(x) - g(x) : x \in K\}$$

Proof. Since π is stationary and g is non-negative, $\pi g = \pi Pg$. Furthermore, since $\pi g < \infty$, evidently $\pi(P - I)g = 0$. So,

$$0 = \int_K ((Pg)(x) - g(x))\pi(dx) + \int_{K^c} ((Pg)(x) - g(x))\pi(dx). \tag{3}$$

On the other hand, (2) guarantees that

$$\int_{K^c} f(x)\pi(dx) \leq - \int_{K^c} ((Pg)(x) - g(x))\pi(dx). \tag{4}$$

Consequently,

$$\begin{aligned} \pi f &= \int_K f(x)\pi(dx) + \int_{K^c} f(x)\pi(dx) \\ &\leq \int_K f(x)\pi(dx) - \int_{K^c} ((Pg)(x) - g(x))\pi(dx), \\ &\hspace{15em} \text{using (4)} \\ &= \int_K f(x)\pi(dx) + \int_K ((Pg)(x) - g(x))\pi(dx), \\ &\hspace{15em} \text{using (3)} \\ &\leq \sup\{f(x) + (Pg)(x) - g(x) : x \in K\}. \end{aligned}$$

■

The function g appearing in Proposition 1 is called a Lyapunov function. By choosing g intelligently so as to exploit the structure of P and f , one can use Proposition 1 to obtain good bounds on πf . For example, quadratic candidates for g often give good bounds on πf , in the setting of functions f that correspond to “queue-lengths” in the queueing network context. Note that Proposition 1 holds for any Markov chain, and makes no recurrence assumptions regarding Z (such as Harris recurrence). Meyn and Tweedie (1993) is a good source for models that are amenable to this type of Lyapunov analysis.

Proposition 1 also extends to continuous-time Markov processes. As an illustration, we offer the following result.

Proposition 2. Let $Z = (Z(t) : t \geq 0)$ be a continuous-time Markov chain on discrete state space S , having generator (i.e. rate matrix) A . Suppose A is uniformizable (i.e. $\|A\| \triangleq \sup\{\sum_{y \in S} |A(x, y)| : x \in S\} < \infty$). Let f and g be two non-negative functions for which $\pi g < \infty$, where π is a probability satisfying $\pi A = 0$. If there exists a set $K \subseteq S$ such that $(Ag)(x) \leq -f(x)$ for $x \in K^c$, then

$$0 \leq \pi f \leq \sup\{f(x) + (Ag)(x) : x \in K\} < \infty.$$

Proof. Let $\lambda(x) = -A(x, x)$ be the jump rate for state $x \in S$, and note that the uniformizability of A ensures that $\sup\{\lambda(x) : x \in S\} = \|A\|/2 < \infty$. Also, the inequality $(Ag)(x) \leq -f(x)$ is equivalent to

$$\sum_y R(x, y)g(y) \leq g(x) - f(x)/\lambda(x),$$

where $R = (R(x, y) : x, y \in S)$ is the stochastic matrix in which $R(x, x) = 0$ and $R(x, y) = A(x, y)/\lambda(x)$ for $x \neq y$. Set $\tilde{\pi}(x) = \pi(x)\lambda(x)/\sum_y \pi(y)\lambda(y)$. The uniformizability of A ensures that $\tilde{\pi}$ is a stationary distribution for R satisfying $\tilde{\pi} g < \infty$. Consequently, we can follow the proof of

Proposition 1, yielding the inequality

$$\begin{aligned} &\sum_x \tilde{\pi}(x) f(x)/\lambda(x) \\ &\leq \sum_{x \in K} \tilde{\pi}(x) \left(\frac{f(x)}{\lambda(x)} + (Rg)(x) - g(x) \right). \end{aligned}$$

So

$$\begin{aligned} \sum_x \pi(x) f(x) &\leq \sum_{x \in K} \pi(x) (f(x) + (Ag)(x)) \\ &\leq \sup\{f(y) + (Ag)(y) : y \in K\}. \end{aligned}$$

■

This bound on πf can be applied similarly as in discrete time. It follows, from Propositions 1 and 2, that bounds on steady-state expectations for Markov chains and processes can easily be generated in practice. For additional discussion of such bounds, see Zeevi and Glynn (2001).

