Adaptive Control Variates for Pricing Multi-Dimensional American Options

Samuel M. T. Ehrlichman   Shane G. Henderson

July 25, 2007

Abstract

We explore a class of control variates for the American option pricing problem. We construct the control variates by using multivariate adaptive linear regression splines to approximate the option's value function at each time step; the resulting approximate value functions are then combined to construct a martingale that approximates a "perfect" control variate. We demonstrate that significant variance reduction is possible even in a high-dimensional setting. Moreover, the technique is applicable to a wide range of both option payoff structures and assumptions about the underlying risk-neutral market dynamics. The only restriction is that one must be able to compute certain one-step conditional expectations of the individual underlying random variables.

1 Introduction

Efficient pricing of American options remains a thorny issue in finance. This is true despite the fact that numerical techniques for solving this problem have been studied for decades – certainly at least since the binomial tree method of Cox et al. (1979). Both tree-based methods and PDE methods are very fast in low dimensions but do not extend well to higher-dimensional problems, arbitrary stochastic processes, or arbitrary payoff structures. In the last decade or so, attention has been turned to simulation techniques to solve such problems.

What makes American options much more difficult to price than their European counterparts, of course, is the embedded optimal stopping problem. Many of the early papers on using simulation to price American options therefore focus on this aspect of the computation. Carriere (1996) uses nonparametric regression techniques to approximate the value of continuing (i.e., not exercising) at every time step, proceeding backwards in time from expiry. This in turn produces a stopping rule: exercise only if the (known) value of exercise exceeds the (approximate) value of continuing. These ideas are developed further in Longstaff and Schwartz (2001) and Tsitsiklis and Van Roy (2001), both of which use linear regression on a fixed set of basis functions to approximate the continuation value.

The continuation value approximations obtained using these methods are not perfect, but they do yield feasible stopping policies. These policies therefore yield lower bounds on the true option price. Recently, Haugh and Kogan (2004), Rogers (2002), and Andersen and Broadie (2004) showed how to compute upper bounds on the option price via a martingale duality. Bolia and Juneja (2005) observed that this same martingale, if computed by function approximation instead of simulation-within-simulation, can serve as a simulation control variate and thereby provide variance reduction. This approach can be viewed as a special case of a class of martingale control variate methods introduced by Henderson and Glynn (2002).

The method introduced by Bolia and Juneja (2005) relies on finding a particular set of basis functions. To avoid internal simulations it is necessary that the basis functions be such that one can easily compute certain one-step conditional expectations. In related work, Rasmussen (2005) computes a control variate for the option price by using a carefully chosen European option (or several such options), evaluated at the exercise time of the American option being priced. Laprise et al. (2006) construct upper and lower piecewise linear approximations of the value function and compute the American option price using a sequence of portfolios of European options. Their method only works in one dimension, though. An earlier use of European options as control variates for American options appears in Broadie and Glasserman (2004), wherein the European options in question expire in a single time step and employed at each step of a stochastic mesh scheme.

The work we present here, like that of Bolia and Juneja (2005), can be thought of as a "primal-dual" method, in the sense of Andersen and Broadie (2004). The martingale-based control variate is used both to improve the quality of the lower bound *and* to derive the upper bound. In our work, as well as that of Andersen and Broadie (2004), the upper bound solution is derived by first considering a suboptimal stopping strategy, and then deriving a corresponding martingale. Thus, a poor choice of stopping strategy will never be "rescued" by the fact that an upper bound is available. However, unlike Andersen and Broadie (2004), our upper bound solutions do not involve any additional simulation trials. As a result, the quality of the upper bound depends not only on the quality of the suboptimal stopping times but *also on a functional approximation* for the martingale from which the upper bound arises. In that sense, our work can also be thought of as primarily a variance reduction technique for lower bound methods, albeit one which produces an upper bound for free.

Our contribution is to identify a technique for computing the control variate that possesses the desired tractability property in a quite general setting. Moreover, construction of the control variate is more or less automatic; once it has been done for one pricing problem it can be extended to other problems without much effort. We demonstrate these extensions in detail for various basket options, barrier options, and Asian options in both a Black-Scholes and stochastic volatility (Heston 1993) model.

Rogers (2002) commented that the selection of the dual martingale may be "more art than science." We contend that our approach takes a bit of the art out of this process and injects, if not science, at least some degree of automation to the procedure.

The remainder of this paper is organized as follows. Section 2 gives some mathematical preliminaries, recalls the pricing algorithm of Longstaff and Schwartz (2001), defines the martingales that we work with and clarifies their linkage with the pricing problem. Section 3 discusses multivariate adaptive regression splines (Friedman 1991), or MARS, and discriminant analysis, which are the techniques we adapt to construct martingales. Section 4 describes the algorithm in detail. Section 5 gives a number of examples, and we offer some conclusions in Section 6.

2 Mathematical Framework

As in most papers that discuss simulation applied to American option pricing, we actually consider the problem of pricing a Bermudan option, which differs from its American counterpart in that it may be exercised only at a finite set of points in time. To simplify notation, we assume that these times are the evenly spaced steps $t = 0, \ldots, T$.

Let $(X_t : t = 0, \ldots, T)$ be an $\mathbb{R}^d$-valued process on a filtered probability space $(\Omega, \mathscr{F}, \mathbb{P})$, where $\mathscr{F} = (\mathscr{F}_t : t = 0, \ldots, T)$ is the natural filtration of (X_t). We assume (X_t) to be Markov, enlarging the state space if necessary to ensure this. We treat X_0 as deterministic, so $\mathscr{F}_0$ is taken to be trivial. Let r be the riskless interest rate which we assume to be constant and, to simplify notation, normalized so that if $s < t$, the time-s dollar value of \$1 to be delivered at time t is $e^{-r(t-s)}$. We assume that the market is arbitrage-free and work exclusively with a risk-neutral (pricing) measure $\mathbb{Q}$ with the same null sets as $\mathbb{P}$. See e.g., Duffie

(2001) or Glasserman (2004) for details on risk-neutral pricing.

