

Adaptive Control Variates for Finite-Horizon Simulation

Sujin Kim and Shane G. Henderson

School of Operations Research and Industrial Engineering
Cornell University

March 27, 2006

Abstract

Adaptive Monte Carlo methods are simulation efficiency improvement techniques designed to adaptively tune simulation estimators. Most of the work on adaptive Monte Carlo methods has been devoted to adaptively tuning importance sampling schemes. We instead focus on adaptive methods based on control variate schemes. We introduce two adaptive control variate methods, and develop their asymptotic properties. The first method uses stochastic approximation to adaptively tune control variate estimators. It is easy to implement, but it requires some non-trivial tuning of parameters. The second method is based on sample average approximation. Tuning is no longer required, but it can be computationally expensive. Numerical results for the pricing of barrier options are presented to demonstrate the methods.

1 Introduction

Suppose that we wish to estimate EX , where X is a real-valued random variable. Suppose also that $\{Y(\theta) : \theta \in \Theta\}$ is a parametric collection of random variables such that $EY(\theta) = 0$ for any θ in the parameter set Θ . Then $X - Y(\theta)$ is an unbiased estimator for EX , where $Y(\theta)$ serves as a control variate, and the parameter θ can be selected so as to minimize the variance of $X - Y(\theta)$. When the parameterization is linear, we can appeal to the standard theory of (linear) control variates, e.g., [Law and Kelton, 2000]. Identifying the θ that minimizes the variance is straightforward in the linear case because the variance is a convex quadratic in θ . When the parameterization is nonlinear, the problem is not so straightforward. We propose two adaptive procedures that tune the parameter θ while estimating EX , and study the large sample properties of the procedures.

It was shown in Glynn and Whitt [1989] that, from an asymptotic perspective (asymptotic in the simulation runlength), there is no advantage gained by nonlinear parameterizations over linear parameterizations. However, the form of the parameterization considered in Glynn and Whitt [1989] is different from the form considered in this paper. Glynn and Whitt [1989] analyzed the asymptotic performance of control variate estimators which have the form $f(\bar{X}_n, \bar{Y}_n)$, where n is the simulation runlength, \bar{X}_n and \bar{Y}_n are sample averages of X and the control Y respectively, and f is a function that combines the two sample averages. They showed that from the asymptotic point of view, we may restrict the choice of f to linear functions. Indeed, in our parametrization, X and the control variate $Y(\theta)$ are linearly related so the negative result does not apply.

Our interest in this problem stems from several application areas. An extended example in this paper (see Sections 2 and 6) shows that one can apply adaptive control variates in pricing certain financial derivatives. A second example arises in the simulation analysis of multiclass processing networks. When these networks are heavily loaded, simulation estimators can suffer from large variance, and so some form of variance reduction is needed. The simulation estimators developed in Henderson and Meyn [1997, 2003] give large variance reductions, but the asymptotic rates of growth in the variance are the same as for the naïve estimator; see Meyn [2005]. One approach to potentially improving these estimators is to develop parameterized estimators. A third class of examples arises in the problem of estimating the “expected cost to absorption” in a Markov chain. This problem has received a great deal of attention because of its applications in radiation transport problems; see, e.g., Kollman et al. [1999].

The common thread underlying these applications is that they involve the simulation of a Markov process. This allows us to construct a parameterized family of control variates using the approximating martingales developed in Henderson and Glynn [2002]. Once we have a parameterized class of control variates at hand, we then need a procedure for selecting a control from within the class.

The first of our procedures is based on a stochastic approximation scheme. At iteration k , several independent replications of $X - Y(\theta_{k-1})$ are generated, conditional on the parameter choice θ_{k-1} from the previous iteration. The sample mean and the gradient (with respect to θ) of the sample variance are then computed, and the parameter θ_{k-1} is updated to θ_k in a stochastic approximation step. This procedure is easily implemented and exhibits good performance with appropriately chosen step sizes. But the selection of the step size is nontrivial and the finite-time performance of the algorithm is strongly impacted by the choice of step sizes.

The second procedure is based on the theory of sample average approximation. In an initial stage, a random sample is generated and a sample variance function, as a function of θ , is defined with respect to the generated sample. The sample variance function can then be viewed as a deterministic function of θ , and the optimal value of θ that minimizes this sample variance function is determined using a nonlinear optimization solver. Then one makes a “production run” using the value of θ returned from the optimization step. The initial optimization can be computationally expensive when compared with one step of the stochastic approximation procedure. However, sample average approximation does not require tuning parameters beyond the choice of runlength, and for very long simulation runs, a vanishingly small fraction of the effort is required in the initial optimization.

Henderson and Simon [2004] also develop an adaptive control variate method for finite-horizon simulations. They give conditions under which adaptive control variate estimators converge at an exponential rate. One of the key assumptions there is the existence of a “perfect” control variate, i.e., a parameter value θ^* such that $\text{var}(X - Y(\theta^*)) = 0$. For the applications we have in mind this assumption is unlikely to hold. Bolia and Juneja [2005] use the martingale control variates developed in Henderson and Glynn [2002], as we do, but in the case of linearly parameterized controls. Maire [2003] expresses the estimation problem as an integration problem over the unit hypercube, and uses the expansion of the integrand for an approximate orthonormal basis as a control variate. An iterative procedure estimates the coefficients of the expansion so that the variance of each estimated coefficient has a polynomial decay. The residual terms are not estimated iteratively so, in general, the convergence rate of the procedure cannot exceed the canonical rate. Henderson et al. [2003] develop adaptive control variate schemes for Markov chains in the steady-state setting. They use a stochastic approximation procedure for tuning control variate estimators developed in Tadić and Meyn [2004] and provide conditions for minimization of an approximation of the steady-state variance.

We view the primary contributions of this paper as follows. We demonstrate how nonlinearly parameterized control variates can arise naturally in Markov chain simulations. We develop 2 algorithms

that are designed to identify a variance-minimizing parameterization. We show how the algorithms should balance the effort spent in searching for good parameterizations and the effort spent in conducting repeated simulation runs at a fixed parameter selection. We give proofs of consistency and central limit theorems. Finally, we discuss implementation issues regarding the practical use of these adaptive methods.

This paper is organized as follows. Section 2 sketches some of the main ideas in pricing barrier options using adaptive control variates, as motivation for the remainder of the paper. Section 3 outlines the general problem and discusses estimation of the gradient of the variance function. Section 4 explores the use of stochastic approximation and Section 5 explores sample average approximation. In Section 6 we describe the results of some limited experiments with the example of Section 2. Finally, Section 7 contains some concluding remarks.

Unless otherwise stated, all vectors are column vectors and all norms are Euclidean. Suffixes can either indicate different instances of a random vector or components of a single vector, with the context clarifying what is intended.

2 A Motivating Example

In this section we describe the problem of pricing barrier options, and explain how parameterized controls can be found. Our goal in this section and Section 6 is not to develop the most efficient known estimators for pricing barrier options, but rather to demonstrate the adaptive control variate methodology in a familiar, but nontrivial, setting, and bring out some of the practical issues involved in applications.

A barrier option is a derivative security that is either activated (knocked-in) or extinguished (knocked-out) when the price of the underlying asset reaches a certain level (barrier) at any time during the lifetime of the option; see, e.g., Glasserman [2004].

The price of the underlying stock at time t is denoted by $S(t)$, for $t \geq 0$. Suppose that the underlying stock price is monitored at discrete times $t_i = i\Delta t, i = 0, 1, 2, \dots, l$, where T is the (deterministic) expiration date of the option and $\Delta t = T/l$ is the time between consecutive monitoring dates. For notational convenience, let S_i denote the underlying stock price at the i th monitoring point (i.e., $S(t_i)$). Assume that the initial stock price S_0 takes a value in an interval H and the barrier is the boundary of H . When the stock price crosses the barrier, the option is knocked out and the payoff is zero. If the option has not been knocked out by time T , then the payoff at time T is $(S_l - K)^+$, where $K > 0$ is the strike price. Hence, the option payoff depends on the complete path $\{S_i, i = 0, \dots, l\}$. Define

$$\begin{aligned}\tau &= \inf\{n \geq 0 : S_n \notin H\} \text{ and} \\ A_i &= 1_{\{\tau > i\}}, i = 0, \dots, l.\end{aligned}$$

Then A_i is the indicator that determines whether the option is alive at time t_i or not. We assume that the market is arbitrage free. Then the price of a knock-out call option is given by

$$e^{-rT} E[A_l(S_l - K)^+],$$

where r is the (assumed constant) risk-free interest rate and the expectation is taken under the risk-neutral measure. Since the discount factor e^{-rT} is constant, pricing the option reduces to estimating the expected payoff with the initial stock price x , i.e., estimating

$$E[A_l(S_l - K)^+ | S_0 = x].$$

Assume that the underlying stock price process $\{S(t) : t \geq 0\}$ is a (time homogeneous) Markov process. Then $\{S_n : n = 0, 1, 2, \dots\}$, where S_n is the stock price at time $t_n = n\Delta t$, is a discrete time Markov chain on the state space $[0, \infty)$. For $i = 0, 1, \dots$, define

$$U^*(x, i) = \begin{cases} E[A_i(S_i - K)^+ | S_0 = x], & \text{if } x \in H, \text{ and} \\ 0 & \text{if } x = 0 \text{ or } x \notin H, \end{cases}$$

so that $U^*(x, i)$ is the expected payoff of the option with the initial stock price x and maturity t_i . Our goal is to estimate $U^*(x, l)$.

We now describe the martingale that serves as a control variate, drawing from the general results of Henderson and Glynn [2002, Section 4]. Let $\tilde{S}_i = S_i A_i$, for $i \geq 0$. Then $\{\tilde{S}_n : n \geq 0\}$ is a Markov process on the state space $\mathcal{S} = H \cup \{0\}$ (assuming that $S_0 \in H \cup \{0\}$). For a real-valued function $f : \mathcal{S} \rightarrow \mathbb{R}$, let $P(x, \cdot)f(\cdot) = E[f(\tilde{S}_1) | \tilde{S}_0 = x]$, provided that the expectation exists. Let $U : \mathcal{S} \times \{0, 1, \dots, l-1\} \rightarrow \mathbb{R}$ be a real-valued function with $U(0, \cdot) = 0$ and for $1 \leq n \leq l$ let

$$M_n(U) = \sum_{i=1}^n [U(\tilde{S}_i, l-i) - P(\tilde{S}_{i-1}, \cdot)U(\cdot, l-i)],$$

provided that the conditional expectations in this expression are finite. Then it is straightforward to show that $(M_n(U) : 1 \leq n \leq l)$ is a martingale and $E_x(M_l(U)) = 0$ for any U , provided that the usual integrability conditions hold, where E_x denotes expectation under the initial condition $\tilde{S}_0 = x$. Therefore, $U^*(x, l)$ can be estimated via i.i.d. replications of

$$(\tilde{S}_l - K)^+ - M_l(U), \tag{1}$$

with $\tilde{S}_0 = x$, where $M_l(U)$ serves as a control variate.

