

FUNCTION-APPROXIMATION-BASED PERFECT CONTROL VARIATES FOR PRICING AMERICAN OPTIONS

Nomesh Bolia
Sandeep Juneja

Tata Institute of Fundamental Research
Mumbai 400005, INDIA

ABSTRACT

Monte Carlo simulation techniques that use function approximations have been successfully applied to approximately price multi-dimensional American options. However, for many pricing problems the time required to get accurate estimates can still be prohibitive, and this motivates the development of variance reduction techniques. In this paper, we describe a zero-variance or ‘perfect’ control variate to price American options. We then discuss how function approximation may be used to approximate this perfect control variate. Empirically, we observe that on simple one dimensional examples, this approximately perfect control variate gives orders of magnitude of variance reduction compared to naive estimation.

1 INTRODUCTION

Accurate estimation of the price of an American option and the optimal exercise policy when the dimensionality of the underlying process is large remains an important problem in option pricing. Typically, this problem is simplified by restricting times at which the option can be exercised to a finite set of values so that it may be modeled as a discrete time Markov decision process (MDP), or more specifically, an optimal stopping problem in discrete time. Unfortunately, the well known numerical techniques for solving such MDP’s suffer from a ‘curse of dimensionality’.

Recently, significant literature has developed that approximately solves the optimal stopping problem associated with the American option using Monte Carlo simulation (see, e.g., the overview in Chapter 8 of Glasserman 2004, Bolia and Juneja 2005). These typically involve generating many sample paths (or more general sample trees) of the underlying assets until the expiry of the option and then using some form of backward induction to obtain an approximate solution. Carrière (1996), Longstaff and Schwartz (2001) and Tsitsiklis and Van Roy (2001) propose regression based function approximations to aid in accurate and quick ex-

ecution of this backward induction step. These methods involve approximating the ‘continuation value’ (the value of the option if not exercised) as a linear combination of certain easily evaluated and well-chosen basis functions. Bolia, Glasserman and Juneja (2004) build upon the function approximation techniques proposed in Longstaff and Schwartz (2001) and Tsitsiklis and Van Roy (2001) and show how they may be used to develop an approximate zero variance or ‘perfect’ importance sampling estimator for pricing American options.

Henderson and Glynn (2001) show that for commonly used performance measures in Markovian settings, a martingale may be constructed that serves as a perfect control variate in the sense that the resultant performance measure estimator has zero variance. This martingale is typically unimplementable as it requires knowledge of measures that are at least as difficult to estimate as the performance measure of interest. For example, in the American option setting, if our interest is in estimating the value of the option at a particular time at a particular value of the underlying assets, then (as we note later), a perfect control variate can be constructed that requires knowledge of the complete option price process, i.e., it requires knowledge of the price of option at each state of the underlying assets, at each time period, up till option maturity.

As mentioned earlier, Bolia, Glasserman and Juneja (2004) used function approximations to develop an approximately perfect importance sampling estimator to price American options. In this paper we extend their analysis, combining it with the results from Henderson and Glynn (2001), to develop an approximately perfect control variate estimator for the same purpose. We note that in simple one dimensional settings, the control variate based estimator that we develop here, provides far greater variance reduction compared to the importance sampling estimator developed by Bolia, Glasserman and Juneja (2004).

In Section 2, we develop our mathematical framework and discuss the perfect zero-variance control variate. Here, we also observe its connection to the well known additive

duality for developing accurate estimates of an upper bound to the option price. In Section 3, we discuss the proposed approximation methodology. Experimental results displaying the effectiveness of the proposed scheme are given in Section 4.

2 MATHEMATICAL FRAMEWORK

We refer the reader to, e.g., Karatzas and Shreve (1991), Duffie (1996), Kallianpur and Karandikar (1999) for a description of continuous time models used for asset price movements and for technical conditions under which the no-arbitrage condition implies the existence of a unique risk neutral measure so that the price of a European option is an expectation of its discounted payoffs under the risk neutral measure. We assume that such conditions hold. Furthermore, we assume that the option under consideration can be exercised at finite number of time points. In practice, options that provide fixed finite number of times at which they can be exercised are referred to as Bermudan options. American options that expire at time T are well approximated by Bermudan options by dividing T into a large number of well spaced intervals. We assume that the probability density functions are available to us so that we can generate samples of asset prices at the finite number of time intervals where they can be exercised. In some cases, this may only be approximately true (see, e.g., Glasserman 2004), however, we do not dwell on this issue in the paper.

Specifically, suppose that the option can be exercised only at $N + 1$ times $0, 1, 2, \dots, N$ (these times need not be integers or equally spaced, we do so for notational simplicity).