3 NONPARAMETRIC MAXIMUM LIKELIHOOD APPROACH TO CONSTRAINED MONTE CARLO

Suppose that we can simulate n iid copies $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n$ of the random vector \mathbf{X} , with the goal of computing the expectation $\alpha = EX_1$ of the first component of \mathbf{X} . Given the knowledge that $EX \in B$, a natural means of dealing with such a constraint is to analyze the simulated data via the principle of (nonparametric) maximum likelihood.

In particular, we look for a probability distribution supported on the n sample points $\mathbf{X}_1, \dots, \mathbf{X}_n$ that maximizes the likelihood of the simulated data set, subject to the constraint that the resulting probability distribution be consistent with the knowledge that $EX \in B$. To be precise, write $\mathbf{X}_i = (X_{i1}, \dots, X_{id})$. The nonparametric maximum likelihood problem involves solving

$$\max_{p_1, \dots, p_n} \prod_{i=1}^n p_i$$

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

$$\sum_{i=1}^n p_i = 1 \tag{5}$$

$$\left(\sum_{i=1}^n p_i X_{ij} : 1 \leq j \leq d \right) \in B. \tag{6}$$

If the convex hull of $\mathbf{X}_1, \dots, \mathbf{X}_n$ intersects B , then the above optimization problem has a feasible point. Furthermore, if B is convex, the above problem involves maximizing a strictly concave function (namely $\sum_{i=1}^n \log p_i$) over a convex set, so that any maximizer must be unique. Let p^* be the associated maximizer. The vector p^* induces the probability

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

(here $\delta_{\mathbf{X}_i}(\cdot) = I(\mathbf{X}_i \in \cdot)$), leading to the point estimator

$$\alpha(n) = \int_{\mathbb{R}^d} x_1 \mu_n(d\mathbf{x}).$$

Proposition 3. Suppose that B has an open interior B° and that $EX \in B^\circ$. Then, $\alpha(n) \rightarrow \alpha$ a.s. as $n \rightarrow \infty$.

Proof. Without the constraint (6) on (p_1, \dots, p_n) , the solution of the nonparametric maximum likelihood estimation problem is $\tilde{p}_i = 1/n$, $1 \leq i \leq n$. Consequently, if $n^{-1} \sum_{i=1}^n \mathbf{X}_i \in B$, the constraint (6) is not active at \tilde{p} , so the “relaxed maximizer” \tilde{p} actually maximizes the fully constrained problem.

By the strong law of large numbers, $n^{-1} \sum_{i=1}^n \mathbf{X}_i \rightarrow EX$ a.s. as $n \rightarrow \infty$. Given $EX \in B^\circ$, it follows that $n^{-1} \sum_{i=1}^n \mathbf{X}_i \in B^\circ \subseteq B$ for n sufficiently large, so that we may conclude that

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

for n large enough. Thus $\alpha(n) = n^{-1} \sum_{i=1}^n X_{i1}$ for n sufficiently large, so that $\alpha(n) \rightarrow EX_1$ a.s. as $n \rightarrow \infty$ by the strong law for iid sequences. ■

Thus, Proposition 3 provides a setting in which $\alpha(n)$ is a (strongly) consistent estimator for α . Note, however, that Proposition 3 does not cover the case in which $Y = (X_2, \dots, X_d)$ satisfies the equality constraint $EY = \mu$. In the remainder of this section, we study this important special case. In particular, we will establish that the estimator $\alpha(n)$ is asymptotically identical to the estimator associated with the method of control variates. This suggests that at least in this special setting, the estimator $\alpha(n)$ behaves sensibly (and, in some sense, optimally).

Before discussing this result, we recall that the method of control variates involves taking advantage of the fact that the random variable $X - \lambda^T(Y - \mu)$ is an unbiased estimator for α , regardless of the value of the “control coefficient” vector λ . (Here we encode λ and Y as column vectors). Assuming that the covariance matrix $\Sigma \triangleq E(Y - \mu)(Y - \mu)^T$ is non-singular, the variance of the estimator is minimized

over λ via the choice

$$\lambda^* = \Sigma^{-1} \cdot EX\tilde{Y},$$

where $\tilde{Y} = Y - \mu$; see Lavenberg and Welch (1981). Since the covariance matrix Σ and the covariance vector $EX\tilde{Y}$ are typically unknown, one estimates λ^* from the n simulated replicates $(X_1, Y_1), \dots, (X_n, Y_n)$ of (X, Y) as follows:

$$\lambda_n = \left(\frac{1}{n} \sum_{i=1}^n \tilde{Y}_i \tilde{Y}_i^T \right)^{-1} \cdot \frac{1}{n} \sum_{i=1}^n X_i \tilde{Y}_i$$

where $\tilde{Y}_i = Y_i - \mu$. This leads to the control variates estimator

$$\alpha_c(n) = \frac{1}{n} \sum_{i=1}^n X_i - \lambda_n^T \frac{1}{n} \sum_{i=1}^n \tilde{Y}_i.$$

