

FUNCTION-APPROXIMATION-BASED IMPORTANCE SAMPLING FOR PRICING AMERICAN OPTIONS

Nomesh Bolia
Sandeep Juneja

Tata Institute of Fundamental Research
Mumbai, INDIA - 400005

Paul Glasserman

Columbia Business School
New York, NY 10027

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 importance sampling measure for American options. We then discuss how function approximation may be used to approximately learn this measure; we test this idea in simple examples. We also note that the zero-variance measure is fundamentally connected to a duality result for American options. While our methodology is geared towards developing an estimate of an accurate lower bound for the option price, we observe that importance sampling also reduces variance in estimating the upper bound that follows from the duality.

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 modelled 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)). 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 execution 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.

In this paper we build upon the function approximation techniques proposed in Longstaff and Schwartz and (2001) and Tsitsiklis and Van Roy (2001) and show how they may be used to develop an effective importance sampling estimator for pricing American options. We view the pricing of an American option as a two-phase procedure: In the first phase the approximately optimal exercise policy is learned, while in the second phase, this policy is evaluated to estimate the American option price. Our experiments suggest that with regression-based methods, the first phase is completed very quickly, even under naive simulation. Thus, our main focus is the second phase, i.e., efficient evaluation of the approximately optimal exercise policy using importance sampling based simulation methods. Note that in this approach, since the policy is evaluated approximately, the price we estimate always lower bounds the true price.

For evaluating the optimal policy, we show that if the original probability density governing the underlying asset movement is biased appropriately by the value functions, we get a zero-variance importance sampling measure. This motivates the use of function approximations to develop an approximate zero-variance estimator. We also briefly discuss how this measure may be learned adaptively, although since the first phase is very quick, no additional benefits are observed from the adaptive approach.

We observe that any importance sampling estimator leads to a 'dual' upper bound on the option price. This is an importance sampling interpretation of the multiplicative dual developed in Jamshidian (2003) (also see Haugh and Kogan (2004) for an 'additive' dual upper bound). Interestingly, this upper bound is tight under the zero-variance estimator, suggesting that as the approximations to the zero-

variance measure improve so does this upper bound. In our experiments, this bound lies within 2–20% of the option price as the number of exercise opportunities ranges from 2 to 50 in a simple one-dimensional example. We observe that importance sampling reduces variance in the estimation of the upper bound.

Regression-based techniques approximate value functions by representing them as linear combinations of a set of basis functions. If the basis functions are not integrable with respect to the original measure and if some of the linear weights assigned are negative then this could lead to implementation difficulties. We resolve this by selecting the basis functions from a family of functions with requisite implementation advantages and we use non-negative least squares to determine the weights assigned to the basis functions. As may be expected, the performance of our approach depends critically on the quality of approximations provided by the basis functions. We explore this further by also experimenting with truncated basis functions.

In Section 2, we develop our mathematical framework and discuss the zero-variance importance sampling distribution. In Section 3, we discuss the proposed importance sampling methodology. Experimental results displaying the effectiveness of the proposed scheme are presented in Section 4.

2 MATHEMATICAL FRAMEWORK

Suppose that the option can be exercised only at $N + 1$ times $0, 1, 2, \dots, N$ (these times need not be equally spaced). 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.

The underlying security prices 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 value in $\{n, n + 1, \dots, N\}$ (recall that τ is a stopping time w.r.t. $\{X_n\}$ if $\{\tau = n\}$ is a measurable function of (X_1, \dots, X_n)). Let

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

the expectation 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.

