ADAPTIVE CONTROL VARIATES IN MONTE CARLO SIMULATION

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Sujin Kim

August 2006

ADAPTIVE CONTROL VARIATES IN MONTE CARLO SIMULATION

Sujin Kim, Ph.D.

Cornell University 2006

Monte Carlo simulation is widely used in many fields. Unfortunately, it usually requires a large amount of computer time to obtain even moderate precision so it is necessary to apply efficiency improvement techniques. Adaptive Monte Carlo methods are specialized Monte Carlo simulation techniques where the methods are adaptively tuned as the simulation progresses. The primary focus of such techniques has been in adaptively tuning importance sampling distributions to reduce the variance of an estimator. We instead focus on adaptive methods based on control variate schemes. In this dissertation we introduce two adaptive control variate methods where a family of parameterized control variates is available, and develop their asymptotic properties.

The first method is based on a stochastic approximation scheme for identifying the optimal choice of control variate. It is easily implemented, but its performance is sensitive to certain tuning parameters, the selection of which is nontrivial. The second method uses a sample average approximation approach. It has the advantage that it does not require any tuning parameters, but it can be computationally expensive and requires the availability of nonlinear optimization software.

We include implementations of the methods and numerical results for two applications. These results suggests that the adaptive methods outperform the naïve approach as long as the parameterization of the control variate is carefully chosen.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1

Introduction

Monte Carlo simulation is widely used in many fields. Unfortunately, it usually needs a large amount of computational effort in order to obtain sufficiently accurate results, especially, for large-scale or complex systems.

The effectiveness of Monte Carlo simulation is closely related to the variance of the simulation estimators. When the variance of the estimators is high, the results may become unacceptably inaccurate. The computational load of simulation has motivated an interest in Monte Carlo methods for reducing the variance of simulation estimators. If we can reduce the variance of an estimator without disturbing its expectation, we can obtain more accurate estimates.

The control variate method is one of the most effective and widely used variance reduction techniques in Monte Carlo simulation [Rubinstein, 1986, Law and Kelton, 2000]. In this dissertation, we develop adaptive Monte Carlo methods for estimating the expected performance measure of stochastic systems based on control variate schemes, and study the asymptotic properties of these procedures. Suppose that one wishes 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 one can estimate EX by a sample average of i.i.d. replications of $X - Y(\theta)$, and the parameter θ can be selected so as to minimize the variance of $X - Y(\theta)$. This method can be viewed as a parameterized variance reduction technique, where $Y(\theta)$ serves as a control variate. We propose adaptive procedures to tune the parameter θ for improving efficiency as the simulation progresses. This idea of adaptive control variates is

also considered in the context of steady-state simulation.

Our interest in this problem stems from several application areas. One of these is the problem of pricing financial derivatives. When the payoff of a derivative security is path dependent, or the model of the dynamics of the underlying assets is complex or high dimensional, it is often necessary to price via simulation [Glasserman, 2004]. An extended example in this paper (see Sections 2.1 and 3.2) shows that one can apply adaptive control variate methods to improve the efficiency of simulations 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 have the same asymptotic rates of growth in the variance as the naïve estimator; see Meyn [2003]. One way to potentially improve on these results 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], Baggerly et al. [2000], Fitzgerald and Picard [2001].

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 "approximating martingales". Henderson and Glynn [2002] show how to define approximating martingales for a variety of performance measures for Markov processes. The idea is to use the simpler approximating process to construct a zero mean martingale for the original process. Once we have a pa-

rameterized class of control variates at hand, we then need a procedure for selecting a control from within the class.

1.1 Adaptive Monte Carlo Methods

Adaptive Monte Carlo methods are designed to adaptively tune simulation estimators as the simulation progresses, with the purpose of improving efficiency. One needs to have a good understanding of the structure of the system being simulated in order to appropriately apply adaptive methods.

Most of the work on adaptive Monte Carlo methods has been devoted to adaptively tuning importance sampling schemes. Importance sampling has been used in various applications to accelerate simulation by minimizing the variance of the simulation estimator; see, e.g. Al-Qaq et al. [1995], Rubinstein and Melamed [1998] for queueing and reliability models and Vázquez-Abad and Dufresne [1998], Su and Fu [2000], Arouna [2003] for pricing financial derivatives, where stochastic approximation is used to tune the change of measure. Another way to tune importance sampling estimators is to select optimal importance sampling distributions via the cross entropy method; see Rubinstein [1999]. Adaptive importance sampling is primarily used for rare event simulation. For a review of its uses in this area, see Hsieh [2002], and for applications to option pricing see Glasserman and Staum [2001]. Kollman et al. [1999] discuss adaptive importance sampling in Markov chains and apply it to radiation transport problems. They provide an adaptive sampling algorithm that converges exponentially to the zero variance solution. Juneja and Shahabuddin [2006] is an excellent and up-to-date reference for importance sampling in general.

A limited amount of work has been done on adaptive control variates. Hen-

derson and Simon [2004] 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 they only work with 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, and therefore, 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.

In this dissertation, we focus on adaptive methods based on control variate schemes. The main contribution of this dissertation is to develop adaptive control variate methods for finite-horizon simulation when non-linearly parameterized control variables are available, and to provide conditions under which the adaptive estimators are consistent. To the best of our knowledge, this is the first application of stochastic approximation and sample average approximation methods in this setting. We also explore adaptive control variate methods for steady-state simulation

based on sample average approximation. We also discuss some implementation issues relevant to the practical use of these adaptive methods.

In general, it will be the case that the optimal variance $v(\theta^*)$ is positive a.s. Consequently, the rates of convergence for our proposed estimators are typically the canonical $n^{-1/2}$ rate, where n is proportional to the computational effort, as evidenced by central limit theorems. This precludes the exponential rates of convergence that are demonstrated in Henderson and Simon [2004]. However, we do briefly consider the case of a perfect control variate in the linearly-parameterized case in Section 2.2. This section sheds further light on the perfect control variate case treated in Henderson and Simon [2004], taking a somewhat different approach to constructing an estimator.

1.2 Review of Simulation Optimization Methodologies

In this section we briefly review some simulation optimization methodologies related to our work. Consider the following optimization problem:

$$\min_{\theta \in \Theta} f(\theta) = E[f(\theta, \xi)], \tag{1.2.1}$$

for some random variable ξ and parameter $\theta \in \Theta$, where $\Theta \in \mathbb{R}^p$ is a set of permissible values of the parameter θ. We assume that the function $f(\theta)$ is differentiable, and can only be evaluated by Monte Carlo simulations. How does one compute an (approximate) minimizer of (1.2.1)? Since (1.2.1) is a stochastic optimization problem, standard stochastic optimization algorithms can be applied. In particular, we consider gradient-based stochastic optimization methods.

There exist several different approaches for estimating the gradient of the function f. The main ones are finite differences, e.g., L'Ecuyer and Perron [1994],

likelihood ratio methods, e.g., Glynn [1990], and conditional Monte Carlo, e.g., [Fu and Hu, 1997]. For our adaptive control variate algorithms, we chose the method of infinitesimal perturbation analysis (IPA). The idea of IPA is simply 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. This yields an unbiased estimator [Glasserman, 1991, L'Ecuyer, 1995]. IPA is usually highly efficient when it is valid (i.e., yields an unbiased gradient estimator). Unfortunately, there are many applications where IPA is not valid. In many cases, $f(\theta, \xi)$ can be replaced by a smoother alternative, and then IPA can be used.

Stochastic approximation (SA) is a class of methods used to solve differentiable simulation optimization problems. The procedure is analogous to the steepest descent gradient search method in deterministic optimization, except here the gradient does not have an analytic expression and must be estimated. Since the basic stochastic algorithms were introduced by Robbins and Monro [1951] and Kiefer and Wolfowitz [1952], a huge amount of work has been devoted to this area. See Kushner and Yin [2003] for asymptotic properties of the various SA algorithms. The SA algorithms are easy to implement, and have been used in many areas. Fu [1990] and L'Ecuyer and Glynn [1994] studied the SA method with IPA gradient estimation and applied it to the optimization of the steady-state mean of a single server queue. The SA method is widely used in adaptive importance sampling to find an optimal importance sampling distribution [Vázquez-Abad and Dufresne, 1998, Su and Fu, 2000, Arouna, 2003].

Another standard method to solve the problem (1.2.1) is that of sample average approximation (SAA). This method approximates the original simulation

optimization problem (1.2.1) with a deterministic optimization problem. One can use the sample average $\frac{1}{N}\sum_{i=1}^{N} f(\theta, \xi_i)$ based on an i.i.d. random sample $\xi_1, .., \xi_N$ as an approximation of the expected value $f(\theta)$ for any θ. Once the sample is fixed, the sample average function becomes deterministic. Consequently, the SAA problem becomes a deterministic optimization problem, and one can solve it using any convenient optimization algorithm. The algorithm can exploit the IPA gradients, which are exact gradients of the sample average $\frac{1}{N}\sum_{i=1}^{N} f(\theta, \xi_i)$. Plambeck et al. [1996] used a SAA method with IPA gradient estimates for solving convex performance functions in stochastic systems and gave extensive computational results. The optimization of SAA problems has also been well studied in simulation [Robinson, 1996, Rubinstein and Shapiro, 1993, Chen and Schmeiser, 2001]. For an introduction to this approach, see Shapiro and Homem-de-Mello [2000], Shapiro [2003].

1.3 Dissertation Outline

In Chapter 2 we study adaptive methods based on control variate schemes in a finite-horizon setting. We assume that a family of mean zero parameterized control variates is available. When the parameterization is linear, we can appeal to the standard theory of (linear) control variates. Identifying the θ that minimizes the variance is straightforward in this case, because the variance is a convex quadratic in θ. It is possible to construct perfect (zero-variance) control variates in certain settings [Henderson and Glynn, 2002, Henderson and Simon, 2004], and we explore the asymptotic behavior of the linear control variate estimators in this case. In general, one can obtain a zero-variance estimator with a finite number of samples, N. The distribution of N has an exponentially decaying tail.

When the parameterization is nonlinear, the problem is not so straightforward. Under some pathwise differentiability and moment conditions, the variance of the control variate estimator $X - Y(\theta)$ becomes a differentiable function in the parameter θ. Once we have a differentiable variance function on hand, we apply simulation optimization algorithms to search for the optimal values of the parameter θ. We propose two adaptive procedures that tune the parameter θ while estimating EX, and study the large-sample properties of these procedures.

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 performs well with appropriately chosen step sizes. But the selection of the step size is nontrivial, and has a strong impact on the finite-time performance of the algorithm.

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 is defined with the generated sample. The sample variance function is deterministic in terms of the parameter θ, and the optimal value of θ that minimizes this sample variance function is determined using a non-linear optimization solver. Then one makes a "production run" using the value of θ returned in the first stage. 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.

In Chapter 3 we examine the performance of the adaptive control variate methods discussed in Chapter 2 applied to two examples. It is important to find a good parameterization for the control variate $Y(\theta)$ to obtain an efficient control variate estimator. The control variate $Y(\theta)$ should approximate the random variable X reasonably well and at the same time the computational expense brought by introducing the control variate should be moderate. We describe how to construct control variate estimators using martingale approximation, and choose good parameterizations for the control variates in the context of our examples. We also discuss the implementation of our methods.

In Chapter 4 we turn our attention to steady-state simulations. We assume that the underlying stochastic process possesses regenerative structure. A wide class of discrete-event simulations is regenerative [Glynn, 1994, Henderson and Glynn, 2001]. The regenerative process enjoys asymptotic properties which provide a clean setting for simulation output analysis. Under mild regularity conditions, a regenerative process satisfies a law of large numbers and a central limit theorem, and consistent estimators for the steady-state mean and time average variance can be obtained [Glynn and Iglehart, 1993, Glynn and Whitt, 2002]

We explore adaptive control variate methods for estimating steady-state performance measures based on the sample average approximation technique. The procedures exploit the regenerative structure of the underlying stochastic processes. The quantities computed over the regenerative cycles are one-dependent identically distributed random variables, so the sample average approximation method for terminating simulations in Chapter 2 can be extended to this setting. To define the sample average approximation problem, we consider time average variance

estimators based on a regenerative method. Under mild regularity assumptions, the control variate estimator based on the regenerative method is consistent and the sample average approximation problem converges to the true problem.

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.

Chapter 2

Adaptive Control Variate Methods for Finite-Horizon Simulation

In this chapter we study adaptive methods based on control variate schemes for the case in which parameterized control variates are available. Suppose that we wish to estimate EX, where X is a real-valued random variable. Suppose also that $EY(\theta) = 0$ for any $\theta \in \Theta$, where Θ is a parameter set. Then $X - Y(\theta)$ is an unbiased estimator for μ, where $Y(\theta)$ serves as a control variate, and one is free to select the parameter θ so as to minimize the variance of $X - Y(\theta)$. When the parameterization is linear, identifying the θ that minimizes the variance is straightforward because the variance is a convex quadratic in θ [Law and Kelton, 2000]. In the nonlinearly parameterized case, the problem is not so straightforward. We propose two adaptive procedures that tune the parameter θ while estimating EX.

Our motivating example for this chapter is the problem of pricing barrier options. Section 2.1 sketches some of the main ideas in pricing barrier options using adaptive control variates. We explore the linearly parameterized case in Section 2.2, which is precisely that of standard control variate theory. We then turn to the more complicated nonlinear-parameterization case. First, in Section 2.3 we outline the general problem and discuss gradient estimation. In Section 2.4 we explore an approach based on stochastic approximation, and then in Section 2.5 we study the sample average approximation approach .

2.1 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 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. We will return to this example in Section 3.2 and describe the results of some simulation experiments.

2.1.1 Pricing Barrier Options

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 $t_i = i\Delta t, i = 0, 1, 2, \ldots, l$, 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 ith 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 $\{S_i, i = 0, \ldots, l\}$. Define

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

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

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

2.1.2 Construction of Martingale Control Variates

Assume that the underlying stock price process $\{S(t) : t \geq 0\}$ is a (time homogeneous) Markov process. Then $\{S_n : n = 0, 1, 2, \ldots\}$, 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 $i = 0, 1, \ldots$, 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 $f : \mathcal{S} \to \mathbb{R}$, let $P(x, \cdot)f(\cdot) = E[f(\tilde{S}_1)|\tilde{S}_0 = x]$, provided that the expectation exists. Let $U : \mathcal{S} \times \{0, 1, \ldots, l-1\} \to \mathbb{R}$ 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

$$(\tilde{S}_l - K)^+ - M_l(U), \tag{2.1.1}$$

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 (2.1.1) 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, $k = 1, \ldots, p$. In this case, for $1 \leq n \leq l$,

$$\begin{aligned} M_n(U) &= \sum_{i=1}^{n} \left[\sum_{k=1}^{p} \theta(k) U_k(\tilde{S}_i, l-i) - P(\tilde{S}_{i-1}, \cdot) \sum_{k=1}^{p} \theta(k) U_k(\cdot, l-i) \right] \\ &= \sum_{k=1}^{p} \theta(k) \left[\sum_{i=1}^{n} U_k(\tilde{S}_i, l-i) - P(\tilde{S}_{i-1}, \cdot) U_k(\cdot, l-i) \right] \\ &= \sum_{k=1}^{p} \theta(k) M_n(U_k), \end{aligned} \tag{2.1.2}$$

so that the control $M_n(U)$ is simply a linear combination of martingales corresponding to the basis functions U_k, $k = 1, \ldots, p$. 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. We further investigate the linear control variate case in Section 2.2.

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 Section 2.3, 2.4 and 2.5. We will return to this barrier option pricing example in Section 3.2.

2.2 The Linear Case

The theory of linear control variates is very well understood; see, for example, Glynn and Szechtman [2002] or Glasserman [2004] for detailed treatments. The standard theory does not cover the perfect (zero-variance) control variate case, so after a brief review of the key ideas we discuss this case in some detail.

2.2.1 Linear Control Variate

Suppose that

$$Y(\theta) = \sum_{i=1}^{p} \theta(i)C(i),$$

where $C(i)$ is a real-valued square-integrable random variable with $EC(i) = 0$ for each $i = 1, \ldots, p$. This is the standard multiple control variates setting. Let θ and C be the corresponding column vectors in $\mathbb{R}^p$, so that $Y(\theta) = \theta^T C$, where

x^T denotes the transpose of the matrix x. Assuming that the covariance matrix $\Lambda = \mathrm{cov}(C, C)$ is nonsingular, the optimal choice of weights θ^* is

$$\theta^* = \Lambda^{-1}\beta,$$

where $\beta = \mathrm{cov}(X, C)$ is a column vector whose ith component is $\mathrm{cov}(X, C(i))$, $i = 1, \ldots, p$. Since θ^* involves moment quantities that are generally unknown, it can be estimated using the sample analogue

$$\theta_n = \Lambda_n^{-1}\beta_n$$

where

$$\begin{aligned} \beta_n &= \frac{1}{n}\sum_{j=1}^{n} X_j C_j - \bar{X}_n \bar{C}_n \text{ and} \\ \Lambda_n &= \frac{1}{n}\sum_{j=1}^{n} C_j C_j^T - \bar{C}_n \bar{C}_n^T. \end{aligned}$$

Here $\{(X_j, C_j) : j \geq 1\}$ are i.i.d. replicates of the vector (X, C), and $\bar{X}_n$ and $\bar{C}_n$ are the usual sample means of the first n observations.

Since Λ is nonsingular and $\Lambda_n \to \Lambda$ as $n \to \infty$ element-wise, it follows that Λ_n is also nonsingular for sufficiently large n, so that the estimator θ_n is well-defined for sufficiently large n. The corresponding estimator for $\mu = EX$ is

$$\mu_n = \bar{X}_n - \theta_n^T \bar{C}_n.$$

One can show that μ_n satisfies a central limit theorem of the form

$$\sqrt{n}(\mu_n - \mu) \Rightarrow \sigma N(0, 1), \tag{2.2.1}$$

where $\Rightarrow$ denotes convergence in distribution, $N(0, 1)$ is a normal random variable with mean 0 and variance 1 and $\sigma^2 = \mathrm{var}(X - Y(\theta^*))$. One can develop an

alternative estimator for θ_n that exploits the fact that $EC = 0$. This will not change the central limit theorem (2.2.1); see Glynn and Szechtman [2002].

Hence, if $\sigma^2 > 0$, the estimator μ_n converges to μ at the canonical rate $n^{-1/2}$ as is well known. In the case where $\sigma^2 = 0$ the central limit theorem (2.2.1) shows that the convergence is faster than the canonical rate, but the exact asymptotic behaviour is not as clear. The next section explores this case in more detail.