The underlying security prices at these times are modeled by a Markov process $(X_n \in \mathcal{X} : n \leq N)$, where $\mathcal{X}$ may be a multi-dimensional state space. The state description may include additional variables such as the value of stochastic interest rates and volatilities, and supplementary path dependent information, so that the resulting process $\{X_n\}$ is Markov. The value of the option at time n if exercised at that time, is denoted by $g_n : \mathcal{X} \rightarrow \mathfrak{R}^+$ (i.e., its exercise value or intrinsic value). Let $\mathcal{T}_n$ denote the set of stopping times taking values in $\{n, n + 1, \dots, N\}$ (recall that a random variable τ is a stopping time w.r.t. $\{X_n\}$ if the occurrence of $\{\tau = n\}$ can be ascertained by observing $(X_1, \dots, X_n)$). Let

$$J_n(x) = \sup_{\tau \in \mathcal{T}_n} E[g_\tau(X_\tau) | X_n = x], \quad x \in \mathcal{X},$$

where the expectation is taken under the risk neutral measure. Then $J_n(x)$ is the value of the option at time n given that the option is not exercised before time n . The initial state $X_0 = x_0$ is fixed and known. So, our pricing problem is to evaluate $J_0(x_0)$. Since no closed form solutions are available for the

above except in the simplest cases, numerical and simulation methods are needed for pricing such multi-dimensional and multiple exercise opportunity options. This formulation is sufficiently general to include discounted payoffs through appropriate definition of the $\{X_n\}$ and $\{g_n\}$ (see Glasserman 2004, p.425), and hence these are not explicitly stated.

Further, we suppose that the pdf of X_{n+1} conditioned on $X_n = x$ evaluated at y is given by $f_n(x, y)$ under the risk-neutral measure. For any function $F : \mathcal{X} \rightarrow \mathfrak{R}$, define

$$\begin{aligned} E_n[F](x) &= E[F(X_{n+1}) | X_n = x] \\ &= \int_{\mathcal{X}} F(y) f_n(x, y) dy. \end{aligned} \quad (1)$$

It is well known that the value functions $J = (J_n(x) : x \in \mathcal{X}, n \leq N)$ satisfy the following backward recursions:

$$\begin{aligned} J_N(x) &= g_N(x) \\ J_n(x) &= \max(g_n(x), E_n[J_{n+1}](x)) \end{aligned} \quad (2)$$

for $n = 0, 1, 2, \dots, N - 1$, for each $x \in \mathcal{X}$. An alternative set of recursions based on the continuation value function $Q = (Q_n(x) : x \in \mathcal{X}, n \leq N - 1)$ may be developed, where

$$Q_n(x) = E_n[J_{n+1}](x).$$

These recursions are:

$$\begin{aligned} Q_{N-1}(x) &= E_{N-1}[g_N](x) \\ Q_n(x) &= E_n[\max(g_{n+1}, Q_{n+1})](x) \end{aligned} \quad (3)$$

for $n = 0, 1, 2, \dots, N - 2$. Note that $J_n(x) = \max(g_n(x), Q_n(x))$.

2.1 Control Variates and a Zero-Variance Estimator

Let $\tau^* = \min\{n : g_n(X_n) = J_n(X_n)\}$. Then, it is well known that τ^* is an optimal stopping time for our problem, i.e., $J_0(x_0) = E[g_{\tau^*}(X_{\tau^*})]$ (see, e.g., Duffie 1996).

As noted, e.g., in Bolia, Glasserman and Juneja (2004), this stopping time is easily and accurately estimated using functional approximations in the sense that if τ approximates τ^* , then $J_0(x_0) - E[g_\tau(X_\tau)]$ is positive, but typically small. Once an approximation τ to τ^* is known, $J_0(x_0)$ may be estimated by the usual Monte Carlo technique of generating independent samples of $g_\tau(X_\tau)$ and taking their average.

Suppose that Y is a random variable that is a function of $(x_0, X_1, \dots, X_\tau)$ whose mean equals zero. Then,

$$g_\tau(X_\tau) - Y$$

provides an unbiased estimator for $E[g_\tau(X_\tau)]$. In the control variate technique, one looks for such a Y that is highly

positively correlated with $g_\tau(X_\tau)$ so that $g_\tau(X_\tau) - Y$ has a much smaller variance compared to $g_\tau(X_\tau)$. One way to arrive at such control variates is through the following martingale: For any collection of functions $(F_n : \mathcal{X} \rightarrow \mathfrak{R}, 1 \leq n \leq N)$ such that $E[|F_n(X_n)|] < \infty$ for all n , the process

$$M_n = \sum_{i=0}^{n-1} (F_{i+1}(X_{i+1}) - E_i[F_{i+1}](X_i)),$$

for $n = 1, \dots, N$ is a zero mean martingale w.r.t. filtration $(\mathcal{F}_n : 1 \leq n \leq N)$, where $\mathcal{F}_n = \sigma(X_1, X_2, \dots, X_n)$ and $X_0 = x_0$. Since, $\tau \leq N$, from the Martingale Stopping Time Theorem, $E[M_\tau] = 0$ (see, e.g., Williams 1991). Thus, M_τ may serve as a control variate for $g_\tau(X_\tau)$. Specializing the results from Henderson and Glynn (2001) to our setting, we note that the random variable $\tilde{M}_{\tau^*}$ associated with the zero mean martingale