For simplicity of exposition, 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. Define

$$(P_n J)(x) = E[J(X_{n+1}) | X_n = x] = \int_{\mathcal{X}} J(y) f_n(x, y) dy. \quad (1)$$

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

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

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

$$Q_n(x) = E[J_{n+1}(X_{n+1}) | X_n = x] = (P_n J_{n+1})(x).$$

These recursions are:

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

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

2.1 Importance Sampling and Zero-Variance Measure

Let τ^* denote an optimal stopping time for our problem, i.e., suppose $J_0(x_0) = E[g_{\tau^*}(X_{\tau^*})]$. Suppose this stopping time is known, and to avoid trivialities suppose $\tau^* > 0$. Then, a naive estimate of $J_0(x_0) = Q_0(x_0)$ is obtained by taking an average of independent identically distributed samples of $g_{\tau^*}(X_{\tau^*})$. Suppose, we generate these samples using the importance sampling pdf's $(\tilde{f}_n(x, \cdot) : x \in \mathcal{X}, n \leq N - 1)$ such that $\tilde{f}_n(x, y) > 0$ whenever $f_n(x, y) > 0$ and vice versa for each x, y and n . Let \tilde{P} denote the resultant measure (let P denote the original measure). Then the unbiased importance sampling (IS) estimator of $J_0(x_0)$ is obtained by taking an average of independent, identically

distributed samples of

$$\frac{f_0(x_0, X_1)}{\tilde{f}_0(x_0, X_1)} \frac{f_1(X_1, X_2)}{\tilde{f}_1(X_1, X_2)} \dots \frac{f_{\tau^*-1}(X_{\tau^*-1}, X_{\tau^*})}{\tilde{f}_{\tau^*-1}(X_{\tau^*-1}, X_{\tau^*})} g_{\tau^*}(X_{\tau^*}). \quad (4)$$

(see, e.g., Glynn and Iglehart 1989).

We assume in our analysis that $Q_n(x) > 0$ for all $x \in \mathcal{X}$ and for $n \leq N - 1$. Now suppose that the importance sampling distribution P^* corresponds to:

$$f_n^*(x, y) = \frac{f_n(x, y) J_{n+1}(y)}{Q_n(x)}$$

for each x, y and n . Equation (1) confirms the validity of $f_n^*(x, \cdot)$ as a pdf. Since $Q_n(X_n) = J_n(X_n)$ when $\tau^* > n$ and $J_{\tau^*}(X_{\tau^*}) = g_{\tau^*}(X_{\tau^*})$, it is easy to see that P^* is a zero-variance measure as (4), with f_n^* replacing \tilde{f}_n , reduces to $Q_0(x_0) = J_0(x_0)$ a.s. (Such zero-variance measure are discussed, e.g., in Kollman et al. 1999; Ahamed, Borkar, and Juneja 2004).

In this paper we use regression based approximations of $J_n(\cdot)$ to develop approximations for the zero-variance measure. First we study some properties associated with this measure.

2.2 Dual Estimator

Let \tilde{L}_n denote the n step likelihood ratio of P w.r.t. \tilde{P} , i.e.,

$$\tilde{L}_n = \frac{f_0(x_0, X_1)}{\tilde{f}_0(x_0, X_1)} \frac{f_1(X_1, X_2)}{\tilde{f}_1(X_1, X_2)} \dots \frac{f_{n-1}(X_{n-1}, X_n)}{\tilde{f}_{n-1}(X_{n-1}, X_n)}.$$

(Define L_n^* similarly). Then the importance sampling estimator for $J_0(x_0)$ equals $\tilde{L}_{\tau^*} g_{\tau^*}(X_{\tau^*})$. Note that this is upper bounded by

$$\max_{n \leq N} \tilde{L}_n g_n(X_n). \quad (5)$$

Thus, the average of independent samples of (5) provides an unbiased estimator for an upper bound on $J_0(x_0)$.

Proposition 1 Under P^* ,

$$\max_{n \leq N} L_n^* g_n(X_n) = J_0(x_0),$$

i.e, the upper bound on the samples is constant and tight.