Let the known function $g : \{0, \dots, T\} \times \mathbb{R}^d$ satisfy $g(t, \cdot) \geq 0$ and $Eg^2(t, X_t) < \infty$ for all $t = 0, \dots, T$. We interpret $g(t, X_t)$ to be the value of exercising the option at time t in state X_t. Let $\mathcal{T}(t)$ be the set of all $\mathscr{F}$-stopping times valued in $\{t, \dots, T\}$. Then the Bermudan option pricing problem is to compute Q_0, where

$$Q_t = \sup_{\tau \in \mathcal{T}(t)} E_t \left[e^{-r(\tau - t)} g(\tau, X_\tau) \right],$$

for $t = 0, \dots, T$. We recall some theory about American and Bermudan options; again see e.g., Duffie (2001) for details. The above optimal stopping problem admits a solution $\tau_t^* \in \mathcal{T}(t)$, so that

$$Q_t = E_t \left[e^{-r(\tau_t^* - t)} g\left(\tau_t^*, X_{\tau_t^*}\right) \right]$$

for each $t = 0, \dots, T$. Moreover, the Q_t's satisfy the backward recursion

$$\begin{aligned} Q_T &= g(T, X_T), \\ Q_t &= \max\left\{ g(t, X_t), e^{-r} E_t Q_{t+1} \right\}, \end{aligned}$$

for $t = 0, \dots, T-1$. Therefore, the optimal stopping times $\tau_0^*, \dots, \tau_T^*$ satisfy

$$\begin{aligned} \tau_T^* &\equiv T, \\ \tau_t^* &= \begin{cases} t & \text{if } g(t, X_t) \geq e^{-r} E_t Q_{t+1}, \\ \tau_{t+1}^* & \text{otherwise,} \end{cases} \end{aligned}$$

for $t = 0, \dots, T-1$. An easy consequence of this is

$$s < t \leq \tau_s^* \Longrightarrow \tau_s^* = \tau_{s+1}^* = \cdots = \tau_t^*. \tag{1}$$

2.1 The Longstaff-Schwartz Method

The least-squares Monte Carlo (LSM) method of Longstaff and Schwartz (2001) provides an approximation to the optimal stopping times (τ_t^*) and hence to the option price process (Q_t). Since the resulting stopping times (τ_t) are suboptimal for the original problem, the value of following such a stopping strategy provides a lower bound on the true price process. Following the notation of Andersen and Broadie (2004), we denote the lower bound process by

$$L_t = E_t e^{-r(\tau_t - t)} g(\tau_t, X_{\tau_t}).$$

We now recall the procedure by which LSM computes the stopping times (τ_t) and hence (L_t). Let $\phi_0, \phi_1, \dots,$ be a collection of functions from $\mathbb{R}^d$ to $\mathbb{R}$ such that $\phi_0 \equiv 1$ and $\{\phi_i(X_t) : i = 0, 1, \dots\}$ form a basis for $L^2(\Omega, \sigma(X_t), \mathbb{Q})$ for all $t = 1, \dots, T$. The algorithm proceeds as follows. Denote $\boldsymbol{\phi} = (\phi_0, \dots, \phi_k)$, for some fixed k. Generate a set of N paths $\{X_t(n) : t = 0, \dots, T; n = 1, \dots, N\}$. Set $\tau_T(n) = T$ and $L_T(n) = g(T, X_T(n))$ for $n = 1, \dots, N$. Then recursively estimate

$$\boldsymbol{\alpha}_t = \operatorname*{argmin}_{\boldsymbol{\alpha}} \sum_{n=1}^{N} 1_{[g(t, X_t(n)) > 0]} \left(\boldsymbol{\alpha}' \boldsymbol{\phi}(X_t(n)) - L_{t+1}(n) \right)^2, \tag{2}$$

$$\begin{aligned} \tau_t(n) &= \begin{cases} t & \text{if } g(t, X_t(n)) > [\boldsymbol{\alpha}_t' \boldsymbol{\phi}(X_t(n))]_+ \\ \tau_{t+1}(n) & \text{otherwise,} \end{cases} \\ L_t(n) &= \begin{cases} g(t, X_t(n)) & \text{if } \tau_t(n) = t, \\ e^{-r} L_{t+1}(n) & \text{otherwise,} \end{cases} \end{aligned}$$

for $t = T-1, \ldots, 0$. The regression in (2) is performed only on those paths which have positive exercise value at time t, thus (we hope) producing a better fit on the paths that actually matter than we would obtain if we performed the regression on the complete set of paths. We shall comment on this point in Section 4.1 when we describe our version of the algorithm with the control variate.

The idea behind the LSM algorithm is that if τ_t is close to the true optimal stopping time τ_t^*, then the lower-bounding value process L_t is close to Q_t. It is shown in Clément et al. (2002) both that the approximations τ_t converge to τ_t^* and that the approximations L_t converge to Q_t as the number of basis functions used $k \to \infty$.

2.2 Martingales and Variance Reduction

As we have already noted, the stopping times (τ_t) obtained in the LSM method are suboptimal, and so the option prices (L_t) implied by the algorithm are lower bounds on the true option prices (Q_t). To obtain an upper bound we employ a martingale duality result developed independently by Haugh and Kogan (2004) and Rogers (2002). Let $\pi = (\pi_t : t = 0, \ldots, T)$ denote a martingale with respect to $\mathscr{F}$. By the optional sampling theorem, for any $t \geq 0$,

$$\begin{aligned} Q_t &= e^{rt} \sup_{\tau \in \mathcal{T}(t)} E_t \left[e^{-r\tau} g(\tau, X_\tau) - \pi_\tau + \pi_\tau \right] \\ &= e^{rt} \sup_{\tau \in \mathcal{T}(t)} E_t \left[e^{-r\tau} g(\tau, X_\tau) - \pi_\tau \right] + e^{rt} \pi_t \\ &\leq e^{rt} E_t \max_{s=t,\ldots,T} \left[e^{-rs} g(s, X_s) - \pi_s \right] + e^{rt} \pi_t \\ &=: U_t. \end{aligned} \tag{3}$$

The martingale π here is arbitrary, and any such choice yields an upper bound. We next give a class of martingales from which to choose.

Let $h_t : \mathbb{R}^d \to \mathbb{R}$ be such that $E|h_t(X_t)| < \infty$ for each $t = 0, 1, \ldots, d$. Define $\pi_0 = 0$, and for $t = 1, \ldots, T$, set

$$\pi_t = \sum_{s=1}^{t} e^{-rs} (h_s(X_s) - E_{s-1} h_s(X_s)). \tag{4}$$

Evidently (π_t) is a martingale, and can be used to obtain an upper bound on the option price as in (3).

Such martingales can also be used to great effect as control variates in estimating the lower bound process. Recall that $L_t = E_t e^{-r(\tau_t - t)} g(\tau_t, X_{\tau_t})$ for each t, and so conditional on $\mathscr{F}_t$, we can compute L_t by averaging conditionally independent replicates of $e^{-r(\tau_t - t)} g(\tau_t, X_{\tau_t})$. Proposition 2.1 shows that if we choose the function h_t so that $h_t(X_t) = L_t$ for each t, then the martingale difference $\pi_{\tau_t} - \pi_t$ is a perfect control variate, in the sense that it is perfectly correlated with $e^{-r(\tau_t - t)} g(\tau_t, X_{\tau_t})$, conditional on $\mathscr{F}_t$. This generalizes a comment in Bolia and Juneja (2005), who show that the proposition holds in the case $t = 0$.