But how should we select the function U ? Our notation suggests that $U = U^*$ would be a good choice, and this is indeed the case. To see why, note that for all $x \in \mathcal{S}$ and $i > 0$,

$$\begin{aligned} U^*(x, i) &= E[A_i(S_i - K)^+ | S_0 = x] \\ &= E[(\tilde{S}_i - K)^+ | \tilde{S}_0 = x], \\ &= E[E[(\tilde{S}_i - K)^+ | \tilde{S}_1, \tilde{S}_0 = x] | \tilde{S}_0 = x] \\ &= E[U^*(\tilde{S}_1, i-1) | \tilde{S}_0 = x] \\ &= \int_{\mathcal{S}} U^*(y, i-1) P(x, dy) \\ &= P(x, \cdot)U^*(\cdot, i-1), \end{aligned}$$

where P is the transition probability kernel of $\{\tilde{S}_n : n \geq 0\}$. It follows that

$$\begin{aligned} M_l(U^*) &= \sum_{i=1}^l [U^*(\tilde{S}_i, l-i) - U^*(\tilde{S}_{i-1}, l-(i-1))] \\ &= U^*(\tilde{S}_l, 0) - U^*(\tilde{S}_0, l) \\ &= (\tilde{S}_l - K)^+ - U^*(x, l). \end{aligned}$$

Hence, if $U = U^*$, then the estimator (1) of $E[A_l(S_l - K)^+ | S_0 = x]$ has zero variance.

So it is desirable that $U \approx U^*$. Suppose that $U(x, i) = U(x, i; \theta)$, where $\theta \in \Theta \subseteq \mathbb{R}^p$ is a p -dimensional vector of parameters.

Remark 1 In our general notational scheme, X is the payoff $(\tilde{S}_l - K)^+$ at time $T = l\Delta t$, EX is the expected payoff $U^*(x, l)$, and $Y(\theta)$ is $M_l(U(\cdot, \cdot; \theta))$.

A linear parameterization arises if

$$U(x, i; \theta) = \sum_{k=1}^p \theta_k U_k(x, i),$$

where $U_k(\cdot, \cdot)$ are given basis functions, $k = 1, \dots, p$. In this case, for $1 \leq n \leq l$,

$$\begin{aligned} M_n(U) &= \sum_{i=1}^n \left[\sum_{k=1}^p \theta_k U_k(\tilde{S}_i, l-i) - P(\tilde{S}_{i-1}, \cdot) \sum_{k=1}^p \theta_k U_k(\cdot, l-i) \right] \\ &= \sum_{k=1}^p \theta_k \left[\sum_{i=1}^n U_k(\tilde{S}_i, l-i) - P(\tilde{S}_{i-1}, \cdot) U_k(\cdot, l-i) \right] \\ &= \sum_{k=1}^p \theta_k M_n(U_k), \end{aligned} \tag{2}$$

so that the control $M_n(U)$ is simply a linear combination of martingales corresponding to the basis functions U_k , $k = 1, \dots, p$. In this sense, the linearly parameterized case leads us back to the theory of linear control variates. Notice that recomputing the control for a new value of θ is straightforward – one simply reweights the previous values of the martingales corresponding to the basis functions. See Kim and Henderson [2004a] for further discussion.

The situation is more complicated when $U(x; \theta)$ arises from a nonlinear parameterization. An example of such a parameterization with $p = 4$ is given by

$$U(x, i; \theta) = \theta_1 x^{\theta_2} + \theta_3 x + \theta_4.$$

Now $Y(\theta)$ is a nonlinear function of a random object Y (the path $(\tilde{S}_i : 0 \leq i \leq l)$) and a parameter vector θ . It is difficult to recompute the value of $X - Y(\theta)$ when θ changes. Essentially one needs to store the sample path of the chain, explicitly or implicitly, in order to be able to do this.

For nonlinear parameterizations, we need a method for selecting a good choice of θ . This is the subject of the next 3 sections. We will return to the barrier option pricing example in Section 6.

3 The Nonlinear Case: Preliminaries

Suppose that $Y(\theta) = h(Y, \theta)$ is a nonlinear function of a random element Y and a parameter vector $\theta \in \Theta \subset \mathbb{R}^p$. Let H denote the support of the probability distribution of (X, Y) , i.e., H is the smallest closed set such that $P((X, Y) \in H) = 1$. Let H_2 be the set

$$\{y : \exists x \text{ with } (x, y) \in H\},$$

i.e., the set of y values that appear in H .

Assumption A1 The parameter set Θ is compact. For all $y \in H_2$, the real-valued function $h(y, \cdot)$ is \mathcal{C}^1 (i.e., continuously differentiable) on \mathcal{U} , where \mathcal{U} is a bounded open set containing Θ .

Assumption A2 The random variable X is square integrable. Also, for all $\theta \in \mathcal{U}$, $EY^2(\theta) < \infty$ and $EY(\theta) = Eh(Y, \theta) = 0$.

For convenience we define $X(\theta) = X - Y(\theta)$. Define

$$v(\theta) = \text{var } X(\theta) = \text{var}(X - Y(\theta))$$

to be the variance of the estimator as a function of θ . As before our overall goal is to estimate EX . Our *intermediate* goal is to identify θ^* which minimizes $v(\theta)$ over $\theta \in \Theta$. In general we cannot expect to find a closed form expression for θ^* as in the linear case, and so we approach this problem from the point of view of stochastic optimization. Regardless of which stochastic optimization method we adopt, we need to impose some structure in order to make progress. We now develop some machinery that will allow us to conclude that $v(\cdot)$ is differentiable.

Assumption A3 For all $y \in H_2$, $h(y, \cdot)$ is Lipschitz on \mathcal{U} , i.e., there exists $C(y) > 0$ such that for all $\theta_1, \theta_2 \in \mathcal{U}$,

$$|h(y, \theta_1) - h(y, \theta_2)| \leq C(y) \|\theta_1 - \theta_2\|,$$

where $\|\cdot\|$ is a metric on \mathbb{R}^p . Therefore,

$$\sup_{\theta \in \mathcal{U}} \left| \frac{\partial h(y, \theta)}{\partial \theta_i} \right| \leq C(y)$$

for all $y \in H_2$ and $i = 1, \dots, p$.

Remark 2 Recall that a \mathcal{C}^1 function is Lipschitz on a compact set. If $h(y, \cdot)$ is \mathcal{C}^1 on \mathbb{R}^p (or on an open set containing the closure of \mathcal{U}), then **A3** is immediate.

To establish the required differentiability, and to obtain an estimator of the gradient of v , we use Infinitesimal Perturbation Analysis (IPA). Let $f(\theta) = Ef(\theta, \xi)$ for some random variable ξ whose distribution does not depend on θ . Recall that the idea in IPA is to take $\nabla_{\theta} f(\theta, \xi)$, the gradient of $f(\theta, \xi)$ for fixed ξ , as an estimate of $\nabla_{\theta} f(\theta)$. If $\nabla_{\theta} f(\theta, \xi)$ is uniformly dominated by an integrable function of ξ , then the gradient and expectation operators can be exchanged and this yields an unbiased estimator. See L'Ecuyer [1995] for detailed sufficient conditions for the interchange to be valid.

An unbiased gradient estimator can be obtained by noting that the sample variance of i.i.d. observations is an unbiased estimator of the variance, so that under **A2**, and for any $m \geq 2$,

$$v(\theta) = EV(m, \theta) := E \frac{1}{m-1} \sum_{i=1}^m (X_i(\theta) - \bar{X}_m(\theta))^2 = E \frac{m}{m-1} \left(\frac{1}{m} \sum_{i=1}^m X_i^2(\theta) - \bar{X}_m^2(\theta) \right), \quad (3)$$

where $(X_1, Y_1), \dots, (X_m, Y_m)$ are i.i.d. replications of (X, Y) and

$$\bar{X}_m(\theta) = \frac{1}{m} \sum_{j=1}^m X_j(\theta),$$

for all $\theta \in \mathcal{U}$. (We include the terms $h(Y_j, \theta)$ in the sample average $\bar{X}_m(\theta)$ even though we know that they have zero mean, because they reduce variance.) Assumption **A1** implies that for each $(x, y) \in H$, $x - h(y, \cdot)$ is a \mathcal{C}^1 function on \mathcal{U} . This provides the pathwise differentiability of $V(m, \theta)$ on \mathcal{U} . We also need some integrability conditions.

Assumption A4 $E \left(C(Y) \left[1 + \sup_{\theta \in \mathcal{U}} |X(\theta)| \right] \right) < \infty$, where $C(Y)$ appears in **A3**.

We can construct an unbiased gradient estimator from (3) as

$$\begin{aligned} g_m(\theta_0) &= \nabla V(m, \theta_0) \\ &= \frac{1}{m-1} \sum_{i=1}^m \nabla_{\theta} (X_i(\theta) - \bar{X}_m(\theta))^2 \Big|_{\theta=\theta_0} \\ &= \frac{-2}{m-1} \sum_{i=1}^m (X_i(\theta) - \bar{X}_m(\theta)) \nabla_{\theta} \left(h(Y_i, \theta) - \frac{1}{m} \sum_{j=1}^m h(Y_j, \theta) \right) \Big|_{\theta=\theta_0}. \end{aligned}$$

Proposition 1 *If **A1** - **A4** hold then $v(\cdot)$ is \mathcal{C}^1 on \mathcal{U} and for $\theta_0 \in \mathcal{U}$,*

$$\begin{aligned} g(\theta_0) &:= \nabla_{\theta} v(\theta) \Big|_{\theta=\theta_0} \\ &= E g_m(\theta_0) \end{aligned} \quad (4)$$

Proof. We apply Proposition 1 in L'Ecuyer [1995] to the sample variance $V(m, \theta)$ component by component. Consider the j th component, for some $j = 1, \dots, p$. The only condition that requires explicit verification is that $\partial V(m, \theta) / \partial \theta_j$ is dominated by an integrable function of $(\mathbf{X}, \mathbf{Y}) = ((X_i, Y_i) : 1 \leq i \leq m)$. We have that

$$\frac{\partial V(m, \theta)}{\partial \theta_j} = \frac{m}{m-1} \left(\frac{-1}{m} \sum_{i=1}^m 2X_i(\theta) \frac{\partial h(Y_i, \theta)}{\partial \theta_j} + 2\bar{X}_m(\theta) \frac{1}{m} \sum_{i=1}^m \frac{\partial h(Y_i, \theta)}{\partial \theta_j} \right). \quad (5)$$

The first term in the parentheses in (5) is integrable by **A4**. For the second term, we apply **A3** and split the sums to obtain

$$\begin{aligned} & \left| \bar{X}_m(\theta) \frac{1}{m} \sum_{i=1}^m \frac{\partial h(Y_i, \theta)}{\partial \theta_j} \right| \\ & \leq \frac{1}{m^2} \sum_{i=1}^m \sup_{\theta \in \mathcal{U}} |X_i(\theta)| C(Y_i) + \frac{1}{m^2} \sum_{i=1}^m \sum_{j \neq i} \sup_{\theta \in \mathcal{U}} |X_i(\theta)| C(Y_j). \end{aligned} \quad (6)$$

If $E \sup_{\theta \in \mathcal{U}} |X_i(\theta)|$ is finite then **A4** implies integrability of this bound and the proof will be complete. Fix $\theta_0 \in \mathcal{U}$. By **A3**,

$$\begin{aligned} |X_1(\theta)| &\leq |X_1| + |h(Y_1, \theta)| \\ &\leq |X_1| + |h(Y_1, \theta_0)| + |h(Y_1, \theta) - h(Y_1, \theta_0)| \\ &\leq |X_1| + |h(Y_1, \theta_0)| + C(Y_1) \|\theta - \theta_0\|. \end{aligned}$$

But $\|\theta - \theta_0\|$ is bounded on the bounded set \mathcal{U} , and so $\sup_{\theta \in \mathcal{U}} |X_1(\theta)|$ is integrable. ■

So under the assumptions **A1** - **A4**, the variance function $v(\theta)$ is continuously differentiable in $\theta \in \mathcal{U}$, and we have an IPA-based unbiased gradient estimator at our disposal. We are now equipped to attempt to minimize $v(\theta)$ over $\theta \in \Theta$.