On the other hand, the optimization problem defining the probability vector $p^* = p_n^*$ is, in the setting of equality constraints,

$$\max_{p_1, \dots, p_n} \sum_{i=1}^n \log p_i$$

subject to $p_i > 0$ for $1 \leq i \leq n$ and

$$\sum_{i=1}^n p_i = 1 \tag{7}$$

$$\sum_{i=1}^n p_i \tilde{Y}_i = 0. \tag{8}$$

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 (7) and $v_n \in \mathbb{R}^{d-1}$ for (8). At the maximizer $p_n^* = (p_{n1}^*, \dots, p_{nn}^*)^T$, the multipliers should satisfy

$$(1/p_{ni}^*) - v_n^T \tilde{Y}_i - \gamma_n = 0 \tag{9}$$

for $1 \leq i \leq n$. Multiplying through (9) by p_{ni}^* and summing over i , we get

$$n - v_n^T \sum_{i=1}^n p_{ni}^* \tilde{Y}_i - \gamma_n \sum_{i=1}^n p_{ni}^* = 0.$$

Hence $\gamma_n = n$ and (9) implies that

$$p_{ni}^* = (n + v_n^T \tilde{Y}_i)^{-1} \tag{10}$$

for $1 \leq i \leq n$. Multiplying through (10) by \tilde{Y}_i and summing over i , we arrive at the equation

$$\sum_{i=1}^n \tilde{Y}_i (n + v_n^T \tilde{Y}_i)^{-1} = 0.$$

In other words, the multiplier v_n that characterizes p_n^* via relation (10) is a root of the equation $f_n(v) = 0$, where $f_n : \mathbb{R}^{d-1} \rightarrow \mathbb{R}^{d-1}$ is given by

$$f_n(v) = \sum_{i=1}^n \tilde{Y}_i (n + v^T \tilde{Y}_i)^{-1}.$$

Given the root $v_n \in \mathbb{R}^{d-1}$, $\alpha(n)$ is then defined by

$$\alpha(n) = \sum_{i=1}^n X_i (n + v_n^T \tilde{Y}_i)^{-1}.$$

If the root v yields p_{ni}^* '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 the next theorem will confirm that the above approach is asymptotically valid (so that the p_{ni}^* 's defining $\alpha(n)$ can be obtained from the root v_n of $f_n(v) = 0$). More importantly, the next theorem establishes that $\alpha(n)$ and $\alpha_c(n)$ become asymptotically identical as $n \rightarrow \infty$.

In order to simplify our exposition, we focus on the case where Y is scalar-valued. We say that a sequence of random vectors $(\chi_n : n \geq 1)$ is $o_p(n^{-1/2})$ if $n^{1/2}\chi_n \Rightarrow 0$ as $n \rightarrow \infty$. Moreover, a sequence $(\chi_n : n \geq 1)$ of random vectors is $O(a_n)$ a.s. ($o(a_n)$ a.s.) if χ_n/a_n is a.s. a bounded sequence ($\chi_n/a_n \rightarrow 0$ a.s. as $n \rightarrow \infty$).

Theorem. Suppose that Y is scalar-valued. If $EX^2 < \infty$ and $EY^4 < \infty$, then

$$\alpha(n) = \alpha_c(n) + O\left(\frac{\log \log n}{n}\right) \text{ a.s.}$$

as $n \rightarrow \infty$, so that

$$n^{1/2}(\alpha(n) - \alpha, \alpha_c(n) - \alpha) \Rightarrow (\sigma N, \sigma N)$$

as $n \rightarrow \infty$, where N is a mean-zero normal r.v. with unit variance, and $\sigma^2 = \text{var } X \cdot (1 - \rho^2)$ with ρ the coefficient of correlation between X and Y .