$$\tilde{M}_n = \sum_{i=0}^{n-1} (J_{i+1}(X_{i+1}) - E_i[J_{i+1}](X_i)), \quad (4)$$

is the perfect control variate for estimating $J_0(x_0)$ in the sense that

$$g_{\tau^*}(X_{\tau^*}) - \tilde{M}_{\tau^*} = J_0(x_0)$$

a.s. This follows by simply observing that along the set $\{\tau^* > n\}$, we have $E_i[J_{i+1}](X_i) = Q_i(X_i) = J_i(X_i)$ for $i \leq n$, so that

$$\tilde{M}_{\tau^*} = J_{\tau^*}(X_{\tau^*}) - J_0(x_0) = g_{\tau^*}(X_{\tau^*}) - J_0(x_0).$$

In this paper we use regression based functional approximations of $J_n(\cdot)$ to approximate the martingale $(M_n : n \leq N)$. First we observe the connection of this martingale with the well known additive duality in American options settings.

2.2 Additive Duality

Haugh and Kogan (2001) note that for any zero mean martingale $(M_n : 1 \leq n \leq N)$ w.r.t. the filtration $(\mathcal{F}_n : 1 \leq n \leq N)$,

$$J_0(x_0) = E[g_{\tau^*}(X_{\tau^*})] \quad (5)$$

$$= E[g_{\tau^*}(X_{\tau^*}) - M_{\tau^*}] \quad (6)$$

$$\leq E[\max_n (g_n(X_n) - M_n)]. \quad (7)$$

Thus, average of independent samples of $\max_n (g_n(X_n) - M_n)$ provide an estimate of an upper bound to the option price. Furthermore, they note

that the inequality in (5) is tight if the martingale $(\tilde{M}_n : 1 \leq n \leq N)$ is considered in the upper bound.

This can be seen by observing that $g_n(X_n) - \tilde{M}_n$ may be re-expressed as:

$$g_n(X_n) - J_n(X_n) + \sum_{i=0}^{n-1} (E_i[J_{i+1}](X_i) - J_i(X_i)) + J_0(x_0).$$

This in turn is dominated by $J_0(x_0)$. To see this, note that $g_n(X_n) - J_n(X_n) \leq 0$, and the well known fact that the process $(J_m(X_m) : 0 \leq m \leq N)$ is a super-martingale, so that for each i , $E_i[J_{i+1}](X_i) - J_i(X_i) \leq 0$ (see, e.g., Duffie 1996).

This suggests that if we have an approximation $\hat{J}_n(\cdot)$ to the true option value $J_n(\cdot)$ (for $1 \leq n \leq N$), then average of independent samples of $[\max_n (g_n - \tilde{M}_n)]$, with $\hat{J}_n(\cdot)$ replacing $J_n(\cdot)$ in the definition of $\tilde{M}_n$, can provide a good upper biased estimate of $J_0(x_0)$. We test this observation empirically in our experiments.

It is also noteworthy that perfect control variates solve the additive dual problem, while, as noted in Bolia, Glasserman and Juneja (2004), the perfect importance sampling distribution solves the multiplicative dual problem.

Before we develop specific approximations to $J_n(\cdot)$, we characterize the errors that may be associated with such approximations in the next section.

2.3 Characterizing Approximation Error

Consider an approximate value function process $\hat{J} = (\hat{J}_n(x) \geq 0, n \leq N, x \in \mathcal{X})$ and a positive $\epsilon > 0$ such that, for all n ,

$$E(\hat{J}_n(X_n) - J_n(X_n))^2 \leq \epsilon^2 \quad (8)$$

and

$$E(E_n[\hat{J}_{n+1}](X_n) - E_n[J_{n+1}](X_n))^2 \leq \epsilon^2. \quad (9)$$

A sufficient condition for this is that the error

$$\sup_{n,x} |\tilde{J}_n(x) - J_n(x)| \leq \epsilon.$$

Here we also assume that $E[J_n(X_n)]^2 < \infty, \forall n$. It follows that $E[E_n[J_{n+1}](X_n)]^2 < \infty, \forall n$. Then, (8) and (9) imply that $E[\hat{J}_{n+1}(X_n)]^2 < \infty$ and $E[E_n[\hat{J}_{n+1}](X_n)]^2 < \infty, \forall n$. Also, from (8) and (9), it follows that

$$E(|\hat{J}_n(X_n) - J_n(X_n)|) \leq \epsilon \quad (10)$$

and

$$E(|E_n[\hat{J}_{n+1}](X_n) - E_n[J_{n+1}](X_n)|) \leq \epsilon. \quad (11)$$

The associated zero mean martingale $\hat{M} = (\hat{M}_n \in \mathfrak{R} : 1 \leq n \leq N)$, as in (4), is then given by

$$\hat{M}_n = \sum_{i=0}^{n-1} \left(\hat{J}_{i+1}(X_{i+1}) - E_i[\hat{J}_{i+1}](X_i) \right). \quad (12)$$

Let $V(Y)$ denote the variance of any rv Y . In the following proposition we note that under the assumption that τ^* is known, the variance of the control variate estimator may not be very large. Thus, if $(J_n : n \leq N)$ can be closely approximated by a known quantity $(\hat{J}_n : n \leq N)$ and we have an accurate approximation for τ , then the corresponding $\hat{M}_\tau$ may provide a good proxy for the zero-variance control variate M_{τ^*} and good simulation efficiency may be expected.