Proposition 2.1. *Suppose that h_t is chosen so that $h_t(X_t) = L_t$ for each $t = 1, \ldots, T$. Then the martingale $\pi = (\pi_t : t = 0, \ldots, T)$ defined in (4) satisfies*

$$\pi_{\tau_t} - \pi_t = e^{-r\tau_t} g(\tau_t, X_{\tau_t}) - e^{-rt} L_t$$

for each $t = 0, 1, \ldots, T$.

Proof. First observe that on the event $[\tau_t = t]$, both sides of the equality we are trying to prove are zero. Hence, it suffices to prove that the result holds in the continuation region, i.e.,

$$1_{[\tau_t > t]} \left(\pi_{\tau_t} - \pi_t \right) = 1_{[\tau_t > t]} \left(e^{-r\tau_t} g\left(\tau_t, X_{\tau_t} \right) - e^{-rt} L_t \right). \tag{5}$$

Now, if $s \in \{t+1, \ldots, T\}$, then

$$\begin{aligned}
1_{[\tau_t \geq s]} E_{s-1} L_s &= 1_{[\tau_t \geq s]} E_{s-1} \left[E_s e^{-r(\tau_s - s)} g\left(\tau_s, X_{\tau_s}\right) \right] \\
&= E_{s-1} 1_{[\tau_t \geq s]} e^{-r(\tau_s - s)} g\left(\tau_s, X_{\tau_s}\right) \\
&= e^r E_{s-1} 1_{[\tau_t \geq s]} e^{-r(\tau_{s-1} - (s-1))} g\left(\tau_{s-1}, X_{\tau_{s-1}}\right) \\
&= e^r 1_{[\tau_t \geq s]} L_{s-1},
\end{aligned}$$

where the penultimate equality uses the fact that

$$s < t \leq \tau_s \Longrightarrow \tau_s = \tau_{s+1} = \cdots = \tau_t,$$

analogous to (1). Therefore,

$$\begin{aligned}
\pi_{\tau_t} - \pi_t &= \sum_{s=t+1}^{T} 1_{[\tau_t \geq s]} e^{-rs} \left(L_s - E_{s-1} L_s\right) \\
&= \sum_{s=t+1}^{T} 1_{[\tau_t \geq s]} e^{-rs} \left(L_s - e^r L_{s-1}\right) \\
&= \sum_{s=t+1}^{T} 1_{[\tau_t \geq s]} e^{-rs} L_s - \sum_{s=t}^{T-1} 1_{[\tau_t \geq s+1]} e^{-rs} L_s \\
&= 1_{[\tau_t = T]} e^{-rT} L_T + \sum_{s=t+1}^{T-1} e^{-rs} L_s \left(1_{[\tau_t \geq s]} - 1_{[\tau_t \geq s+1]}\right) - 1_{[\tau_t > t]} e^{-rt} L_t \\
&= \sum_{s=t+1}^{T} e^{-rs} L_s 1_{[\tau_t = s]} - 1_{[\tau_t > t]} e^{-rt} L_t.
\end{aligned}$$

So

$$\begin{aligned}
1_{[\tau_t > t]} \left(\pi_{\tau_t} - \pi_t\right) &= 1_{[\tau_t > t]} \sum_{s=t+1}^{T} 1_{[s = \tau_t]} e^{-rs} L_s - 1_{[\tau_t > t]} e^{-rt} L_t \\
&= 1_{[\tau_t > t]} \left(e^{-r\tau_s} L_{\tau_s} - e^{-rt} L_t\right),
\end{aligned}$$

proving (5). $\square$

Proposition 2.1 shows that conditional on $\mathscr{F}_t$, we can estimate L_t with zero (conditional) variance by

$$e^{-r(\tau_t - t)} g(\tau_t, X_{\tau_t}) - e^{rt}(\pi_{\tau_t} - \pi_t).$$

Since $\mathscr{F}_0$ is the trivial sigma field, by taking $t = 0$ we get a zero variance estimator of L_0, the lower bound on the option price at time 0. In other words, $e^{rt}\left(\pi_{\tau_t} - \pi_t\right)$ is the "perfect" additive control variate for estimating L_t from $e^{-r(\tau_t - t)} g(\tau_t, X_{\tau_t})$.

Of course, we cannot set $h_t(X_t) \equiv L_t$, since we are trying to compute L_t in the first place. But this observation motivates us to search for a set of functions $\{\hat{h}_t\}$ such that

$$\hat{h}_t(X_t) \approx L_t$$

for each t. (In this paper the approximation is in the mean-square sense.) Let us write $\hat{L}_t$ for $\hat{h}_t(X_t)$, and let the induced martingale be $\hat{\pi} = (\hat{\pi}_t : t = 0, 1, \ldots, T)$, where $\hat{\pi}_0 = 0$ and

$$\hat{\pi}_t = \sum_{s=1}^{t} e^{-rs} \left(\hat{L}_s(X_s) - E_{s-1} \hat{L}_s(X_s)\right).$$

We use the approximately optimal martingale $\hat{\pi}$ evaluated at time τ_0 as a control variate in estimating L_0, as indicated by the remark following Proposition 2.1; details are given in Section 4.

Observe that there are two distinct approximations being performed. The one described in the preceding paragraph approximates the value of the option at time t by a (more tractable) function of the underlying state X_t. In contrast, (2) projects the realized value of the option at time $t+1$ onto a certain space of random variables measurable with respect to $\mathcal{F}_t$. In the language of Glasserman and Yu (2004), the approximation used to compute the martingale is "regression later," whereas the approximation (2) used for the stopping strategy is "regression now."

In addition to serving as a control variate, the martingale $\hat{\pi}$ begets an upper bound on the true option price, as in (3). Andersen and Broadie (2004) show that the martingale π is the optimal one to use in computing the upper bound, and indeed that the inequality in (3) would actually be an equality if we had $\tau_t = \tau_t^*$ almost surely, for $t = 0, \ldots, T$. This motivates the use of the martingale $\hat{\pi}$ to compute the upper bound. We note that the same observation is made in Bolia and Juneja (2005).