2.2.2 Exponential Convergence

It is possible to construct perfect (zero-variance) control variates in certain settings [Henderson and Glynn, 2002, Henderson and Simon, 2004]. Of course, as mentioned in the introduction, the perfect-control-variate case is unlikely to arise in the applications we have in mind. Nonetheless, partly to provide another perspective on the results of Henderson and Simon [2004] and partly for completeness, we outline the asymptotic behavior of μ_n in this case.

Let

$$\mathbf{X}_n = \begin{bmatrix} X_1 \\ X_2 \\ \vdots \\ X_n \end{bmatrix} \text{ and } \mathbf{C}_n = \begin{bmatrix} 1 & C_1(1) & C_1(2) & \cdots & C_1(p) \\ 1 & C_2(1) & C_2(2) & \cdots & C_2(p) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & C_n(1) & C_n(2) & \cdots & C_n(p) \end{bmatrix}$$

be the column vector of observations of X and the matrix with jth row containing a 1 together with C_j^T.

Define $N = \inf\{n \geq 1 : \mathbf{C}_n \text{ has full column rank}\}$. Proposition 2.2.2 below shows that N is almost surely finite when Λ is nonsingular and

$$\mu_N = \bar{X}_N - \theta_N^T \bar{C}_N = \mu$$

almost surely. Hence, if we know that a perfect control exists, then we can continue

the simulation until time N and report $\bar{X}_N - \theta_N^T \bar{C}_N$ as an estimate of μ that is almost-surely correct. Therefore, in the case when a perfect control variate exists, *the controlled estimator gives the exact answer in finite time.*

It will typically be the case that $N = p + 1$ a.s. However, in certain situations N may be random.

Example 1. Suppose that with probability 0.5, $C(1)$ is uniformly distributed on the interval $(-1, 1)$ and $C(2) = C(1) - 1$, and with probability 0.5, $C(1)$ and $C(2)$ are independent uniform random variables on $(-1, 1)$ and $(0, 2)$ respectively. Suppose further that $X = 2C(1) + C(2) + \mu$. Then with probability 0.5^n, $C_i(2) = C_i(1) - 1$ for $i = 1, \ldots, n$. Hence, $P(N = 3) = 7/8$ and for $n \geq 4$, $P(N = n) = (1/2)^n$. At time N, and not before, we learn the exact coefficients of the linear function that defines X. This then gives μ. If $X = 2C(1) + C(2) + \mu$ except at, say, $C = (1, 1)$ then the linear relationship still holds with probability 1. However, now μ_N equals μ only with probability 1, and not on all sample paths.

In this example N has an exponential tail. This observation is true in general assuming only second moments on X and C. Before stating this result precisely we need a lemma.

Lemma 2.2.1. *The matrix $\mathbf{C}_n$ has full column rank if and only if Λ_n is positive definite.*

Proof. It is well-known (e.g., Rice [1988, p. 477]) that $\mathbf{C}_n$ has full column rank if and only if $\mathbf{C}_n^T \mathbf{C}_n$ is nonsingular. Define

$$\Sigma_n = \frac{1}{n} \sum_{i=1}^{n} C_i C_i^T.$$

Then

$$\begin{aligned}\mathbf{C}_n^T\mathbf{C}_n &= \begin{bmatrix} 1 & 1 & \dots & 1 \\ C_1 & C_2 & \dots & C_n \end{bmatrix} \begin{bmatrix} 1 & C_1^T \\ 1 & C_2^T \\ \vdots & \vdots \\ 1 & C_n^T \end{bmatrix} \\ &= n \begin{bmatrix} 1 & \bar{C}_n^T \\ \bar{C}_n & \Sigma_n \end{bmatrix}. \end{aligned} \tag{2.2.2}$$

Premultiplying $\mathbf{C}_n^T\mathbf{C}_n$ by the nonsingular elementary matrix

$$B = \begin{bmatrix} 1 & 0 \\ -\bar{C}_n & I \end{bmatrix}$$

where I is the $p \times p$ identity matrix, we obtain

$$B\mathbf{C}_n^T\mathbf{C}_n = n \begin{bmatrix} 1 & \bar{C}_n^T \\ 0 & \Lambda_n \end{bmatrix},$$

which is nonsingular if and only if Λ_n is nonsingular. □

We can now state the main result of this section.

Proposition 2.2.2. *Suppose that $X \in \mathbb{R}$ and $C \in \mathbb{R}^p$ have finite second moments, $EC = 0$, $\Lambda = \operatorname{cov}(C, C)$ is positive definite and $X = C^T\theta^* + \mu$ a.s. Then N, as defined above, is finite a.s., $\mu_N = \mu$ a.s., and N has an exponentially decaying tail, i.e., $P(N > n) \leq ar^n$ for some $a > 0$ and $r < 1$.*

Proof. From Lemma 2.2.1, N can be alternatively defined as

$$\inf\{n \geq 1 : \Lambda_n \text{ is nonsingular}\}. \tag{2.2.3}$$

Since Λ_n converges elementwise to Λ under the second moment assumption almost surely, it follows that N is finite almost surely.

Next, $X = C^T\theta^* + \mu$ a.s., and so

$$\mathbf{X}_n = \mathbf{C}_n \begin{bmatrix} \mu \\ \theta^* \end{bmatrix} \tag{2.2.4}$$

almost surely, for any $n \geq 1$. The relation (2.2.4) also holds at time N, since

$$\begin{aligned} P\left(\mathbf{X}_N \neq \mathbf{C}_N \begin{bmatrix} \mu \\ \theta^* \end{bmatrix}\right) &= \sum_{n=1}^{\infty} P\left(\mathbf{X}_n \neq \mathbf{C}_n \begin{bmatrix} \mu \\ \theta^* \end{bmatrix}, N = n\right) \\ &\leq \sum_{n=1}^{\infty} P\left(\mathbf{X}_n \neq \mathbf{C}_n \begin{bmatrix} \mu \\ \theta^* \end{bmatrix}\right) \\ &= 0. \end{aligned}$$

Taking (2.2.4) at time N and premultiplying by $\mathbf{C}_N^T$, we then get

$$\mathbf{C}_N^T\mathbf{X}_N = \mathbf{C}_N^T\mathbf{C}_N \begin{bmatrix} \mu \\ \theta^* \end{bmatrix} \text{ a.s.}$$

If we use the representation (2.2.2) to expand out this relation, we find that

$$\bar{X}_N = \mu + \bar{C}_N^T\theta^* \text{ and} \tag{2.2.5}$$

$$\frac{1}{N}\sum_{i=1}^{N} C_i X_i = \bar{C}_N\mu + \Sigma_N\theta^* \tag{2.2.6}$$

almost surely. From (2.2.5), $\bar{C}_N^T\theta^* = \bar{X}_N - \mu$ a.s., so that

$$\bar{C}_N\bar{C}_N^T\theta^* = \bar{C}_N\bar{X}_N - \bar{C}_N\mu \text{ a.s.} \tag{2.2.7}$$

Adding (2.2.6) and (2.2.7) and rearranging, we then see that

$$\Lambda_N\theta^* = \beta_N \text{ a.s.},$$

so that

$$\theta^* = \Lambda_N^{-1}\beta_N = \theta_N \text{ a.s.}$$

It follows from this relation and (2.2.5) that

$$\mu_N = \bar{X}_N - \bar{C}_N^T \theta_N = \mu \text{ a.s.}$$

as claimed.

To prove the exponentially decaying tail property, note that $\mathbf{C}_n$ has full column rank if and only if at least $p+1$ of the vectors $C_1, \ldots, C_n$ are affinely independent [Bazaraa et al., 1993, p. 36]. Since Λ is nonsingular, it follows that there exist $p+1$ affinely-independent points $c_1, \ldots, c_{p+1}$ contained in the support of C_1. Now let $\epsilon > 0$ be such that the open balls $B(c_i, \epsilon)$ centered at c_i with radius ϵ are disjoint, and moreover if $x_i \in B(c_i, \epsilon)$ for all $i = 1, \ldots, p+1$, then $\{x_1, \ldots, x_{p+1}\}$ are affinely independent. Let $\tau_i = \inf\{k : C_k \in B(c_i, \epsilon)\}$, and let $N' = \max_i \tau_i$. Then at least $p+1$ of $C_1, \ldots, C_{N'}$ are affinely independent, and so $\mathbf{C}_{\mathbf{N}'}$ is nonsingular. It follows that $N \leq N'$. Furthermore, $P(C_1 \in B(c_i, \epsilon)) > 0$ since c_i is contained in the support of C_1. Hence, each τ_i is a geometric random variable and therefore N' has a geometric tail. Since $N \leq N'$ this gives the result. □

2.3 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. We assume the following:

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

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

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,

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

for all $y \in H_2$ and $j = 1, \ldots, p$.

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 we use the following result on Infinitesimal Perturbation Analysis (IPA) from L'Ecuyer [1995]. Let $f(\theta) = Ef(\theta, \xi)$ for some random variable ξ whose distribution does not depend on θ. The basic idea in IPA is to take $\nabla_\theta f(\theta, \xi)$, the gradient of $f(\theta, \xi)$ for fixed ξ, as an estimate of $\nabla_\theta f(\theta)$. This yields an unbiased estimator if the gradient and expectation can be exchanged. The following theorem gives sufficient conditions for the interchange to be valid. Since each component of the gradient can be dealt with separately, there is no loss of generality if we assume for the purposes of this theorem that $p = 1$.

Theorem 2.3.1. *[L'Ecuyer, 1995] Let $\theta_0 \in \Upsilon$, where Υ is an open interval, and let H be a measurable set such that $P(\xi \in H) = 1$. Suppose that for every $z \in H$, there is a $D(z)$, where $D(z)$ is at most countable, such that*

(i) $\forall z \in H$, $f(\cdot, z)$ is continuous everywhere in Υ,

(ii) $\forall z \in H$, $f(\cdot, z)$ is differentiable everywhere in $\Upsilon \backslash D(z)$,

(iii) there exists a function $\phi : H \to [0, \infty)$ such that

$$\sup_{\theta \in \Upsilon \backslash D(z)} |f'(\theta, z)| \leq \phi(z)$$

$\forall z \in H$ with $E\phi(\xi) < \infty$, and

(iv) $f(\theta, \xi)$ is almost surely differentiable at $\theta = \theta_0$, i.e.,

$$P\left(\xi \in \left\{z : f'(\theta_0, z) = \lim_{\delta \to 0} \frac{f(\theta_0 + \delta, z) - f(\theta_0, z)}{\delta}\right\}\right) = 1.$$

Then $f(\cdot)$ is differentiable at $\theta = \theta_0$, and

$$f'(\theta_0) = Ef'(\theta_0, \xi).$$

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 $m \geq 2$,

$$\begin{aligned} v(\theta) = EV(m,\theta) &:= E\frac{1}{m-1}\sum_{i=1}^{m}(X_i(\theta) - \bar{X}_m(\theta))^2 \\ &= E\frac{m}{m-1}\left(\frac{1}{m}\sum_{i=1}^{m}X_i^2(\theta) - \bar{X}_m^2(\theta)\right), \end{aligned} \tag{2.3.1}$$

where $(X_1, Y_1), \ldots, (X_m, Y_m)$ 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 (2.3.1) 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 2.3.2. *If **A1 - A4** hold then $v(\cdot)$ is $\mathcal{C}^1$ on $\mathcal{U}$ and for $\theta_0 \in \mathcal{U}$,*

$$\begin{aligned} g(\theta_0) &:= \nabla_\theta v(\theta)|_{\theta=\theta_0} \\ &= Eg_m(\theta_0) \end{aligned} \tag{2.3.2}$$

Proof. We apply Theorem 2.3.1 to the sample variance $V(m,\theta)$ component by component. Consider the jth component, for some $j = 1, \ldots, p$. The only condition that requires explicit verification is that $\partial V(m,\theta)/\partial\theta(j)$ is dominated by an integrable function of $(\mathbf{X}, \mathbf{Y}) = ((X_i, Y_i) : 1 \le i \le m)$. We have that

$$\frac{\partial V(m,\theta)}{\partial\theta(j)} = \frac{m}{m-1}\left(\frac{-1}{m}\sum_{i=1}^{m} 2X_i(\theta)\frac{\partial h(Y_i,\theta)}{\partial\theta(j)} + 2\bar{X}_m(\theta)\frac{1}{m}\sum_{i=1}^{m}\frac{\partial h(Y_i,\theta)}{\partial\theta(j)}\right). \tag{2.3.3}$$

The first term in the parentheses in (2.3.3) is integrable by **A4**. For the second term, we apply **A3** and split the sums to obtain

$$\begin{aligned}&\left|\bar{X}_m(\theta)\frac{1}{m}\sum_{i=1}^{m}\frac{\partial h(Y_i,\theta)}{\partial\theta(j)}\right| \\ &\quad\le \frac{1}{m^2}\sum_{i=1}^{m}\sup_{\theta\in\mathcal{U}}|X_i(\theta)|C(Y_i) + \frac{1}{m^2}\sum_{i=1}^{m}\sum_{k\neq i}\sup_{\theta\in\mathcal{U}}|X_i(\theta)|C(Y_k). \end{aligned}\tag{2.3.4}$$

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**,

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

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 $\theta \in \Theta$.

2.4 The Stochastic Approximation Method

Stochastic approximation is a class of stochastic optimization methods used to solve problems with differentiable objective functions. 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

$$\theta_{n+1} = \Pi_\Theta(\theta_n - a_n g_n(\theta_n)), \tag{2.4.1}$$

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

$$\sum_{n=1}^{\infty} a_n = \infty \text{ and } \sum_{n=1}^{\infty} a_n^2 < \infty. \tag{2.4.2}$$

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 as follows. Let $m \geq 2$ be a fixed positive integer.

In Section 2.4.1, we give conditions under which the stochastic approximation estimator μ_n is consistent and a central limit theorem is satisfied. We propose several estimators for the asymptotic variance in the central theorem, which provides a way to estimate a confidence interval for μ. Section 2.4.2 shows that under additional conditions θ_n converges to some random variable θ^* a.s. as $n \to \infty$.

Initialization: Choose θ_0.

For $k = 1$ to n

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

Compute

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

Next k

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

Figure 2.1: The stochastic approximation algorithm

2.4.1 Asymptotic Properties of the Stochastic Approximation Estimator

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

Theorem 2.4.1 (Liptser and Shiryayev 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 $L_\infty = \lim_{n\to\infty} L_n$. *Then*

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

Let $\mathcal{F}_n = \sigma\{(X_{k,i}, Y_{k,i}) : 1 \le k \le n, 1 \le i \le m\}$ 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 2.4.2. *Assume* ***A1-A4****. Then* $\mu_n \to \mu$ *a.s. as* $n \to \infty$.

Proof. For $k \ge 1$ and $n \ge 1$, 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 $M_n/n \to 0$ a.s. as $n \to \infty$.

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

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

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 2.4.1, $M_n/n \to 0$ a.s. as $n \to \infty$. □

Remark 3. The proof of Proposition 2.4.2 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 $\theta_k \to \theta^*$ a.s. as $k \to \infty$ 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 Shiryayev [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 $(\mathcal{F}_{k,n} : k = 1, \ldots, n)$ 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 $k = 1, \ldots, n$. 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 2.4.3 (Liptser and Shiryayev 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*

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

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) $\sum_{k=1}^{n} E(\xi_{k,n}^2 I(|\xi_{k,n}| > \delta)|\mathcal{F}_{k-1,n}) \to 0$ *in probability, for all* $\delta \in (0, 1]$,

(ii) $\sum_{k=1}^{n} E(\xi_{k,n}^2|\mathcal{F}_{k-1,n}) \to \eta^2$ *in probability, and*

(iii) $\sum_{k=1}^{\lfloor nc_n \rfloor} E(\xi_{k,n}^2|\mathcal{F}_{k-1,n}) \to 0$ *in probability*

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

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

as $n \to \infty$, 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

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

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

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

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 2.4.4. *Assume **A1-A5** and that $\theta_n \to \theta^*$ for some random variable θ^* a.s. as $n \to \infty$. 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 $n \to \infty$.

Proof. To show the central limit theorem we apply Theorem 2.4.3. 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 2.4.2, $(\xi_{k,n}, \mathcal{F}_{k,n} : n \geq 1, 1 \leq k \leq n)$ 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 $(X_1,Y_1),\ldots,(X_m,Y_m)$ are i.i.d. replications of (X,Y), independent of $(X_{k,i},Y_{k,i})$, $i = 1,\ldots,m$, $k \geq 1$. 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

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

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

The second and third conditions of Theorem 2.4.3 are easily dealt with. We see that

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

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

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

as $n \to \infty$. The central limit theorem is therefore a consequence of Theorem 2.4.3.

□

Hence we see that the stochastic approximation estimator μ_n satisfies a strong law and central limit theorem as $n \to \infty$. It will almost invariably be the case that $v(\theta^*) > 0$ a.s. so that the rate of convergence of μ_n is the canonical rate $n^{-1/2}$. This is the best that can be hoped for with the Monte Carlo nature of the estimation procedure we used.

Recall that our motivation for choosing $m > 1$ was to obtain an unbiased gradient estimator with low variance. This additional averaging of m terms in each step of the algorithm does not slow convergence, at least to first order, in the sense that the variance of the estimator and the limiting variance that appear in the central limit theorem are each reduced by a factor of m. Therefore the choice of $m \geq 2$ is essentially immaterial from the central-limit-theorem point of view. Of course, these are large sample results, and it may be beneficial to carefully choose m in small samples. We do not explore that possibility here.

In the rather special case where $v(\theta^*) = 0$ a.s. the central limit theorem above still holds in the sense that $\sqrt{n}(\mu_n - \mu) \Rightarrow 0$ as $n \to \infty$. The rate of convergence is then faster than $n^{-1/2}$, and its exact nature depends on the rate at which $\theta_n \to \theta^*$ a.s. We do not explore this case further here, because we believe that the case $v(\theta^*) = 0$ a.s. is unlikely to arise in the applications we have in mind. See Henderson and Simon [2004] for an exploration of increased convergence rates

when θ^* is constant and $v(\theta^*) = 0$.