4 Stochastic Approximation

Stochastic approximation is a class of stochastic optimization methods used to solve problems with differentiable objective functions. See Spall [2003] for an accessible introduction and references. In the presence of nonconvexity the algorithm may only converge to a local minimum. The general form of the algorithm is a recursion where an approximation θ_n for the optimal solution is updated to θ_{n+1} using an estimator $g_n(\theta_n)$ of the gradient $g(\theta_n)$ of the objective function evaluated at θ_n . For a minimization problem, the recursion is of the form

$$\theta_{n+1} = \Pi_{\Theta}(\theta_n - a_n g_n(\theta_n)), \quad (7)$$

where Π_{Θ} denotes a projection of points outside Θ back into Θ , and $\{a_n\}$ is a sequence of positive real numbers such that

$$\sum_{n=1}^{\infty} a_n = \infty \text{ and } \sum_{n=1}^{\infty} a_n^2 < \infty. \quad (8)$$

We use IPA to obtain $g_n(\theta_n)$, as discussed in the previous section.

Our stochastic approximation algorithm for finding θ^* and estimating EX is given in Figure 1. Let $m \geq 2$ be a fixed positive integer.

Initialization: Choose θ_0 .

For $k = 1$ to n

 Generate the i.i.d. sample $(X_{k,i}, Y_{k,i}) \sim (X, Y)$, $i = 1, \dots, m$, independent of all else.

 Compute

$$\begin{aligned} A_k(\theta_{k-1}) &= \frac{1}{m} \sum_{i=1}^m [X_{k,i} - h(Y_{k,i}, \theta_{k-1})], \\ g_{k-1}(\theta_{k-1}) &= \frac{-2}{m-1} \sum_{i=1}^m [X_{k,i} - h(Y_{k,i}, \theta_{k-1}) - A_k(\theta_{k-1})] \\ &\quad \left. \nabla_{\theta} \left[h(Y_{k,i}, \theta) - \frac{1}{m} \sum_{j=1}^m h(Y_{k,j}, \theta) \right] \right|_{\theta=\theta_{k-1}} \\ \theta_k &= \Pi_{\Theta}(\theta_{k-1} - a_{k-1} g_{k-1}(\theta_{k-1})). \end{aligned}$$

Next k

Set $\mu_n = n^{-1} \sum_{k=1}^n A_k(\theta_{k-1})$.

Figure 1: The stochastic approximation algorithm

We first show consistency of the estimator μ_n of $\mu = EX$. We apply the following martingale strong law of large numbers which can be found in Liptser and Shiriyayev [1989, p. 144]. Let $(\mathcal{F}_n : n \geq 0)$ be a filtration, i.e. an increasing sequence of σ -fields.

Theorem 2 (Liptser and Shiriyayev 1989) *Let $(M_n, \mathcal{F}_n : n \geq 0)$ be a square-integrable martingale with $M_0 = 0$. Let $(L_n : n \geq 0)$ be nondecreasing in n with $L_n \in \mathcal{F}_n$ for all n . Define*

$$V_n = \sum_{k=1}^n E((M_k - M_{k-1})^2 | \mathcal{F}_{k-1})$$

and assume that

$$\sum_{n=1}^{\infty} \frac{V_{n+1} - V_n}{(1 + L_n)^2} < \infty \text{ a.s. and } P(L_\infty = \infty) = 1,$$

where $L_\infty = \lim_{n \rightarrow \infty} L_n$. Then

$$\frac{M_n}{L_n} \rightarrow 0 \text{ a.s.}$$

Let $\mathcal{F}_n = \sigma\{(X_{k,i}, Y_{k,i}) : 1 \leq k \leq n, 1 \leq i \leq m\}$ be the sigma field containing the information from the first n steps of the stochastic approximation algorithm. Let \mathcal{F}_0 be the trivial sigma field and θ_0 be any deterministic guess for θ^* . (If θ_0 is not deterministic then we can extend \mathcal{F}_0 appropriately, so there is no loss of generality in this convention.)

Proposition 3 *Assume A1-A4. Then $\mu_n \rightarrow \mu$ a.s. as $n \rightarrow \infty$.*

Proof. For $k \geq 1$ and $n \geq 1$, define

$$\begin{aligned} \zeta_k(\theta_{k-1}) &= A_k(\theta_{k-1}) - \mu \text{ and} \\ M_n &= \sum_{k=1}^n \zeta_k(\theta_{k-1}). \end{aligned}$$

Then

$$\mu_n = \mu + \frac{M_n}{n},$$

and hence it suffices to show that $M_n/n \rightarrow 0$ a.s. as $n \rightarrow \infty$.

Define $M_0 = 0$. Since $E(\zeta_k(\theta_{k-1})|\mathcal{F}_{k-1}) = 0$ for all $k \geq 1$, $(M_n, \mathcal{F}_n : n \geq 0)$ is a martingale. Moreover, for all $n \geq 1$,

$$\begin{aligned} E(M_n^2) &= \sum_{k=1}^n \text{var}(A_k(\theta_{k-1})) \\ &= \sum_{k=1}^n \frac{1}{m} E(v(\theta_{k-1})) < \infty, \end{aligned}$$

where the finiteness follows from the fact that $v(\cdot)$ is continuous on the compact set Θ and therefore bounded. Define $L_n = n$ for all $n \geq 0$ and

$$V_n = \sum_{k=1}^n E((M_k - M_{k-1})^2 | \mathcal{F}_{k-1}) = \sum_{k=1}^n E(\zeta_k^2(\theta_{k-1}) | \mathcal{F}_{k-1}) = \frac{1}{m} \sum_{k=1}^n v(\theta_{k-1}).$$

Then $P(L_\infty = \infty) = 1$ and

$$\sum_{n=1}^{\infty} \frac{V_{n+1} - V_n}{(1 + L_n)^2} = \frac{1}{m} \sum_{n=1}^{\infty} \frac{v(\theta_n)}{(1 + n)^2} \leq \frac{\sup_{\theta \in \Theta} v(\theta)}{m} \sum_{n=1}^{\infty} \frac{1}{(1 + n)^2} < \infty \text{ a.s.}$$

Therefore, by Theorem 2, $M_n/n \rightarrow 0$ a.s. as $n \rightarrow \infty$. ■

Remark 3 The proof of Proposition 3 is based on the square integrability of $X_1(\cdot)$ and the continuity of $v(\cdot)$ on Θ . The square-integrability condition may seem too strong. But if $\theta_k \rightarrow \theta^*$ a.s. as $k \rightarrow \infty$ for some random variable θ^* that takes on countably many values, then under the Lipschitz continuity of $h(y, \cdot)$ and finite first moment conditions, μ_n is still strongly consistent.

We now assess the rate of convergence of μ_n through a central limit theorem. We use the following martingale central limit theorem which can be found in Liptser and Shiriyayev [1989, p. 444]. A martingale difference sequence $(\xi_{k,n}, \mathcal{F}_{k,n} : n \geq 1, 1 \leq k \leq n)$ is a collection of mean-zero random variables $\xi_{k,n}$ and filtrations $(\mathcal{F}_{k,n} : k = 1, \dots, n)$ such that $\xi_{k,n}$ is measurable with respect to $\mathcal{F}_{k,n}$ for all $n \geq 1$ and $1 \leq k \leq n$, and $E(\xi_{k,n} | \mathcal{F}_{k-1,n}) = 0$ for all $n \geq 1$ and $k = 1, \dots, n$. Here we have adopted the convention that $\mathcal{F}_{0,n}$ is the trivial sigma field for all $n \geq 1$, so that θ_0 is a deterministic approximation for θ^* .

Theorem 4 (Liptser and Shiriyayev 1989) Assume that $(\mathcal{F}_{k,n} : 1 \leq k \leq n, n \geq 1)$ is nested i.e. $\mathcal{F}_{k,n} \subseteq \mathcal{F}_{k,n+1}$, for all $k \leq n, n \geq 1$. Let η^2 be a \mathcal{G} -measurable random variable where

$$\mathcal{G} \subseteq \sigma(\cup_{n \geq 1} \mathcal{F}_{n,n}).$$

Let Z be a random variable with characteristic function

$$E(e^{itZ}) = E \exp\left(-\frac{t^2}{2}\eta^2\right), t \in \mathbb{R},$$

so that Z is a mixture of mean-zero normal random variables. Let $(\xi_{k,n}, \mathcal{F}_{k,n} : n \geq 1, 1 \leq k \leq n)$ be a martingale difference sequence with $E(\xi_{k,n}^2) < \infty$, for all $n \geq 1, 1 \leq k \leq n$. Assume that

- (i) $\sum_{k=1}^n E(\xi_{k,n}^2 I(|\xi_{k,n}| > \delta) | \mathcal{F}_{k-1,n}) \rightarrow 0$ in probability, for all $\delta \in (0, 1]$,
- (ii) $\sum_{k=1}^n E(\xi_{k,n}^2 | \mathcal{F}_{k-1,n}) \rightarrow \eta^2$ in probability, and
- (iii) $\sum_{k=1}^{\lfloor nc_n \rfloor} E(\xi_{k,n}^2 | \mathcal{F}_{k-1,n}) \rightarrow 0$ in probability

for a certain sequence $(c_n)_{n \geq 1}$ with $c_n \downarrow 0, nc_n \rightarrow \infty$ as $n \rightarrow \infty$. Then

$$S_n = \sum_{k=1}^n \xi_{k,n} \Rightarrow Z$$

as $n \rightarrow \infty$, where \Rightarrow denotes convergence in distribution.

The central limit theorem below assumes that θ_n converges to some random variable θ^* a.s. Establishing this result requires some care, so we state our main results assuming that this convergence holds and then give sufficient conditions for the convergence of θ_n . The theory does not require that θ^* be a minimizer of $v(\theta)$ over Θ although we would certainly prefer this to be the case. Before stating the central limit theorem we need another assumption. Let

$$E = \{\omega : \theta_k(\omega) \rightarrow \theta^*(\omega) \text{ as } k \rightarrow \infty\}$$

so that $P(E) = 1$ and let

$$\Gamma = \{\theta^*(\omega) = \lim_{k \rightarrow \infty} \theta_k(\omega) : \omega \in E\} \subseteq \Theta$$

be the set of limiting values of θ_k .

Assumption A5 For any $\gamma \in \Gamma$, there is a neighbourhood $\mathcal{N}(\gamma)$ of γ such that the collection $\{X^2(\theta) : \theta \in \mathcal{N}(\gamma)\}$ is uniformly integrable.

Remark 4 A set of sufficient conditions for **A5** is **A1-A3** and $EK^2(Y) < \infty$.

Theorem 5 Assume **A1-A5** and that $\theta_n \rightarrow \theta^*$ for some random variable θ^* a.s. as $n \rightarrow \infty$. Let Z be a random variable with characteristic function

$$E(e^{itZ}) = E \exp\left(-\frac{t^2}{2}v(\theta^*)\right), t \in \mathbb{R},$$

i.e., $Z = v^{1/2}(\theta^*)N(0, 1)$ is a mixture of mean-zero normal random variables. Then

$$\sqrt{mn}(\mu_n - \mu) \Rightarrow Z$$

as $n \rightarrow \infty$.