To compute the martingale $\hat{\pi}$, we need to be able to compute the conditional expectation $E_{s-1}h_s(X_s)$ efficiently. We restrict the class of functions $\{h_t\}$ considered so that these conditional expectations can be evaluated without the need to resort to further simulation, in the same spirit as Bolia and Juneja (2005) and Rasmussen (2005). Bolia and Juneja (2005) use a particular parametric form for h_t which is easily fit by least squares, but is tightly coupled with the specific stochastic process considered. Rasmussen (2005) chooses h_t to be the value of a European option, or a combination of several European options, that are highly correlated with the American option being priced. Indeed, in many examples Rasmussen (2005) simply chooses h to be given by $h_t(X_t) = E_t g(T, X_T)$ so that $\pi_t = e^{-rt}E_t g(T, X_T) - Eg(T, X_T)$. The success of their method, therefore, depends on the ability to find particular European options which can be easily priced and which correlate well with the American option in question. Our method also involves the pricing of European options in a sense, although not necessarily options on traded assets. Like Broadie and Glasserman (2004), the European options we use as control variates each expire after a single time step. These options are automatically selected using the MARS fitting procedure, and in general are priced easily. We now explore MARS.

3 MARS and Extensions

Multivariate adaptive regression splines (Friedman 1991), or MARS, is a nonparametric regression technique that has enjoyed widespread use in a variety of applications since its introduction. For example, Chen et al. (1999) use MARS to approximate value functions of a stochastic dynamic programming problem, although for a different purpose than we do here.

Given observed responses $y(1), \ldots, y(N) \in \mathbb{R}$ and predictors $x(1), \ldots, x(N) \in \mathbb{R}^d$, MARS fits a model of the form

$$y \approx \hat{f}(x) = \alpha_0 + \sum_{m=1}^{M_1} \alpha_{1,m} f_{1,m}(x) + \sum_{m=1}^{M_2} \alpha_{2,m} f_{2,m}(x) + \cdots + \sum_{m=1}^{M_p} \alpha_{p,m} f_{p,m}(x).$$

Each function $f_{1,m}$ takes one of two forms,

$$f_{1,m}(x) \in \left\{ \left(\left[x^{(i)} - x^{(i)}(n)\right]\right)_+, \left(\left[x^{(i)}(n) - x^{(i)}\right]\right)_+ \right\},$$

for some $i = 1, \ldots, d$, and some $n = 1, \ldots, N$. Here, $x^{(i)}$ denotes the i'th coordinate of x. Each function $f_{j,m}$ for $j > 1$ is a product of functions used in previous sums so that the total degree is j. In our setting, we take $p = 1$ so the fitted model can be written

$$y \approx \hat{f}(x) = \alpha_0 + \sum_{i=1}^{d} \sum_{j=1}^{J_i} \alpha_{i,j} \left(q_{i,j} \left[x^{(i)} - k_{i,j} \right] \right)_+ , \tag{6}$$

where $q_{i,j} \in \{-1, +1\}$ and the knots $k_{i,j}$ are chosen from the data: $k_{i,j} \in \{x^{(i)}(n) : n = 1, \ldots, N\}$, for each $i = 1, \ldots, d$ and each $j = 1, \ldots, J_i$. A function with the form (6) may be called an additive linear spline.

We present a simplified version of the MARS fitting algorithm here, as we are only concerned with the $p = 1$ case. Full details are given in Friedman (1991), and a summary can be found in Hastie et al. (2001). MARS produces a fitted model by proceeding in a stepwise manner. At each step, the algorithm attempts to add each possible pair of basis functions[1]

$$\left\{ \left(x^{(i)} - x^{(i)}(n) \right)_+ , \left(x^{(i)}(n) - x^{(i)} \right)_+ \right\}$$

in turn for $n = 1, \ldots, N$ and $i = 1, \ldots, d$. It adds a basis function if the improvement in fit from adding that function exceeds a given threshold, up to a specified number of basis functions M_{max}. Upon completion of this procedure, the algorithm prunes some of the basis functions it has selected if doing so will improve the weighted mean-square error criterion

$$\frac{\frac{1}{N} \sum_{n=1}^{N} \left(y_n - \hat{f}(x_n) \right)^2}{\left(1 - \frac{CM_1+1}{N}\right)^2},$$

where C is a specified penalty parameter.

Friedman (1991) argues that the computation time of the fitting algorithm has an upper bound proportional to $dNM_{\max}^4$. Our implementation of MARS takes $M_{\max} = 21 \vee (2d+1)$, so for $d < 10$ the computational time is simply proportional to dN; for higher dimensions, it is proportional to $d^5 N$. However, in our experiments, we have found that the threshold criterion is often met before $M_{\max}$ basis functions are even considered, so even though the upper bounds discussed above are valid, they may be quite pessimistic.

3.1 Computing the Approximating Martingale

Suppose we have used MARS to fit

$$\hat{h}_t(x) = \alpha_0 + \sum_{i=1}^{d} \sum_{j=1}^{J_i} \alpha_{i,j} \left(q_{i,j} \left[x^{(i)} - k_{i,j} \right] \right)_+$$

for each time step $t = T, \ldots, 1$. Then for each $t = 1, \ldots, T$, the t'th increment of the resulting martingale $(\hat{\pi}_t)$ is given by

$$\hat{\pi}_t - \hat{\pi}_{t-1} = e^{-rt} \sum_{i=1}^{d} \sum_{j=1}^{J_i} \alpha_{i,j} \left(\left(q_{i,j} \left[X_t^{(i)} - k_{i,j} \right] \right)_+ - E_{t-1} \left[\left(q_{i,j} \left[X_t^{(i)} - k_{i,j} \right] \right)_+ \right] \right), \tag{7}$$

where we have suppressed the dependence of the fitted parameters on the time step t in the notation. Having simulated, say, $X_s(1), \ldots, X_s(N')$, $s = 1, \ldots, T$, it is evident how to compute the first term inside the sum in (7). The second term can be computed explicitly as long as we can compute expressions of the form

$$E_{t-1} \left[\left(X_t^{(i)} - k \right)_+ \right] . \tag{8}$$

[1] In fact, the algorithm sorts the $x_n(i)$'s and skips a small number of observations between each knot it considers. This helps to prevent over-fitting and offers some computational benefits as well.

But this is nothing but the expected value of a vanilla European call option on a *single* underlying random variable. Such conditional expectations can often be computed very easily. Even if the underlying random variables have complex dynamics, such as arises in a stochastic volatility model, we may be able to simplify the problem enough by selecting our discretization scheme carefully so that an answer is within reach. Typically, this will involve replacing the state variable X_t with some transformation of $\log X_t$. See Section 5 for specific examples of how we compute the conditional expectation.