Proof. We will start by establishing that for n sufficiently large, the equation $f_n(v) = 0$ has a root v_n for which all the quantities p_{ni}^* , $1 \leq i \leq n$, defined by (10) are positive. It will then follow by convexity and Theorem

4.38 of Avriel (1976) that $p_n^* = (p_{n1}^*, \dots, p_{nn}^*)^T$ is the global solution of the fully constrained problem.

Let $\tilde{v}_n = \sum_{i=1}^n \tilde{Y}_i / (\sum_{i=1}^n \tilde{Y}_i^2 / n)$. By the strong law of large numbers (as applied to $\sum_{i=1}^n \tilde{Y}_i^2 / n$) and the law of the iterated logarithm (as applied to $\sum_{i=1}^n \tilde{Y}_i$), it follows that $\tilde{v}_n = O((n \log \log n)^{1/2})$ a.s. as $n \rightarrow \infty$. Put $I_n = [\tilde{v}_n - n^\delta, \tilde{v}_n + n^\delta]$ for $\delta \in (0, 1/2)$.

We claim that $f_n(\cdot)$ is continuous on I_n for n large. Given the form of f_n , this will follow if we can show that $n + v \tilde{Y}_i > 0$ for $1 \leq i \leq n$ and $v \in I_n$, for n large. To verify this, observe that the Borel-Cantelli lemma and the hypothesis $E\tilde{Y}^4 < \infty$ imply that $\tilde{Y}_n = o(n^{1/4})$ a.s. as $n \rightarrow \infty$. Consequently, for $v \in I_n$, $v \tilde{Y}_i = o(n)$ a.s. as $n \rightarrow \infty$, proving the uniform positivity of the terms $n + v \tilde{Y}_i$, for n large.

We will now show that $f_n(\tilde{v}_n - n^\delta)$ and $f_n(\tilde{v}_n + n^\delta)$ are of opposite sign, for n large. The continuity of $f_n(\cdot)$ then guarantees the existence of a root $v_n \in I_n$ at which the quantities $(n + v_n \tilde{Y}_i)^{-1}$ are uniformly positive in $i \in \{1, \dots, n\}$. As argued earlier, this implies that v_n is the desired global maximizer.

Observe that for $|x| \leq 1/2$, $|(1+x)^{-1} - (1-x)^{-1}| \leq 2x^2$. Then, for n sufficiently large,

$$\sup_{v \in I_n} \left| \left(1 + \frac{v}{n} \tilde{Y}_i\right)^{-1} - \left(1 - \frac{v}{n} \tilde{Y}_i\right)^{-1} \right| \leq 2 \left(\frac{\tilde{v}_n}{n}\right)^2 \tilde{Y}_i^2 \quad (11)$$

So,

$$\begin{aligned} f_n(v) &= \frac{1}{n} \sum_{i=1}^n \tilde{Y}_i \left(1 + \frac{v}{n} \tilde{Y}_i\right)^{-1} \\ &= \frac{1}{n} \sum_{i=1}^n \tilde{Y}_i \left(1 - \frac{v}{n} \tilde{Y}_i\right) + O\left(\left(\frac{\tilde{v}_n}{n}\right)^2 \sum_{i=1}^n \frac{|\tilde{Y}_i|^3}{n}\right) \\ &= \frac{1}{n} \sum_{i=1}^n \tilde{Y}_i \left(1 - \frac{\tilde{v}_n}{n} \tilde{Y}_i\right) + (\tilde{v}_n - v) \cdot \frac{1}{n^2} \sum_{i=1}^n \tilde{Y}_i^2 \\ &\quad + O\left(\frac{\log \log n}{n}\right) \text{ a.s.} \\ &= (\tilde{v}_n - v) \cdot \frac{1}{n} E\tilde{Y}^2 + O\left(\frac{\log \log n}{n}\right) \text{ a.s.} \end{aligned} \quad (12)$$

uniformly in $v \in I_n$. (The definition of \tilde{v}_n was used to eliminate the first term in the third equality above.) Hence, $f_n(\tilde{v}_n - n^\delta) = n^{\delta-1} E\tilde{Y}^2 + O((\log \log n)/n)$ a.s. and $f_n(\tilde{v}_n + n^\delta) = -n^{\delta-1} E\tilde{Y}^2 + O((\log \log n)/n)$ a.s., proving that the endpoints of I_n are of opposite sign.