Proof. To show the central limit theorem we apply Theorem 4. Let

$$\xi_{k,n} = \frac{\sqrt{m}(A_k(\theta_{k-1}) - \mu)}{\sqrt{n}}$$

so that

$$\sqrt{mn}(\mu_n - \mu) = \sum_{k=1}^n \xi_{k,n}.$$

As in Proposition 3, $(\xi_{k,n}, \mathcal{F}_{k,n} : n \geq 1, 1 \leq k \leq n)$ is a martingale difference sequence with $E\xi_{k,n}^2 = Ev(\theta_{k-1})/n < \infty$, where $\mathcal{F}_{k,n} = \mathcal{F}_k$ for all n . Fix $\delta > 0$ and let

$$W_n = \sum_{k=1}^n E(\xi_{k,n}^2 I(|\xi_{k,n}| > \delta) | \mathcal{F}_{k-1,n}).$$

If $\zeta_k(\theta_{k-1}) = A_k(\theta_{k-1}) - \mu$, then

$$\begin{aligned} W_n &= \frac{m}{n} \sum_{k=1}^n E[\zeta_k^2(\theta_{k-1}) I(\zeta_k^2(\theta_{k-1}) > n\delta^2/m) | \mathcal{F}_{k-1,n}] \\ &= \frac{m}{n} \sum_{k=1}^n E[\zeta_k^2(\theta_{k-1}) I(\zeta_k^2(\theta_{k-1}) > n\delta^2/m) | \theta_{k-1}]. \end{aligned}$$

For any $\theta \in \Theta$, let $\zeta(\theta) = \frac{1}{m} \sum_{j=1}^m (X_j - h(Y_j, \theta) - \mu)$, where $(X_1, Y_1), \dots, (X_m, Y_m)$ are i.i.d. replications of (X, Y) , independent of $(X_{k,i}, Y_{k,i})$, $i = 1, \dots, m$, $k \geq 1$. Then

$$W_n = \frac{m}{n} \sum_{k=1}^n f(\theta_{k-1}, n\delta^2/m),$$

where

$$f(\theta, b) = E[\zeta^2(\theta) I(\zeta^2(\theta) > b)].$$

Let $\omega \in E$ be fixed, and let $\gamma = \theta^*(\omega)$. Assumption **A5** ensures that the collection $(\zeta^2(\theta) : \theta \in \mathcal{N}(\gamma))$ is uniformly integrable and so for all $\epsilon > 0$, there exists $K_\epsilon > 0$ such that $f(\theta, K_\epsilon) \leq \epsilon$ for all $\theta \in \mathcal{N}(\gamma)$. Fix $\epsilon > 0$. Let $n_1 = n_1(\omega) \geq 1$ be such that $\theta_n(\omega) \in \mathcal{N}(\gamma)$ for all $n \geq n_1$ and let $n_2 \geq 1$ be such that $n\delta^2/m \geq K_\epsilon$ for all $n \geq n_2$. Let $n^* = \max\{n_1, n_2\} + 1$. Then

$$\begin{aligned} W_n &= \frac{m}{n} \sum_{k=1}^n f(\theta_{k-1}, n\delta^2/m) \\ &= \frac{m}{n} \sum_{k=1}^{n^*} f(\theta_{k-1}, n\delta^2/m) + \frac{m}{n} \sum_{k=n^*+1}^n f(\theta_{k-1}, n\delta^2/m) \\ &\leq \frac{m}{n} \sum_{k=1}^{n^*} f(\theta_{k-1}, 0) + \frac{m}{n} \sum_{k=n^*+1}^n f(\theta_{k-1}, K_\epsilon). \end{aligned}$$

Hence

$$0 \leq \limsup_{n \rightarrow \infty} W_n \leq 0 + \limsup_{n \rightarrow \infty} \frac{m}{n} \sum_{k=n^*+1}^n \epsilon = m\epsilon.$$

Since ϵ and $\omega \in E$ were arbitrary, we conclude that $W_n \rightarrow 0$ as $n \rightarrow \infty$ a.s.

The second and third conditions of Theorem 4 are easily dealt with. We see that

$$\sum_{k=1}^n E(\xi_{k,n}^2 | \mathcal{F}_{k-1}) = \sum_{k=1}^n \frac{m}{n} E((A_k(\theta_{k-1}) - \mu)^2 | \mathcal{F}_{k-1}) = \frac{1}{n} \sum_{k=1}^n v(\theta_{k-1}) \rightarrow v(\theta^*)$$

as $n \rightarrow \infty$ a.s., since $\{\theta_k\}$ converges a.s., and v is continuous. For the third condition, let $c_n = n^{-1/2}$. Then

$$\sum_{k=1}^{\lfloor nc_n \rfloor} E(\xi_{k,n}^2 | \mathcal{F}_{k-1}) = \frac{1}{n} \sum_{k=1}^{\lfloor n^{1/2} \rfloor} v(\theta_{k-1}) \leq \frac{n^{1/2} \sup_{\theta \in \Theta} v(\theta)}{n} \rightarrow 0$$

as $n \rightarrow \infty$. The central limit theorem is therefore a consequence of Theorem 4. ■

Hence we see that the stochastic approximation estimator μ_n satisfies a strong law and central limit theorem as $n \rightarrow \infty$. In general, it will be the case that the optimal variance $v(\theta^*)$ is positive a.s. so that the rate of convergence of μ_n is the canonical $n^{-1/2}$ rate. This is the best that can be hoped for with the Monte Carlo nature of the estimation procedure we used.

The central limit theorem above still holds when $v(\theta^*) = 0$ a.s. In this case, $\sqrt{n}(\mu_n - \mu) \Rightarrow 0$ as $n \rightarrow \infty$ so the rate of convergence is faster than $n^{-1/2}$. The actual rate of convergence depends on the rate at which $\theta_n \rightarrow \theta^*$ a.s. We do not explore this case further here. See Henderson and Simon [2004] for an exploration of increased convergence rates when θ^* is deterministic and $v(\theta^*) = 0$.

Recall that we chose $m \geq 2$ to obtain an unbiased gradient estimator. The theorem above shows that the estimator μ_n converges at the rate $(mn)^{-1/2}$, so the additional averaging of m terms in each step of the algorithm does not slow down the rate of convergence, at least to first order. Therefore the choice of $m \geq 2$ is not significant from a large sample theory point of view. In small samples, there may be some benefit to carefully choosing m . We do not explore that possibility here.

The central limit theorem suggests a confidence interval procedure, provided that the variance can be estimated. Suppose that $\theta_k \rightarrow \theta^*$ a.s. for some deterministic $\theta^* \in \Theta$, so that the variance appearing in the central limit theorem is deterministic and equal to $v(\theta^*)$. Several estimators immediately suggest themselves. One could use (a) the sample variance of all mn samples, (b) the average of the sample

variances of the m samples in each iteration, or (c) the sample variance of the n averages. Kim and Henderson [2004b] show that each of these estimators is consistent and satisfies a central limit theorem (under appropriate conditions), and compares the estimators based on their limiting variance. The conclusion is that estimators (a) and (b) behave almost identically when m is large, but when m is small, (a) is slightly preferred to (b). In view of these results, we used (b) because it seems easier to compute than (a). More precisely, our variance estimator is

$$S_n^2 = \frac{1}{n} \sum_{k=1}^n \left(\frac{1}{m-1} \sum_{i=1}^m (X_{k,i}(\theta_{k-1}) - A_k(\theta_{k-1}))^2 \right). \quad (9)$$

We now give conditions under which θ_n converges to some random variable θ^* a.s. as $n \rightarrow \infty$. Theorem 6 below is an immediate specialization of Kushner and Yin [2003, Theorem 2.1, p. 127]. We first need some definitions.

A box $B \subset \mathbb{R}^p$ is a set of the form

$$B = \{x \in \mathbb{R}^p : a(i) \leq x(i) \leq b(i), i = 1, \dots, p\},$$

where $a(i), b(i) \in \mathbb{R}$ and $a(i) \leq b(i), i = 1, \dots, p$. For $x \in B$ define the set $\mathcal{C}(x)$ as follows. For x in the interior of B , $\mathcal{C}(x) = \{0\}$. For x on the boundary of B , $\mathcal{C}(x)$ is the convex cone generated by the outward normals of the faces on which x lies. A *first-order critical point* x of a \mathcal{C}^1 function $f : B \rightarrow \mathbb{R}$ satisfies

$$-\nabla f(x) = z \text{ for some } z \in \mathcal{C}(x).$$

A first-order critical point is either a point where the gradient $\nabla f(x)$ is zero, or a point on the boundary of B where the gradient “points towards the interior of B ”. Let $S(f, B)$ be the set of first-order critical points of f in B . We define the distance from a point x to a set S to be

$$d(x, S) = \inf_{y \in S} \|x - y\|.$$

The projection $y = \Pi_B x$ is a pointwise projection defined by

$$y(i) = \begin{cases} a(i) & \text{if } x(i) < a(i), \\ x(i) & \text{if } a(i) \leq x(i) \leq b(i), \text{ and} \\ b(i) & \text{if } b(i) < x(i), \end{cases}$$

for each $i = 1, \dots, p$.

Let $(\mathcal{G}_n : n \geq 0)$ be a filtration, where the initial guess θ_0 is measurable with respect to \mathcal{G}_0 , and G_n (an estimate for the gradient of f at θ_n) is measurable with respect to \mathcal{G}_{n+1} for all $n \geq 0$.

Theorem 6 *Let B be a box in \mathbb{R}^p and $f : \mathbb{R}^p \rightarrow \mathbb{R}$ be \mathcal{C}^1 . Suppose that for $n \geq 0$, $\theta_{n+1} = \Pi_B(\theta_n - a_n G_n)$ with the following additional conditions.*

- (i) *The conditions (8) hold.*
- (ii) $\sup_n E \|G_n\|^2 < \infty$.
- (iii) $E[G_n | \mathcal{G}_n] = \nabla f(\theta_n)$ for all $n \geq 0$.

Then,

$$d(\theta_n, S(f, B)) \rightarrow 0$$

as $n \rightarrow \infty$ a.s. Moreover, suppose that $S(f, B)$ is a discrete set. Then, on almost all sample paths, θ_n converges to a unique point in $S(f, B)$ as $n \rightarrow \infty$.

Notice that the point in $S(f, B)$ that θ_n converges to can be random. We can apply Theorem 6 in our context, but first we need one more assumption.

Assumption A6 The random variables X , $K(Y)$ and $Y(\theta_0)$, for some fixed $\theta_0 \in \Theta$, all have finite 4th moments.

Remark 5 When **A1-A3** and **A6** hold, $EY^4(\theta)$ is bounded in $\theta \in \Theta$.

Corollary 7 Let Θ be a box in \mathbb{R}^p and suppose **A1 - A4**, **A6** hold. Then $d(\theta_n, S(v, \Theta)) \rightarrow 0$ as $n \rightarrow \infty$ a.s. Moreover, suppose that $S(v, \Theta)$ is a discrete set. Then, on almost all sample paths, θ_n converges to a unique point in $S(v, \Theta)$ as $n \rightarrow \infty$.

Proof. The only condition of Theorem 6 that needs verification is the condition $\sup_n E\|G_n\|^2 < \infty$. In our case, $G_n = g_n(\theta_n)$, and

$$\|g_n(\theta_n)\|^2 \leq \sup_{\theta \in \Theta} \|g_n(\theta)\|^2.$$

But the distribution of $g_n(\theta)$ does not depend on n , so the result follows if

$$\sup_{\theta \in \Theta} E\|g_1(\theta)\|^2 < \infty.$$

The argument is similar to one used in Proposition 1 and is omitted. It is this argument that requires the stronger moment assumption **A6**. ■

Corollary 7 does not ensure that θ_n converges to a deterministic θ^* as $n \rightarrow \infty$. For that we need to impose further conditions. One simple condition is that the set of first-order critical points $S(v, \Theta)$ consists of a single element θ^* . This condition is unlikely to be easily verified in practice.

We will see in Section 6 that the stochastic approximation procedure works well so long as the parameters of the procedure are chosen appropriately. However, as with any stochastic approximation procedure, it can be difficult to select good values for these parameters. For this reason we also consider a second estimator based on quite a different approach.