Proof Note that

$$L_n^* g_n(X_n) = \frac{Q_0(x_0)}{J_1(X_1)} \frac{Q_1(X_1)}{J_2(X_2)} \dots \frac{Q_{n-1}(X_{n-1})}{J_n(X_n)} g_n(X_n).$$

It follows that this is upper bounded by $Q_0(x_0) = J_0(x_0)$ since $J_i(x) = \max(g_i(x), Q_i(x))$. Thus, $\max_{n \leq N} L_n^* g_n(X_n) \leq J_0(x_0)$. The result follows as

$$\max_{n \leq N} L_n^* g_n(X_n) \geq L_{\tau^*}^* g_{\tau^*}(X_{\tau^*}) = J_0(x_0).$$

□

This result is an IS formulation of a multiplicative duality result in Jamshidian (2003). It suggests that if \tilde{P} is close to P^* , then $\tilde{E}(\max_{n \leq N} \tilde{L}_n g_n(X_n))$ provides a close upper bound to $J_0(x_0)$ (where \tilde{E} denotes the expectation operator under \tilde{P}). Also, note that

$$\tilde{E}(\max_{n \leq N} \tilde{L}_n g_n(X_n)) = E(\max_{n \leq N} \tilde{L}_n g_n(X_n) / \tilde{L}_N)$$

This provides two ways of estimating the upper bound associated with \tilde{P} : simulating under the IS distribution \tilde{P} and using the expression on the left, or simulating under the original distribution and using the expression on the right. Jamshidian (2003) does the latter. In our experiments, using IS (with the approximate zero-variance distribution) to estimate the upper bound yields lower variance.

2.3 Characterizing Approximation Error

Suppose that we have a $\tilde{J} = (\tilde{J}_n(x) > 0, n \leq N - 1, \tilde{J}_N(x) \geq 0, x \in X)$ and a positive $\epsilon > 0$ such that,

$$(1 - \epsilon) \leq \frac{\tilde{J}_n(x)}{J_n(x)} \leq (1 + \epsilon) \quad (6)$$

for all n and x . For each $n \leq N - 1$ and $x \in \mathcal{X}$, consider the pdf

$$\tilde{f}_n(x, y) = \frac{f_n(x, y) \tilde{J}_{n+1}(y)}{(P_n \tilde{J}_{n+1})(x)}, \quad y \in \mathcal{X},$$

and denote the associated measure by \tilde{P} . Here we assume that $(P_n \tilde{J}_{n+1})(x) < \infty$, also $\tilde{J}_N(x) > 0$ if $g_N(x) > 0$. Let \tilde{L}_n denote the corresponding likelihood ratio of P w.r.t. \tilde{P} restricted to time 0 to n . Further, for any random variable Y , let $V(Y)$ denote its variance.

Proposition 2 Under (6),

$$\max_{n \leq N} \tilde{L}_n g_n(X_n) \leq J_0(x_0) \left(\frac{1 + \epsilon}{1 - \epsilon} \right)^N. \quad (7)$$

Furthermore,

$$V(\tilde{L}_{\tau^*} g_{\tau^*}(X_{\tau^*})) \leq \left(\left(\frac{1 + \epsilon}{1 - \epsilon} \right)^N - 1 \right) J_0(x_0)^2 = NO(\epsilon). \quad (8)$$

Proof From (6) it follows that

$$(1 - \epsilon)J_n(x) \leq (P_n \tilde{J}_{n+1})(x) \leq (1 + \epsilon)J_n(x).$$

Note that

$$\begin{aligned} \tilde{L}_n g_n(X_n) &= \\ &= \frac{(P_0 \tilde{J}_1)(x_0)}{\tilde{J}_1(X_1)} \frac{(P_1 \tilde{J}_2)(X_1)}{\tilde{J}_2(X_2)} \dots \frac{(P_{n-1} \tilde{J}_n)(X_{n-1})}{\tilde{J}_n(X_n)} g_n(X_n) \\ &\leq J_0(x_0) \left(\frac{1 + \epsilon}{1 - \epsilon} \right)^n. \end{aligned}$$

From this, (7) follows. In particular,

$$\tilde{L}_{\tau^*} g_{\tau^*}(X_{\tau^*}) \leq J_0(x_0) \left(\frac{1 + \epsilon}{1 - \epsilon} \right)^N. \quad (9)$$

Note that

$$V(\tilde{L}_{\tau^*} g_{\tau^*}(X_{\tau^*})) = \tilde{E}(\tilde{L}_{\tau^*}^2 g_{\tau^*}(X_{\tau^*})^2) - J_0(x_0)^2.$$