The central limit theorem suggests a confidence interval procedure, provided that the variance can be estimated. Suppose that $\theta_k \to \theta^*$ a.s. for some fixed $\theta^* \in \Theta$, so that the variance appearing in the central limit theorem is deterministic and equal to $v(\theta^*)$. To estimate $v(\theta^*)$ we can use any one of the three estimators

$$\begin{aligned} S_n^2 &= \frac{1}{mn-1}\sum_{k=1}^{n}\sum_{i=1}^{m}\left(X_{k,i}(\theta_{k-1}) - \mu_n\right)^2, \\ \hat{S}_n^2 &= \frac{1}{n}\sum_{k=1}^{n}\left(\frac{1}{m-1}\sum_{i=1}^{m}(X_{k,i}(\theta_{k-1}) - A_k(\theta_{k-1}))^2\right), \text{ and} \\ \tilde{S}_n^2 &= \frac{m}{n-1}\sum_{k=1}^{n}(A_k(\theta_{k-1}) - \mu_n)^2. \end{aligned} \tag{2.4.3}$$

The estimator S_n^2 is the sample variance using all mn samples, $\hat{S}_n^2$ is the average of the sample variances of m terms in each iteration, and $\tilde{S}_n^2$ is m times the sample variance of the averages computed at each iteration. The following proposition shows that all three estimators are strongly consistent, so they can be used to construct asymptotically valid confidence intervals.

Proposition 2.4.5. *Assume **A1-A4** and that θ_n converges to some fixed $\theta^* \in \Theta$ a.s. Then*

(i) $S_n^2, \hat{S}_n^2, \tilde{S}_n^2 \to v(\theta^*)$ as $n \to \infty$ a.s.

*(ii) Assume also **A5**, and that $v(\theta^*) > 0$. Then*

$$\frac{\sqrt{nm}(\mu_n - \mu)}{\eta_n} \Rightarrow N(0,1)$$

as $n \to \infty$, where η_n can be $S_n, \hat{S}_n$ or $\tilde{S}_n$.

Proof. For part (i), write

$$\begin{aligned} S_n^2 &= \frac{1}{nm-1}\sum_{k=1}^{n}\sum_{i=1}^{m} X_{k,i}^2(\theta_{k-1}) - \frac{nm}{nm-1}\mu_n^2 \\ &= \frac{1}{nm-1}\sum_{k=1}^{n}\sum_{i=1}^{m} X_{k,i}^2(\theta^*) - \frac{nm}{nm-1}\mu_n^2 \qquad (2.4.4) \\ &\quad + \frac{1}{nm-1}\sum_{k=1}^{n}\sum_{i=1}^{m}(X_{k,i}^2(\theta_{k-1}) - X_{k,i}^2(\theta^*)) \qquad (2.4.5) \end{aligned}$$

By the SLLN and Proposition 2.4.2,

$$\frac{1}{nm-1}\sum_{k=1}^{n}\sum_{i=1}^{m} X_{k,i}^2(\theta^*) - \frac{nm}{nm-1}\mu_n^2 \to E(X_1^2(\theta^*)) - \mu^2 = v(\theta^*)$$

as $n \to \infty$ a.s. Therefore it suffices to show that the last term in (2.4.5) converges to 0 a.s. as $n \to \infty$.

Since $\theta_k \to \theta^*$ as $k \to \infty$ a.s., for any given $\epsilon > 0$, there exists a random N such that for all $k \geq N$, $\|\theta^* - \theta_k\| < \epsilon$ a.s. Then

$$\begin{aligned} &\frac{1}{nm-1}\sum_{k=1}^{n}\sum_{i=1}^{m}(X_{k,i}^2(\theta_{k-1}) - X_{k,i}^2(\theta^*)) \\ &\leq \frac{1}{nm-1}\sum_{k=1}^{n}\sum_{i=1}^{m} |X_{k,i}(\theta_{k-1}) - X_{k,i}(\theta^*)|\ |X_{k,i}(\theta_{k-1}) + X_{k,i}(\theta^*)| \\ &\leq \frac{2}{nm-1}\sum_{k=1}^{n}\sum_{i=1}^{m} C(Y_{k,i}) \sup_{\theta\in\mathcal{U}} |X_{k,i}(\theta)| \|\theta_{k-1} - \theta^*\| \\ &\leq \frac{2}{nm-1}\sum_{k=1}^{N}\sum_{i=1}^{m} C(Y_{k,i}) \sup_{\theta\in\mathcal{U}} |X_{k,i}(\theta)| \|\theta_{k-1} - \theta^*\| \qquad (2.4.6) \\ &\quad + \frac{2}{nm-1}\sum_{k=N+1}^{n}\sum_{i=1}^{m} C(Y_{k,i}) \sup_{\theta\in\mathcal{U}} |X_{k,i}(\theta)|\epsilon \qquad (2.4.7) \end{aligned}$$

Now, (2.4.6) converges to 0 a.s. as $n \to \infty$ since N is finite. **A4** implies that $C(Y_1)\sup_{\theta\in\mathcal{U}}|X_1(\theta)|$ is integrable and hence the SLLN ensures that

$$\frac{2}{nm-1}\sum_{k=N+1}^{n}\sum_{i=1}^{m} C(Y_{k,i}) \sup_{\theta\in\mathcal{U}} |X_{k,i}(\theta)|\epsilon \to 2\epsilon E\left(C(Y_1)\sup_{\theta\in\mathcal{U}}|X_1(\theta)|\right)$$

as $n \to \infty$ a.s. Since ϵ is arbitrary, (2.4.7) converges to 0 a.s. as $n \to \infty$.

Essentially the same argument can be applied to $\hat{S}_n$ and $\tilde{S}_n$. We omit the details.

Part (ii) is an immediate consequence of Part (i) and the converging together lemma (e.g., Chung [1974, p. 93]). □

Under the conditions of Proposition 2.4.5(ii), an asymptotic $100(1-\alpha)\%$ confidence interval for μ is

$$\left[\mu_n - z\frac{\eta_n}{\sqrt{nm}}, \mu_n + z\frac{\eta_n}{\sqrt{nm}}\right],$$

where η_n can be $S_n, \hat{S}_n$ or $\tilde{S}_n$ and z is chosen such that $P(-z \leq N(0,1) \leq z) = 1-\alpha$.

But which variance estimator should we use? Some insight into this question can be obtained by assuming that $\theta_k = \theta^*$ for all k, and then considering the second-order behavior of the variance estimators as given by central limit theorems. This case is easier to analyze than the general case because the $X_{k,i}(\theta^*)$s are i.i.d. random variables.

Proposition 2.4.6. *Suppose that $\theta_k = \theta^*$ for all $k \geq 0$. Suppose that $EX^4(\theta^*) < \infty$. Then*

$$\begin{aligned}
\sqrt{mn}(S_n^2 - v(\theta^*)) &\Rightarrow \sigma N(0,1), \\
\sqrt{mn}(\hat{S}_n^2 - v(\theta^*)) &\Rightarrow \hat{\sigma} N(0,1), \text{ and} \\
\sqrt{mn}(\tilde{S}_n^2 - v(\theta^*)) &\Rightarrow \tilde{\sigma} N(0,1)
\end{aligned}$$

as $n \to \infty$*, where*

$$\begin{aligned}
\sigma^2 &= E[X_1(\theta^*) - \mu]^4 - v^2(\theta^*), \\
\hat{\sigma}^2 &= E[X_1(\theta^*) - \mu]^4 - \frac{m-3}{m-1} v^2(\theta^*), \text{ and} \\
\tilde{\sigma}^2 &= E[X_1(\theta^*) - \mu]^4 + (2m-3) v^2(\theta^*).
\end{aligned}$$

Proof. First consider S_n^2. Notice that the $X_{k,i}(\theta^*)$s are i.i.d. Therefore

$$\begin{aligned}
&\sqrt{nm}\left(S_n^2(\theta^*) - v(\theta^*)\right) \\
&= \sqrt{nm}\left(\frac{1}{nm-1}\sum_{k=1}^{n}\sum_{i=1}^{m} X_{k,i}^2(\theta^*) - \frac{nm}{nm-1}\mu_n^2 - v(\theta^*)\right) \\
&= \sqrt{nm}\left(\frac{1}{nm}\sum_{k=1}^{n}\sum_{i=1}^{m} X_{k,i}^2(\theta^*) - \mu_n^2 - v(\theta^*) + o_p((nm)^{-1/2})\right).
\end{aligned}$$

Let $g(x, y) = x - y^2$. Then

$$\frac{1}{nm}\sum_{k=1}^{n}\sum_{i=1}^{m} X_{k,i}^2(\theta^*) - \mu_n^2 - v(\theta^*) = g(\frac{1}{nm}\sum_{k=1}^{n}\sum_{i=1}^{m} X_{k,i}^2(\theta^*), \mu_n) - g(E(X_1^2(\theta^*)), \mu).$$

By the delta method,

$$\sqrt{nm}\left(g(\frac{1}{nm}\sum_{k=1}^{n}\sum_{i=1}^{m} X_{k,i}^2(\theta^*), \mu_n) - g(E(X_1^2(\theta^*)), \mu)\right) \Rightarrow \sigma N(0,1),$$

where

$$\begin{aligned}
\sigma^2 &= \nabla g\left(E[X_1(\theta^*)]^2, \mu\right)^T \operatorname{cov}(X_1^2(\theta^*), X_1(\theta^*)) \nabla g\left(E(X_1^2(\theta^*)), \mu\right) \\
&= E(X_1^4(\theta^*)) - 4\mu E(X_1^3(\theta^*)) + 8\mu^2 E(X_1^2(\theta^*)) - [E(X_1^2(\theta^*))]^2 - 4\mu^4 \\
&= E(X_1(\theta^*) - \mu)^4 - v^2(\theta^*).
\end{aligned}$$

The central limit theorem for $\hat{S}_n^2$ follows from the ordinary central limit theorem. We get

$$\sqrt{nm}(\hat{S}_n^2(\theta^*) - v(\theta^*)) \Rightarrow \hat{\sigma} N(0,1),$$

where

$$\begin{aligned}\hat{\sigma}^2 &= m \operatorname{var}\left(\frac{1}{m-1}\sum_{i=1}^{m}(X_{1,i}(\theta^*) - A_1(\theta^*))^2\right) \\ &= m\frac{1}{m}\left(E(X_1(\theta^*)-\mu)^4 - \frac{m-3}{m-1}E(X_1(\theta^*)-\mu)^2\right) \\ &= E(X_1(\theta^*)-\mu)^4 - \frac{m-3}{m-1}v^2(\theta^*).\end{aligned}$$

(The second equality above requires some algebra.)

The proof of the central limit theorem for $\tilde{S}_n^2$ follows essentially the same argument that we used for S_n^2 and is omitted. □

Notice that $\tilde{\sigma}^2 > \sigma^2, \hat{\sigma}^2$ for $m \geq 2$, so on that basis we prefer either S_n^2 or $\hat{S}_n^2$ to $\tilde{S}_n^2$. The difference between σ^2 and $\hat{\sigma}^2$ is much smaller and vanishes as m grows. So the choice between these estimators essentially comes down to computational convenience, so long as m is large enough. We used $\hat{S}_n^2$ in our experiments.

2.4.2 Convergence of the Stochastic Approximation Algorithm

We now give conditions under which θ_n converges to some random variable θ^* a.s. as $n \to \infty$. Theorem 2.4.7 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

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

where $a(i), b(i) \in \mathbb{R}$ and $a(i) \leq b(i), i = 1, \ldots, p$. 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 $f : B \to \mathbb{R}$ 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

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

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

Let $(\mathcal{G}_n : n \ge 0)$ 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 $n \ge 0$.

Theorem 2.4.7. *Let B be a box in $\mathbb{R}^p$ and $f : \mathbb{R}^p \to \mathbb{R}$ be $\mathcal{C}^1$. Suppose that for $n \ge 0$, $\theta_{n+1} = \Pi_B(\theta_n - a_n G_n)$ with the following additional conditions.*

(i) The conditions (2.4.2) hold.

(ii) $\sup_n E\|G_n\|^2 < \infty$.

(iii) $E[G_n|\mathcal{G}_n] = \nabla f(\theta_n)$ *for all* $n \ge 0$.

Then,

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

as $n \to \infty$ 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 $n \to \infty$.

Notice that the point in $S(f, B)$ that θ_n converges to can be random. We can apply Theorem 2.4.7 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 fourth moments.

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

Corollary 2.4.8. *Let Θ be a box in $\mathbb{R}^p$ and suppose **A1 - A4, A6** hold. Then $d(\theta_n, S(v, \Theta)) \to 0$ as $n \to \infty$ 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 $n \to \infty$.*

Proof. The only condition of Theorem 2.4.7 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 the one used in Proposition 2.3.2 and is omitted. It is this argument that requires the stronger moment assumption **A6**. □

Corollary 2.4.8 does not ensure that θ_n converges to a deterministic θ^* as $n \to \infty$. 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 Chapter 3 that the stochastic approximation procedure works well so long as the step size parameters of the procedure are chosen appropriately. However, the selection of the parameters is a nontrivial problem. Various procedures have been developed where the step size parameters are adaptively updated as the number of iterations n grows [Ruppert, 1985]. But with any stochastic approximation procedure, it can be still difficult to select good values for these parameters. For this reason we also consider a second estimator based on quite a different approach.

2.5 The Sample Average Approximation Method

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

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

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

$$\hat{\mu}_n = \frac{1}{n} \sum_{i=1}^{n} [X_i - h(Y_i, \hat{\theta})]. \tag{2.5.1}$$

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. This is the case in the sample average approximation (SAA) method that we now adopt [Shapiro, 2003].

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

$$V(m, \theta) = \frac{m}{m-1} \left(\frac{1}{m} \sum_{i=1}^{m} \tilde{X}_i^2(\theta) - \bar{X}_m^2(\theta) \right), \tag{2.5.2}$$

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 (2.5.2) 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 the problem $\mathcal{P}_m$ obtained from the first phase. In the second phase, we then estimate μ via the sample average (2.5.1), using $\hat{\theta}_m$ in place of $\hat{\theta}$. Our sample average approximation algorithm for estimating μ is given in Figure 2.2.

In Section 2.5.1 we show that the sample average approximation estimator $\hat{\mu}_n$ satisfies a strong law and central limit theorem. These results require a little care, because $\hat{\theta}_m$ is a random variable. We show in Section 2.5.2 that the set of first-order critical points for the SAA problem $\mathcal{P}_m$ converges to the set of first-order critical points for the original problem with probability 1 as the sample size m gets large. The optimal choice of m is an important issue from an implementation

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

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

For a fixed θ, define

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

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

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

Figure 2.2: The sample average approximation algorithm

standpoint. Section 2.5.3 provides an approximate form for the optimal m when the computational budget is fixed. The behaviour of the optimal m depends on the characteristics of the original optimization problem.

2.5.1 Asymptotic Properties of the Sample Average Approximation Estimator

The results in this section are based on a uniform version of the strong law of large numbers (ULLN). The following proposition, which appears as Proposition 7 in Shapiro [2003], provides conditions for ULLN. 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 2.5.1 (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*

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

as $n \to \infty$ 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 2.5.2. *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 $\hat{\mu}_n(\hat{\theta}) \to \mu$ as $n \to \infty$ a.s., and*

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

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

Proof. For the strong law note that

$$\begin{aligned} |\hat{\mu}_n(\hat{\theta}) - \mu| &\leq \left|\frac{1}{n}\sum_{i=1}^{n}(X_i-\mu)\right| + \left|\frac{1}{n}\sum_{i=1}^{n} h(Y_i,\hat{\theta})\right| \\ &\leq \left|\frac{1}{n}\sum_{i=1}^{n}(X_i-\mu)\right| + \sup_{\theta\in\Theta}\left|\frac{1}{n}\sum_{i=1}^{n} h(Y_i,\theta)\right|. \end{aligned} \tag{2.5.3}$$

The first term in (2.5.3) converges to 0 as $n \to \infty$ by the strong law of large numbers. The second term converges to 0 by an application of Proposition 2.5.1.

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 $x \in \mathbb{R}$,

$$P\left(\sqrt{n}(\hat{\mu}_n(\hat{\theta}) - \mu) \leq x \,|\, \hat{\theta}\right) \to \Phi\left(\frac{x}{v^{1/2}(\hat{\theta})}\right) 1_{\{v(\hat{\theta})>0\}} + 1_{\{x\geq 0\}} 1_{\{v(\hat{\theta})=0\}} \qquad (2.5.4)$$

as $n \to \infty$, where Φ is the distribution function of a normal random variable with mean 0 and variance 1, and $1_{\{\cdot\}}$ is an indicator function. The dominated convergence theorem ensures that we can take expectations through (4.2.14), and so

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

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 $m(n) \to \infty$ as $n \to \infty$. In this case, $\hat{\mu}_n(\hat{\theta}_{m(n)})$ behaves the same as $\hat{\mu}_n(\theta^*)$ as $n \to \infty$, at least to first order.

Theorem 2.5.3. *Suppose that $\hat{\theta}_{m(n)} \to \theta^*$ as $n \to \infty$ 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 $\hat{\mu}_n(\hat{\theta}_{m(n)}) \to \mu$*

as $n \to \infty$ *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 $n \to \infty$*.*

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

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

Notice that θ^* is independent of the samples used to compute $\hat{\mu}_n$ for every n. By Theorem 4.2.4, 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 $n \to \infty$.

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

$$\begin{aligned}P(|D_{2,n}| > \epsilon) &\leq \epsilon^{-2} E D_{2,n}^2 \\ &= \frac{1}{n\epsilon^2} \sum_{j=1}^{n} E[h(Y_j, \hat{\theta}_{m(n)}) - h(Y_j, \theta^*)]^2 \qquad (2.5.5) \\ &= \frac{1}{\epsilon^2} E[h(Y_1, \hat{\theta}_{m(n)}) - h(Y_1, \theta^*)]^2. \qquad (2.5.6)\end{aligned}$$

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

$$[h(Y_1, \hat{\theta}_{m(n)}) - h(Y_1, \theta^*)]^2 \leq K^2(Y_1) \, \|\hat{\theta}_{m(n)} - \theta^*\|^2. \qquad (2.5.7)$$

The normed term in (2.5.7) is bounded, and so the dominated convergence theorem implies that (2.5.6) converges to 0 as $n \to \infty$. □

2.5.2 Convergence of the Solutions of the Sample Average Approximation Problem

Theorem 2.5.3 requires that the sequence of the first-stage solutions $\{\hat{\theta}_m\}$ converges to a random variable θ^* as $m \to \infty$ a.s. If the problem $\mathcal{P}$ has a unique optimal solution θ^*, and $\hat{\theta}_m$ solves the problem $\mathcal{P}_m$ exactly, then, as in Shapiro [2003], this requirement would follow using standard arguments and an extension of a uniform law of large numbers to nonlinear functions of means. (Recall from (2.5.2) 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. [2007] 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 the 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 2.5.4. *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 $\xi_1, \ldots, \xi_m$ are i.i.d. and*

(i) for all $\xi \in H$, $f(\cdot, \xi) = (f_1(\cdot, \xi), \ldots, f_d(\cdot, \xi))$ is $\mathcal{C}^1$ on an open set containing

Θ,

(ii) the component functions $f_i(\theta, \xi)$ $(i = 1, \ldots, d)$ *are dominated by an integrable function, and*

(iii) the gradient components $\partial f_i(\theta, \xi)/\partial\theta(j)$ *are dominated by an integrable function* $(i = 1, \ldots, d,\ j = 1, \ldots, p)$.