5 Sample Average Approximation

The stochastic approximation method above estimates the parameter θ^* that solves the optimization problem

$$\mathcal{P} : \min_{\theta \in \Theta} v(\theta)$$

and the target mean μ simultaneously. An alternative is a two-phase approach where we first compute an estimate $\hat{\theta}$ of θ^* , and in a second phase estimate μ using

$$\hat{\mu}_n = \frac{1}{n} \sum_{i=1}^n [X_i - h(Y_i, \hat{\theta})]. \quad (10)$$

If $\hat{\theta}$ is a deterministic approximation for θ^* , then the ordinary strong law and central limit theorem immediately apply. In general, however, $\hat{\theta}$ will be a random variable that depends on sampling in the initial phase, as occurs with the sample average approximation (SAA) method that we now adopt. (See Shapiro [2004] for an introduction to this approach.)

Let $m \geq 2$ be given and suppose that we generate, and then fix, the random sample $(\tilde{X}_1, \tilde{Y}_1), (\tilde{X}_2, \tilde{Y}_2), \dots, (\tilde{X}_m, \tilde{Y}_m)$. For a fixed θ , the sample variance of $(\tilde{X}_i(\theta) : 1 \leq i \leq m)$ is

$$V(m, \theta) = \frac{m}{m-1} \left(\frac{1}{m} \sum_{i=1}^m \tilde{X}_i^2(\theta) - \bar{X}_m^2(\theta) \right), \quad (11)$$

where

$$\bar{X}_m(\theta) = \frac{1}{m} \sum_{i=1}^m \tilde{X}_i(\theta).$$

The SAA problem corresponding to \mathcal{P} is

$$\mathcal{P}_m : \quad \min_{\theta \in \Theta} V(m, \theta),$$

i.e., we minimize the *sample* variance. Once the sample is fixed, the SAA problem can be solved using any convenient optimization algorithm. The algorithm can exploit the IPA gradients derived earlier, which are exact gradients of $V(m, \theta)$. In our implementation we used a quasi-Newton procedure that exploits the IPA gradients.

The term “sample average approximation” may seem inappropriate because the function $V(m, \cdot)$ in (11) is not a sample average. It is, instead, a nonlinear function of sample averages. But the standard theory for sample average approximation is readily extended to this setting, and we give the extensions that we require below. So the term is not unreasonable and we retain it.

Let $\hat{\theta}_m$ be a first-order critical point for problem \mathcal{P}_m obtained from the first phase. In the second phase, we then estimate μ via the sample average (10), using $\hat{\theta}_m$ in place of $\hat{\theta}$. Our sample average approximation algorithm for estimating μ is given in Figure 2.

The first stage: Choose a positive integer $m \geq 2$.

Generate the i.i.d. sample $(\tilde{X}_i, \tilde{Y}_i) \sim (X, Y)$, $i = 1, \dots, m$.

For a fixed θ , define

$$V(m, \theta) = \frac{m}{m-1} \left(\frac{1}{m} \sum_{i=1}^m \tilde{X}_i^2(\theta) - \left(\frac{1}{m} \sum_{i=1}^m \tilde{X}_i(\theta) \right)^2 \right),$$

where $\tilde{X}_i(\theta) = \tilde{X}_i - h(\tilde{Y}_i, \theta)$.

Find $\hat{\theta}_m$, a first order critical point for the problem

$$\min_{\theta \in \Theta} V(m, \theta).$$

The second stage:

Generate the i.i.d. sample $(X_j, Y_j) \sim (X, Y)$, $j = 1, \dots, n$, independent of the sample $(\tilde{X}_i, \tilde{Y}_i)$, $i = 1, \dots, m$.

Compute $\hat{\mu}_n = n^{-1} \sum_{j=1}^n X_j - h(Y_j, \hat{\theta}_m)$.

Figure 2: The sample average approximation algorithm

We next show that $\hat{\mu}_n$ satisfies a strong law and central limit theorem. These results require a little care, because $\hat{\theta}_m$ is a random variable. We first need a uniform version of the strong law. Proposition 8 appears as Proposition 7 in Shapiro [2004]. We say that $f(y, \theta)$ is *dominated* by an integrable function $f(\cdot)$ if $Ef(Y) < \infty$ and for every $\theta \in \Theta$, $|f(Y, \theta)| \leq f(Y)$ a.s.

Proposition 8 (Shapiro 2003) *Suppose that for every $y \in H_2$, the function $f(y, \cdot)$ is continuous on (the compact set) Θ , and $f(y, \theta)$ is dominated by an integrable function. Then $Ef(Y, \theta)$ is continuous as a function of $\theta \in \Theta$ and*

$$\sup_{\theta \in \Theta} \left| \frac{1}{n} \sum_{i=1}^n f(Y_i, \theta) - Ef(Y, \theta) \right| \rightarrow 0$$

as $n \rightarrow \infty$ a.s.

We can now state a version of the strong law and central limit theorem for the case where $\hat{\theta}$ is random. There is no need for $\hat{\theta}$ to be a solution of \mathcal{P}_m ; it can be any random variable taking values in Θ . To emphasize the dependence of $\hat{\mu}_n$ on θ we write $\hat{\mu}_n(\theta)$.

Theorem 9 *Suppose that **A1-A3** hold, that $EK(Y) < \infty$, and that the samples used in constructing $\hat{\theta}$ are independent of those used in computing $\hat{\mu}_n$. Then $\hat{\mu}_n(\hat{\theta}) \rightarrow \mu$ as $n \rightarrow \infty$ a.s., and*

$$\sqrt{n}(\hat{\mu}_n(\hat{\theta}) - \mu) \Rightarrow v^{1/2}(\hat{\theta})N(0, 1)$$

as $n \rightarrow \infty$, where $N(0, 1)$ is independent of $\hat{\theta}$.

Proof. For the strong law note that

$$\begin{aligned} |\hat{\mu}_n(\hat{\theta}) - \mu| &\leq \left| \frac{1}{n} \sum_{i=1}^n (X_i - \mu) \right| + \left| \frac{1}{n} \sum_{i=1}^n h(Y_i, \hat{\theta}) \right| \\ &\leq \left| \frac{1}{n} \sum_{i=1}^n (X_i - \mu) \right| + \sup_{\theta \in \Theta} \left| \frac{1}{n} \sum_{i=1}^n h(Y_i, \theta) \right|. \end{aligned} \quad (12)$$

The first term in (12) converges to 0 as $n \rightarrow \infty$ by the strong law of large numbers. The second term converges to 0 by an application of Proposition 8.

For the central limit theorem, first note that conditional on $\hat{\theta}$, μ_n is an average of i.i.d. random variables with finite variance. Hence the ordinary central limit theorem ensures that for each fixed $x \in \mathbb{R}$,

$$P\left(\sqrt{n}(\hat{\mu}_n(\hat{\theta}) - \mu) \leq x \mid \hat{\theta}\right) \rightarrow \Phi\left(\frac{x}{v^{1/2}(\hat{\theta})}\right) \mathbf{1}_{\{v(\hat{\theta}) > 0\}} + \mathbf{1}_{\{x \geq 0\}} \mathbf{1}_{\{v(\hat{\theta}) = 0\}} \quad (13)$$

as $n \rightarrow \infty$, where Φ is the distribution function of a normal random variable with mean 0 and variance 1 and $\mathbf{1}_{\{\cdot\}}$ is an indicator function. The dominated convergence theorem ensures that we can take expectations through (13), and so

$$\begin{aligned} &P(\sqrt{n}(\hat{\mu}_n(\hat{\theta}) - \mu) \leq x) \\ &\rightarrow E \left[\Phi\left(\frac{x}{v^{1/2}(\hat{\theta})}\right) \mathbf{1}_{\{v(\hat{\theta}) > 0\}} + \mathbf{1}_{\{x \geq 0\}} \mathbf{1}_{\{v(\hat{\theta}) = 0\}} \right] \\ &= P(v^{1/2}(\hat{\theta})N(0, 1) \leq x) \end{aligned}$$

for all $x \in \mathbb{R}$, which is the desired central limit theorem. ■

Hence the strong law and central limit theorem continue to hold in the case where $\hat{\theta}$ is random. In particular, if we first solve, or approximately solve, \mathcal{P}_m to get $\hat{\theta}_m$, and then compute $\mu_n(\hat{\theta}_m)$, then the resulting estimator is “well behaved” as the number of samples n gets large.

Now, as the computational budget gets large, one would naturally want to eventually zero in on a fixed θ^* that solves \mathcal{P} using some vanishing fraction of the budget, and use the remainder of the budget to estimate μ . This can be modelled by assuming that $m = m(n)$ is a function of n such that $m(n) \rightarrow \infty$ as $n \rightarrow \infty$. In this case, $\hat{\mu}_n(\hat{\theta}_{m(n)})$ behaves the same as $\hat{\mu}_n(\theta^*)$ as $n \rightarrow \infty$, at least to first order.

Theorem 10 *Suppose that $\hat{\theta}_{m(n)} \rightarrow \theta^*$ as $n \rightarrow \infty$ a.s., for some random variable θ^* . Suppose further that **A1** - **A3** hold and the samples used in computing $\hat{\theta}_{m(n)}$ are independent of those used to compute $\hat{\mu}_n$ for every n . Then $\hat{\mu}_n(\hat{\theta}_{m(n)}) \rightarrow \mu$ as $n \rightarrow \infty$ a.s. If, in addition, $EK^2(Y) < \infty$, then*

$$\sqrt{n}(\hat{\mu}_n(\hat{\theta}_{m(n)}) - \mu) \Rightarrow v^{1/2}(\theta^*)N(0, 1)$$

as $n \rightarrow \infty$.

Proof. The proof of the strong law is very similar to the analogous result in the previous section and is therefore omitted. To prove the central limit theorem, note that

$$\begin{aligned} \sqrt{n}(\hat{\mu}_n(\hat{\theta}_{m(n)}) - \mu) &= \sqrt{n}(\hat{\mu}_n(\theta^*) - \mu) + \sqrt{n}(\hat{\mu}_n(\hat{\theta}_{m(n)}) - \hat{\mu}_n(\theta^*)) \\ &= D_{1,n} - D_{2,n}, \text{ say.} \end{aligned}$$

Notice that θ^* is independent of the samples to compute $\hat{\mu}_n$ for every n . By Theorem 9, it suffices to show that

$$D_{2,n} = \frac{1}{\sqrt{n}} \sum_{j=1}^n [h(Y_j, \hat{\theta}_{m(n)}) - h(Y_j, \theta^*)] \Rightarrow 0$$

as $n \rightarrow \infty$.

Chebyshev's inequality ensures that for any fixed $\epsilon > 0$

$$\begin{aligned} P(|D_{2,n}| > \epsilon) &\leq \epsilon^{-2} E D_{2,n}^2 \\ &= \frac{1}{n\epsilon^2} \sum_{j=1}^n E[h(Y_j, \hat{\theta}_{m(n)}) - h(Y_j, \theta^*)]^2 \end{aligned} \tag{14}$$

$$= \frac{1}{\epsilon^2} E[h(Y_1, \hat{\theta}_{m(n)}) - h(Y_1, \theta^*)]^2. \tag{15}$$

Now, $[h(Y_1, \hat{\theta}_{m(n)}) - h(Y_1, \theta^*)]^2 \rightarrow 0$ as $n \rightarrow \infty$ a.s. Moreover,

$$[h(Y_1, \hat{\theta}_{m(n)}) - h(Y_1, \theta^*)]^2 \leq K^2(Y_1) \|\hat{\theta}_{m(n)} - \theta^*\|^2. \tag{16}$$

The normed term in (16) is bounded, and so the dominated convergence theorem implies that (15) converges to 0 as $n \rightarrow \infty$. ■

It remains to give conditions under which $\hat{\theta}_m \rightarrow \theta^*$ as $m \rightarrow \infty$ a.s. If problem \mathcal{P} has a unique optimal solution θ^* and $\hat{\theta}_m$ solves problem \mathcal{P}_m exactly then, as in Shapiro [2004], this would follow using standard arguments and an extension of a uniform law of large numbers to nonlinear functions of means. (Recall from (11) that $V(m, \theta)$ is essentially a nonlinear function of sample means, rather than a sample mean itself.) However, the best that we can hope for from a computational point of view is that $\hat{\theta}_m$ is a first-order critical point for the problem \mathcal{P}_m . So, to obtain convergence to a fixed θ^* , we first prove convergence of first-order critical points to those of the true problem \mathcal{P} . Our next result extends Theorem 3.1 in Bastin et al. [2004] for sample averages to nonlinear functions of sample averages.