Proposition 1 Under (8) and (9),

$$V\left(g_{\tau^*}(X_{\tau^*}) - \hat{M}_{\tau^*}\right) \leq (8N^2 + N)\epsilon^2. \quad (13)$$

Proof Throughout the proof, we suppress the use of X_n in expressing $J_n(X_n)$ and $g_n(X_n)$ for ease of notation. Set $\hat{J}_0(x_0) = J_0(x_0)$. Then by rearranging terms, $Y = g_{\tau^*} - \hat{M}_{\tau^*}$ can be re-expressed as

$$Y = g_{\tau^*} - \hat{J}_{\tau^*} + \sum_{k=0}^{\tau^*-1} (E_k[\hat{J}_{k+1}] - \hat{J}_k) + J_0(x_0). \quad (14)$$

Subtracting and adding $E_k[J_{k+1}]$ to each term in the sum in (14), and noting from (2) that $J_k = E_k[J_{k+1}]$ for $k \leq (\tau^* - 1)$ and $J_{\tau^*} = g_{\tau^*}$, we get

$$\begin{aligned} Y &= (J_{\tau^*} - \hat{J}_{\tau^*}) + \sum_{k=0}^{\tau^*-1} (E_k[\hat{J}_{k+1}] - E_k[J_{k+1}]) \\ &\quad + \sum_{k=0}^{\tau^*-1} (J_k - \hat{J}_k) + J_0(x_0). \end{aligned}$$

We may re-express the variance $E(Y - J_0(x_0))^2$ as $E[Z_1 + Z_2 + Z_3]^2$, where $Z_1 = J_{\tau^*} - \hat{J}_{\tau^*}$, $Z_2 = \sum_{k=0}^{\tau^*-1} (E_k[\hat{J}_{k+1}] - E_k[J_{k+1}])$, and $Z_3 = \sum_{k=0}^{\tau^*-1} (J_k - \hat{J}_k)$. Let $I_{\{A\}}$ denote the indicator function of any event A , i.e., $I_{\{A\}} = 1$ if A occurs and 0 otherwise. Then,

$$(E[Z_1])^2 \leq E[Z_1^2] \leq \sum_{k=1}^N E\left[(J_k - \hat{J}_k)^2 I_{\{\tau^*=k\}}\right] \leq N\epsilon^2. \quad (15)$$

The Cauchy-Schwartz inequality implies that

$$\begin{aligned} E[Z_2^2] &\leq \sum_{k=0}^{N-1} E(E_k[\hat{J}_{k+1}] - E_k[J_{k+1}])^2 + 2 \sum_{i < j} \\ &\quad [E(E_i[\hat{J}_{i+1}] - E_i[J_{i+1}])^2 \cdot E(E_j[\hat{J}_{j+1}] - E_j[J_{j+1}])^2]^{1/2}. \end{aligned}$$

It therefore follows from (9) that $E[Z_2^2] \leq N^2\epsilon^2$. Similarly, $E[Z_3^2] \leq N^2\epsilon^2$. Further,

$$\begin{aligned} E[Z_1 Z_2] &\leq [E(J_{\tau^*} - \hat{J}_{\tau^*})^2]^{1/2} \sum_{k=0}^{N-1} [E(E_k[\hat{J}_{k+1}] - E_k[J_{k+1}])^2]^{1/2} \\ &\leq N^2\epsilon^2, \end{aligned}$$

where the first inequality follows from the Cauchy-Schwartz inequality, and the second one from (15) and (11). Similarly, $E[Z_1 Z_3] \leq N^2\epsilon^2$ and $E[Z_2 Z_3] \leq N^2\epsilon^2$. Therefore,

$$E[Z_1 + Z_2 + Z_3]^2 \leq N\epsilon^2 + 8N^2\epsilon^2$$

and the proposition follows. $\square$

3 PROPOSED METHODOLOGY

As in Bolia, Glasserman and Juneja (2004), we adopt a two-phase approach to pricing the option. In the first phase, we approximately estimate the optimal stopping policy. For this purpose we develop approximations $\hat{Q}$ for $Q = (Q_n(x), n \leq N - 1, x \in \mathcal{X})$. Tsitsiklis and Van Roy (2001) find that methods based on least squares regressions are better suited for estimating Q rather than $J = (J_n(x), n \leq N, x \in \mathcal{X})$, as the estimates of the former have better bias characteristics.