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

Proof. If $d(\hat{\theta}_m, S(g \circ \bar{f}, \Theta)) \not\to 0$, 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 2.5.1, $\bar{f}_m(\hat{\theta}_m) \to \bar{f}(\theta^*)$ and $\nabla_\theta \bar{f}_m(\hat{\theta}_m) \to \nabla_\theta \bar{f}(\theta^*)$ as $m \to \infty$ 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 $m \to \infty$, 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 2.5.5. *Suppose that* ***A1-A4*** *hold,* Θ *is convex and* $EK^2(Y) < \infty$. *Then* $d(\hat{\theta}_m, S(v, \Theta)) \to 0$ *as* $m \to \infty$ *a.s.*

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

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

Notice that

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

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

By the proof of Proposition 2.3.2 and Remark 4,

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

are all dominated by an integrable function ($i = 1,\ldots,p$). By Theorem 2.5.4, it follows that

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

as $m \to \infty$. □

Corollary 2.5.5 shows that $\hat{\theta}_m$ converges to the set of first-order critical points of v as $m \to \infty$. 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.

2.5.3 Allocation of Computational Budget

The limiting results in the previous two sections 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 $m(c), n(c) \to \infty$ as $c \to \infty$ to ensure that $\hat{\theta}_{m(c)} \to \theta^*$ and $\hat{\mu}_{n(c)}(\hat{\theta}_{m(c)}) \to \mu$. 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 behaviour 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 $\nabla^2_\theta h(y, \cdot)$ 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, the 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

$$v(\theta) - v(\theta^*) \leq \lambda \|\theta - \theta^*\|^2, \tag{2.5.8}$$

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

$$\begin{aligned} m^* &\approx \operatorname{argmin} \left\{ \frac{1}{n} E(v(\hat{\theta}_m)) = \frac{v(\theta^*)m + \gamma}{m(c - \alpha m)} : 1 \leq m \leq \frac{c}{\alpha} \right\} \\ &= \frac{\sqrt{(\gamma\alpha)^2 + \gamma\alpha v(\theta^*)c} - \gamma\alpha}{\alpha v(\theta^*)}. \end{aligned} \tag{2.5.9}$$

The expression (2.5.9) is asymptotically (i.e., as $c \to \infty$) 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 (2.5.8), 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^*.

Chapter 3

Numerical Results

In this chapter we examine the performance of the adaptive control variate methods discussed in Chapter 2 on two examples. In Section 3.1 we consider a problem of estimating accrued costs till absorption for discrete time Markov chains on a finite state space. We describe how to construct control variate estimators for the Markov chains and discuss implementation details of the stochastic approximation and sample average approximation algorithms. In Section 3.2 we return to the barrier call option example presented in Section 2.1. We discuss how to choose a good parameterization for the control variate estimators and address the properties and implementation issues of our methods. Section 3.3 contains some concluding remarks.

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.

3.1 Accrued Costs Prior to Absorption

The objective of this example is to demonstrate the feasibility of our methods rather than to provide a comprehensive comparison. In Section 3.1.1 we describe our example, estimating accumulated cost till absorption, and in Section 3.1.2 we discuss the implementation of our methods. The results in Section 3.1.3 show that both adaptive methods outperform a naïve approach.

3.1.1 Construction of Martingale Control Variates

Let $Z = (Z_n : n \geq 0)$ be a discrete time Markov chain on the finite state space $S = \{0, 1, \ldots, d\}$. Suppose that Z reaches the absorbing state 0 almost surely starting from any $Z_0 > 0$, and let $T = \inf\{n \geq 0 : Z_n = 0\}$ be the time till absorption. Let $f : S \to \mathbb{R}$ be a given cost function. Define

$$\mu(x) = E\left(\sum_{k=0}^{T-1} f(Z_k) | Z_0 = x\right) \tag{3.1.1}$$

for all $x \in S - \{0\}$ and set $\mu(0) = 0$, so that μ is the expected cost accrued until absorption. If we view f and μ as column vectors, then μ satisfies

$$\mu = f + P\mu,$$

where P is the transition matrix of Z, and we take $f(0) = 0$. Suppose that μ is unknown and that we wish to estimate it.

Now we show how to define an approximating martingale control variate for μ. Let $u : S \to \mathbb{R}$ be a real-valued function on the state space S with $u(0) = 0$, and for $n \geq 0$ let

$$M_n(u) = u(Z_n) - u(Z_0) - \sum_{j=0}^{n-1} [(P - I)u](Z_j),$$

where I is the identity matrix. Then $(M_n(u) : n \geq 0)$ is the well-known Dynkin martingale; see, e.g., Karlin and Taylor [1981, p. 308]. The optional sampling theorem ensures that $E_x M_T(u) = 0$ for any u, where E_x denotes expectation under the initial condition $Z_0 = x$. Therefore, one can estimate $\mu(x)$ via i.i.d. replications of

$$\left[\sum_{k=0}^{T-1} f(Z_k)\right] - M_T(u),$$

where $Z_0 = x$ and $M_T(u)$ serves as a parameterized control variate. In our general notational scheme, X is the accrued cost till absorption and $Y(\theta)$ is $M_T(u)$, where

u depends on a parameter θ as described below. By (3.1.1),

$$\sum_{k=0}^{T-1} f(Z_k) - M_T(\mu) = \mu(x)$$

and hence, if $u = \mu$, then we have a zero-variance estimator.

So it is desirable to find a good choice of the function u. Suppose that $u(x) = u(x;\theta)$, where $\theta \in \Theta \subseteq \mathbb{R}^p$ is a $p-$dimensional vector of parameters. A linear parameterization arises if

$$u(x;\theta) = \sum_{i=1}^{p} \theta(i) u_i(x),$$

where $u_i(\cdot)$ are given basis functions, $i = 1, \ldots, p$. In this case,

$$\begin{aligned} M_n(u) &= u(Z_n;\theta) - u(Z_0;\theta) - \sum_{j=0}^{n-1} [(P-I)u](Z_j;\theta) \\ &= \sum_{i=1}^{p} \theta(i) u_i(Z_n) - \sum_{i=1}^{p} \theta(i) u_i(Z_0) - \sum_{j=0}^{n-1} [(P-I) \sum_{i=1}^{p} \theta(i) u_i](Z_j) \\ &= \sum_{i=1}^{p} \theta(i) \left[u_i(Z_n) - u_i(Z_0) - \sum_{j=0}^{n-1} [(P-I)u_i](Z_j) \right] \\ &= \sum_{i=1}^{p} \theta(i) M_n(u_i) \end{aligned} \tag{3.1.2}$$

so that $M_n(u)$ is simply a linear combination of martingales corresponding to the basis functions u_i, $i = 1, \ldots, p$. Therefore, the control variate

$$Y(\theta) = \sum_{i=1}^{p} \theta(i) M_T(u_i)$$

is simply a linear combination of zero-mean random variables. In this sense, the linearly parameterized case leads us back to the theory of linear control variates.

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

$$u(x;\theta) = \theta(1) x^{\theta(2)},$$

where $p = 2$. In this case, $Y(\theta)$ is a nonlinear function of a random object Y (the Markov chain Z) and a parameter vector θ.

In Section 3.1.2 and Section 3.1.3, we focus on the case where $u(\cdot;\theta)$ is nonlinearly parameterized and apply our adaptive methods to select a good choice of θ and estimate μ.

3.1.2 Implementation

We first verify that **A1-A6** in Chapter 2 are satisfied. Let $u(\cdot;\theta)$ be given, where $u(0;\theta) = 0$ for all $\theta \in \Theta$. Let $M_T(u(\theta)) = -u(x;\theta) - \sum_{j=0}^{T-1}(P - I)u(Z_j;\theta)$ under some fixed initial state $Z_0 = x$. Then $X(\theta) = X - M_T(u(\theta))$ is an estimator of $\mu(x)$. Let $V = (0, V(1), \ldots, V(d))^\top$, where $V(j) = \sum_{k=0}^{T-1} I(Z_k = j)$ is the number of visits to state j before absorption. Then

$$\begin{aligned} X(\theta) &= \sum_{j=0}^{T-1} f(Z_j) + u(x;\theta) + \sum_{j=0}^{T-1}[(P - I)u(\theta)](Z_j) \\ &= u(x;\theta) + \sum_{j=0}^{T-1}[(P - I)(u(\theta) - \mu)](Z_j) \\ &= u(x;\theta) + \sum_{k=0}^{d} V(k)[(P - I)(u(\theta) - \mu)](k) \\ &= u(x;\theta) + V^\top(P - I)(u(\theta) - \mu). \end{aligned}$$

To verify that **A1-A6** are satisfied we proceed as follows. First suppose that Θ is convex and compact, that there exists a bounded open set $\mathcal{U}$ such that $\Theta \subset \mathcal{U}$, and that $u(y;\cdot) : \mathcal{U} \to \mathbb{R}$ is $\mathcal{C}^1$ and Lipschitz for all $y \in S$ (these assumptions are all satisfied in our particular example below). Since S is finite and $\mathcal{U}$ is bounded, there exists a $K > 0$ such that for all $\theta_1, \theta_2 \in \mathcal{U}$ and $y \in S$,

$$|u(y;\theta_1) - u(y;\theta_2)| \leq K\|\theta_1 - \theta_2\|,$$

and $\{u(y;\theta), \frac{\partial}{\partial\theta(i)}u(y;\theta) : \theta \in \mathcal{U}, y \in S, i = 1, \ldots, p\}$ are uniformly bounded, i.e.

$$C = \sup_{\theta\in\mathcal{U}, y\in S, i=1,\ldots,p} \left\{ |u(y;\theta)|, \left| \frac{\partial u(y;\theta)}{\partial\theta(i)} \right| \right\} < \infty.$$

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

$$\begin{aligned} |M_T(u(\theta_1)) - M_T(u(\theta_2))| &\leq |u(x;\theta_1) - u(x;\theta_2)| \\ &+ \sum_{j=0}^{T-1} |[(P-I)(u(\theta_1) - u(\theta_2))](Z_j)| \\ &\leq K\|\theta_1 - \theta_2\| + T\|P - I\| \, \|u(\theta_1) - u(\theta_2)\| \\ &\leq K\|\theta_1 - \theta_2\| + T\|P - I\| \cdot dK\|\theta_1 - \theta_2\|. \end{aligned}$$

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

$$\begin{aligned} |X(\theta)| &\leq |u(x;\theta)| + |V^\top(P-I)(u(\theta) - \mu)| \\ &\leq |u(x;\theta)| + \|V^\top(P-I)\| \, \|(u(\theta) - \mu)\| \\ &\leq C + dT\|(P-I)\|(dC + \|\mu\|), \end{aligned}$$

and similarly,

$$\begin{aligned} |\frac{\partial}{\partial\theta(i)}X(\theta)| &\leq |\frac{\partial}{\partial\theta(i)}u(x;\theta)| + |V^\top(P-I)(\frac{\partial}{\partial\theta(i)}u(\theta) - \mu)| \\ &\leq C + dT\|(P-I)\|(dC + \|\mu\|). \end{aligned}$$

Since all of these bounds depend only on the random variable T, which has a finite moment generating function in a neighborhood of 0, we can easily verify that assumptions **A1-A6** are satisfied.

For the simulation experiment, we use the "random walk" transition matrix P

given by

$$P = \begin{pmatrix} 1 & 0 & 0 & 0 & \dots & 0 & 0 & 0 \\ q(1) & 0 & p(1) & 0 & \dots & 0 & 0 & 0 \\ 0 & q(2) & 0 & p(2) & \dots & 0 & 0 & 0 \\ \vdots & \ddots & & \ddots & & & & \vdots \\ 0 & 0 & 0 & 0 & \dots & q(d-1) & 0 & p(d-1) \\ 0 & 0 & 0 & 0 & \dots & 0 & 1 & 0 \end{pmatrix},$$

where $q(i) > 0$ for all $i = 1, \dots, d-1$. We take

$$u(y; \theta) = \theta(1) y^{\theta(2)},$$

where $\theta = (\theta_1, \theta_2) \in \Theta$, $\Theta = \{x \in \mathbb{R}^2 : a(j) \leq x(j) \leq b(j), j = 1, 2\}$ and $a(j) \geq 0, j = 1, 2$. Then $u(y; \cdot)$ is $\mathcal{C}^1$ for all $y \in S$ and the moment generating function of T is defined in a neighborhood of 0. We took $d = 30$ and $f(x) = 1$, so that the random variable $X = T$ is the time till absorption in state 0.

In the stochastic approximation algorithm, we took $m = 100$ and

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

where e, $A > 0$ and $\alpha \in (1/2, 1]$ are tunable constants. This form of the gain sequence is advocated in Spall [2003]. We used the average of the sample variances of m terms in each step as an estimator of $v(\theta^*)$. For the SAA estimator, we first replicated $m = 100$ samples. We obtained $\hat{\theta}_m$ by applying a quasi-Newton method with a linesearch (supplied as part of the MATLAB$^{\text{TM}}$ package) using IPA gradients to solve the sample average approximation problem $\mathcal{P}_m$. 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 4.2.4.

3.1.3 Simulation Results

The values $V_{\text{naïve}}$, V_{SA} and V_{SAA} are, respectively, the estimated variances obtained from the naïve, SA and SAA estimators. We used the same CPU time for all three estimators for a given initial state x to allow a fair comparison.

Example 2. In this example, we let $p(x) = .25$ and $\theta_0 = (1, 1)$. Table 3.1 shows that the SAA estimators outperform the SA estimators, and the SA estimators outperform the naive estimator. 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. A contour plot of the variance surface as a function of θ for initial state $x = 15$ appears in Figure 3.1. We see that the function is not convex, but appears to have a unique first-order critical point, so that we can expect convergence of the parameter estimates to θ^*. In the plot, this appears to be the point $(2, 1)$.

Table 3.1: Estimated squared standard errors in Example 2

x	CPU time (sec)	$V_{\text{naïve}}/V_{SA}$	$V_{\text{naïve}}/V_{SAA}$
5	16.8	19	2.6 E+10
10	20.2	21	2.9 E+10
15	21.8	32	8.9 E+10
20	25.8	23	6.4 E+11
25	28.6	5.0	3.6 E+3
30	29.8	1.9	91

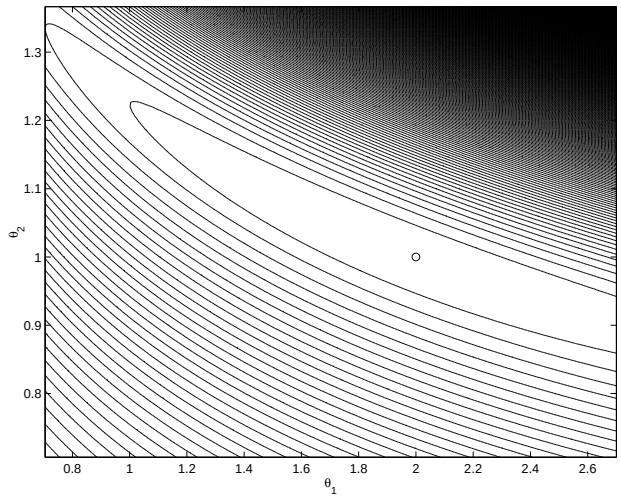

Figure 3.1: Contour Plot of $v(\cdot)$ for Example 2 with initial state $x = 15$ and runlength 1000

Remark 6. If the simulation run length n is long enough, then from Theorems 2.4.4 and 2.5.3 we would expect the SA and SAA estimators to be fairly similar in performance.

Example 3. In this example, $p(x) = .0001 + .4998/x$ and $\theta_0 = (2, 1)$. The results are given in Table 3.2 and are similar to those of Example 2. The SAA estimator outperforms the other estimators, but not by as large a margin.

3.2 Pricing Barrier Options

In this section we return to the discretely monitored barrier call option example presented in Section 2.1. Selecting a good parameterization is crucial to obtaining efficient control variate estimators. We discuss how to find a good parameterization in practice with this example and address issues which arise in the application of our methods.

Table 3.2: Estimated squared standard errors in Example 3

x	CPU time (sec)	$V_{\text{naïve}}/V_{SA}$	$V_{\text{naïve}}/V_{SAA}$
5	15.5	6.4	3.4 E+3
10	17.0	9.5	85
15	17.6	14	57
20	19.5	2.1	44
25	21.2	7.3	37
30	21.8	2.6	36

3.2.1 Implementation

We assume that under the risk-neutral measure, the underlying stock price $\{S(t) : t \geq 0\}$ is governed by the dynamics

$$\frac{dS(t)}{S(t)} = rdt + \sigma dW(t), \tag{3.2.1}$$

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

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

where $Z_1, \ldots, Z_l$ 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 3.2 displays surface plots of the estimated expected payoff function $U^*(x,i)$. For any fixed $i=1,..,l-1$, $U^*(x,i)$ 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 3.2. Let our parameterization have the form

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

where $\theta=(\theta_1,\theta_2,\ldots,\theta_{4(l-1)})\in\Theta$, $\Theta=\{y\in\mathbb{R}^{4(l-1)}: a(j)\leq y(j)\leq b(j), j=1,2,\ldots,4(l-1)\}$ and $a(j)\geq 0$, $j=1,2,\ldots,4(l-1)$. (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 $U(x,i;\cdot):\mathbb{R}^{4(l-1)}\to\mathbb{R}$ is $\mathcal{C}^1$ for all $(x,i)\in\mathcal{S}\times\{0,1,\ldots,l-1\}$ and $U(\cdot,i;\cdot):(0\ \infty)\times\mathbb{R}^{4(l-1)}\to\mathbb{R}$ is

$\mathcal{C}^1$ for all $i \in \{0, 1, \ldots, l-1\}$. Details on both the verification of **A1-A6**, and the computation of the control variate $M_l(U(\theta))$ are given in Appendix.