Our last task is to establish the relationship with control variates. Since (12) holds uniformly in $v \in I_n$ and $v_n \in I_n$,

it follows that

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

as $n \rightarrow \infty$. So, (11) shows that

$$\begin{aligned} \alpha(n) &= \frac{1}{n} \sum_{i=1}^n X_i \left(1 + \frac{v_n}{n} \tilde{Y}_i\right)^{-1} \\ &= \frac{1}{n} \sum_{i=1}^n X_i \left(1 - \frac{v_n}{n} \tilde{Y}_i\right) \\ &\quad + O\left(\left(\frac{\tilde{v}_n}{n}\right)^2 \cdot \frac{1}{n} \sum_{i=1}^n |X_i \tilde{Y}_i^2|\right) \\ &= \frac{1}{n} \sum_{i=1}^n X_i \left(1 - \frac{\tilde{v}_n}{n} \tilde{Y}_i\right) + (\tilde{v}_n - v_n) \cdot \frac{1}{n^2} \sum_{i=1}^n X_i \tilde{Y}_i^2 \\ &\quad + O\left(\frac{\log \log n}{n}\right) \\ &= \alpha_c(n) + (\tilde{v}_n - v_n) \cdot \frac{1}{n^2} \sum_{i=1}^n X_i \tilde{Y}_i^2 \\ &\quad + O\left(\frac{\log \log n}{n}\right) \\ &= \alpha_c(n) + O\left(\frac{\log \log n}{n}\right) \text{ a.s.} \end{aligned}$$

as $n \rightarrow \infty$. (Note that Cauchy-Schwarz inequality implies that $E|X\tilde{Y}^2| < \infty$.) The central limit theorem follows from a converging together argument and the existing central limit theory for $\alpha_c(n)$. ■

The above theorem makes concrete the assertion that nonparametric maximum likelihood, in the presence of equality constraints, basically coincides with the method of control variates.

4 A NUMERICAL EXAMPLE

In this section, we offer a numerical example to complement the theory developed in Section 3. Specifically, we consider the problem of numerically computing the price of an Asian option via simulation. Such a price can be expressed as the expectation of a random variable X (where the expectation is computed under the so-called “equivalent martingale measure”; see Duffie (1996) for details).

In particular,

$$X = \max\left(\int_0^t \xi(s) ds - k, 0\right)$$

where $\xi = (\xi(s) : s \geq 0)$ is a stochastic process describing the price of the underlying security and k is the “strike

price”. A common specification for ξ is to assume that it is a geometric Brownian motion. We will follow this convention, and will assume that

$$\xi(t) = \exp(B(t))$$

where $B = (B(t) : t \geq 0)$ is standard Brownian motion. For our example, we choose $t = 1$ and $k = 1.25$.

A commonly used control variate in the setting of Asian options is

$$Y = \int_0^t \xi(s) ds.$$

For our geometric Brownian motion example

$$\begin{aligned} EY &= \int_0^1 E \exp(B(t)) dt \\ &= \int_0^1 \exp(t/2) dt \\ &= 2(\exp(1/2) - 1). \end{aligned}$$

For the purposes of this numerical study, we consider three estimators. The first estimator we consider, denoted $\tilde{\alpha}(n)$, is the conventional Monte Carlo estimator based on computing a sample mean formed from n iid replications of the random variable X . Our second estimator for $\alpha = EX$ is the control variates estimator $\alpha_c(n)$ described in Section 3. Finally, the third estimator is $\alpha(n)$, the nonparametric maximum likelihood estimator.

Since we are interested in comparing the variability of the three estimators just defined, we repeat $m = 1000$ times the simulation described in the previous paragraph. The estimators obtained by averaging over m the α ’s are $\tilde{\alpha}(n, m)$, $\alpha_c(n, m)$, and $\alpha(n, m)$, in correspondence with previous notation. The (sample) standard deviations of these three estimators are $\tilde{s}(m)$, $s_c(m)$, and $s(m)$.

The results are summarized in Tables 1 and 2; the true value of α here is 0.3247. Observe that the nonparametric maximum likelihood estimator $\alpha(n, m)$ and the control variates estimator $\alpha_c(n, m)$ give similar results, in accordance with our theorem.