Let $f(\theta, \xi)$ be a \mathbb{R}^d -valued function of $\theta \in \Theta \subset \mathbb{R}^p$ and a random vector ξ and let $\bar{f}(\theta) = Ef(\theta, \xi)$. Let

$$\bar{f}_m(\cdot) = \frac{1}{m} \sum_{i=1}^m f(\cdot, \xi_i)$$

denote a sample average of m i.i.d. realizations of the function $f(\cdot, \xi)$. Suppose that $g(x)$ is a real-valued \mathcal{C}^1 function of $x \in D \subset \mathbb{R}^d$, where D is an open set containing the range of \bar{f} and \bar{f}_m for all m . We seek conditions under which first-order critical points of $g \circ \bar{f}_m = g(\bar{f}_m(\cdot))$ on Θ converge to those of $g \circ \bar{f}$.

Theorem 11 *Consider the functions defined immediately above. Let H denote the support of the probability distribution of ξ . Suppose that Θ is convex and compact, the samples ξ_1, \dots, ξ_m are i.i.d. and*

- (i) *for all $\xi \in H$, $f(\cdot, \xi) = (f_1(\cdot, \xi), \dots, f_d(\cdot, \xi))$ is \mathcal{C}^1 on an open set containing Θ ,*
- (ii) *the component functions $f_j(\theta, \xi)$ ($j = 1, \dots, d$) are dominated by an integrable function, and*
- (iii) *the gradient components $\partial f_j(\theta, \xi) / \partial \theta_i$ are dominated by an integrable function ($i = 1, \dots, p$, $j = 1, \dots, d$).*

Let $\hat{\theta}_m \in S(g \circ \bar{f}_m, \Theta)$ be the set of first-order critical points of $g \circ \bar{f}_m$ on Θ . Then $d(\hat{\theta}_m, S(g \circ \bar{f}, \Theta)) \rightarrow 0$ as $m \rightarrow \infty$ a.s.

Proof. If $d(\hat{\theta}_m, S(g \circ \bar{f}, \Theta)) \not\rightarrow 0$, then by passing to a subsequence if necessary, we can assume that for some $\epsilon > 0$, $d(\hat{\theta}_m, S(g \circ \bar{f}, \Theta)) \geq \epsilon$ for all $m \geq 1$. Since Θ is compact, by passing to a further subsequence if necessary, we can assume that $\hat{\theta}_m$ converges to a point $\theta^* \in \Theta$. It follows that $\theta^* \notin S(g \circ \bar{f}, \Theta)$. On the other hand, by Proposition 8, $\bar{f}_m(\hat{\theta}_m) \rightarrow \bar{f}(\theta^*)$ and $\nabla_{\theta} \bar{f}_m(\hat{\theta}_m) \rightarrow \nabla_{\theta} \bar{f}(\theta^*)$ as $m \rightarrow \infty$ a.s.

Since Θ is convex, each $\hat{\theta}_m$ satisfies the first order condition

$$\langle g'(\bar{f}_m(\hat{\theta}_m)) \nabla_{\theta} \bar{f}_m(\hat{\theta}_m), u - \hat{\theta}_m \rangle \geq 0, \text{ for all } u \in \Theta, \text{ a.e.}$$

Taking the limit as $m \rightarrow \infty$, we obtain that

$$\langle g'(\bar{f}(\theta^*)) \nabla_{\theta} \bar{f}(\theta^*), u - \theta^* \rangle \geq 0, \text{ for all } u \in \Theta, \text{ a.e.}$$

Therefore, $\theta^* \in S(g \circ \bar{f}, \Theta)$ and we obtain a contradiction. ■

We now obtain the following corollary.

Corollary 12 *Suppose that **A1-A4** hold, Θ is convex and $EK^2(Y) < \infty$. Then $d(\hat{\theta}_m, S(v, \Theta)) \rightarrow 0$ as $m \rightarrow \infty$ a.s.*

Proof. If $g(x, y) = x - y^2$, then

$$V(m, \theta) = \frac{m}{m-1} \left(\frac{1}{m} \sum_{i=1}^m X_i^2(\theta) - \bar{X}_m^2(\theta) \right) = \frac{m}{m-1} g \left(\frac{1}{m} \sum_{i=1}^m X_i^2(\theta), \frac{1}{m} \sum_{i=1}^m X_i(\theta) \right).$$

Notice that

$$S(V(m, \cdot), \Theta) = S \left(g \left(\frac{1}{m} \sum_{i=1}^m X_i^2(\cdot), \frac{1}{m} \sum_{i=1}^m X_i(\cdot) \right), \Theta \right),$$

i.e., the sets of first-order critical points of these two functions coincide.

By the proof of Proposition 1 and Remark 4,

$$X(\theta), X^2(\theta), \frac{\partial h(Y, \theta)}{\partial \theta_i} \text{ and } 2X(\theta) \frac{\partial h(Y, \theta)}{\partial \theta_i}$$

are all dominated by an integrable function ($i = 1, \dots, p$). By Theorem 11, it follows that

$$d(\hat{\theta}_m, S(g(EX^2(\cdot), EX(\cdot)), \Theta)) = d(\hat{\theta}_m, S(v, \Theta)) \rightarrow 0$$

as $m \rightarrow \infty$. ■

Corollary 12 shows that $\hat{\theta}_m$ converges to the set of first-order critical points of v as $m \rightarrow \infty$. This does not guarantee that the sequence $\{\hat{\theta}_m\}$ converges almost surely, as was the case for stochastic approximation. In general we cannot guarantee this because when there are multiple critical points, the particular critical point chosen depends, among other things, on the optimization algorithm that is used. Of course, a simple sufficient condition that ensures convergence is the existence of a unique first-order critical point. This condition is clearly difficult to verify in practice.

The limiting results above establish that our procedure is a sensible one. However, these results do not shed light on how much effort to devote to searching for θ^* versus how much to allocate to the “production run” that estimates μ . The computational effort required to compute $\hat{\theta}_m$ and $\hat{\mu}_n(\hat{\theta}_m)$ for a given $\hat{\theta}_m$ is approximately proportional to m and n . Letting $m = m(c)$ and $n = n(c)$ be functions of the total computational budget c we therefore have

$$\alpha_1 m(c) + \alpha_2 n(c) \approx c,$$

for some constants α_1 and α_2 . Without loss of generality we assume that $\alpha_1 = \alpha > 1$ and $\alpha_2 = 1$.

Now, $m(c)$ and $n(c)$ must satisfy $m(c), n(c) \rightarrow \infty$ as $c \rightarrow \infty$ to ensure that $\hat{\theta}_{m(c)} \rightarrow \theta^*$ and $\hat{\mu}_{n(c)}(\hat{\theta}_{m(c)}) \rightarrow \mu$. The mean squared error of $\hat{\mu}_n(\hat{\theta}_m)$ is then

$$\text{mse}(\hat{\mu}_n(\hat{\theta}_m)) = \text{var}(\hat{\mu}_n(\hat{\theta}_m)) = \frac{1}{n} E v(\hat{\theta}_m).$$

We wish to determine m that minimizes $n^{-1} E v(\hat{\theta}_m)$, where $n = c - \alpha m$. We proceed heuristically as follows.

The asymptotic behavior of the optimal solution $\hat{\theta}_m$ of the approximation problem (\mathcal{P}_m) provides a guideline for determining the optimal m . Suppose that assumptions **A1-A4** hold. In addition, assume that for all $y \in H_2$, $h(y, \cdot)$ is \mathcal{C}^2 (i.e., twice continuously differentiable) on \mathcal{U} and that $\nabla_{\theta}^2 h(y, \cdot)$ is uniformly dominated by an integrable function. Then, under some uniform integrability conditions, $v(\cdot)$ is a \mathcal{C}^2 function on \mathcal{U} . If Θ is convex, problem \mathcal{P} has a unique optimal solution θ^* , and the Hessian matrix $\nabla^2 v(\theta^*)$ is positive definite, then $\hat{\theta}_m$ tends to θ^* at a stochastic rate of order $m^{-1/2}$ [Shapiro, 1993]. Under additional uniform integrability conditions, $E \|\hat{\theta}_m - \theta^*\| = O(m^{-1/2})$. From the second order Taylor approximation to $v(\hat{\theta}_m)$ and using the continuity of the Hessian matrix $\nabla^2 v(\cdot)$, we obtain

$$v(\theta) - v(\theta^*) \leq \lambda \|\theta - \theta^*\|^2, \tag{17}$$

for all θ in a convex compact neighborhood \mathcal{W} of θ^* and for some constant λ which depends on the eigenvalues of $\nabla^2 v(\theta), \theta \in \mathcal{W}$. Therefore, we expect that

$$E[v(\hat{\theta}_m)] - v(\theta^*) = O(m^{-1}).$$

Assume that $E[v(\hat{\theta}_m)] - v(\theta^*) \sim \frac{\gamma}{m}$, for some constant γ . Then the asymptotically optimal m^* is

$$\begin{aligned} m^* &\approx \operatorname{argmin} \left\{ \frac{1}{n} E(v(\hat{\theta}_m)) = \frac{v(\theta^*)m + \gamma}{m(c - \alpha m)} : 1 \leq m \leq \frac{c}{\alpha} \right\} \\ &= \frac{\sqrt{(\gamma\alpha)^2 + \gamma\alpha v(\theta^*)c} - \gamma\alpha}{\alpha v(\theta^*)}. \end{aligned} \quad (18)$$

Expression (18) is asymptotically (i.e., as $c \rightarrow \infty$) of the form $m \approx R\sqrt{c}$, where

$$R = \sqrt{\frac{\gamma}{\alpha v(\theta^*)}}.$$

Thus we see that the optimal choice of m^* is of the order \sqrt{c} . The coefficient R provides some insight. When α is large, solving the approximation problem (\mathcal{P}_m) is expensive, so we trade off some computational accuracy for more production runs. Similarly, if the optimal variance $v(\theta^*)$ is large, then it is not worth spending too much effort finding θ^* . From (17), we can view γ as a measure of the curvature of $v(\cdot)$ at θ^* . Therefore, if the curvature is high, then we invest more effort in finding θ^* .

6 Numerical Results

We now return to the discretely monitored barrier call option example presented in Section 2. We assume that under the risk-neutral measure, the underlying stock price $\{S(t) : t \geq 0\}$ is governed by the dynamics

$$\frac{dS(t)}{S(t)} = rdt + \sigma dW(t), \quad (19)$$

where $(W(t) : t \geq 0)$ is a standard Brownian motion, the risk-free interest rate r and volatility σ are constants and S_0 is fixed; see Glasserman [2004] for more about this model. In order to simulate the price process, we generate independent replications of the stock price using the form

$$S_i = S_{i-1} \exp \left((r - \sigma^2/2)\Delta t + \sigma\sqrt{\Delta t}Z_i \right), i = 1, \dots, l,$$

where Z_1, \dots, Z_l are i.i.d. standard (mean 0 and variance 1) normal random variables.