3.2 An Extension of MARS

The function approximation (6) is separable in $\{x^{(i)} : i = 1, \ldots, d\}$. Of course, the function $L_t = L_t(X_t)$ we are trying to approximate will not be separable in general. Indeed, even if the payoff function g is separable, we cannot expect that $L_t(X_t)$ will be separable except for the case $t = T$. For example, consider the case $t = T - 1$. Here,

$$Q_{T-1} = Q_{T-1}(X_{T-1}) = \max\left\{g\left(T-1, X_{T-1}\right), e^{-r}E_{T-1}\left[g(T, X_T)\right]\right\}.$$

So even if g is separable, and even if $E_{T-1}g(T, X_T)$ is separable (which it may not be if there is dependence in the components of X_T), Q_{T-1} will typically not be separable, as the maximum of two separable functions need not be separable.

Intuitively, separability of g is equivalent to the European version of the option being decomposable into options on the individual components of X. Separability of Q_t for $t < T$, on the other hand, would mean that *the decision of whether to exercise early could be made separately for these options*, which is not the case. Since L_t can be made arbitrarily close to Q_t by employing sufficiently many basis functions, it follows that L_t will not be separable either. Thus, the best we can ever hope for with the approximation (6) is to obtain an approximation to the projection of $L_t = L_t(\cdot)$ on the space of separable functions. In particular, $E\left(\hat{L}_t - L_t\right)^2$ may be large no matter how much effort is spent on computing $\hat{L}_t$. This fact suggests that MARS may produce inadequate approximations to the optimal martingale.

In order to at least partially address this issue, we consider a more general form of the approximating multivariate linear spline,

$$y \approx \hat{f}(x) = \alpha_0 + \sum_{j=1}^{J} \alpha_j \left[a_j' x - k_j\right]_+, \tag{9}$$

where we have additional parameters $a_j \in \mathbb{R}^d$, $j = 1, \ldots, d$, to estimate. This is quite similar to the form (6), except that now we consider linear combinations of the x's as predictors. One can think of the a_j vectors as giving a reparameterization of the state variables. If we a priori choose the a_j's, then the problem essentially reduces to the previous one. But this would require user intervention. We prefer an automated procedure, although one can certainly reparameterize manually before invoking our approach.

The following proposition indicates that it is possible to achieve good function approximations with expressions of the form (9).

Proposition 3.1. *Suppose X is an $\mathbb{R}^d$-valued random variable, and $f : \mathbb{R}^d \to \mathbb{R}$ satisfies $Ef^2(X) < \infty$. Then for any $\epsilon > 0$ there is a function $\hat{f}$ of the form (9) such that $E(f(X) - \hat{f}(X))^2 < \epsilon$.*

Proof. Jones (1987) shows that there exists a sequence of vectors $\left(a_m \in \mathbb{R}^d : m = 1, 2, \ldots\right)$ such that

$$E\left(f(X) - \sum_{j=1}^{m} g_j\left(a_j' X\right)\right)^2 \longrightarrow 0$$

as $m \to \infty$. Here, the functions $(g_m : \mathbb{R} \to \mathbb{R} : m = 1, 2, \ldots,)$ are given recursively by

$$g_m(z) = E\left[f(X) - \sum_{j=1}^{m-1} g_j(a_j' X) \;\middle|\; a_m' X = z \right].$$

Accordingly, choose m sufficiently great such that $E\left(f(X) - \sum_{j=1}^{m} g_j(a_j' X) \right)^2 < \epsilon/2$. By induction, $E g_j^2(a_j' X) < \infty$ for $j = 1, \ldots, m$. Since continuous functions are dense in L^2 (Rudin 1987, Theorem 3.14), we conclude from the Stone-Weierstrass Theorem that there exist linear splines $\hat{g}_j : \mathbb{R} \to \mathbb{R}$ such that

$$E\left(g_j(a_j' X) - \hat{g}_j(a_j' X) \right)^2 < \epsilon/2^{j+1},$$

for each $j = 1, \ldots, m$. The result now follows from the triangle inequality and the observation that, in one dimension, any linear spline can be written in the form (9). $\square$

The expression (9) can be thought of as a specific example of a projection pursuit regression fit (Friedman and Stuetzle 1981; Friedman et al. 1983) using truncated linear splines as its univariate basis functions. Projection pursuit regression methods typically estimate the linear directions $a_1, \ldots, a_J$ and the remaining parameters simultaneously. This requires numerical optimization and can be slow. We instead adopt the simpler approach of Zhang et al. (2003), who first identify candidate a_j's and then run the MARS fitting algorithm with $x = (x(1), \ldots, x(d))$ replaced by $(a_1' x, \ldots, a_J' x)$. Zhang et al. (2003) provide two methods for selecting the a_r's. We consider the method that uses linear discriminant analysis, or LDA (Fisher 1936). Given responses $y(1), \ldots, y(N) \in \mathbb{R}$ and predictors $x(1), \ldots, x(N) \in \mathbb{R}^d$, choose some $\tilde{n} \in \{1, \ldots, N\}$ and define the corresponding LDA direction to be

$$S_x^{-1} \left[\frac{1}{|\{n : y(n) < y(\tilde{n})\}|} \sum_{y(n) < y(\tilde{n})} x(n) - \frac{1}{N - |\{n : y(n) < y(\tilde{n})\}|} \sum_{y(n) \geq y(\tilde{n})} x(n) \right],$$

where S_x denotes the sample covariance matrix of the $x(n)$'s. Observe that the bracketed term is nothing but the vector connecting the centroids of the two subpopulations of predictors.

This idea can be extended by performing LDA on the second moments of the predictor variables, leading to directions given by the eigenvectors of

$$S_x^{-1} \left[S_{x(n)|[y(n) < y(\tilde{n})]} - S_{x|[y(n) \geq y(\tilde{n})]} \right] S_x^{-1}, \tag{10}$$

where $S_{x|A}$ is the conditional sample covariance matrix of x given A. Zhang et al. (2003) argue that one typically only needs the eigenvectors of (10) corresponding to the two or three greatest magnitude eigenvalues.

We have found that including linear combinations of the components of X_t when estimating the approximation $\hat{L}_t$ as described in this section provides a dramatic improvement over the "vanilla" MARS fit in terms of reducing variance. This is most notable in the case of basket options (see Section 5.2), where the value functions are highly non-separable.

4 The Algorithm

We now describe how these pieces are put together to compute lower and upper bounds for the Bermudan option price. Like Bolia and Juneja (2005), we use a two phase procedure. In phase one, we compute the suboptimal stopping times $\tau_0, \ldots, \tau_{T-1}$ and the approximate value functions $\hat{L}_1, \ldots, \hat{L}_T$, working backwards from time T. This is done with a small number of simulation trials. In phase two, we run a large number of simulation trials to estimate the lower bound of the option price,

$$E\left[e^{-r\tau_0} g\left(\tau_0, X_{\tau_0}\right) - \hat{\pi}_{\tau_0}\right],$$

and the upper bound,

$$E \max_{t=0,\dots,T} \left[e^{-rt} g(t, X_t) - \hat{\pi}_t\right].$$

Before presenting the algorithm, we point out how the control variate can be used not only to estimate L_0 but also to improve our estimates of the stopping times (τ_t).