Now,

$$\begin{aligned} \tilde{E}(\tilde{L}_{\tau^*}^2 g_{\tau^*}(X_{\tau^*})^2) \\ = E(\tilde{L}_{\tau^*} g_{\tau^*}(X_{\tau^*})^2) \leq J_0(x_0)^2 \left(\frac{1 + \epsilon}{1 - \epsilon} \right)^N, \end{aligned}$$

where the last step follows from (9). \square

The above analysis suggests that if J_n can be closely approximated by \tilde{J}_n a known quantity (e.g., a function of European options whose value is known), then the corresponding \tilde{P} may provide a good proxy for the zero-variance measure and good simulation efficiency may be expected. However, there may still remain issues concerning efficient generation of random variables from the distributions \tilde{f}_n .

3 PROPOSED METHODOLOGY

As mentioned earlier, 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 importance sampling. However, for importance sampling we also need to find estimators \hat{J} for J . Furthermore, to generate samples and compute the likelihood ratio, it is desirable that the integral $(P_n \tilde{J}_{n+1})(x)$ be known explicitly, and that we may be able to sample cheaply from the associated densities

\tilde{f} . We now discuss how this is achieved using functional approximations.

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 $r \in \mathfrak{R}^K$ 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 $r_0^*, r_2^*, \dots, r_{N-1}^*$ 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 conduct importance sampling we consider a parameterized value function $\hat{J} : \mathcal{X} \times \mathfrak{R}^K \rightarrow \mathfrak{R}$ that assigns values $\hat{J}(x, s)$ to state x where again, $s \in \mathfrak{R}^K$, and

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

We choose each $\phi_k(\cdot)$ so that $(P_n \phi_k)(x)$ can be explicitly evaluated and it is easy to generate samples from the probability density functions

$$\frac{f_n(x, y) \phi_k(y)}{(P_n \phi_k)(x)}.$$

We estimate parameters $s_1^*, s_2^*, \dots, s_N^*$, under non-negativity constraints, 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 first discuss the procedure outlined above in detail. Then, we extend these to adaptive methodologies that combine the two phases.

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 M sample paths $(x_{m,n} : n \leq N, m \leq M)$ of the process $(X_n : n \leq N)$ using the original measure P .

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

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

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:

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

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). This step differs slightly from the method in Longstaff and Schwartz (2001), in that we do not restrict ourselves to in-the-money paths above. Since phase 1 requires relatively less time, we do not focus on these issues.

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

$$s_N^* = \arg \min_{s \geq 0} \sum_{m=1}^M [g_N(x_{m,N}) - \sum_{k=1}^K \phi_k(x_{m,N})s(k)]^2.$$

Note that we now use the non-negative least squares method (as in Lawson and Hanson 1974).

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

$$s_n^* = \arg \min_{s \geq 0} \sum_{m=1}^M [g_{\tau_{m,n-1}}(x_{m,\tau_{m,n-1}}) - \sum_{k=1}^K \phi_k(x_{m,n})s(k)]^2.$$

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 $J_n(x_{m,n})$.

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 involving importance sampling to evaluate the price of the option. The importance sampling probability densities are given by

$$\hat{f}_n(x, y) = \frac{f_n(x, y) \hat{J}(y, s_{n+1}^*)}{\int_{y \in \mathcal{X}} f_n(x, y) \hat{J}(y, s_{n+1}^*) dy}.$$

This may be re-expressed as

$$\hat{f}_n(x, y) = \sum_{k \leq K} p_k^*(x) \frac{f_n(x, y) \phi_k(y)}{(P_n \phi_k)(x)},$$

where

$$p_k^*(x) = \frac{s_{n+1}^*(k) (P_n \phi_k)(x)}{\sum_{k \leq K} s_{n+1}^*(k) (P_n \phi_k)(x)}.$$

Note that $p_k^*(x) \geq 0$ and $\sum_{k \leq K} p_k^*(x) = 1$. Hence, if we can easily generate a sample from the pdf $\frac{f_n(x, \cdot) \phi_k(\cdot)}{(P_n \phi_k)(x)}$, then this makes generation from $\hat{f}_n(x, \cdot)$ also straightforward.