Table 1: Asian Option Pricing Simulation

Parameter	Sample size n		
	100	1000	10000
$\tilde{\alpha}(n, m)$	0.3231	0.3263	0.3280
$\alpha_c(n, m)$	0.3230	0.3225	0.3243
$\alpha(n, m)$	0.3236	0.3225	0.3250
$\tilde{s}(m)$	0.0740	0.0755	0.0759
$s_c(m)$	0.0233	0.0241	0.0228
$s(m)$	0.0237	0.0234	0.0226

Table 2: Asian Option Pricing Confidence Intervals

95 % Confidence Interval	Sample size n		
	100	1000	10000
$\tilde{\alpha}(n, m) - z_{0.95}\tilde{s}(m)/\sqrt{m}$	0.3084	0.3216	0.3265
$\tilde{\alpha}(n, m) + z_{0.95}\tilde{s}(m)/\sqrt{m}$	0.3378	0.3310	0.3295
$\alpha_c(n, m) - z_{0.95}s_c(m)/\sqrt{m}$	0.3184	0.3210	0.3238
$\alpha_c(n, m) + z_{0.95}s_c(m)/\sqrt{m}$	0.3276	0.3240	0.3248
$\alpha(n, m) - z_{0.95}s(m)/\sqrt{m}$	0.3190	0.3210	0.3246
$\alpha(n, m) + z_{0.95}s(m)/\sqrt{m}$	0.3282	0.3240	0.3254

REFERENCES

- Asmussen, S. 1987. *Applied Probability and Queues*. New York: Wiley.
- Avellaneda, M., R. Buff, C. Friedman, N. Grandchamp, L. Kruk, and J. Newman. 2000. Weighted Monte Carlo: A new technique for calibrating asset-pricing models. *International J. of Theoretical and Applied Finance* 4: 1: 1–29.
- Avellaneda, M., and R. Gamba. 2000. Conquering the greeks in Monte Carlo: Efficient calculation of the market sensitivities and hedge-ratios of financial assets by direct numerical simulation.
- Avriel, M. 1976. *Nonlinear Programming: Analysis and Methods*. Englewood Cliffs: Prentice-Hall.
- Bertsimas, D., D. Gamarnik, and J. N. Tsitsiklis. 1998. Geometric bounds for stationary distributions of infinite Markov chains via Lyapunov functions. *Laboratory for Information and Decision Systems and Operations Research Center, MIT*.
- Bertsimas, D., D. Gamarnik, and J. N. Tsitsiklis. 2000. Performance of multiclass markovian queueing networks via piecewise linear Lyapunov functions. To appear in *The Annals of Applied Probability*.
- Duffie, D. 1996. *Dynamic Asset Pricing Theory*. 2nd ed. Princeton: Princeton University Press.
- Glynn, P. W., and D. L. Iglehart. 1989. Importance sampling for stochastic systems. *Management Science* 35: 1367–1393.
- Kumar, S., and P. R. Kumar. 1994. Performance bounds for queueing networks and scheduling policies. *IEEE Transactions on Automatic Control* 8: 1600–1611.
- Lavenberg, S. S., and P. D. Welch. 1981. A perspective on the use of control variables to increase the efficiency of Monte Carlo simulations. *Management Science* 27: 322–335.
- Meyn, S. P., and R. L. Tweedie. 1993: *Markov Chains and Stochastic Stability*. New York: Springer Verlag.
- Szechtman, R., and P. W. Glynn. 2001. Efficient Monte Carlo computation in the presence of constraints. In preparation.

Zeevi, A., and P. W. Glynn. 2001. Simple explicit upper and lower bounds for steady-state expectations of Markov processes . In preparation.

AUTHOR BIOGRAPHIES

ROBERTO SZECHTMAN is a Ph.D. student in the Department of Management Science and Engineering at Stanford University. His research interests include applied probability and large deviations theory.

PETER W. GLYNN received his Ph.D. from Stanford University, after which he joined the faculty of the Department of Industrial Engineering at the University of Wisconsin-Madison. In 1987, he returned to Stanford, where he is the Thomas Ford Professor of Engineering in the Department of Management Science and Engineering. He was a co-winner of the 1993 Outstanding Simulation Publication Award sponsored by the TIMS College on Simulation. His research interests include discrete-event simulation, computational probability, queueing, and general theory for stochastic systems.