4.1 Using the Control Variate to Estimate the Stopping Times

Since for each $t = 1, \dots, T$, the random variable $\hat{\pi}_t$ can be computed without knowing $\tau_0, \dots, \tau_{t-1}$, we can in fact use a modified version of the control variate for computing stopping times in phase one as well as for estimating the bounds on the price in phase two. (A similar idea is explored in Rasmussen 2005.) Specifically, we replace the approximation (2) by

$$\boldsymbol{\alpha}_t = \operatorname*{argmin}_{\boldsymbol{\alpha}} \sum_{n=1}^{N} 1_{[g(t, X_t(n))>0]} \left(\boldsymbol{\alpha}' \boldsymbol{\phi}(X_t(n)) - \left[L_{t+1}(n) - e^{r(t+1)}(\hat{\pi}_{\tau_{t+1}} - \hat{\pi}_t)\right]\right)^2. \tag{11}$$

For the purposes of this regression, our estimate of L_{t+1} comes from samples of $e^{-r(\tau_{t+1}-(t+1))} g\left(\tau_{t+1}, X_{\tau_{t+1}}\right)$, so we can write the predictors in the regression (11) as

$$Z_t \left(e^{-r\tau_{t+1}} g\left(\tau_{t+1}, X_{\tau_{t+1}}\right) - \left(\hat{\pi}_{\tau_{t+1}} - \hat{\pi}_t\right)\right),$$

where $Z_t := 1_{[g(t,X_t)>0]} e^{r(t+1)}$.

We now provide evidence that this actually improves (or at least, does not worsen) our stopping time estimates. For the time being, let us ignore the term Z_t. Evidently,

$$E_t e^{-r\tau_{t+1}} g\left(\tau_{t+1}, X_{\tau_{t+1}}\right) = E_t \left(e^{-r\tau_{t+1}} g\left(\tau_{t+1}, X_{\tau_{t+1}}\right) - \left(\hat{\pi}_{\tau_{t+1}} - \hat{\pi}_t\right)\right), \tag{12}$$

since $\hat{\pi}$ is a martingale. Moreover, by Proposition 2.1,

$$\pi_{\tau_{t+1}} - \pi_t = (\pi_{\tau_{t+1}} - \pi_{t+1}) + (\pi_{t+1} - \pi_t) = e^{-r\tau_{t+1}} g\left(\tau_{t+1}, X_{\tau_{t+1}}\right) - e^{-r(t+1)} E_t L_{t+1},$$

so

$$\mathrm{Var}_t \left[e^{-r\tau_{t+1}} g\left(\tau_{t+1}, X_{\tau_{t+1}}\right) - \left(\pi_{\tau_{t+1}} - \pi_t\right)\right] = \mathrm{Var}_t \left[E_t L_{t+1}\right] = 0,$$

where $\mathrm{Var}_t[\cdot]$ denotes the conditional variance $E_t(\cdot)^2 - E_t^2(\cdot)$. Therefore,

$$\mathrm{Var}_t \left[e^{-r\tau_{t+1}} g\left(\tau_{t+1}, X_{\tau_{t+1}}\right) - \left(\hat{\pi}_{\tau_{t+1}} - \hat{\pi}_t\right)\right] = \mathrm{Var}_t \left[\left(\pi_{\tau_{t+1}} - \pi_t\right) - \left(\hat{\pi}_{\tau_{t+1}} - \hat{\pi}_t\right)\right].$$

Taking expectations and invoking the variance decomposition formula gives

$$\begin{aligned} &E \,\mathrm{Var}_t \left[e^{-r\tau_{t+1}} g\left(\tau_{t+1}, X_{\tau_{t+1}}\right) - \left(\hat{\pi}_{\tau_{t+1}} - \hat{\pi}_t\right)\right] \\ &\qquad = \mathrm{Var} \left[\left(\pi_{\tau_{t+1}} - \pi_t\right) - \left(\hat{\pi}_{\tau_{t+1}} - \hat{\pi}_t\right)\right] - \mathrm{Var}\, E_t \left[\left(\pi_{\tau_{t+1}} - \pi_t\right) - \left(\hat{\pi}_{\tau_{t+1}} - \hat{\pi}_t\right)\right] \\ &\qquad = \mathrm{Var} \left[\left(\pi_{\tau_{t+1}} - \pi_t\right) - \left(\hat{\pi}_{\tau_{t+1}} - \hat{\pi}_t\right)\right]. \end{aligned}$$

On the other hand,

$$\begin{aligned} E \operatorname{Var}_t e^{-r\tau_{t+1}} g\left(\tau_{t+1}, X_{\tau_{t+1}}\right) &= E \operatorname{Var}_t \left[\pi_{\tau_{t+1}} - \pi_t + e^{-r(t+1)} E_t L_{t+1}\right] \\ &= E \operatorname{Var}_t \left[\pi_{\tau_{t+1}} - \pi_t\right] \\ &= \operatorname{Var}\left[\pi_{\tau_{t+1}} - \pi_t\right]. \end{aligned}$$

But $\hat{\pi}$ is a projection of π, and so

$$E \operatorname{Var}_t \left[e^{-r\tau_{t+1}} g\left(\tau_{t+1}, X_{\tau_{t+1}}\right) - \left(\hat{\pi}_{\tau_{t+1}} - \hat{\pi}_t\right)\right] \leq E \operatorname{Var}_t \left[e^{-r\tau_{t+1}} g\left(\tau_{t+1}, X_{\tau_{t+1}}\right)\right]. \tag{13}$$

Equation (12) says that the regressor has the same conditional bias regardless of the presence of the control variate; equation (13) says that on average, the regressor with the control variate has lower conditional variance than the one without. Therefore, the control variate should improve the quality of the stopping times.