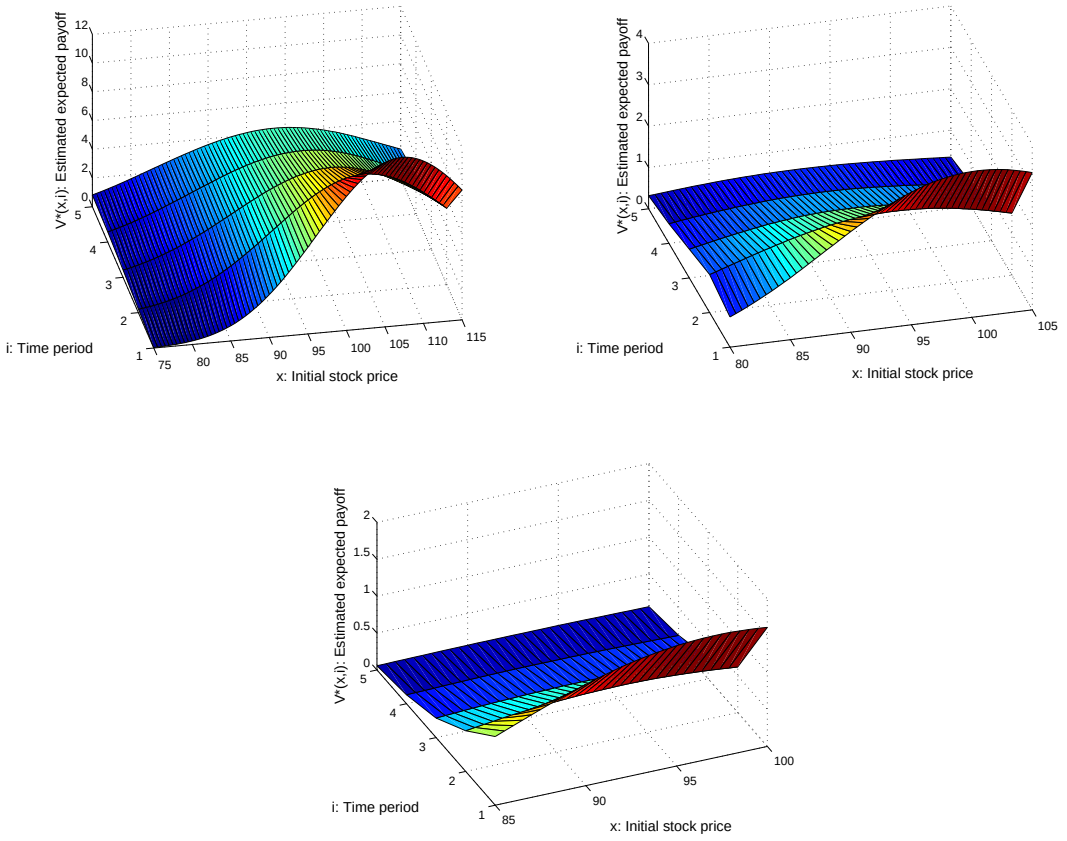

Figure 3.2: Surface plots of the estimated expected payoff $U^*(x, i)$. Upper left: $\sigma = .4$, $l = 6$ and barriers at $H_l = 75$ and $H_u = 115$. Upper right: $\sigma = .6$, $l = 6$ and barriers at $H_l = 80$ and $H_u = 105$. Lower: $\sigma = .6$, $l = 6$ and barriers at $H_l = 85$ and $H_u = 100$.

3.2.2 Simulation Results

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 3.3 - 3.5 report numerical results for options with various volatilities, monitoring dates and barriers. In the stochastic approximation algorithm, we took $m = 500$ and used (2.4.3) 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 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.

In Tables 3.3 - 3.5, 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 $V_{\text{naïve}}$, 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 2.4.4 and 2.5.3 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 3.2, 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 monitor-

ing 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.

Table 3.3: Estimated variance reduction ratio: $H_l = 75$ and $H_u = 115$

Volatility	Frequency of monitoring	CPU time (sec)	SA ratio	$V_{\text{naïve}}/V_{SA}$	SAA ratio	$V_{\text{naïve}}/V_{SAA}$
$\sigma = .4$	$l = 3$	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

3.3 Concluding Remarks

The two adaptive estimation procedures developed in Chapter 2 have somewhat complementary characteristics. The stochastic approximation scheme has a low computational effort per replication, but typically requires some tuning of the gain sequence to achieve satisfactory performance. The sample average approximation method is more robust, but can be computationally expensive in the initial optimization phase.

The simulation experiments in this chapter should be viewed as a demonstration of the feasibility of the two methods rather than a comprehensive comparison. The sample average approximation method outperforms the stochastic approximation scheme and the naïve approach. In most cases the stochastic approximation scheme

Table 3.4: Estimated variance reduction ratio: $H_l = 80$ and $H_u = 105$

Volatility	Frequency of monitoring	CPU time (sec)	SA ratio	$V_{\text{naïve}}/V_{SA}$	SAA ratio	$V_{\text{naïve}}/V_{SAA}$
$\sigma = .4$	$l = 3$	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 3.5: Estimated variance reduction ratio: $H_l = 85$ and $H_u = 100$

Volatility	Frequency of monitoring	CPU time (sec)	SA ratio	$V_{\text{naïve}}/V_{SA}$	SAA ratio	$V_{\text{naïve}}/V_{SAA}$
$\sigma = .4$	$l = 3$	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

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.

Chapter 4

Adaptive Control Variate Methods for Steady-State Simulation

In Chapter 2 we developed adaptive control variate techniques for finite-horizon simulations. We now turn our attention to the steady-state case. In this chapter we discuss adaptive control variate methods for estimating steady-state performance measures when the underlying stochastic processes possess regenerative structure. The procedure is similar to the one used for terminating simulations in Chapter 2. We confine our attention to the sample average approximation technique for tuning the control variate parameter, and we develop adaptive estimators based on a regenerative method.

Let $X = (X_n : n \geq 0)$ be a discrete time stochastic process on a state space S, and let $f : S \to \mathbb{R}$ be a real-valued function defined on the state space S. Under very general conditions, $\{f(X_n) : n \geq 0\}$ satisfies a law of large numbers (LLN), so there exists a constant α for which

$$\frac{1}{n}\sum_{i=0}^{n-1} f(X_i) \to \alpha \tag{4.0.1}$$

as $n \to \infty$ a.s. Our task is to estimate the constant α, which is called the *steady state mean* of f. The natural estimator for α is the time average

$$\alpha_n = \frac{1}{n}\sum_{i=0}^{n-1} f(X_i).$$

Under additional conditions, the process X and the function f satisfy a central limit theorem (CLT), so there exists a positive constant σ for which

$$n^{1/2}(\alpha_n - \alpha) \Rightarrow \sigma N(0,1)$$

as $n \to \infty$, where $N(0,1)$ denotes a normal random variable having mean 0 and variance 1. The constant σ^2 is called the *time-average variance constant* (TAVC).

Suppose that $h(\cdot;\theta) : S \to \mathbb{R}$ is a real-valued function of $x \in S$ for any parameter vector $\theta \in \Theta$, where Θ is a parameter set. Suppose also that for all $\theta \in \Theta$,

$$\frac{1}{n}\sum_{i=0}^{n-1} h(X_i;\theta) \to 0 \tag{4.0.2}$$

as $n \to \infty$ a.s. Then, the time average

$$\alpha_n(\theta) = \frac{1}{n}\sum_{i=0}^{n-1}[f(X_i) - h(X_i;\theta)]$$

is a strongly consistent estimator for α. Here $n^{-1}\sum_{i=0}^{n-1} h(X_i;\theta)$ serves as a control variate. Let $\sigma^2(\theta)$ denote the TAVC arising in the assumed-to-hold CLT for $\alpha_n(\theta)$, i.e.,

$$n^{1/2}(\alpha_n(\theta) - \alpha) \Rightarrow \sigma(\theta)N(0,1).$$

Now we are free to choose the parameter θ that minimizes the TAVC $\sigma^2(\theta)$.

A natural question is where the control variate comes from. Many applications involve the simulation of an appropriate Markov process. We briefly describe how to define control variates in Markov process simulation, drawing from the general results of Henderson and Glynn [2002, Section 6]. Suppose that $X = (X_n : n \geq 0)$ is a positive Harris recurrent discrete-time Markov chain with stationary probability measure π. Suppose that $\pi|f| \triangleq \int_S |f(x)|\pi(dx) < \infty$. By the strong law $\alpha_n \to \alpha = \pi f$ a.s. as $n \to \infty$.

Let $u : S \to \mathbb{R}$ be a real-valued function on the state space S, and for $n \geq 1$ let

$$\begin{aligned}\alpha_n(u) &= \frac{1}{n}\sum_{i=0}^{n-1}[f(X_i) + \int_S u(y)P(X_i,dy) - u(X_i)] \\ &= \frac{1}{n}\sum_{i=0}^{n-1}[f(X_i) + (P-I)u(X_i)],\end{aligned}$$

where I is the identity operator and P is the transition probability kernel of X. Observe that if u is π integrable, then $(P-I)u$ is π integrable, and $\pi[(P-I)u] = [\pi P - \pi]u = [\pi - \pi]u = 0$. So by the strong law, $\alpha_n(u) \to \alpha$ a.s. as $n \to \infty$.

Then how do we select the function u? Consider *Poisson's equation*,

$$(P-I)u(x) = -(f(x) - \alpha), \forall x \in S. \tag{4.0.3}$$

Suppose that u^* is a solution to (4.0.3). Then $\alpha_n(u^*)$ is a zero-variance estimator of α. So it is desirable that $u \approx u^*$. Suppose that $u(x) = u(x, \theta)$, where $\theta \in \Theta$ is a parameter vector. Then

$$\alpha_n(\theta) = \frac{1}{n}\sum_{i=0}^{n-1}[f(X_i) + (P-I)u(X_i;\theta)]$$

and $\frac{1}{n}\sum_{i=0}^{n-1}(P-I)u(X_i;\theta)$ serves as a parameterized control variate.

When the parameterization is linear, we can appeal to the theory of linear control variates in steady-state simulation. The variance of $\alpha_n(\theta)$ is then a convex quadratic in θ. We can identify the value of θ that minimizes this variance, and then estimate it from the simulation. Loh [1994] examines this problem using both regenerative and batch means methods. In the nonlinearly parameterized case, the problem is not so straightforward, as we already saw in the finite-horizon setting. Fortunately, the steady-state case involves many of the same ideas that we used for finite-horizon simulation in Chapter 2. An adaptive control variate scheme for the steady-state setting using stochastic approximation was developed in Henderson et al. [2003]. We instead focus on the sample average approximation method to estimate the optimal value of θ. We assume that the underlying stochastic process is regenerative, and exploit its regenerative structure.

The asymptotic properties of the regenerative method provide a clean setting for simulation output analysis. The idea is to identify random times at which the

process probabilistically restarts, and use these regeneration points to obtain valid point and interval estimates for the steady-state mean. For an overview of this method, see Shedler [1993].

Regenerative structure is present in a wide class of discrete-event simulations. It is well known that regeneration times can be easily identified in the setting of irreducible positive recurrent discrete state space Markov chains, for example. It has also been shown that "well-posed" simulations of general state-space Markov chains (GSSMC) exhibit regenerative structure [Glynn, 1994]. However, identification of the corresponding regeneration times is non-trivial in general.

A widely used model for a discrete-event simulation is a generalized semi-Markov process (GSMP) [Shedler, 1993, Haas, 1999, Henderson and Glynn, 2001]. One can rigorously define the GSMP through a related GSSMC. Therefore, a GSMP with a well-posed GSSMC naturally possesses regenerative structure. If a GSMP has a "single state," then the regeneration times can be easily identified and the standard regenerative method can be applied to analyze the simulation output. However, most discrete-event simulations do not have a single state, and then it becomes very difficult to determine the regenerative cycle boundaries. Haas [1999] provides conditions on the building blocks of a GSMP under which a strong law and a functional central limit theorem hold. Then the batch means method can be used to obtain a point estimate and a confidence interval for the steady-state mean even when the GSMP does not have a single state. Henderson and Glynn [2001] discuss the state of the art of regenerative methods for general discrete-event simulation, and examine the issue of identifying regeneration times from a practical standpoint.

The remainder of this chapter is organized as follows. We start in Section

4.1 by describing the regenerative structure of stochastic processes, and reviewing the regenerative method. Under mild regularity conditions, a regenerative process satisfies a LLN and a CLT. Also, we can obtain strongly consistent estimators for the steady-state mean and TAVC. In Section 4.2, we provide conditions under which the the TAVC function $\sigma^2(\theta)$ is differentiable and explore the use of sample average approximation to identify the optimal parameter value θ^* that minimizes $\sigma^2(\theta)$. In Section 4.3 we consider the regenerative and batch means approaches to estimate $\sigma^2(\theta)$, and then provide conditions under which the sample average approximation problem converges to the true problem.

4.1 Regenerative Processes

Let $X = (X_n : n \geq 0)$ be a discrete time stochastic process on a state space S. For a strictly increasing sequence of non negative finite random times $T = \{T(k) : k \geq 0\}$, define the random cycles of the process X by $Y_k = \{X_n : T(k-1) \leq n < T(k)\}$, $k \geq 1$. We say that X is *classically regenerative* if there exists such a sequence T with the property that the sequence of cycles $\{Y_k : k \geq 1\}$ is independent and identically distributed. The random times $T = \{T(k) : k \geq 0\}$ are said to be *regeneration times* for the process X, where $T(0)$ is the first regeneration time. If the sequence of cycles $\{Y_k : k \geq 1\}$ is identically distributed but 1-dependent, then X is called *1-dependent regenerative*. Any well-posed steady-state simulation has either independent or 1-dependent regenerative structure [Glynn, 1994]. The theory of classical regenerative processes has been extensively studied, and much of this theory can be extended to the setting of 1-dependent regenerative processes.

Throughout this chapter, we assume that the stochastic process X is a 1-dependent regenerative process with regeneration times $T(0) = 0 < T(1) < \cdots$.

Since $T(0) = 0$, the first cycle starts at time 0. This initial condition can be relaxed, but for convenience we restrict our attention to this initial case. Denote the kth cycle length by $\tau_k = T(k) - T(k-1)$, $k \geq 1$ and let

$$F_k = \sum_{i=T(k-1)}^{T(k)-1} f(X_i), k \geq 1,$$

for some real-valued function f defined on the state space S. For example, $f(s)$ can be viewed as a reward (or cost) when in state s. Then F_k is the reward (or cost) accumulated over the kth cycle and $\{F_k : k \geq 1\}$ is 1-dependent and identically distributed. The following theorem gives a LLN and CLT for regenerative processes.

Theorem 4.1.1. *Let $X = (X_n : n \geq 0)$ be a 1-dependent regenerative process with state space S and let $f : S \to \mathbb{R}$. Suppose that $E(|F_1| + \tau_1) < \infty$, and define $Z_k = F_k - \alpha\tau_k$ for $k \geq 1$. Then*

(i) the strong law

$$\alpha_n \to \alpha = \frac{EF_1}{E\tau_1} \tag{4.1.1}$$

holds as $n \to \infty$ a.s.

(ii) Moreover, if $EZ_1^2 < \infty$, then the central limit theorem (CLT)

$$n^{1/2}(\alpha_n - \alpha) \Rightarrow \sigma N(0,1) \tag{4.1.2}$$

holds, in which case,

$$\sigma^2 = \frac{EZ_1^2 + 2EZ_1Z_2}{E\tau_1}.$$

The proof is nearly identical to the classically regenerative case given in Glynn and Iglehart [1993], and so we omit it here. Theorem 4.1.1 shows that the behavior

of a regenerative process in a cycle determines the asymptotic behavior of the process.

The expression for σ^2 in Theorem 4.1.1 suggests the following variance estimator for σ^2 :

$$\hat{\sigma}_n^2 = \frac{1}{n} \sum_{k=1}^{l(n)-1} [Z_k^2(n) + 2Z_k(n)Z_{k+1}(n)],$$

where $Z_k(n) = F_k - \alpha_n \tau_k$ and $l(n) = \sup\{k \geq 0 : T(k) \leq n\}$ is the number of completed regenerative cycles by time n. The next theorem shows that $\hat{\sigma}_n^2$ is a consistent estimator for σ^2 when the CLT (4.1.2) holds.

Theorem 4.1.2. *Assume that the conditions in Theorem 4.1.1 hold. Then*

$$\hat{\sigma}_n^2 \Rightarrow \sigma^2$$

as $n \to \infty$, where $\Rightarrow$ denotes weak convergence. Furthermore, if $E(F_1^2 + \tau_1^2) < \infty$, then $\hat{\sigma}_n^2$ is strongly consistent, that is, $\hat{\sigma}_n^2 \to \sigma^2$ as $n \to \infty$ a.s.

The proof follows as in Glynn and Iglehart [1993], and so is omitted. In Section 4.3, we will use the regenerative variance estimator $\hat{\sigma}_n^2$ in the development of the sample average approximation method.

4.2 Sample Average Approximation Method for Steady-State Simulation

Suppose that $h(\cdot;\theta) : S \to \mathbb{R}$ is a real-valued function of $x \in S$ for any parameter vector $\theta = (\theta(1), \theta(2), \ldots, \theta(p)) \in \Theta \subset \mathbb{R}^p$, where Θ is a parameter set. Suppose also that for all $\theta \in \Theta$,

$$\frac{1}{n} \sum_{i=0}^{n-1} h(X_i;\theta) \to 0 \tag{4.2.1}$$

as $n \to \infty$ a.s. Define

$$\alpha_n(\theta) = \frac{1}{n}\sum_{i=0}^{n-1}[f(X_i) - h(X_i;\theta)].$$

If the corresponding moment conditions in Theorem 4.1.1 hold, then

$$n^{1/2}(\alpha_n(\theta) - \alpha) \Rightarrow \sigma(\theta)N(0,1)$$

as $n \to \infty$, for each $\theta \in \Theta$. Our goal is to find θ^* that solves the optimization problem

$$\mathcal{P}: \qquad \min_{\theta\in\Theta} \sigma^2(\theta),$$

where the TAVC $\sigma^2(\theta)$ is well-defined for any $\theta \in \Theta$.