3.3 Adaptively Learning the Parameters

We briefly outline a methodology to adaptively find parameters $r_0^*, r_1^*, \dots, r_{N-1}^*$ so that

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

for each x and n . Again suppose that M sample paths $(x_{m,n} : n \leq N, m \leq M)$ of the process $(X_n : n \leq N)$ are generated. Suppose that path m is generated using measure $P(m)$ that is equivalent to P and under which the process $\{X_n\}$ remains Markov. Let L_{m,n_1,n_2} ($0 \leq n_1 < n_2 \leq N$) denote the ratio of the product of the original pdf's of transitions $(x_{m,n_1}, x_{m,n_1+1}), \dots, (x_{m,n_2-1}, x_{m,n_2})$ and the product of new pdf's of these transitions under $P(m)$. Then, we may again set

$$r_{N-1}^* = \arg \min_r \sum_{m=1}^M [g_N(x_{m,N}) L_{m,N-1,N} - \sum_{k=1}^K \phi_k(x_{m,N}) r(k)]^2 L_{m,0,N-1}$$

and, recursively, set

$$r_n^* = \arg \min_r \sum_{m=1}^M [g_{\tau_{m,n}}(x_{m,\tau_{m,n}}) L_{m,n,\tau_{m,n}} - \sum_{k=1}^K \phi_k(x_{m,n}) r(k)]^2 L_{m,0,n-1}.$$

Note that these parameters are determined using the weighted least squares method (see, e.g., Bertsekas and

Tsitsiklis 1996). Again, if $\tau_{m,n}$ is a realization of the optimal stopping time, then $g_{\tau_{m,n}}(x_{m,\tau_{m,n}})L_{m,n,\tau_{m,n}}$ above is an unbiased sample of $Q_n(x_{m,n})$. The term $L_{m,0,n-1}$ is used to unbias the effect of generating the state $x_{m,n}$ using importance sampling. The parameters $s_1^*, s_2^*, \dots, s_N^*$ may be determined similarly, with a non-negativity constraint.

In practice, this adaptive approach may be implemented as follows: one may initially generate a fixed number of samples under the naive measure and learn the optimal stopping rule and the new importance sampling measure. This measure may then be used to generate more sample paths to refine the optimal stopping rule and update the importance sampling measure. This process may then be repeated many times. These paths (after the naive simulation phase) may be used to provide samples of payoffs from the approximately optimal policy, whose average then estimates the option price.

In numerical experiments (not reported), we find that the adaptive method leads to a slight reduction in variance compared to the two-step approach, especially when the number of exercise opportunities is large. However, due to increase in the number of updates of the weight parameters, the average per sample effort increases.

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

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

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

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

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

Then,

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

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 $\beta_n^*(k) = s_{n+1}^*(k) \frac{\tilde{\sigma}}{\hat{\sigma}} \exp\left[\frac{\tilde{\mu}^2}{2\tilde{\sigma}^2} - \frac{\hat{\mu}^2}{2\hat{\sigma}^2}\right]$, and set $p_n^*(k) = \frac{\beta_n^*(k)}{\sum_{k \leq K} \beta_n^*(k)}$. Then our importance sampling pdf's have the form

$$f_n^*(x, y) = \sum_{k \leq K} p_n^*(k) d_{\tilde{\mu}, \tilde{\sigma}^2}^{(k)}(y),$$

where each $p_n^*(k)$ depends upon x . Thus, for each x , $f_n^*(x, \cdot)$ corresponds to a non-negative mixture of log-normal distributions and it is easy to generate samples from it.

In an attempt to improve the approximations in the case of non-negative least squares method, we also consider truncated basis functions of the form

$$\phi_k(y) = \exp[\alpha_{1k} \log^2 y + \alpha_{2k} \log y] I(a_k < y < b_k),$$

where $I(\cdot)$ is an indicator function. In this case, $f_n^*(x, y)$ corresponds to an appropriate non-negative mixture of truncated log-normal distributions.