In the second phase we evaluate this policy using control variates. However, for control variates we also need to find estimators $\hat{J}$ for J . Furthermore, to compute the value of the control variate, it is desirable that the integral $E_n[\hat{J}_{n+1}](x)$ be known explicitly. We now discuss how this is achieved using functional approximations. (the discussion here is essentially that in Bolia, Glasserman and Juneja 2004, we repeat it to make this presentation self contained).

Let $\phi_k : \mathcal{X} \rightarrow \mathfrak{R}$ for $1 \leq k \leq K$ denote a set of basis functions. Consider a parameterized value function $\hat{Q} : \mathcal{X} \times \mathfrak{R}^K \rightarrow \mathfrak{R}$ that assigns values $\hat{Q}(x, r)$ to state x , where $r = (r(1), \dots, r(K)) \in \mathfrak{R}^K$ is a vector of free parameters and

$$\hat{Q}(x, r) = \sum_{k=1}^K \phi_k(x) r(k).$$

Using simulated paths, we find parameters $r_0^*, r_2^*, \dots, r_{N-1}^*$ (each $\in \mathbb{R}^K$) so that

$$\hat{Q}(x, r_n^*) \approx Q_n(x)$$

for each x and n . These approximations are then used to estimate the optimal stopping policy.

To compute the control variate, consider a parameterized value function $\hat{J} : \mathcal{X} \times \mathbb{R}^K \rightarrow \mathbb{R}$ that assigns values $\hat{J}(x, s)$ to state x , where again, $s = (s(1), \dots, s(K)) \in \mathbb{R}^K$, and

$$\hat{J}(x, s) = \sum_{k=1}^K \phi_k(x) s(k).$$

We choose each $\phi_k(\cdot)$ so that $E_n[\phi_k](x)$ can be explicitly evaluated. We estimate parameters $s_1^*, s_2^*, \dots, s_N^*$ (each in $\mathbb{R}^K$) so that $\hat{J}(x, s_n^*) \approx J_n(x)$ for each x and n .

Different basis functions can be used for $\hat{Q}$ and $\hat{J}$. These could further depend upon the time period n . The number of basis functions used could also be a function of n . However, to keep the notation simple we avoid these generalizations. We now discuss the procedure outlined above in detail

3.1 Approximately Evaluating the Optimal Policy

As mentioned earlier, we follow the type of approach used in Longstaff and Schwartz (2001) and Tsitsiklis and Van Roy (2001) to approximate the optimal stopping time τ^* . This involves generating L sample paths $(x_{m,n} : n \leq N, m \leq L)$ of the process $(X_n : n \leq N)$ using the densities $f_n(\cdot, \cdot)$.

The parameters $r_0^*, \dots, r_{N-1}^*$ are found recursively:

$$r_{N-1}^* = \arg \min_r \sum_{m=1}^L [g_N(x_{m,N}) - \sum_{k=1}^K \phi_k(x_{m,N-1}) r(k)]^2.$$

Here, $g_N(x_{m,N})$ is an unbiased sample of $Q_{N-1}(x_{m,N-1}) = E_{N-1}[g_N](x_{m,N-1})$. Then, using the approximation $\hat{Q}(\cdot, r_{N-1}^*)$ for $Q_{N-1}(\cdot)$ along each generated path m we can approximately evaluate when to exercise the option, given that we have not exercised it till time $N-2$. Call this time $\tau_{m,N-2}$.

Recursively, consider time n . Suppose that we know $\tau_{m,n}$, the time to exercise the option along path m , given that we have not exercised it till time n . Then, parameters r_n^* are found as a solution to the least squares problem:

$$r_n^* = \arg \min_r \sum_{m=1}^L [g_{\tau_{m,n}}(x_{m,\tau_{m,n}}) - \sum_{k=1}^K \phi_k(x_{m,n}) r(k)]^2.$$

Note that if $\tau_{m,n}$ is a realization of the optimal stopping time, then $g_{\tau_{m,n}}(x_{m,\tau_{m,n}})$ above is an unbiased sample of the continuation value $Q_n(x_{m,n})$ (and hence a reasonable proxy).

We modify this approach to determine the parameters $s_1^*, \dots, s_N^*$. Set

$$s_N^* = \arg \min_s \sum_{m=1}^L [g_N(x_{m,N}) - \sum_{k=1}^K \phi_k(x_{m,N}) s(k)]^2.$$

The parameters s_n^* for $n \leq N-1$ are found after parameters r_n^* have been determined. Knowing r_n^* allows us to determine whether to exercise the option at state $x_{m,n}$ or not by comparing $g_n(x_{m,n})$ and $\hat{Q}(x_{m,n}, r_n^*)$ for each m . Then, $\tau_{m,n-1}$ is known for each m . Set

$$s_n^* = \arg \min_s \sum_{m=1}^L [g_{\tau_{m,n-1}}(x_{m,\tau_{m,n-1}}) - \sum_{k=1}^K \phi_k(x_{m,n}) s(k)]^2.$$