We follow a similar procedure to the one used in the finite-horizon setting, first, imposing some structure on the problem to obtain a differentiable TAVC function $\sigma^2(\cdot)$. We define the key random variables associated with regenerative cycles. For any fixed parameter value θ and $k \geq 1$, define

$$\begin{aligned}
H_k(\theta) &= \sum_{i=T(k-1)}^{T(k)-1} h(X_i;\theta), \\
W_k(\theta) &= \sum_{i=T(k-1)}^{T(k)-1} |h(X_i;\theta)|, \\
Z_k(\theta) &= (F_k - H_k(\theta)) - \alpha\tau_k, \text{ and} \\
\sigma^2(\theta) &= \frac{EZ_1^2(\theta) + 2EZ_1(\theta)Z_2(\theta)}{E\tau_1}.
\end{aligned}$$

We assume the following conditions.

Assumption B1 The parameter set Θ is compact and for all $x \in S$, the function $h(x,\cdot)$ is $\mathcal{C}^1$ on $\mathcal{U}$, where $\mathcal{U}$ is a bounded open set containing Θ. Moreover, $EH_1(\theta) = 0$ for any $\theta \in \mathcal{U}$.

Assumption B2 The moment conditions $E(\tau_1^2 + F_1^2) < \infty$, and $EW_1^2(\theta_0) < \infty$ for some fixed $\theta_0 \in \mathcal{U}$, hold.

Assumption B3 For all $x \in S$, $h(x;\cdot)$ is Lipschitz on $\mathcal{U}$, i.e., $\exists c(x) > 0$ such that for all $\theta_1, \theta_2 \in \mathcal{U}$,

$$|h(x;\theta_1) - h(x;\theta_2)| \leq c(x)\|\theta_1 - \theta_2\|,$$

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

$$\sup_{\theta\in\mathcal{U}} \left| \frac{\partial h(x;\theta)}{\partial \theta(j)} \right| \leq c(x)$$

for all $x \in S$ and $j = 1, \ldots, p$. Define $C_k = \sum_{i=T(k-1)}^{T(k)-1} c(X_i), k \geq 1$ and assume that $EC_1^2 < \infty$.

Remark 7. Suppose that $U(\theta)$ is a random variable for each $\theta \in \mathcal{U}$. We say that $U(\cdot)$ is *dominated* by an integrable random variable $\tilde{U}$ if $E\tilde{U} < \infty$ and for every $\theta \in \mathcal{U}$, $|U(\theta)| \leq \tilde{U}$ a.s. Under **B2** and **B3**, $H_1(\cdot)^2$ is dominated by an integrable random variable, and hence so is $Z_1^2(\cdot)$. To see why, note that for any $\theta \in \mathcal{U}$,

$$\begin{aligned} H_1^2(\theta) &= [H_1(\theta_0) + (H_1(\theta) - H_1(\theta_0))]^2 \\ &\leq 2H_1^2(\theta_0) + 2(H_1(\theta) - H_1(\theta_0))^2 \\ &\leq 2W_1^2(\theta_0) + 2C_1^2\|\theta - \theta_0\|^2. \end{aligned}$$

But $\mathcal{U}$ is bounded, and hence $\|\theta - \theta_0\|^2$ is bounded.

Proposition 4.2.1. *Assume that **B1-B3** hold. Then $\sigma^2(\cdot)$ is $\mathcal{C}^1$ on $\mathcal{U}$ and*

$$\nabla_\theta \sigma^2(\theta) = \frac{E\nabla_\theta Z_1^2(\theta) + 2E\nabla_\theta(Z_1(\theta)Z_2(\theta))}{E\tau_1}.$$

Proof. It suffices to show that $EZ_1(\cdot)$ and $EZ_1(\cdot)Z_2(\cdot)$ are $\mathcal{C}^1$ on $\mathcal{U}$ and the gradient and expectation can be exchanged. We apply Theorem 2.3.1 to $Z_1(\theta)$ and $Z_1(\theta)Z_2(\theta)$ component by component. Consider the jth component, for some $j \in \{1, \ldots, p\}$. The only condition that requires explicit verification is that

$\partial Z_1(\theta)/\partial\theta(j)$ and $\partial Z_1(\theta)Z_2(\theta)/\partial\theta(j)$ are dominated by an integrable function of X. With probability 1,

$$\frac{\partial Z_k(\theta)}{\partial\theta(j)} = -\frac{\partial}{\partial\theta(j)}\left(\sum_{i=T(k-1)}^{T(k)-1} h(X_i;\theta)\right) = -\sum_{i=T(k-1)}^{T(k)-1}\frac{\partial h(X_i;\theta)}{\partial\theta(j)}, k \geq 1.$$

Hence,

$$\begin{aligned}\left|\frac{\partial Z_1^2(\theta)}{\partial\theta(j)}\right| &= 2\left|\frac{\partial}{\partial\theta(j)}\left(\sum_{i=0}^{T(1)-1} h(X_i;\theta)\right) Z_1(\theta)\right| \\ &\leq 2C_1|Z_1(\theta)| \text{ and} \end{aligned} \tag{4.2.2}$$

$$\begin{aligned}\left|\frac{\partial(Z_1(\theta)Z_2(\theta))}{\partial\theta(j)}\right| &= \left|\frac{\partial}{\partial\theta(j)}\left(\sum_{i=0}^{T(1)-1} h(X_i;\theta)\right) Z_2(\theta)\right| \\ &+ \left|\frac{\partial}{\partial\theta(j)}\left(\sum_{i=T(1)}^{T(2)-1} h(X_i;\theta)\right) Z_1(\theta)\right| \\ &\leq C_1|Z_2(\theta)| + C_2|Z_1(\theta)|. \end{aligned} \tag{4.2.3}$$

By **B3** and **Remark 7**, the right hand sides of equations (4.2.2) and (4.2.3) are dominated by an integrable function. □

We now introduce the sample average approximation (SAA) procedure to solve the problem $\mathcal{P}$. The procedure exploits the regenerative structure of the underlying stochastic processes. The quantities computed over the regenerative cycles are one-dependent identically distributed random variables, so the sample average approximation method for terminating simulations in Chapter 2 can be extended to this setting. Suppose that we have a sample path $\tilde{X}_0, \tilde{X}_1, \ldots, \tilde{X}_m$ and that $V(m,\theta)$ is an estimate for $\sigma^2(\theta)$ based on $\tilde{X}_0, \tilde{X}_1, \ldots, \tilde{X}_m$. We will discuss specific choices of variance estimator $V(m,\theta)$ in Section 4.3. Then the SAA to problem $\mathcal{P}$ is

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

Let $\hat{\theta}_m$ be a solution for problem $\mathcal{P}_m$, perhaps obtained by using some deterministic nonlinear optimization algorithm. Then in a second phase, α is estimated using

$$\alpha_n(\hat{\theta}_m) = \frac{1}{n}\sum_{i=0}^{n-1}[f(X_i) - h(X_i;\hat{\theta}_m)],$$

where the samples $X_0, X_1, \ldots, X_{n-1}$ are independent of $\tilde{X}_0, \tilde{X}_1, \ldots, \tilde{X}_m$.

The asymptotic theory for finite-horizon simulation can be extended to this setting by analyzing simulation outputs via regenerative cycles. We split the sample average $\alpha_n(\hat{\theta}_m)$ into two parts: A random sum of 1-dependent identically distributed random variables, and a remainder term that converges to zero a.s. To show that $\alpha_n(\hat{\theta}_m)$ satisfies a SLLN and CLT, we first need the following Proposition 4.2.2, a uniform version of the strong law for a random number of samples and a following lemma.

Proposition 4.2.2. *Suppose that $\{U_i(\theta) : i \geq 1\}$ is a κ-dependent stationary sequence of random variables for any $\theta \in \Theta$, where Θ is a compact parameter set. Let $\{l(n) : n \geq 1\}$ be a family of random indices such that $l(n)/n \to \lambda$ a.s as $n \to \infty$ for some $\lambda < \infty$. Suppose that $U_i(\cdot)$ is continuous on Θ w.p.1 and dominated by an integrable random variable for all $i \geq 1$. Then*

$$\sup_{\theta\in\Theta}\left|\frac{1}{n}\sum_{i=1}^{l(n)} U_i(\theta) - \lambda E U_1(\theta)\right| \to 0 \tag{4.2.4}$$

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

Proof. Observe that

$$\begin{aligned} & \sup_{\theta\in\Theta}\left|\frac{1}{n}\sum_{i=1}^{l(n)} U_i(\theta) - \lambda E U_1(\theta)\right| \\ = \; & \sup_{\theta\in\Theta}\left|\frac{1}{n}\sum_{i=1}^{l(n)}[U_i(\theta) - E U_1(\theta)] - R_n(\theta)\right| \end{aligned}$$

$$\leq \quad \frac{l(n)}{n} \sup_{\theta \in \Theta} \left| \frac{1}{l(n)} \sum_{i=1}^{l(n)} [U_i(\theta) - EU_1(\theta)] \right| + \sup_{\theta \in \Theta} |R_n(\theta)|, \tag{4.2.5}$$

where

$$R_n(\theta) = E[U_1(\theta)] \left(\lambda - \frac{l(n)}{n} \right).$$

Since $U_1(\theta)$ is dominated by an integrable random variable, the second term in (4.2.5) converges to 0 a.s. as $n \to \infty$. So it suffices to show that the first term in (4.2.5) converges to 0 a.s. as $n \to \infty$.

Now, $EU_1(\cdot)$ is continuous by Lebesque's Dominated Convergence Theorem (LDCT). Moreover, for all $\theta \in \Theta$, $\frac{1}{m} \sum_{i=1}^{m} U_i(\theta) \to EU_1(\theta)$ a.s. as $m \to \infty$, and $l(n) \to \infty$ a.s. as $n \to \infty$. First, we will show that $\sup_{\theta \in \Theta} \left| \frac{1}{m} \sum_{i=1}^{m} U_i(\theta) - EU_1(\theta) \right|$ $\to 0$ as $m \to \infty$ a.s. Then by Theorem 2.1 in [Gut, 1988] (p.10),

$$\sup_{\theta \in \Theta} \left| \frac{1}{l(n)} \sum_{i=1}^{l(n)} U_i(\theta) - EU_1(\theta) \right| \to 0$$

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

We follow the proof of Proposition 7 in Shapiro [2003]. Choose a point $\bar{\theta} \in \Theta$, a sequence γ_k of positive numbers converging to zero, and define $V_k := \{\theta \in \Theta : ||\theta - \bar{\theta}|| \leq \gamma_k\}$ and

$$\delta_k^i := \sup_{\theta \in V_k} |U_i(\theta) - U_i(\bar{\theta})|, \text{ for } i \geq 1.$$

Note that $\delta_k^i, i \geq 1$ are κ-dependent identically distributed random variables. Since $U_1(\cdot)$ is continuous w.p.1 and dominated by an integrable random variable on Θ, by LDCT we have that

$$\lim_{k \to \infty} E[\delta_k^1] = E[\lim_{k \to \infty} \delta_k^1] = 0. \tag{4.2.6}$$

Note that

$$\sup_{\theta \in V_k} \left| \frac{1}{m} \sum_{i=1}^{m} U_i(\theta) - \frac{1}{m} \sum_{i=1}^{m} U_i(\bar{\theta}) \right| \leq \frac{1}{m} \sum_{i=1}^{m} \delta_k^i. \tag{4.2.7}$$

By the LLN, the right hand side of (4.2.7) converges to $E[\delta_k^1]$ a.s. as $m \to \infty$. Together with (4.2.6) this implies that for any given $\varepsilon > 0$, there exist a neighborhood W of $\bar{\theta}$ such that w.p.1 for sufficiently large M,

$$\sup_{\theta \in W \cap \Theta} \left| \frac{1}{M} \sum_{i=1}^{M} U_i(\theta) - \frac{1}{M} \sum_{i=1}^{M} U_i(\bar{\theta}) \right| < \varepsilon.$$

Since Θ is compact, there exists a finite number of points $\theta_1, \theta_2, \ldots, \theta_J \in \Theta$ and corresponding neighborhoods $W_1, \ldots, W_J$ covering Θ such that w.p.1 for sufficiently large M,

$$\sup_{\theta \in W_j \cap \Theta} \left| \frac{1}{M} \sum_{i=1}^{M} U_i(\theta) - \frac{1}{M} \sum_{i=1}^{M} U_i(\theta_j) \right| < \varepsilon, j = 1, \ldots, J. \tag{4.2.8}$$

Furthermore, since $EU_1(\cdot)$ is continuous on Θ, these neighborhoods can be chosen in such a way that

$$\sup_{\theta \in W_j \cap \Theta} |E[U_1(\theta)] - E[U_1(\theta_j)]| < \varepsilon, j = 1, \ldots, J. \tag{4.2.9}$$

Again by the LLN, w.p.1 for M large enough,

$$\left| \frac{1}{M} \sum_{i=1}^{M} U_i(\theta_j) - E[U_1(\theta_j)] \right| < \varepsilon, j = 1, \ldots, J. \tag{4.2.10}$$

By (4.2.8)-(4.2.10), w.p.1 for M large enough

$$\sup_{\theta \in \Theta} \left| \frac{1}{M} \sum_{i=1}^{M} U_i(\theta) - E[U_1(\theta)] \right| < 3\varepsilon.$$

□

Lemma 4.2.3. *Suppose that $U_1, U_2, \ldots$ are independent and identically distributed nonnegative random variables with $EU_1 < \infty$. Then $U_n/n \to 0$ as $n \to \infty$ a.s. and hence*

$$\max_{1 \le i \le n} U_i/n \to 0.$$

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

The proof is as in Durrett [1999], so is omitted.

We can now state a version of the strong law and central limit theorem for the case where $\hat{\theta}$ is random.

Theorem 4.2.4. *Suppose that **B1-B3** hold, and that the samples used in constructing $\hat{\theta}$ are independent of those used in computing $\alpha_n(\hat{\theta})$ for every n. Then $\alpha_n(\hat{\theta}) \to \alpha$ as $n \to \infty$ a.s., and*

$$\sqrt{n}(\alpha_n(\hat{\theta}) - \alpha) \Rightarrow \sigma(\hat{\theta})N(0,1)$$

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

Proof. For the strong law note that

$$\begin{aligned} |\alpha_n(\hat{\theta}) - \alpha| &\leq \left|\frac{1}{n}\sum_{i=0}^{n-1}(f(X_i) - \alpha)\right| + \left|\frac{1}{n}\sum_{i=0}^{n-1} h(X_i, \hat{\theta})\right| \\ &\leq |\alpha_n - \alpha| + \sup_{\theta\in\Theta}\left|\frac{1}{n}\sum_{i=0}^{n-1} h(X_i, \theta)\right|. \end{aligned} \tag{4.2.11}$$

The first term in (4.2.11) converges to 0 as $n \to \infty$ by the LLN (4.1.1). For the second term, we have

$$\sup_{\theta\in\Theta}\left|\frac{1}{n}\sum_{i=0}^{n-1} h(X_i, \theta)\right| \leq \sup_{\theta\in\Theta}\left|\frac{1}{n}\sum_{k=1}^{l(n)} H_k(\theta)\right| + R_n, \tag{4.2.12}$$

where

$$R_n = \sup_{\theta\in\Theta}\left[\frac{1}{n}\sum_{i=T(l(n))}^{T(l(n)+1)-1} |h(X_i, \theta)|\right].$$

The first term in (4.2.12) converges to 0 by Proposition 4.2.2. Note that

$$\begin{aligned} R_n &\leq \sup_{\theta\in\Theta}\left[\frac{1}{n}\sum_{i=T(l(n))}^{T(l(n)+1)-1} |h(X_i, \theta_0)| + c(X_i)||\theta - \theta_0||.\right] \\ &\leq \frac{W_{l(n)+1}(\theta_0)}{n} + \frac{C_{l(n)+1}}{n}\sup_{\theta\in\Theta}||\theta - \theta_0||. \end{aligned} \tag{4.2.13}$$

The results in Lemma 4.2.3 also hold for a 1-dependent sequence of stationary random variables. Therefore,

$$\max_{1\le k\le n+1} W_k(\theta_0)/(n+1) \text{ and } \max_{1\le k\le n+1} C_k/(n+1) \to 0$$

as $n \to \infty$ a.s. Since $l(n)+1 \le n+1$ and Θ is bounded, (4.2.13) converges to 0. This completes the proof of the strong law. We now turn to the central limit theorem.

By Theorem 4.1.1, for each fixed $t \in \mathbb{R}$,

$$P\left(\sqrt{n}(\alpha_n(\hat{\theta})-\alpha) \le t \,|\, \hat{\theta}\right) \to \Phi\left(\frac{t}{\sigma(\hat{\theta})}\right) I(\sigma(\hat{\theta})>0) + I(t\ge 0)I(\sigma(\hat{\theta})=0) \tag{4.2.14}$$

as $n \to \infty$, where Φ is the distribution function of a normal random variable with mean 0 and variance 1 and $I(\cdot)$ is an indicator function. LDCT ensures that we can take expectations through (4.2.14), and so

$$\begin{aligned} & P(\sqrt{n}(\alpha_n(\hat{\theta})-\alpha) \le t) \\ & \quad \to \quad E\left[\Phi\left(\frac{t}{\sigma(\hat{\theta})}\right) I(\sigma(\hat{\theta})>0) + I(t\ge 0)I(\sigma(\hat{\theta})=0)\right] \\ & \quad = \quad P(\sigma(\hat{\theta})N(0,1) \le t) \end{aligned}$$

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

Next, we study the asymptotic behavior of $\alpha_n(\hat{\theta}_m)$ as the computational budget gets large. Assume that $m = m(n)$ is a function of n such that $m(n) \to \infty$ as $n \to \infty$. If $\hat{\theta}_{m(n)} \to \theta^*$ in probability as $n \to \infty$, then $\alpha_n(\hat{\theta}_{m(n)})$ behaves the same as $\alpha_n(\theta^*)$, asymptotically as $n \to \infty$.

Theorem 4.2.5. *Suppose that $\hat{\theta}_{m(n)} \to \theta^*$ as $n \to \infty$ in probability, for some random variable θ^*. Suppose further that **B1 - B3** hold and the samples used in*

computing $\hat{\theta}_{m(n)}$ are independent of those used to compute $\alpha_n(\hat{\theta}_{m(n)})$ for every n. Then $\alpha_n(\hat{\theta}_{m(n)}) \to \alpha$ as $n \to \infty$ a.s., and

$$\sqrt{n}(\alpha_n(\hat{\theta}_{m(n)}) - \alpha) \Rightarrow \sigma(\theta^*)N(0,1)$$

as $n \to \infty$, where $N(0,1)$ is independent of θ^*.