Specifically, we focus on the problem of pricing an American put with a strike price of 40, $T = 1$ year, $r = 0.06$ and $\sigma = 0.20$. For this option for the initial stock price $X_0 = 36$, Longstaff and Schwartz (2001) use finite difference methods to find that for time periods $N = 50$ the price equals 4.478.

First, we empirically establish that the time taken to learn the optimal parameters $r^* = (r_0^*, \dots, r_{N-1}^*)$ is quite small. We do this by generating a fixed number of initial phase paths to determine an approximately optimal stopping policy using the r^* computed from these paths. Then we evaluate the option price under this policy by generating 100,000 samples using naive simulation. The results for $X_0 = 36$, $N = 50$ are shown in Table 1. Four untruncated

basis functions are used. Their parameters (α_1, α_2) are set at $(-1, 0)$, $(-1, 1)$, $(-1, 2)$ and $(-1, 3)$. The results show that only a small number of paths are needed in the initial phase to approximately learn the optimal policy for this example. Similar results were observed (not reported) for initial stock price equal to 50 and 30 (at this price it is optimal to exercise immediately, so we focussed on estimating the continuation value at time zero).

Table 1: Comparison of Price Estimates as the Number of Paths in Initial Phase Varies

Paths in Initial Phase	Estimated Mean (100,000 paths)	95% C. I.
2,000	4.471	(4.454, 4.488)
5,000	4.481	(4.463, 4.499)
10,000	4.485	(4.468, 4.502)
50,000	4.475	(4.457, 4.492)

In the second experiment we compare importance sampling with naive simulation on our example using the r^* and the s^* learned from 10,000 generated paths using naive simulation. We conduct simulations for $N = 2, 10, 20$ and 50 (T remains equal to one year) and for initial stock price $X_0 = 36$ and 50. For each simulation 100,000 paths are generated to evaluate the option price. The results are shown in Table 2. The variance reduction factor (VR) corresponds to the ratio of the estimate of naive variance and the estimate of importance sampling variance. We also estimate the dual upper bound via importance sampling simulation. As illustrated in Table 3, importance sampling yields more precise estimates of the upper bound than does naive simulation. The basis functions used are the same four used in the first experiment. We also experimented with varying the parameters of the four basis functions (to neighboring values) and varying the number of basis functions to 2, 3 and 6 similar functions. The results (not reported) are more-or-less insensitive to these perturbations. Similar results were seen for initial stock price 30 (not reported), however this is a less interesting case as here with a large probability it is optimal to exercise quickly.

A problem with the experiments reported in Table 2 is that in most cases for each n only one or two of the $s_n^*(k)$ assume a positive value; thus only one or two basis functions are used in conducting importance sampling, suggesting that better approximations of J may lead to improved performance. With this in mind, we test truncated basis functions. However, more analysis and experimentation is needed to develop insights for improving approximations.

Parameters for the truncated basis functions are specified in Table 4. These are chosen so that near the initial stock price $X_0 = 36$, we may be able to approximate the terminal value J_N reasonably accurately. Again we compare importance sampling with naive simulation on our example using the r^* and the s^* learned from 10,000 generated paths using naive

Table 2: Point Estimates and 95% CI Halfwidths Using Naive Simulation and Importance Sampling

N	X_0	Naive Est.	IS Est.	VR
2	36	4.189 (.023)	4.206 (.009)	6.2
10	36	4.426 (.018)	4.438 (.008)	5.4
20	36	4.451 (.018)	4.458 (.008)	5.3
50	36	4.455 (.017)	4.475 (.008)	5.3
2	50	0.310 (.007)	0.311 (.001)	16
10	50	0.316 (.007)	0.319 (.001)	15.9
20	50	0.325 (.007)	0.321 (.002)	14.7
50	50	0.322 (.007)	0.325 (.002)	15.2

Table 3: Upper Bound Estimates (with 95% CI Halfwidths) Using IS and Naive Simulation