When the term Z_t is reintroduced, it is not clear that these properties are maintained. Although the $\mathscr{F}_t$-measurability of Z_t implies

$$E_t Z_t e^{-r\tau_{t+1}} g\left(\tau_{t+1}, X_{\tau_{t+1}}\right) = E_t Z_t \left(e^{-r\tau_{t+1}} g\left(\tau_{t+1}, X_{\tau_{t+1}}\right) - \left(\hat{\pi}_{\tau_{t+1}} - \hat{\pi}_t\right)\right),$$

so the conditional bias is still unchanged, the inequality corresponding to (13) may not hold and so we may not actually reduce variance by including the control variate. The point is that even though $\hat{\pi}$ is a projection of π, we cannot conclude that variance reduction occurs when we are restricted to the subset $[g(t, X_t) > 0]$. Nevertheless, it is reasonable to assume that there is variance reduction except perhaps when the option is deep out of the money so that the event $[g(t, X_t) > 0]$ occurs with very low probability. We have found in our numerical experiments that there is a (modest) improvement in using the control variate to estimate the stopping time; we continue to perform the regression using (11), retaining the indicator function as a heuristic.

4.2 The Algorithm

We establish the notation we use in the description of the algorithm. Let $\boldsymbol{\phi} = (\phi_0 = 1, \phi_1, \ldots, \phi_k)$ be a fixed set of basis functions which will be used in estimating the stopping time, as in (2). The fitted coefficients of this regression will be denoted $\boldsymbol{\alpha}_t \in \mathbb{R}^{k+1}$, for $t = 0, \ldots, T-1$. We will denote by $\boldsymbol{\theta}_t$ the complete set of parameters specifying the fitted extended MARS model (9) for the approximate value function $\hat{h}_t(\cdot) = \hat{h}_t(\cdot; \boldsymbol{\theta}_t)$, for $t = 1, \ldots, T$. The variables $Y(n)$, $n = 1, \ldots, N_1$, keep track of the cash flow along each path; the variables $cv(n)$ are the corresponding values of the control variate described in Section 4.1.

Algorithm 1 is the first phase of the pricing method where the stopping times and the control variate parameters are fit. After simulating the price paths and setting the value of the option at expiry to equal the payoff in lines 1-3, the algorithm proceeds backwards in time. Line 5 is where the MARS fitting algorithm is invoked; note that the "minimization" in this line is not a true minimization, due to the adaptive nature of the MARS fitting procedure. Line 6 updates the control variate for the stopping time as described in Section 4.1. Lines 7-14 are the Longstaff-Schwartz algorithm. Note that the fit in line 7 is trivial when $t = 0$ since we have assumed that X_0 is constant.

Algorithm 2 is the second phase wherein the lower and upper bounds on the option price are computed. New price paths are simulated (line 1), and the realized values are plugged in to the expressions for the martingale control variate (line 2) and the stopping strategy (line 3), both of which were computed during the first phase. Finally, the lower and upper bounds are computed on lines 4-5.

Algorithm 1 Phase One: Fit Stopping Times and Control Variate

simulate $(X_0(n), \ldots, X_T(n) : n = 1, \ldots, N_1)$
$Y(n) \leftarrow g(T, X_T(n))$, for $n = 1, \ldots, N_1$
$cv(n) \leftarrow 0$, for $n = 1, \ldots, N_1$
for $t = T-1, \ldots, 0$ **do**
  $\boldsymbol{\theta}_t \leftarrow \operatorname{argmin}_{\boldsymbol{\theta}} \sum_{n=1}^{N_1} \left(Y(n) - \hat{h}(X_{t+1}(n); \boldsymbol{\theta})\right)^2$
  $cv(n) \leftarrow cv(n) + \hat{h}(X_{t+1}(n); \boldsymbol{\theta}_t) - E\left[\hat{h}(X_{t+1}; \boldsymbol{\theta}_t) \mid X_t(n)\right]$, for $n = 1, \ldots, N_1$
  $\boldsymbol{\alpha}_t \leftarrow \operatorname{argmin}_{\boldsymbol{\alpha}} \sum_{n=1}^{N_1} 1_{[g(t, X_t(n))>0]} e^{-r} \left(\boldsymbol{\alpha}' \boldsymbol{\phi}(X_t(n)) - (Y(n) - cv(n))\right)^2$
  for $n = 1, \ldots, N_1$ **do**
    if $g(t, X_t(n)) > [\boldsymbol{\alpha}_t' \boldsymbol{\phi}(X_t(n))]_+$ **then**
      $Y(n) \leftarrow g(t, X_t(n)); cv(n) \leftarrow 0$
    else
      $Y(n) \leftarrow e^{-r} Y(n); cv(n) \leftarrow e^{-r} cv(n)$
    end if
  end for
end for

Algorithm 2 Phase Two: Compute Option Price Lower and Upper Bounds

simulate *new* paths $(X_0(n), \ldots, X_T(n) : n = 1, \ldots, N_2)$
$\hat{\pi}_t(n) \leftarrow \sum_{s=1}^{t} e^{-rs} \left(\hat{h}(X_s(n); \boldsymbol{\theta}_s) - E\left[\hat{h}(X_s; \boldsymbol{\theta}_s) \mid X_{s-1}(n)\right]\right)$, for $n = 1, \ldots, N_2$, $t = 1, \ldots, T$
$\tau_0(n) \leftarrow T \wedge \min\{t = 0, \ldots, T-1 : g(t, X_t(n)) > \boldsymbol{\alpha}_t' \boldsymbol{\phi}(X_t(n))\}$, for $n = 1, \ldots, N_2$
$L_0 \leftarrow \frac{1}{N_2} \sum_{n=1}^{N_2} e^{-r\tau_0(n)} g(\tau_0(n), X_{\tau_0(n)}(n)) - \hat{\pi}_{\tau_0(n)}(n)$
$U_0 \leftarrow \frac{1}{N_2} \sum_{n=1}^{N_2} \max_{t=0,\ldots,T} \left(g(t, X_t(n)) - \hat{\pi}_t(n)\right)$

5 Numerical Examples

In this section, we describe how we have applied this algorithm to several multidimensional American option pricing problems, and we provide numerical results. In particular, we show how the conditional expectations (8) are computed.

All computations were performed using the R language (R Development Core Team 2005). The MARS fitting algorithm was originally developed in the S language by Hastie and Tibshirani; it was ported to R by Leisch et al. (2005). R is an interpreted language and thus can be fairly slow. Additionally, raw computation times may reflect details of implementation (e.g., R's garbage collection routines) and mask information that would be relevant in evaluating our algorithm. For this reason, we report the *ratio* of the computation time of the naïve estimator with that of our estimator for a fixed degree of accuracy, which we now explain.

In all experiments we fix the run lengths for Phase 1 and Phase 2 to 10,000 and 20,000 respectively, using common random numbers across experiments. We record the following quantities.

r_1 The time required in Phase 1 for both the LSM method and for MARS to fit the $\hat{L}_t$ functions.

r_2 The time required in Phase 2 to compute the MARS-based estimators of the lower and upper bounds.

$\tilde{r}_1$ The time required in Phase 1 for the LSM method alone.