Proof. The strong law can be proved exactly as in the proof of Theorem 4.2.4. To prove the central limit theorem, note that

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

Notice that θ^* is independent of the samples used to compute α_n for every n. By Theorem 4.2.4, $D_{1,n} \Rightarrow \sigma(\theta^*)N(0,1)$ as $n \to \infty$. Thus, it suffices to show that

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

as $n \to \infty$. Let us write

$$\begin{aligned}D_{2,n} &= \frac{1}{\sqrt{n}} \sum_{k=1}^{l(n)+2} [H_k(\hat{\theta}_{m(n)}) - H_k(\theta^*)] \\ &- \frac{1}{\sqrt{n}} \sum_{i=n+1}^{T(l(n)+2)-1} [h(X_i, \hat{\theta}_{m(n)}) - h(X_i, \theta^*)] \\ &= D_{3,n} - R_n.\end{aligned}$$

Observe that

$$\begin{aligned}|R_n| &\leq \frac{1}{\sqrt{n}} \sum_{i=n+1}^{T(l(n)+2)-1} |h(X_i, \hat{\theta}_{m(n)}) - h(X_i, \theta^*)| \\ &\leq \frac{1}{\sqrt{n}} \sum_{i=n+1}^{T(l(n)+2)-1} c(X_i)||\hat{\theta}_{m(n)} - \theta^*|| \\ &\leq \frac{1}{\sqrt{n}}(C_{l(n)+1} + C_{l(n)+2})||\hat{\theta}_{m(n)} - \theta^*||. \qquad (4.2.15)\end{aligned}$$

Apply Lemma 4.2.3 to $\{C_k^2, k \geq 1\}$. Then $C_{l(n)+1}^2/n$ and $C_{l(n)+2}^2/n$ converge to 0 a.s. as $n \to \infty$ and hence $(C_{l(n)+1} + C_{l(n)+2})/\sqrt{n} \to 0$ a.s. as $n \to \infty$. Since $||\hat{\theta}_{m(n)} - \theta^*||$ is bounded, (4.2.15) converges to 0 as $n \to \infty$ a.s.

To show that $D_{3,n} \Rightarrow 0$, we will adapt techniques from Janson [1983] and Henderson and Glynn [2001]. For any fixed θ and $\theta^* \in \mathcal{U}$, define

$$\begin{aligned}
\Delta H_k(\theta, \theta^*) &= H_k(\theta) - H_k(\theta^*), k \geq 1, \\
\mathcal{F}_N(\theta, \theta^*) &= \sigma(\Delta H_1(\theta, \theta^*), \tau_1, \Delta H_2(\theta, \theta^*), \tau_2, \ldots, \Delta H_N(\theta, \theta^*), \tau_N), \\
S_N(\theta, \theta^*) &= \sum_{k=1}^{N} \Delta H_k(\theta, \theta^*) \text{ and} \\
W_N(\theta, \theta^*) &= E(S_{N+1}^2(\theta, \theta^*) - \sum_{k=1}^{N+1} \Delta H_k^2(\theta, \theta^*) \\
&- 2\sum_{k=1}^{N} \Delta H_k(\theta, \theta^*) \Delta H_{k+1}(\theta, \theta^*) | \mathcal{F}_N(\theta, \theta^*)).
\end{aligned}$$

Note that $E[H_k(\theta) - H_k(\theta^*)] = 0$ so $W(\theta, \theta^*) = (W_N(\theta, \theta^*) : N \geq 1)$ is a martingale with respect to the filtration $\mathcal{F}(\theta, \theta^*) = (\mathcal{F}_N(\theta, \theta^*) : N \geq 1)$. Define $a \wedge b = \min\{a, b\}$. If T is a randomized stopping time with respect to $\mathcal{F}(\theta, \theta^*)$, then $EW_{T \wedge N}(\theta, \theta^*) = EW_1(\theta, \theta^*) = 0$. Hence,

$$\begin{aligned}
ES_{1+T\wedge N}^2(\theta, \theta^*) &= E[\sum_{k=1}^{1+T\wedge N} \Delta H_k^2(\theta, \theta^*) + 2\sum_{k=1}^{T\wedge N} \Delta H_k(\theta, \theta^*) \Delta H_{k+1}(\theta, \theta^*)] \\
&\leq E[\sum_{k=1}^{1+(T\wedge N)} \Delta H_k^2(\theta, \theta^*) + \sum_{k=1}^{T\wedge N} (\Delta H_k^2(\theta, \theta^*) + \Delta H_{k+1}^2(\theta, \theta^*))] \\
&\leq 3E[\sum_{k=1}^{1+(T\wedge N)} \Delta H_k^2(\theta, \theta^*)].
\end{aligned}$$

Noting that $l(n)+1$ is a randomized stopping time, we obtain

$$\begin{aligned} ES^2_{l(n)+2}(\theta,\theta^*) &\leq \liminf_{n\to\infty} ES^2_{1+[l(n)+1)\wedge N]}(\theta,\theta^*) && (4.2.16)\\ &\leq 3\liminf_{n\to\infty} E\sum_{k=1}^{1+[(l(n)+1)\wedge N]}\Delta H_k^2(\theta,\theta^*)\\ &= 3E\sum_{k=1}^{l(n)+2}\Delta H_k^2(\theta,\theta^*) && (4.2.17)\\ &= 3E\Delta H_1^2(\theta,\theta^*)E(l(n)+2), && (4.2.18)\end{aligned}$$

where (4.2.16) follows from Fatou's lemma, (4.2.17) follows from the monotone convergence theorem, and (4.2.18) is a variant of Wald's equation for 1-dependent random variables. Note that

$$D_{3,n} = \frac{1}{\sqrt{n}} S_{l(n)+2}(\hat{\theta}_{m(n)},\theta^*).$$

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

$$\begin{aligned} P(|D_{3,n}|>\epsilon) &\leq \epsilon^{-2}ED^2_{3,n}\\ &= \frac{1}{n\epsilon^2}ES^2_{l(n)+2}(\hat{\theta}_{m(n)},\theta^*)\\ &= \frac{1}{n\epsilon^2}E[E[S^2_{l(n)+2}(\hat{\theta}_{m(n)},\theta^*)|\hat{\theta}_{m(n)},\theta^*]]\\ &\leq \frac{3}{n\epsilon^2}E[E[\Delta H_1^2(\hat{\theta}_{m(n)},\theta^*)|\hat{\theta}_{m(n)},\theta^*]E[l(n)+2|\hat{\theta}_{m(n)},\theta^*]]\\ &= 3\frac{E[l(n)+2]}{n\epsilon^2}E[\Delta H_1^2(\hat{\theta}_{m(n)},\theta^*)]\\ &\leq 3\frac{E[l(n)+2]}{n\epsilon^2}EC_1^2E||\hat{\theta}_{m(n),\theta^*}-\theta^*||^2. && (4.2.19)\end{aligned}$$

By LDCT, $E||\hat{\theta}_{m(n)}-\theta^*||^2 \to 0$, and by Theorem 3.1 of Janson [1983], $E(l(n)+2)/n \to \lambda$ as $n\to\infty$. Therefore, (4.2.19) converges to 0 as $n\to\infty$. □

It remains to give conditions under which $\hat{\theta}_m \to \theta^*$. 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 $\mathcal{P}_m$. If the gradient of the variance estimator converges to the

gradient of $\sigma^2(\theta)$ uniformly on Θ a.s, then by sample-path analysis we can prove the convergence of first-order critical points to those of the true problem $\mathcal{P}$.

Theorem 4.2.6. *Suppose that*

(i) Θ *is convex and compact,*

(ii) $\sigma^2(\cdot)$ *is* $\mathcal{C}^1$ *on an open set containing* Θ,

(ii) $V(m,\cdot)$ *is* $\mathcal{C}^1$ *on an open set containing* Θ *w.p.1, for all* $m \geq 1$, *and*

(iv) $\sup_{\theta\in\Theta} ||\nabla_\theta V(m,\theta) - \nabla_\theta \sigma^2(\theta)|| \to 0$ *a.s. as* $n \to \infty$.

Let $\hat{\theta}_m$ *be a first-order critical point of* $V(m,\cdot)$ *on* Θ *and* $S(\sigma^2(\cdot),\Theta)$ *be the set of first order critical points of* $\sigma^2(\cdot)$ *on* Θ. *Then* $d(\hat{\theta}_m, S(\sigma^2(\cdot),\Theta)) \to 0$ *as* $m \to \infty$ *a.s.*

Proof. If $d(\hat{\theta}_m, S(\sigma^2(\cdot),\Theta)) \not\to 0$, then by passing to a subsequence if necessary, we can assume that for some $\epsilon > 0$, $d(\hat{\theta}_m, S(\sigma^2(\cdot),\Theta)) \geq \epsilon$ 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 $\theta^* \notin S(\sigma^2(\cdot),\Theta)$. On the other hand, $\sigma^2(\cdot)$ is $\mathcal{C}^1$ and $\nabla_\theta V(m,\hat{\theta}_m) \to \nabla_\theta \sigma^2(\theta^*)$ as $m \to \infty$ a.s. Since Θ is convex, each $\hat{\theta}_m$ satisfies the first order condition

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

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

$$\langle \nabla_\theta \sigma^2(\theta^*), u - \theta^* \rangle \geq 0, \text{ for all } u \in \Theta, \text{ a.e.}$$

Therefore, $\theta^* \in S(\sigma^2(\cdot),\Theta)$ and we obtain a contradiction. □

Theorem 4.2.6 gives conditions under which $\hat{\theta}_m$ converges to the set of first-order critical points of σ^2 as $m \to \infty$. As discussed in Chapter 2, a simple sufficient condition that ensures convergence to a fixed θ^* is the existence of a unique first-order critical point, but this condition is not easy to verify in practice.

In general, the condition (iv) in Theorem 4.2.6 is hard to verify. A condition that may be more easily verified is that if the variance estimator $V(m, \cdot)$ converges to $\sigma^2(\cdot)$ uniformly on Θ in probability and the problem $\mathcal{P}$ has a unique solution which is well separated, then $\hat{\theta}_m$ is a (weakly) consistent estimator for θ^*. A well separated optimal solution θ^* is an optimal solution that is unique and such that for any neighborhood of θ^* the gap between the optimal value and the value at any point outside the neighborhood is bounded away from zero. A sequence of random variables $\{\xi_n : n \geq 1\}$ is $o_p(1)$ if and only if $\xi_n \to 0$ in probability as $n \to \infty$.

Theorem 4.2.7. *Suppose that* $(\mathcal{P})$ *has a unique optimal solution* θ^* *and assume that*

(i) $\sup_{\theta \in \Theta} |V(m, \theta) - \sigma^2(\theta)| = o_p(1)$ *and*

(ii) for any $\varepsilon > 0$,

$$\inf_{\theta \in \{\tilde{\theta} : d(\tilde{\theta}, \theta^*) \geq \varepsilon\}} \sigma^2(\theta) > \sigma^2(\theta^*). \tag{4.2.20}$$

Let $\hat{\theta}_m$ *be an optimal solution for problem* $\mathcal{P}_m$. *Then* $\hat{\theta}_m \to \theta^*$ *as* $m \to \infty$ *in probability.*

Proof. Since $V(m, \hat{\theta}_m) \leq V(m, \theta^*)$ and $V(m, \theta^*) = \sigma^2(\theta^*) + o_p(1)$,

$$V(m, \hat{\theta}_m) \leq \sigma^2(\theta^*) + o_p(1).$$

Then

$$\begin{aligned} 0 \leq \sigma^2(\hat{\theta}_m) - \sigma^2(\theta^*) &\leq \sigma^2(\hat{\theta}_m) - V(m, \hat{\theta}_m) + o_p(1) \\ &\leq \sup_{\theta \in \Theta} |\sigma^2(\theta) - V(m, \theta)| + o_p(1). \end{aligned}$$

By (i), $\sigma^2(\hat{\theta}_m) \to \sigma^2(\theta^*)$ in probability as $m \to \infty$.

By (ii), for all $\varepsilon > 0$ and some $\eta > 0$,

$$\sigma^2(\theta) - \sigma^2(\theta^*) > \eta, \text{ for all } \theta \in \{\theta : d(\theta, \theta^*) \geq \varepsilon\}.$$

Therefore,

$$P(d(\hat{\theta}_m, \theta^*) \geq \varepsilon) \leq P(\sigma^2(\hat{\theta}_m) - \sigma^2(\theta^*) > \eta). \tag{4.2.21}$$

The righthand side of (4.2.21) goes to 0 and hence $\hat{\theta}_m \to \theta^*$ as $m \to \infty$ in probability. □

4.3 Variance Estimators

In this section, we consider two estimators for the TAVC $\sigma^2(\theta)$ and provide conditions under which the sample average approximation problems converge to the true problem. The first estimator is based on the regenerative method. The use of regenerative structure allows us to easily analyze the asymptotic properties of this estimator. The second estimator is based on the batch means method, which is currently more applicable than the regenerative method. It is strongly consistent under the harder-to-verify assumption that the output process obeys a *strong invariance principle* (also called *strong approximation*; see Damerdji [1994]).

4.3.1 Regenerative Method

The variance estimator of $\sigma^2(\theta)$ derived using the regenerative method is

$$V_{RG}(m;\theta) = \frac{1}{m}\sum_{k=1}^{l(m)-1}[Z_k^2(m;\theta) + 2Z_k(m;\theta)Z_{k+1}(m;\theta)],$$

where $Z_k(m;\theta) = (F_k - H_k(\theta)) - \alpha_m(\theta)\tau_k$. The SAA problem based on this estimator is

$$\mathcal{P}_{RG}(m): \qquad \min_{\theta\in\Theta} V_{RG}(m;\theta).$$

Let $\hat{\theta}_{RG}(m)$ be a first-order critical point for problem $\mathcal{P}_{RG,m}$.

Theorem 4.1.2 implies that under the conditions of Theorem 4.2.5, the estimator $V_{RG}(m,\theta)$ is a strongly consistent estimator of $\sigma^2(\theta)$ for every $\theta \in \Theta$. We can prove that under the same conditions, $\hat{\theta}_{RG}(m)$ converges to the set of first-order critical points of the true problem $\mathcal{P}$. Our next result extends Theorem 2.5.4 in Chapter 2 for finite-horizon simulation to steady-state simulation.

Let $u(\theta) = (u^{(1)}(\theta),\ldots,u^{(d)}(\theta))$ be a $\mathbb{R}^d$-valued function of $\theta \in \Theta \subset \mathbb{R}^p$ and $\{U_m(\theta) = (U_m^{(1)}(\theta),\ldots,U_m^{(d)}(\theta)) : m \geq 1\}$ be a family of $\mathbb{R}^d$-valued random variables parameterized by θ such that $U_m(\theta) \to u(\theta)$ a.s. as $m \to \infty$ for all $\theta \in \Theta$. Suppose that $\Upsilon(x)$ 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 u and U_m for all m. We seek conditions under which first-order critical points of $\Upsilon \circ U_m = \Upsilon(U_m(\cdot))$ on Θ converge to those of $\Upsilon \circ u$ on Θ.

Theorem 4.3.1. *Consider the family of random variables $\{U_m(\cdot) : m \geq 1\}$ and the function $u(\cdot)$ defined immediately above. Suppose that Θ is convex and compact and*

(i) $U_m(\cdot) = (U_m^{(1)}(\cdot),\ldots,U_m^{(d)}(\cdot))$ is $\mathcal{C}^1$ on an open set containing Θ w.p.1, for all $m \geq 1$,

(ii) $u(\cdot)$ *is* $\mathcal{C}^1$ *on an open set containing* Θ,

(iii) $\sup_{\theta\in\Theta}\left|U_m^{(r)}(\theta)-u^{(r)}(\theta)\right|\to 0$, *a.s. as* $m\to\infty$ $(r=1,\ldots,d)$ *and*

(iv) $\sup_{\theta\in\Theta}\left|\partial U_m^{(r)}(\theta)/\partial\theta(j)-\partial u^{(r)}(\theta)/\partial\theta(j)\right|\to 0$ *a.s. as* $m\to\infty$. $(r=1,\ldots,d,$ $j=1,\ldots,p)$.

Let $\hat{\theta}_m\in S(\Upsilon\circ U_m,\Theta)$ *be the set of first-order critical points of* $\Upsilon\circ U_m$ *on* Θ. *Then* $d(\hat{\theta}_m,S(\Upsilon\circ u,\Theta))\to 0$ *as* $m\to\infty$ *a.s.*

This is a corollary of Theorem 4.2.6, so the proof is omitted. We now obtain the following corollary.