N	X_0	Dual Upper Bound (IS)	Dual Upper Bound (Naive)
2	36	4.280 (.008)	4.292 (.03)
10	36	4.912 (.004)	4.926 (.03)
20	36	5.007 (.003)	4.985 (.03)
50	36	5.063 (.003)	5.048 (.03)
2	50	0.315 (.002)	0.311 (.01)
10	50	0.358 (.002)	0.363 (.01)
20	50	0.381 (.002)	0.386 (.01)
50	50	0.395 (.002)	0.395 (.01)

simulation. For each simulation 100,000 paths are generated to evaluate the option price. The results are shown in Table 5 where we show point estimates under naive simulation and IS along with 95% CI halfwidths, and variance reduction factor VR. The performance of importance sampling compared with naive simulation was better for $X_0 = 30$ with these basis functions. For $X_0 = 50$, with these basis functions, even the mean estimates were significantly off the mark with the bias increasing with the increase in time periods, both under naive and importance sampling simulation.

Table 4: Truncated Basis Functions

Basis fn. i	a_i	b_i	α_{1i}	α_{2i}
1	0	16	0	-0.05
2	16	25	-0.1	1.5
3	25	30	-1.4	5.2
4	30	35	-1.4	4.7
5	35	40	-1.4	4.5
6	40	50	-0.3	0.1
7	50	60	-0.2	0.1
8	60	∞	-0.1	0.0

Table 5: Comparison between Naive Simulation and IS Using Truncated Basis Functions

N	Naive Est.	IS Est.	VR
2	4.229 (.023)	4.200 (.004)	33.2
3	4.372 (.019)	4.378 (.006)	12.8
10	4.424 (.017)	4.415 (.008)	5.0
20	4.431 (.016)	4.418 (.010)	2.3
50	4.457 (.017)	4.457 (.022)	0.6

5 CONCLUDING REMARKS

This paper develops an importance sampling 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 importance sampling distribution. We use approximations to the value function to approximate this optimal distribution. Similar ideas lead to an upper bound on the option price that is tight under the zero variance measure; our IS method appears to reduce variance in estimating this upper bound as well.

The main obstacle to applying this method lies in the choice of basis functions. This is always an issue in developing function approximations for dynamic programming. But the IS method introduces a further consideration: it must be practical to sample from the distributions obtained by biasing the original transition law by each of the basis functions. Our examples accomplish this in the case of the lognormal distribution, but this remains a challenge in more complex models.

ACKNOWLEDGMENTS

This research was supported in part by NSF grant DMI0300044.

REFERENCES

Ahamed, T. P. I., V. S. Borkar and S. Juneja. 2004. Adaptive Importance Sampling Technique for Markov Chains Using Stochastic Approximation. Submitted for publication.

Bertsekas, D.P. and J.N. Tsitsiklis. 1996. *Neuro-Dynamic Programming*, Athena Scientific, Belmont, Massachusetts.

Carrière, J. 1996. Valuation of Early-Exercise Price of Options Using Simulations and Non-Parametric Regression. *Insurance: Mathematics and Economics* 19, 19-30.

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

Glynn, P.W. and D.L. Iglehart. 1989. Importance Sampling for Stochastic Simulations. *Management Science* 35, 1367-1392.

Haugh, M. B. and L. Kogan. 2004. Pricing American Options: A Duality Approach. *Operations Research* 52-2, 258-270.

Jamshidian, F. 2003. Minimax Optimality of Bermudan and American Claims and their Monte-Carlo Upper Bound Approximation. Preprint.

Kollman, C., K. Baggerly, D. Cox and R. Picard. 1999. Adaptive Importance Sampling on Discrete Markov Chains. *Annals of Applied Probability* 9, 391-412.

Lawson, C.L. and R.J.Hanson. 1974. Solving Least Squares Problems, Prentice-Hall, Englewood Cliffs, New Jersey.

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.

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>.

PAUL GLASSERMAN is the Jack R. Anderson Professor in the Decision, Risk, and Operations division of Columbia Business School. His research interests include variance reduction techniques and modelling and computational issues in risk management and the pricing of derivative securities. He is author of the book *Monte Carlo Methods in Financial Engineering*, published by Springer in 2004. His email address is <pg20@columbia.edu>, and his web page is <www.paulglasserman.net>.

SANDEEP JUNEJA is an academic member 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>.