The parameters r_n^* and s_n^* above are determined using the least squares method (see, e.g., Bertsekas and Tsitsiklis 1996). Again, if $\tau_{m,n-1}$ is a realization of the optimal stopping time, then $g_{\tau_{m,n-1}}(x_{m,\tau_{m,n-1}})$ above is an unbiased sample of $J_n(x_{m,n})$. Note that in Bolia, Glasserman and Juneja (2004) we had restricted s_n^* to be non-negative, hence non-negative least squares regression was used. This was useful in implementing importance sampling. However this restriction is no longer necessary in control variate settings suggesting that $\hat{J}_n$ estimated here is more accurate than that in Bolia, Glasserman and Juneja (2004).

3.2 Estimating the Option Price

Once $(\hat{J}(x, s_n^*) : x \in \mathcal{X}, n \leq N)$ are known, we start the second phase of the algorithm using the control variate to evaluate the price of the option.

As in (4), we define the martingale $\hat{M} = (\hat{M}_n \in \mathbb{R} : 0 \leq n \leq N)$ as

$$\hat{M}_n = \sum_{i=0}^{n-1} (\hat{J}_{i+1}(X_{i+1}) - E_i[\hat{J}_{i+1}](X_i)).$$

We generate another set independent paths of the stock price that is also independent of the paths generated earlier to estimate the optimal policy. The martingale realization at time step n along the path m ($x_{m,n} : 1 \leq n \leq N$) is then given by $\hat{M}_{m,n}$. The estimate of the option price is the average of $g_{\tau_{m,0}}(x_{\tau_{m,0}}) - \hat{M}_{m,\tau_{m,0}}$ over all m . The estimate of the upper bound on the option price is the average of $\max_n(g_n(x_{m,n}) - \hat{M}_{m,n})$ over all m . The stopping time $\tau_{m,0}$

along path m is found by comparing the $\hat{Q}_n$ approximated earlier with g_n at every time step n .

4 NUMERICAL RESULTS

We conduct simulation experiments to price a one dimensional American put assuming that the risk-neutral stock price process follows the stochastic differential equation:

$$dS = rSdt + \sigma SdZ$$

where r and σ are constants, Z is standard Brownian motion, and the stock does not pay dividends. The expiration time for the option is denoted by T . Let X_n denote $S(n\Delta t)$ and set $N = T/\Delta t$. Then $(X_n : n \leq N)$ is a Markov process. Let $N(a, b)$ denote a Gaussian random variable with mean a and variance b . Note that X_{n+1} has the distribution

$$X_n \exp[(r - \sigma^2/2)\Delta t + \sqrt{\Delta t}N(0, \sigma^2)]$$

or, given $X_n = x$, we may set

$$X_{n+1} = \exp[N(\hat{\mu}, \hat{\sigma}^2)],$$

where $\hat{\mu} = (r - \sigma^2/2)\Delta t + \log x$ and $\hat{\sigma}^2 = \Delta t\sigma^2$ (see, e.g., Glasserman 2004, p.94).

Thus, $f_n(x, y)$ (or $f(x, y)$ as the densities are time-homogeneous) equals

$$\frac{1}{\sqrt{2\pi\hat{\sigma}^2}} \exp\left[-\frac{1}{2\hat{\sigma}^2}(\log y - \hat{\mu})^2\right].$$

For each $k \leq K$, and constants $(\alpha_{1k}, \alpha_{2k})$ we select our basis function

$$\phi_k(y) = \exp[\alpha_{1k} \log^2 y + \alpha_{2k} \log y]. \quad (16)$$

We keep $\alpha_{1k} \leq 0$. Note that if $\alpha_{1k} > \frac{1}{2\hat{\sigma}^2}$, $E_n[\phi_k](x)$ blows up and if $\alpha_{1k} = 0$, then $\phi_k(y) = y^{\alpha_{2k}}$.

Then,

$$\begin{aligned} \phi_k(y)f(x, y) &= \frac{1}{\sqrt{2\pi\hat{\sigma}^2}} \times \\ &\exp\left[-\frac{1}{2\hat{\sigma}^2}(\log y - \hat{\mu})^2\right] \exp[\alpha_{1k} \log^2 y + \alpha_{2k} \log y]. \end{aligned}$$