We consider a double barrier knock-out call option. Let H_l and H_u denote the lower and upper barrier levels, respectively, and $\mathcal{S} = [H_l, H_u] \cup \{0\}$. Then \tilde{S}_i is defined as $\tilde{S}_i = 1_{\{\tau > i\}} S_i$, where $\tau = \inf\{n \geq 0 : S_n < H_l \text{ or } S_n > H_u\}$. Suppose that $U(\cdot, \cdot; \theta)$ is given, where $U(0, \cdot; \theta) = 0$ for all $\theta \in \Theta$. Let

$$M_l(U(\theta)) = \sum_{i=1}^l U(\tilde{S}_i, l - i; \theta) - P(\tilde{S}_{i-1}, \cdot)U(\cdot, l - i; \theta)$$

under some fixed initial state $\tilde{S}_0 = x$. Then $X(\theta) = (\tilde{S}_l - K)^+ - M_l(U(\theta))$ is an estimator of $U^*(x, l) = E_x[(\tilde{S}_l - K)^+]$.

In order to obtain an efficient estimator $X(\theta)$ it is important to find a good parameterization for the function $U(x, i; \theta)$. The function should approximate the expected payoff $U^*(x, i)$ reasonably well and at the same time should enable the computation of the control variate $M_l(U(\theta))$ with a moderate amount of computational effort. To get a sense of how to choose the parameterization, we estimated the

expected payoff function $U^*(\cdot, \cdot)$. (In general, one needs at least some idea of how this function behaves in order to choose an effective parameterization.) Figure 3 displays surface plots of the estimated expected payoff function $U^*(x, i)$. For any fixed $i = 1, \dots, l-1$, $U^*(x, i)$ initially increases as x increases. But as x approaches the upper barrier H_u , $U^*(x, i)$ reaches a maximum and then decreases. For each fixed i , $U^*(\cdot, i)$ is nearly concave, at least for the higher levels of volatility in Figure 3. Let our parameterization have the form

$$U(x, i; \theta) = \begin{cases} 0 & \text{if } x = 0, \\ (x - K)^+ & \text{if } i = 0, \text{ and} \\ \theta_{4(i-1)+1}x^{\theta_{4(i-1)+2}} + \theta_{4(i-1)+3}x + \theta_{4i} & \text{if } i = 1, 2, \dots, l-1 \text{ and } x \neq 0, \end{cases}$$

where $\theta = (\theta_1, \theta_2, \dots, \theta_{4(l-1)}) \in \Theta$, $\Theta = \{y \in \mathbb{R}^{4(l-1)} : a(j) \leq y(j) \leq b(j), j = 1, 2, \dots, 4(l-1)\}$ and $a(j) \geq 0, j = 1, 2, \dots, 4(l-1)$. (Parameterizations that better fit the true value function are certainly possible, but we wanted to get a sense of how well we could do with very simple parameterizations.) Then $U(x, i; \cdot) : \mathbb{R}^{4(l-1)} \rightarrow \mathbb{R}$ is \mathcal{C}^1 for all $(x, i) \in \mathcal{S} \times \{0, 1, \dots, l-1\}$ and $U(\cdot, i; \cdot) : (0, \infty) \times \mathbb{R}^{4(l-1)} \rightarrow \mathbb{R}$ is \mathcal{C}^1 for all $i \in \{0, 1, \dots, l-1\}$. Details on both the verification of **A1-A6**, and the computation of the control variate $M_i(U(\theta))$ are given in the appendix.




Figure 3: Surface plots of the estimated expected payoff $U^*(x, i)$. Left: $\sigma = .4$, $l = 6$ and barriers at $H_l = 75$ and $H_u = 115$. Center: $\sigma = .6$, $l = 6$ and barriers at $H_l = 80$ and $H_u = 105$. Right: $\sigma = .6$, $l = 6$ and barriers at $H_l = 85$ and $H_u = 100$.

We examine the performance of the proposed estimators relative to the standard Monte Carlo technique. We assume that the annual drift ν is 5% and the initial stock price S_0 is 90. The option has $K = S_0$ and maturity $T = .25$. Tables 1-3 report numerical results for options with various volatilities, monitoring dates and barriers. We use the terms naïve, SA and SAA to represent the estimators obtained through naïve Monte Carlo estimation, the stochastic approximation method and the sample average approximation method, respectively. In the stochastic approximation algorithm, we took $m = 500$ and

$$a_k = \frac{e}{C + k^\alpha},$$

where $e, C > 0$ and $\alpha \in (1/2, 1]$ are tunable constants. This form of the gain sequence is advocated in Spall [2003]. We used (9) as an estimator of $v(\theta^*)$. For the SAA estimator, we first replicated $m = 500$ samples. We obtained $\hat{\theta}_m$ by applying a quasi-Newton method with a line search (supplied as part of the MATLABTM package) using IPA gradients to solve the sample average approximation problem \mathcal{P}_m . We

allocate 10% of the CPU time on this optimization stage. As an estimator of the variance $v(\hat{\theta})$, we used the sample variance of $X(\hat{\theta})$ over n replicates, where $\hat{\theta}$ is viewed as fixed, in the sense of Theorem 9.

In Tables 1-3, the ‘‘SA ratio’’ denotes the ratio of the sample variance of $(\tilde{S}_l - K)^+$ to the estimated variance obtained from the SA estimator, both based on mn samples. Similarly, the ‘‘SAA ratio’’ is the ratio of the sample variance $(\tilde{S}_l - K)^+$ to that of $X(\hat{\theta}_m)$ for given $\hat{\theta}_m$, both over n replicates. So ‘‘SA ratio’’ and ‘‘SAA ratio’’ present the variance reduction ratios without considering the computational effort of computing the control variates and estimating θ^* . The fourth and sixth columns in Tables 1-3 show that both the SA and SAA estimators produce a significant variance reduction. Comparing the two columns, we see that the SAA estimators outperform the SA estimators. A problem with the SA estimator is that it is very sensitive to the step size parameters a_k and the initial point θ_0 . We performed preliminary simulations with this method, tuning the parameters heuristically until reasonable performance was observed.

The values V_{Naive} , V_{SA} and V_{SAA} are, respectively, the estimated variances obtained from the naïve, SA and SAA estimators using the same CPU time. These estimated variances provide a fair comparison among the three estimators. The fifth and seventh columns in Tables 1-3 show that in most cases the SAA estimators outperform both the SA and naïve estimators. The SA estimators outperform the naïve estimators in the cases with barriers at $H_l = 80$ and $H_u = 105$, and with barriers at $H_l = 85$ and $H_u = 100$. But when $H_l = 75$ and $H_u = 115$, we do not observe an apparent advantage in variance reduction with the SA estimators compared to the naïve estimators. In this last case, under a fixed computational budget, the SA estimators do not achieve a sufficient variance reduction to outweigh the computational effort to compute the control variates and estimate θ^* . However, if the simulation run length n is long enough then from Theorems 5 and 10 we would expect the SA and SAA estimators to be fairly similar in performance.

We see that our adaptive methods work better for $\sigma = .6$ than for $\sigma = .4$. In fact, the best performance for the SAA method is obtained with $\sigma = .6$ and barriers at $H_l = 80$ and $H_u = 105$. Both the SA and SAA methods show the worst performance with $\sigma = .4$ and barriers at $H_l = 75$ and $H_u = 115$. These results show that finding a good parameterization is crucial to obtaining an efficient estimator. As observed in Figure 3, for each fixed i , $U^*(\cdot, i)$ is nearly concave for high volatilities so our parameterization works well. (When $\theta_{4(l-i)+1} < 0$ and $\theta_{4(l-i)+2} > 1$, $U(\cdot, i; \theta)$ is concave.) However, when the gap between the two barriers is wide and the volatility is low, the option has low knock-out probability and hence as i decreases the shape of the function $U^*(x, i)$ closely resembles the shape of the payoff $(x - K)^+$. Therefore our parameterization does not approximate the expected payoff function well, and as a consequence our methods do not show satisfactory performance in this case.

In most cases the variance reduction ratio decreases as the number of monitoring dates l increases. One explanation for this is that as l increases, the number of parameters in the control variate increases, and so more effort is required in the optimization stage.

7 Final Remarks

We have developed two adaptive Monte Carlo methods for the case in which parameterized control variates are available. The first method is based on a stochastic approximation technique. It is easy to implement and the computational effort per replication is low, but its performance is sensitive to the tuning parameters. The second method uses a sample average approximation approach. This method is robust, but it involves solving a non-linear optimization problem in the initial optimization phase, which can be computationally expensive.

Volatility	Frequency of monitoring	CPU time (sec)	SA ratio	V_{Naive}/V_{SA}	SAA ratio	V_{Naive}/V_{SAA}
$\sigma = .4$	$l = 3$	374	23	3.5	96	19
	$l = 6$	1839	3.3	0.26	25	2.7
	$l = 12$	7846	4.2	0.20	3.6	0.24
$\sigma = .6$	$l = 3$	189	49	8.2	543	112
	$l = 6$	1599	5.4	0.51	76	10
	$l = 12$	5716	6.3	0.36	9.1	0.72

Table 1: Estimated variance reduction ratio with barriers at $H_l = 75$ and $H_u = 115$.

Volatility	Frequency of monitoring	CPU time (sec)	SA ratio	V_{Naive}/V_{SA}	SAA ratio	V_{Naive}/V_{SAA}
$\sigma = .4$	$l = 3$	90	33	5.7	179	39
	$l = 6$	731	13	1.3	65	8.5
	$l = 12$	5651	2.0	0.12	9.8	0.83
$\sigma = .6$	$l = 3$	94	170	30	1058	238
	$l = 6$	1180	45	5.8	158	27
	$l = 12$	1611	12	1.1	25	3.0

Table 2: Estimated variance reduction ratio with barriers at $H_l = 80$ and $H_u = 105$.

Volatility	Frequency of monitoring	CPU time (sec)	SA ratio	V_{Naive}/V_{SA}	SAA ratio	V_{Naive}/V_{SAA}
$\sigma = .4$	$l = 3$	129	84	15	142	33
	$l = 6$	915	63	8.7	146	27
	$l = 12$	1337	15	1.5	27	3.6
$\sigma = .6$	$l = 3$	87	119	23	174	45
	$l = 6$	238	245	42	387	83
	$l = 12$	508	28	3.8	9.0	1.6

Table 3: Estimated variance reduction ratio with barriers at $H_l = 85$ and $H_u = 100$.

The simulation experiments in Section 6 should not be viewed as comprehensive comparisons, but rather a demonstration of the feasibility of the two methods. The sample average approximation method outperforms the stochastic approximation scheme and the naïve approach. In most cases the stochastic approximation scheme outperforms the naïve approach, but not always. The computational expense per replication brought by introducing the adaptive control variate is justified only when a sufficient reduction in variance is achieved. A good parameterization is essential in this regard. In choosing parameterizations, it is helpful to have some knowledge or intuition about the form of the true value functions.

Acknowledgments

We would like to thank the editors and referees for helpful comments that improved the paper. This research was supported by National Science Foundation grants DMI 0230528, DMI 0224884 and DMI 0400287.

Appendix: Additional Details of the Numerical Example

We first discuss the verification of our assumptions for a general class of martingales, and then specialize to the particular parameterization we used.