Corollary 4.3.2. *Suppose that* ***B1-B3*** *hold and* Θ *is convex. Then*

$$d(\hat{\theta}_{RG}(m),S(\sigma^2,\Theta))\to 0$$

as $m\to\infty$ *a.s.*

Proof. If $\Upsilon(\zeta_1,\ldots,\zeta_8)=\zeta_1-2\zeta_8\zeta_2+\zeta_8^2\zeta_3+2\zeta_4-2\zeta_8\zeta_5-2\zeta_8\zeta_6+2\zeta_8\zeta_7$, then

$$\begin{aligned} V_{RG,m}(\theta) &= \Upsilon(U_m(\theta)) \text{ and} \\ \sigma^2(\theta) &= \Upsilon(u(\theta)), \end{aligned}$$

where

$$Y_k(\theta)=F_k-H_k(\theta),\theta\in\Theta,k\geq 1,$$

$$
\begin{aligned}
U_m^{(1)}(\theta) &= \frac{\sum_{k=1}^{l(m)-1} Y_k^2(\theta)}{m}, & u^{(1)}(\theta) &= \frac{EY_1^2(\theta)}{E\tau_1},\\
U_m^{(2)}(\theta) &= \frac{\sum_{k=1}^{l(m)-1} Y_k(\theta)\tau_k}{m}, & u^{(2)}(\theta) &= \frac{EY_1(\theta)\tau_1}{E\tau_1},\\
U_m^{(3)}(\theta) &= \frac{\sum_{k=1}^{l(m)-1} \tau_k^2}{m}, & u^{(3)}(\theta) &= \frac{E\tau_1^2}{E\tau_1},\\
U_m^{(4)}(\theta) &= \frac{\sum_{k=1}^{l(m)-1} Y_k(\theta)Y_{k+1}(\theta)}{m}, & u^{(4)}(\theta) &= \frac{EY_1(\theta)Y_2(\theta)}{E\tau_1},\\
U_m^{(5)}(\theta) &= \frac{\sum_{k=1}^{l(m)-1} Y_k(\theta)\tau_{k+1}}{m}, & u^{(5)}(\theta) &= \frac{EY_1(\theta)\tau_2}{E\tau_1},\\
U_m^{(6)}(\theta) &= \frac{\sum_{k=1}^{l(m)-1} \tau_k Y_{k+1}(\theta)}{m}, & u^{(6)}(\theta) &= \frac{E\tau_1 Y_2(\theta)}{E\tau_1},\\
U_m^{(7)}(\theta) &= \frac{\sum_{k=1}^{l(m)-1} \tau_k\tau_{k+1}}{m}, & u^{(7)}(\theta) &= \frac{E\tau_1\tau_2}{E\tau_1} \text{ and}\\
U_m^{(8)}(\theta) &= \alpha_m(\theta), & u^{(8)}(\theta) &= \alpha.
\end{aligned}
$$

Note that $Y_1(\cdot)$ is $\mathcal{C}^1$ on $\mathcal{U}$ and

$$Y_1(\theta), Y_1^2(\theta), \frac{\partial Y_1(\theta)}{\partial\theta(j)} \text{ and } \frac{\partial Y_1^2(\theta)}{\partial\theta(j)}$$

are all dominated by an integrable function ($j = 1, \ldots, p$). By Proposition 4.2.2,

$$U_m^{(r)}(\theta) \to u^{(r)}(\theta) \text{ and } \partial U_m^{(r)}(\theta)/\partial\theta(j) \to \partial u^{(r)}(\theta)/\partial\theta(j),\ r = 1, \ldots, 7, j = 1, \ldots, p$$

uniformly on Θ as $m \to \infty$ a.s. It remains to show that $\alpha_m(\theta) \to \alpha$ and $\partial\alpha_m(\theta)/\partial\theta(j) \to 0$, $j = 1, \ldots p$ uniformly on Θ as $m \to \infty$ a.s. Then by Theorem 4.3.1,

$$d(\hat{\theta}_{RG}(m), S(\sigma^2, \Theta)) = d(\hat{\theta}_{RG}(m), S(\Upsilon \circ u, \Theta)) \to 0$$

as $m \to \infty$. The proof of the uniform convergence of $\alpha_m(\theta)$ is very similar to the proof of the similar result in Theorem 4.2.4. Nothing that $\partial H_1(\theta)/\partial\theta(j), j = 1, \ldots, p$ is dominated by C_1, by Theorem 2.3.1, we obtain

$$E\left[\frac{\partial H_1(\theta)}{\partial\theta(j)}\right] = \frac{\partial}{\partial\theta(j)}\left[EH_1(\theta)\right] = 0,\ \theta \in \Theta, j = 1, \ldots, p.$$

Now,

$$\begin{aligned}\sup_{\theta\in\Theta}\left|\frac{\partial\alpha_m(\theta)}{\partial\theta(j)}\right| &\leq \sup_{\theta\in\Theta}\left|\frac{1}{m}\sum_{k=1}^{l(m)}\frac{\partial H_k(\theta)}{\partial\theta(j)}\right| + \sup_{\theta\in\Theta}\left[\frac{1}{m}\sum_{k=T(l(m))}^{T(l(m)+1)-1}\left|\frac{\partial h(X_k,\theta)}{\partial\theta(j)}\right|\right] \\ &\leq \sup_{\theta\in\Theta}\left|\frac{1}{m}\sum_{k=1}^{l(m)}\frac{\partial H_k(\theta)}{\partial\theta(j)}\right| + \frac{C_{l(m)+1}}{m}. \qquad (4.3.1)\end{aligned}$$

By Proposition 4.2.2 and Lemma 4.2.3, (4.3.1) converges to 0. □

It is difficult to apply the regenerative method when the regeneration times cannot be easily identified. For this reason we consider a second estimator based on the batch means method, which does not require identification of the regeneration times.

4.3.2 Batch Means Method

Suppose that we have a sample path $X_0, X_1, \ldots, X_{m-1}$. Divide this sample path into b_m adjacent batches, each of size k_m. For simplicity, we assume that $m = k_m b_m$. The ith batch consists of observations $X_{(i-1)k_m}, \ldots, X_{ik_m-1}$. The sample mean $M_i(k_m;\theta)$ for the ith batch is

$$M_i(k_m;\theta) = \frac{1}{k_m}\sum_{j=(i-1)k_m}^{ik_m-1}\left[f(X_j) - h(X_j;\theta)\right], i \geq 1.$$

The grand mean of the individual batch means is

$$\bar{M}(m;\theta) = \frac{1}{b_m}\sum_{i=1}^{b_m} M_i(k_m;\theta).$$

Then we can estimate $\sigma^2(\theta)$ using

$$V_{BM}(m;\theta) = \frac{k_m}{b_m - 1}\sum_{i=1}^{b_m}(M_i(k_m;\theta) - \bar{M}(m;\theta))^2.$$

Then the SAA problem based on this estimator is

$$\mathcal{P}_{BM,m}: \qquad \min_{\theta\in\Theta} V_{BM}(m;\theta).$$

Let $\theta_{BM}(m)$ be the resulting estimator of a minimizer of $\sigma^2(\cdot)$.

Due to its simplicity, batch means is one of the most widely used methods in steady-state output analysis. In the classical batch means method, the number of batches b_m is fixed and the resulting batch means are approximately i.i.d. normal random variables for sufficiently large batch size k_m. The TAVC $\sigma^2(\theta)$ is not estimated but is instead canceled out. If on the other hand the number of batches b_m increases as the sample size m increases, under some conditions the estimator $V_{BM,m}(\theta)$ becomes a consistent estimator for $\sigma^2(\theta)$.

We would like to use Theorem 4.3.3 to establish that $\theta_{BM}(m) \to \theta^*$ as $m \to \infty$. To do so, we need to establish uniform convergence in probability of the batch means estimator $V_{BM}(m, \theta)$. Unfortunately, we have not been able to do so. However, we conjecture that a result of this form should hold. First, we state one more assumption, and a useful result, that we believe are needed.

Assumption B4 Assume that $\sigma^2(\theta) > 0$ for all $\theta \in \Theta$, and there exists $\delta \in (0, 2)$ such that for all $\theta \in \Theta$,

$$E\left[\sum_{n=T_1}^{T_2-1} |f(X_n) - h(X_n;\theta)|\right]^{2+\delta} < \infty, \text{ and } E\left[\tau_1^{1+\delta/2}\right] < \infty.$$

Remark 8. A set of sufficient conditions for **B4** is that there exists $\delta \in (0, 2)$ such that for all $\theta \in \Theta$,

$$E\tilde{F}_1^{2+\delta}(\theta), EW_1^{2+\delta}(\theta) < \infty,$$

where $\tilde{F}_i = \sum_{k=T(i-1)}^{T(i)-1} |f(X_k)|$ and

$$E\left[\tau_1^{1+\delta/2}\right], EC_1^{2+\delta} < \infty.$$

The following theorem provides conditions under which a sequence of random functions is uniformly convergent in probability.

Theorem 4.3.3. *[Newey, 1991] Let $\Theta \subset \mathbb{R}^p$ be a compact set, $Q_n(\theta)$ be a random function of $\theta \in \Theta$ and the sample size n, and $q(\theta)$ be a non-random function of $\theta \in \Theta$. Suppose that*

(i) *for each $\theta \in \Theta$, $Q_n(\theta) - q(\theta) = o_p(1)$,*

(ii) *there is B_n such that $B_n = O_p(1)$, that is, B_n is bounded in probability and for all $\theta, \tilde{\theta} \in \Theta$,*

$$|Q_n(\tilde{\theta}) - q(\theta)| \leq B_n||\tilde{\theta} - \theta|| \text{ and}$$

(iii) *$q(\cdot)$ is continuous on Θ.*

Then $\sup_{\theta \in \Theta} |Q_n(\theta) - q(\theta)| = o_p(1)$.

Conjecture *Suppose that $(\mathcal{P})$ has a unique optimal solution θ^* and assume that (4.2.20) is satisfied. Let $\hat{\theta}_{BM}(m)$ be an optimal solution for problem $\mathcal{P}_{BM,m}$. Then under **B1-B4** and some conditions for b_m and k_m, $\hat{\theta}_{BM}(m) \to \theta^*$ as $m \to \infty$ in probability.*

Appendix A

Additional Details of the Barrier Option 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 $\Theta \subset \mathcal{U}$, $U(x,i;\cdot) : \mathcal{U} \to \mathbb{R}$ is $\mathcal{C}^1$ for all $(x,i) \in \mathcal{S} \times \{0,1,\ldots,l-1\}$, and $U(\cdot,i;\cdot) : [H_l, H_u] \times \mathcal{U} \to \mathbb{R}$ is Lipschitz for all $i \in \{0,1,\ldots,l-1\}$. (These assumptions are all satisfied in our particular example.) Since $\{0,1,\ldots,l-1\}$ is finite and $\mathcal{U}$ is bounded, there exists a $C > 0$ such that for all $\theta_1, \theta_2 \in \mathcal{U}$ and $(x,i) \in \mathcal{S} \times \{0,1,\ldots,l\}$,

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

and

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

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

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

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

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

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 $x = 0$,

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

For $H_l \leq x \leq H_u$,

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

Therefore,

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

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

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

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

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

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

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

where

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

Then

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

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 dissertation.

BIBLIOGRAPHY

W. A. Al-Qaq, M. Devetsikiotis, and J.-K. Townsend. Stochastic gradient optimization of importance sampling for the efficient simulation of digital communication systems. *IEEE Transactions on Communications*, 43:2975–2985, 1995.

B. Arouna. Robbins-Monro algorithms and variance reduction in finance. *Journal of Computational Finance*, 7(2):1245–1255, 2003.

K. Baggerly, D. Cox, and R. Picard. Exponential convergence of adaptive importance sampling for markov chains. *Journal of Applied Probability*, 37(2), 2000.

F. Bastin, C. Cirillo, and P. L. Toint. Convergence theory for nonconvex stochastic programming with an application to mixed logit. *Mathematical Programming*, 108:207–234, 2007.

M. S. Bazaraa, H. D. Sherali, and C. M. Shetty. *Nonlinear Programming: Theory and Algorithms*. Wiley, New York, 2nd edition, 1993.

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.

H. Chen and B. W. Schmeiser. Stochastic root finding via retrospective approximation. *IIE Transactions*, 33:259–275, 2001.

K. L. Chung. *A Course in Probability Theory*, volume 21 of *Probability and Mathematical Statistics*. Academic Press, San Diego, 2nd edition, 1974.

H. Damerdji. Strong consistency of the variance estimator in steady-state simulation output analysis. *Mathematics of Operations Research*, 19(2):494 – 512, 1994.

R. Durrett. *Essentials of Stochastic Processes*. Springer-Verlag, New York, 1999.

M. Fitzgerald and R. Picard. Accelerated Monte Carlo for particle dispersion. *Communications in Statistics, Part A – Theory and Methods*, 30(11):2459–2471, 2001.

M. C. Fu. Convergence of the GI/G/1 queue usinginfinitesimal perturbation analysis. *Journal of Optimization Theory and Applications*, 65:149–160, 1990.

M. C. Fu and J.-Q. Hu. *Conditional Monte Carlo: Gradient Estimation and Optimization Applications*. Kluwer, Boston, 1997.

P. Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer-Verlag, New York, 2004.

P. Glasserman. *Gradient Estimation Via Perturbation Analysis*. Kluwer, The Netherlands, 1991.

P. Glasserman and J. Staum. Conditioning on one-step survival for barrier option simulations. *Operations Research*, 49:923–937, 2001.

P. W. Glynn. Likelihood ratio gradient estimation for stochastic systems. *Communications of the ACM*, 33:75–84, 1990.

P. W. Glynn. Some topics in regenerative steady-state simulation. *Acta Applicandae Mathematicae*, 34:225–236, 1994.

P. W. Glynn and D. L. Iglehart. Conditions for the applicability of the regenerative method. *Management Science*, 39:1108–1111, 1993.

P. W. Glynn and R. Szechtman. Some new perspectives on the method of control variates. In K. T. Fang, F.J.Hickernell, and H. Niederreiter, editors, *Monte Carlo and Quasi-Monte Carlo Methods 2000*, pages 27–49, Berlin, 2002. Springer-Verlag.

P. W. Glynn and W. Whitt. Necessary conditions in limit theorems for cumulative processes. *Stochastic Processes and Their Applications*, 98:199–209, 2002.

A. Gut. *Stopped Random Walks: Limit Theorems and Applications*. Springer-Verlag, New York NY, 1st edition, 1988.

P. J. Haas. On simulation output analysis for generalized semi-Markov processes. *Communications in Statistics: Stochastic Models*, 15:53–80, 1999.

S. G. Henderson and P. W. Glynn. Regenerative steady-state simulation of discrete event systems. *ACM Transactions on Modeling and Computer Simulation*, 11: 313–345, 2001.

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. Variance reduction for simulation in multiclass queueing networks. *IIE Transactions*, 2003. Submitted.

S. G. Henderson and S. P. Meyn. Efficient simulation of multiclass queueing networks. In S. Andradóottir, 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 B. Simon. Adaptive simulation using perfect control variates. *Journal of Applied Probability*, 41(3):859–876, 2004.

S. G. Henderson, S. P. Meyn, and V. B. Tadić. Performance evaluation and policy selection in multiclass networks. *Discrete Event Dynamic Systems*, 13:149–189, 2003. Special issue on learning and optimization methods.

M. Hsieh. Adaptive Monte Carlo methods for rare event simulations. In E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, editors, *Proceedings of the 2002 Winter Simulation Conference*, pages 108–115, Piscataway NJ, 2002. IEEE.

S. Janson. Renewal theory for m-dependent variables. *Annals of Probability*, 11: 558–568, 1983.

S. Juneja and P. Shahabuddin. Rare-event simulation techniques: An introduction and recent advances. In S. G. Henderson and B. L. Nelson, editors, *Handbook of Simulation*, volume 13 of *Handbooks in Operations Research and Management Science*. Elsevier, 2006.

S. Karlin and H. M. Taylor. *A Second Course in Stochastic Processes*. Academic Press, Boston, 1981.

J. Kiefer and J. Wolfowitz. Stochastic estimation of the maximum of a regression function. *Annals of Mathematical Statistics*, 23:462–466, 1952.

C. Kollman, K. Baggerly, D. Cox, and R. Picard. Adaptive importance sampling

on discrete Markov chains. *Ann. Appl. Probab.*, 9(2):391–412, 1999. ISSN 1050-5164.

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.

P. L'Ecuyer and P. W. Glynn. Stochastic optimization by simulation: Convergence proofs for the GI/G/1 queue in steady-state. *Management Science*, 40:1562–1578, 1994.

P. L'Ecuyer and G. Perron. On the convergence rates of IPA and FDC derivative estimators. *Operations Research*, 42(4):643–656, 1994.

R. S. Liptser and A. N. Shiryayev. *Theory of Martingales*. Kluwer Academic, Boston, 1989.

W. W. Loh. *On the Method of Control Variates*. PhD thesis, Department of Operations Research, Stanford University, Stanford, CA, 1994.

S. Maire. Reducing variance using iterated control variates. *Journal of Statistical Computation and Simulation*, 73(1):1–29, 2003.

S. P. Meyn. Value functions, optimization and performance evaluation in stochastic network models. *IEEE Transactions on Automatic Control*, 2003. Submitted.

W. K. Newey. Uniform convergence in probability and stochastic equicontinuity. *Econometrica*, 59:1161–1167, 1991.

E. L. Plambeck, B.-R. Fu, S. M. Robinson, and R. Suri. Sample-path optimization of convex stochastic performance functions. *Mathematical Programming*, 75: 137–176, 1996.

J. A. Rice. *Mathematical Statistics and Data Analysis*. Wadsworth & Brooks/Cole, Pacific Grove, California, 1988.

H. Robbins and S. Monro. A stochastic approximation method. *Annals of Mathematical Statistics*, 22:400–407, 1951.

S. M. Robinson. Analysis of sample-path optimization. *Mathematics of Operations Research*, 21:513–528, 1996.

R. Y. Rubinstein. *Monte Carlo Optimization, Simulation and Sensitivity of Queueing Networks*. Wiley, New York, 1986.

R. Y. Rubinstein. The cross-entropy method for combinatorial and continuous optimization. *Methodology and Computing in Applied Probability*, 1:127–190, 1999.

R. Y. Rubinstein and B. Melamed. *Modern Simulation and Modeling*. Wiley, 1998.

R. Y. Rubinstein and A. Shapiro. *Discrete Event Systems: Sensitivity Analysis and Stochastic Optimization by the Score Function Method*. Wiley, Chichester, 1993.

D. Ruppert. A newton-raphson version of the multivariate robbins-monro procedure. *Annals of Statistics*, 13:236–245, 1985.

A. Shapiro. Monte Carlo sampling methods. In A. Ruszczynski and A. Shapiro, editors, *Stochastic Programming*, Handbooks in Operations Research and Management Science. Elsevier, 2003.

A. Shapiro. Asypmtotic behavior of optimal solutions in stochastic programming. *Mathematics of Operations Research*, 18:829–845, 1993.

A. Shapiro and T. Homem-de-Mello. On rate of convergence of monte carlo approximations of stochastic programs. *SIAM Journal on Optimization*, 11:70–86, 2000.

G. S. Shedler. *Regenerative Stochastic Simulation*. Academic Press, Boston, 1993.

J. C. Spall. *Introduction to Stochastic Search and Optimization: Estimation, Simulation and Control*. Wiley, Hoboken, New Jersey, 2003.

Y. Su and M. C. Fu. Importance sampling in derivative securities pricing. In J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick, editors, *Proceedings of the 2000 Winter Simulation Conference*, pages 587–596, Piscataway NJ, 2000. IEEE.

V. B. Tadić and S. P. Meyn. Adaptive Monte Carlo algorithms using control variates. *Manuscript*, 2004.

F. Vázquez-Abad and D. Dufresne. Accelerated simulation for pricing Asian options. In D. Medeiros, E. Watson, J. S. Carson, and M. S. Manivannan, editors, *Proceedings of the 1998 Winter Simulation Conference*, pages 1493–1500, Piscataway, NJ, 1998. IEEE.