After simple algebraic manipulations this can be seen to equal

$$\frac{\tilde{\sigma}}{\hat{\sigma}} \exp\left[\frac{\tilde{\mu}^2}{2\tilde{\sigma}^2} - \frac{\hat{\mu}^2}{2\hat{\sigma}^2}\right] d_{\tilde{\mu}, \tilde{\sigma}^2}^{(k)}(y),$$

where $\tilde{\mu} = \frac{\hat{\mu} + \hat{\sigma}^2 \alpha_{2k}}{1 - 2\hat{\sigma}^2 \alpha_{1k}}$ and $\tilde{\sigma}^2 = \frac{\hat{\sigma}^2}{1 - 2\hat{\sigma}^2 \alpha_{1k}}$ and $d_{\tilde{\mu}, \tilde{\sigma}^2}^{(k)}(y)$ is the pdf of the $\exp[N(\tilde{\mu}, \tilde{\sigma}^2)]$ distributed random variable. Let $\beta_n^*(k) = s_{n+1}^*(k) \frac{\tilde{\sigma}}{\hat{\sigma}} \exp\left[\frac{\tilde{\mu}^2}{2\tilde{\sigma}^2} - \frac{\hat{\mu}^2}{2\hat{\sigma}^2}\right]$, then,

$$E_n[\hat{J}_{n+1}](x) = \sum_{k \leq K} \beta_n^*(k).$$

Thus, we have closed form expressions for $E_n[\hat{J}_{n+1}](\cdot)$ and hence an implementable control variate.

Specifically, in all experiments, we price an American put with a strike price of 40, $T = 1$ year and $r = 0.06$. We consider three values of σ - 0.1, 0.2 and 0.3. In all experiments, we use 4 basis function having the structure as in (16). The respective parameters used are $(\alpha_{11}, \alpha_{21}) = (-1, 0)$, $(\alpha_{12}, \alpha_{22}) = (-1, 1)$, $(\alpha_{13}, \alpha_{23}) = (-1, 2)$ and $(\alpha_{14}, \alpha_{24}) = (-1, 3)$. Bolia, Glasserman and Juneja (2004) empirically establish that the time taken to learn the optimal parameters $r^* = (r_0^*, \dots, r_{N-1}^*)$ is quite small. Therefore, in all the further experiments, we report results with the r^* and the s^* learnt from 10,000 paths. Longstaff and Schwartz (2001) use finite difference methods to find that for time periods $N = 50$, $\sigma = 0.2$ and initial stock price $X_0 = 36$, the price equals 4.478.

In our first experiment, we compare the control variate estimator with naive simulation on our example using the r^* and the s^* as learnt above. We conduct simulations for $N = 10, 20$ and 50 (T remains equal to one year) and for initial stock price $X_0 = 36$ (in-the-money). We do this for 3 values of σ : 0.1, 0.2 and 0.3. For each simulation 10^6 paths are generated to evaluate the option price. The results are shown in Table 1. The variance reduction factor (VR) corresponds to the ratio of the estimate of naive variance and the estimate of control variate variance. If we define time factor as the ratio of time taken per path with the control variate to time taken per path in the naive method], the computational effort reduction factor (CR) in the table (calculated as the ratio of VR and time factor) indicates the overall reduction in the computational effort.

In our second experiment, we price the same option, but with $X_0 = 50$ (all other parameters remain the same). The results are reported in Table 1. We observe that the control variate estimator performs better in all cases. The least improvement is when the option is out-of-the-money with $X_0 = 50$ and $\sigma = 0.1$. This is as expected since starting from $X_0 = 50$, because of a low volatility of 0.1, fewer paths go below the strike price of 40 (for a put option, a path can contribute to the price only if the stock price is less than the strike price) than in other cases, and hence $\hat{J}_n$'s are poor approximations of J_n in this region.

Finally, we report the upper bound for the various cases of volatility in Table 2. The last column in the table is the corresponding (lower biased) estimate of the option price obtained using the control variate method described above.

5 CONCLUDING REMARKS

This paper develops a control variate based method for the pricing of American options. The method is based on the observation that the value function of the pricing problem provides a zero-variance control variate based estimator. We use approximations to the value function to approximate this optimal estimator. Similar ideas lead to an upper bound on the option price that is tight under the control variate based on the true option price.

REFERENCES

- Bertsekas, D.P. and J.N. Tsitsiklis. 1996. *Neuro-Dynamic Programming*. Belmont, Massachusetts: Athena Scientific.
- Bolia, N., P. Glasserman, and S. Juneja. 2004. Function-Approximation-based Importance Sampling for Pricing American Options. In *Proceedings of the 2004 Winter Simulation Conference*, ed. R.G. Ingalls, M.D. Rossetti, J. S. Smith, and B. A. Peters, 604-611. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers.
- Bolia, N. and S. Juneja. 2005. Monte Carlo Methods for Pricing Financial Options. *Sadhana*, 30, 347-386.
- Carrière, J. 1996. Valuation of Early-Exercise Price of Options Using Simulations and Non-Parametric Regression. *Insurance: Mathematics and Economics* 19, 19-30.
- Duffie, D. 1996. *Dynamic Asset Pricing Theory*. Princeton, New Jersey: Princeton University Press.
- Glasserman, P. 2004. *Monte Carlo Methods in Financial Engineering*. New York: Springer-Verlag.
- Haugh, M. B. and L. Kogan. 2004. Pricing American Options: A Duality Approach. *Operations Research* 52(2), 258-270.
- Henderson, S. and P. Glynn. 2001. Approximating Martingales for Variance Reduction in Markov Process Simulation. *Mathematics of Operations Research* 27, 253-271.
- Kallianpur, G. and R. L. Karandikar. 1999. *Introduction to Option Pricing Theory*. Birkhauser Verlag.
- Karatzas, I. and S. Shreve. 1991. *Brownian Motion and Stochastic Calculus*. New York: Springer Verlag.
- Longstaff, F. A., and E. S. Schwartz. 2001. Valuing American Options by Simulation: A Simple Least-Squares Approach. *Review of Financial Studies* 14, 113-147.
- Tsitsiklis, J., and B. Van Roy. 2001. Regression Methods for Pricing Complex American-Style Options. *IEEE Transactions on Neural Networks* 12, 694-703.
- Williams, D. 1991. *Probability with Martingales*, U.K.: Cambridge University Press.