First assume that Θ is convex and compact. Suppose that there exists a bounded open set \mathcal{U} such that $\Theta \subset \mathcal{U}$, $U(x, i; \cdot) : \mathcal{U} \rightarrow \mathbb{R}$ is \mathcal{C}^1 for all $(x, i) \in \mathcal{S} \times \{0, 1, \dots, l-1\}$, and $U(\cdot, i; \cdot) : [H_l, H_u] \times \mathcal{U} \rightarrow \mathbb{R}$ is Lipschitz for all $i \in \{0, 1, \dots, l-1\}$. (These assumptions are all satisfied in our particular example.) Since $\{0, 1, \dots, l-1\}$ is finite and \mathcal{U} is bounded, there exists a $C > 0$ such that for all $\theta_1, \theta_2 \in \mathcal{U}$ and $(x, i) \in \mathcal{S} \times \{0, 1, \dots, l\}$,

$$|U(x, i; \theta_1) - U(x, i; \theta_2)| \leq C \|\theta_1 - \theta_2\|,$$

and

$$D = \sup_{\theta \in \mathcal{U}, (x, i) \in \mathcal{S} \times \{0, 1, \dots, l-1\}, k=1, \dots, p} \left\{ |U(x, i; \theta)|, \left| \frac{\partial U(x, i; \theta)}{\partial \theta_k} \right| \right\} < \infty.$$

Moreover, for any $\theta_1, \theta_2 \in \mathcal{U}$,

$$\begin{aligned} & |M_l(U(\theta_1)) - M_l(U(\theta_2))| \\ & \leq \sum_{i=1}^l |U(\tilde{S}_i, l-i; \theta_1) - U(\tilde{S}_i, l-i; \theta_2)| \\ & \quad + \sum_{i=1}^l |P(\tilde{S}_{i-1}, \cdot)U(\cdot, l-i; \theta_1) - P(\tilde{S}_{i-1}, \cdot)U(\cdot, l-i; \theta_2)| \\ & \leq lC \|\theta_1 - \theta_2\| + \sum_{i=1}^l P(\tilde{S}_{i-1}, \cdot)C \|\theta_1 - \theta_2\| \\ & \leq 2lC \|\theta_1 - \theta_2\|. \end{aligned}$$

For any $\theta \in \mathcal{U}$,

$$\begin{aligned} |X(\theta)| &\leq (\tilde{S}_l - K)^+ + \sum_{i=1}^l \left(|U(\tilde{S}_i, l-i; \theta)| + P(\tilde{S}_{i-1}, \cdot) |U(\cdot, l-i; \theta)| \right) \\ &\leq H_u + 2lD, \end{aligned}$$

and similarly,

$$\left| \frac{\partial}{\partial \theta_i} X(\theta) \right| \leq 2lD.$$

Since all of these bounds are finite we can easily verify that assumptions **A1-A6** are satisfied.

Next we discuss the computation of the martingale for the particular parameterization we chose. First, we compute the transition kernel $P(x, \cdot)$ for $x \in \mathcal{S}$. If $x = 0$,

$$P(0, y) = \mathbb{P}(\tilde{S}_1 = y | \tilde{S}_0 = 0) = \begin{cases} 1 & \text{if } y = 0 \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

For $H_l \leq x \leq H_u$,

$$P(x, [-\infty, y]) = \mathbb{P}(\tilde{S}_1 \leq y, | \tilde{S}_0 = x) = \begin{cases} 0 & \text{if } y < 0, \\ \mathbb{P}(S_1 < H_l \text{ or } S_1 > H_u | S_0 = x) & \text{if } 0 \leq y < H_l, \\ \mathbb{P}(S_1 \leq y | S_0 = x) + P(S_1 > H_u | S_0 = x) & \text{if } H_l \leq y \leq H_u, \text{ and} \\ 1 & \text{if } y > H_u. \end{cases}$$

Therefore,

$$P(x, y) = \begin{cases} 0 & \text{if } y \notin \mathcal{S}, \text{ and} \\ \mathbb{P}(S_1 < H_l \text{ or } S_1 > H_u | S_0 = x) & \text{if } y = 0. \end{cases}$$

If $H_l \leq y \leq H_u$, then, letting

$$\Gamma = \left(r - \frac{1}{2}\sigma^2 \right) \Delta t + \ln x, \quad C = \frac{1}{\sigma\sqrt{\Delta t}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{\Gamma^2}{2\sigma^2\Delta t}\right),$$

and Φ be the distribution function of a standard normal random variable, we have that

$$\begin{aligned} \frac{dP(x, [-\infty, y])}{dy} &= \frac{d\mathbb{P}(x \exp((r - \frac{1}{2}\sigma^2)\Delta t + \sigma\sqrt{\Delta t}Z) \leq y)}{dy} \\ &= \frac{d\Phi\left(\frac{\ln(\frac{y}{x}) - (r - \frac{1}{2}\sigma^2)\Delta t}{\sigma\sqrt{\Delta t}}\right)}{dy} \\ &= \frac{1}{\sigma\sqrt{2\pi\Delta t}y} \exp\left(-\frac{1}{2\sigma^2\Delta t}(\ln y - (r - \frac{1}{2}\sigma^2)\Delta t - \ln x)^2\right) \\ &= \frac{1}{\sigma\sqrt{2\pi\Delta t}y} \exp\left(-\frac{\Gamma^2}{2\sigma^2\Delta t} - \frac{(\ln y)^2}{2\sigma^2\Delta t} + \frac{\Gamma}{\sigma^2\Delta t} \ln y\right) \\ &= C \exp\left(-\frac{(\ln y)^2}{2\sigma^2\Delta t} + \left(\frac{\Gamma}{\sigma^2\Delta t} - 1\right) \ln y\right). \end{aligned}$$

To compute $M_l(U(\theta))$, it suffices to compute $P(x, \cdot)U(\cdot, i)$ for $x \in \mathcal{S}$ and $i = 0, \dots, l-1$. For $H_l \leq A < B \leq H_u$ and $p \geq 0$, let

$$\begin{aligned} \Psi(x; p, A, B) &:= \int_A^B y^p C \exp\left(-\frac{(\ln y)^2}{2\sigma^2\Delta t} + \left(\frac{\Gamma}{\sigma^2\Delta t} - 1\right)\ln y\right) dy \\ &= C \int_{\ln A}^{\ln B} \exp\left(-\frac{u^2}{2\sigma^2\Delta t} + \left(\frac{\Gamma}{\sigma^2\Delta t} + p\right)u\right) du \\ &= C \exp\left(\frac{\beta^2}{4\alpha}\right) \sqrt{\frac{\pi}{\alpha}} \left[\Phi\left(\sqrt{2\alpha}\left(\ln B - \frac{\beta}{2\alpha}\right)\right) - \Phi\left(\sqrt{2\alpha}\left(\ln A - \frac{\beta}{2\alpha}\right)\right) \right] \\ &= \exp\left(p\Gamma + \frac{p^2\sigma^2\Delta t}{2}\right) \left[\Phi\left(\frac{\ln B - \Gamma}{\sigma\sqrt{\Delta t}} - p\sigma\sqrt{\Delta t}\right) - \Phi\left(\frac{\ln A - \Gamma}{\sigma\sqrt{\Delta t}} - p\sigma\sqrt{\Delta t}\right) \right], \end{aligned}$$

where

$$\alpha = \frac{1}{2\sigma^2\Delta t} \text{ and } \beta = \frac{\Gamma}{\sigma^2\Delta t} + p.$$

Then

$$P(x, \cdot)U(\cdot, i; \theta) = \begin{cases} 0 & \text{if } x = 0, \\ \Psi(x; 1, K, H_u) - K\Psi(x; 0, K, H_u) & \text{if } i = 0 \text{ and } x \neq 0, \text{ and} \\ \theta_{4(i-1)+1}\Psi(x; \theta_{4(i-1)+2}, H_l, H_u) \\ + \theta_{4(i-1)+3}\Psi(x; 1, H_l, H_u) + \theta_{4i}\Psi(x; 0, H_l, H_u) & \text{if } i = 1, 2, \dots, l-1 \text{ and } x \neq 0. \end{cases}$$

Computing the control variate $M(U(\theta))$ therefore involves the evaluation of the distribution function of a normal random variable. The error in the approximation to the normal distribution function used in our simulation experiment is of the order 10^{-6} and it may therefore very slightly bias our adaptive control variate estimators. We do not explore this issue further in this paper.

References

- F. Bastin, C. Cirillo, and P. L. Toint. Convergence theory for nonconvex stochastic programming with an application to mixed logit. *Mathematical Programming B*, 2004. To appear.
- N. Bolia and S. Juneja. Function-approximation-based perfect control variates for pricing American options. In M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, editors, *Proceedings of the 2005 Winter Simulation Conference*, Piscataway, New Jersey, 2005. IEEE.
- P. Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer-Verlag, New York, 2004.
- P. W. Glynn and W. Whitt. Indirect estimation via $l = \lambda w$. *Operations Research*, 37(1):82–103, 1989.
- S. G. Henderson and P. W. Glynn. Approximating martingales for variance reduction in Markov process simulation. *Mathematics of Operations Research*, 27:253–271, 2002.
- S. G. Henderson and S. P. Meyn. Efficient simulation of multiclass queueing networks. In S. Andradóttir, K. J. Healy, D. H. Withers, and B. L. Nelson, editors, *Proceedings of the 1997 Winter Simulation Conference*, pages 216–223, Piscataway NJ, 1997. IEEE.

- S. G. Henderson and S. P. Meyn. Variance reduction for simulation in multiclass queueing networks. *IEEE Transactions*, 2003. To appear.
- S. G. Henderson, S. P. Meyn, and V. B. Tadić. Performance evaluation and policy selection in multi-class networks. *Discrete Event Dynamic Systems*, 13:149–189, 2003. Special issue on learning and optimization methods.
- S. G. Henderson and B. Simon. Adaptive simulation using perfect control variates. *Journal of Applied Probability*, 41(3):859–876, 2004.
- S. Kim and S. G. Henderson. Adaptive control variates. In R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, editors, *Proceedings of the 2004 Winter Simulation Conference*, pages 621–629, Piscataway NJ, 2004a. IEEE.
- S. Kim and S. G. Henderson. Adaptive control variates. *School of Operations Research and Industrial Engineering, Cornell University*, Technical Report(No. 1417), 2004b.
- C. Kollman, K. Baggerly, D. Cox, and R. Picard. Adaptive importance sampling on discrete Markov chains. *Ann. Appl. Probab.*, 9(2):391–412, 1999.
- H. J. Kushner and G. G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer-Verlag, New York, 2nd edition, 2003.
- A. M. Law and W. D. Kelton. *Simulation Modeling and Analysis*. McGraw-Hill, New York, 3rd edition, 2000.
- P. L’Ecuyer. On the interchange of derivative and expectation for likelihood ratio derivative estimators. *Management Science*, 41:738–748, 1995.
- R. S. Liptser and A. N. Shiriyayev. *Theory of Martingales*. Kluwer Academic, Boston, 1989.
- S. Maire. Reducing variance using iterated control variates. *Journal of Statistical Computation and Simulation*, 73(1):1–29, 2003.
- S. P. Meyn. Workload models for stochastic networks: value functions and performance evaluation. *IEEE Transactions on Automatic Control*, 2005. To appear.
- A. Shapiro. Asymtotic behavior of optimal solutions in stochastic programming. *Mathematics of Operations Research*, 18:829–845, 1993.
- A. Shapiro. Monte Carlo sampling methods. In A. Ruszczyński and A. Shapiro, editors, *Stochastic Programming*, Handbooks in Operations Research and Management Science. Elsevier, 2004.
- J. C. Spall. *Introduction to Stochastic Search and Optimization: Estimation, Simulation and Control*. Wiley, Hoboken, New Jersey, 2003.
- V. B. Tadić and S. P. Meyn. Adaptive Monte Carlo algorithms using control variates. *Manuscript*, 2004.