AUTHOR BIOGRAPHIES

NOMESH BOLIA is a project assistant in the School of Technology and Computer Science in the Tata Institute of Fundamental Research. He completed his B. Tech. from the Indian Institute of Technology Bombay in 2003 in mechanical engineering. His e-mail address is <nomesh@tcs.tifr.res.in>.

SANDEEP JUNEJA is a faculty in the School of Technology and Computer Science in the Tata Institute of Fundamental Research. His research interests include variance reduction techniques in performance analysis of stochastic systems. His e-mail address is <juneja@tifr.res.in> and his web page is <www.tcs.tifr.res.in/~sandeepj>.

Table 1: Point Estimates and 95% CI Halfwidths using Naive Simulation and Control Variate

$\sigma = 0.1, X_0 = 36$				
N	Naive Est.	CV Est.	VR	CR
10	3.769 (0.002)	3.770 (0.000)	80.2	39.5
20	3.882 (0.002)	3.881 (0.000)	169.0	79.0
50	3.952 (0.001)	3.952 (0.000)	4391.3	1928.2
$\sigma = 0.2, X_0 = 36$				
N	Naive Est.	CV Est.	VR	CR
10	4.440 (0.006)	4.442 (0.001)	60.2	42.0
20	4.458 (0.005)	4.460 (0.001)	61.1	44.8
50	4.472 (0.005)	4.471 (0.001)	59.1	45.7
$\sigma = 0.3, X_0 = 36$				
N	Naive Est.	CV Est.	VR	CR
10	5.678 (0.009)	5.673 (0.001)	59.5	41.8
20	5.675 (0.008)	5.678 (0.001)	82.8	62.3
50	5.681 (0.008)	5.677 (0.001)	63.5	47.7
$\sigma = 0.1, X_0 = 50$				
N	Naive Est.	CV Est.	VR	CR
10	0.003 (0.000)	0.003 (0.000)	4.7	3.5
20	0.003 (0.000)	0.003 (0.000)	4.2	3.2
50	0.003 (0.000)	0.003 (0.000)	2.2	1.7
$\sigma = 0.2, X_0 = 50$				
N	Naive Est.	CV Est.	VR	CR
10	0.318 (0.002)	0.317 (0.001)	17.4	13.1
20	0.319 (0.002)	0.319 (0.000)	22.5	17.4
50	0.317 (0.002)	0.317 (0.000)	27.7	22.0
$\sigma = 0.3, X_0 = 50$				
N	Naive Est.	CV Est.	VR	CR
10	1.230 (0.005)	1.233 (0.001)	23.6	17.5
20	1.237 (0.005)	1.237 (0.001)	36.8	28.4
50	1.241 (0.005)	1.238 (0.001)	40.8	32.1

Table 2: Upper Bound Estimates and 95% CI Halfwidths

$\sigma = 0.1, X_0 = 36$		
N	Dual Upper Bound	Est. Price
10	4.352 (0.001)	4.000
20	4.286 (0.001)	4.000
50	4.215 (0.001)	4.000
$\sigma = 0.2, X_0 = 36$		
N	Dual Upper Bound	Est. Price
10	4.919 (0.001)	4.442
20	4.793 (0.001)	4.461
50	4.703 (0.001)	4.471
$\sigma = 0.3, X_0 = 36$		
N	Dual Upper Bound	Est. Price
10	6.120 (0.002)	5.678
20	6.096 (0.001)	5.682
50	6.175 (0.002)	5.666
$\sigma = 0.1, X_0 = 50$		
N	Dual Upper Bound	Est. Price
10	0.014 (0.000)	0.003
20	0.015 (0.000)	0.003
50	0.016 (0.000)	0.003
$\sigma = 0.2, X_0 = 50$		
N	Dual Upper Bound	Est. Price
10	0.394 (0.001)	0.320
20	0.404 (0.001)	0.319
50	0.414 (0.001)	0.318
$\sigma = 0.3, X_0 = 50$		
N	Dual Upper Bound	Est. Price
10	1.377 (0.001)	1.239
20	1.413 (0.001)	1.242
50	1.445 (0.001)	1.239