$\tilde{r}_2$ The time required in Phase 2 to compute the naïve estimator of the lower bound.

s^2 An estimate of the variance of the MARS-based estimator of the lower bound.

$\tilde{s}^2$ An estimate of the variance of the naïve estimator of the lower bound.

$\hat{L}_0$ The MARS-based estimate of the lower bound.

We then compute the Phase 2 run lengths ($\tilde{n}$ and n for the naïve and MARS-based estimators respectively) required to achieve a confidence interval half-width for the lower bound that is approximately 0.1% of the lower bound estimate. Hence

$$\tilde{n} = \frac{1.96^2 \tilde{s}^2}{0.001^2 \hat{L}_0^2} \text{ and}$$
$$n = \frac{1.96^2 s^2}{0.001^2 \hat{L}_0^2}.$$

We then compute approximations for the computational time corresponding to these run lengths, viz

$$\tilde{R} = \tilde{r}_1 + \frac{\tilde{n}}{20,000}\tilde{r}_2 \text{ and}$$
$$R = r_1 + \frac{n}{20,000} r_2.$$

Finally, we report

$$\text{TR} = \tilde{R}/R$$

as an estimate of the speed-up factor (or time reduction) of the MARS-based estimator over the naïve estimator. We also report

$$\text{VR} = \tilde{s}^2/s^2$$

as the variance reduction factor. The former measure represents the true improvement in efficiency of the MARS-based estimator over the naïve estimator, while the latter measure indicates the variance reduction without adjustment for computation time.

In all examples, VR and TR are reported to two significant figures.

5.1 Asian Options

We begin by pricing Bermudan-Asian put options, under both the Black-Scholes and Heston (1993) models. In the Black-Scholes case, we have $(X_t : t = 0, \ldots, T) = ((S_t, A_t) : t = 0, \ldots, T)$, where S_0 is given and $S_1, \ldots, S_T$ are generated according to

$$S_t = S_{t-1} \exp\left(r - \frac{1}{2}\sigma^2 + \sigma W_t\right),$$

for independent standard normal variates $W_1, \ldots, W_T$. The average process $(A_t : t = 0, \ldots, T)$ is given by $A_0 = 0$ and, for $t \geq 1$,

$$A_t = \frac{1}{t}\sum_{s=1}^{t} S_s = \frac{1}{t}S_t + \frac{t-1}{t}A_{t-1} = \frac{1}{t}S_{t-1}\exp\left(r - \frac{1}{2}\sigma^2 + \sigma W_t\right) + \frac{t-1}{t}A_{t-1}. \tag{14}$$

The averaging dates are assumed to coincide with the possible exercise dates, excluding the date $t = 0$.

In continuous time, the Heston (1993) model is given by

$$\begin{aligned} dS_t &= \mu S_t dt + \sqrt{V_t} dW_t^{(1)}, \\ dV_t &= \kappa(\theta - V_t)dt + \sqrt{V_t}\sigma\left(\rho dW_t^{(1)} + \sqrt{1-\rho^2} dW_t^{(2)}\right), \end{aligned} \tag{15}$$

where $W^{(1)}$ and $W^{(2)}$ are independent Brownian motions. We approximate (15) in discrete time by applying the first-order Euler discretization to the logarithms of S_t and V_t. See Glasserman (2004, pp. 339–376) for details. This gives $(X_t : t = 0, \ldots, T) = (S_t, V_t, A_t : t = 0, \ldots, T)$ where S_0 and V_0 are given, and

$$\begin{aligned} V_t &= V_{t-1}\exp\left(\frac{\kappa\theta}{V_{t-1}} - \kappa - \frac{1}{2V_{t-1}}\sigma^2 + \frac{\rho\sigma W_t^{(1)} + \sqrt{1-\rho^2}\sigma W_t^{(2)}}{\sqrt{V_{t-1}}}\right) \\ S_t &= S_{t-1}\exp\left(r - \frac{1}{2}V_{t-1} + \sqrt{V_{t-1}} dW_t^{(1)}\right), \end{aligned} \tag{16}$$

where $W_1^{(1)}, \ldots, W_T^{(1)}, W_1^{(2)}, \ldots, W_T^{(2)}$ are independent standard normal variates. (The process $(A_t : t = 0, \ldots, T)$ is still given by (14).) The scheme (16) is not an exact discretization of (15); we ignore the discretization error and henceforth consider (16) to be the *true* dynamics of the underlying.

The payoff function of the Bermudan-Asian put is given by $g(0, \cdot) \equiv 0$ and

$$g(t, X_t) = (K - A_t)_+$$

for $t \geq 1$.

Let us consider the Heston case, as the Black-Scholes case is an easy specialization thereof. As mentioned in Section 3.1, we apply the MARS algorithm not to X_t but to a transformation of X_t which replaces S_t and V_t by their logarithms and A_t by the logarithm of the *geometric* average

$$\tilde{A}_t = \exp\frac{1}{t}\sum_{s=1}^{t}\log S_s.$$

We do not include the LDA directions in the Asian case. This yields an approximation

$$\hat{L}_t = \sum_{j=1}^{J_S} \alpha_{S,j} \Big(q_{S,j} \big[\log S_t - k_{S,j} \big] \Big)_+ + \sum_{j=1}^{J_V} \alpha_{V,j} \Big(q_{V,j} \big[\log V_t - k_{V,j} \big] \Big)_+ + \sum_{j=1}^{J_A} \alpha_{A,j} \Big(q_{A,j} \big[\log \tilde{A}_t - k_{A,j} \big] \Big)_+ .$$

The marginal conditional distributions of $\log S_t, \log V_t$, and $\log \tilde{A}_t$ given $\mathscr{F}_{t-1}$ are Gaussian, with mean and variance given by

$$\begin{aligned} E_{t-1} \log \begin{bmatrix} S_t \\ V_t \\ \tilde{A}_t \end{bmatrix} &= \begin{bmatrix} \log S_{t-1} + r - \frac{1}{2} V_{t-1} \\ \log V_{t-1} + \kappa\theta / V_{t-1} - \kappa - \frac{1}{2}\sigma^2 / V_{t-1} \\ \frac{1}{t}\Big((t-1) \log \tilde{A}_{t-1} + \log S_{t-1} + r - \frac{1}{2} V_{t-1} \Big) \end{bmatrix}, \\ \mathrm{Var}_{t-1} \log \begin{bmatrix} S_t \\ V_t \\ \tilde{A}_t \end{bmatrix} &= \begin{bmatrix} V_{t-1} \\ \sigma^2 / V_{t-1} \\ (1/t^2)\, V_{t-1} \end{bmatrix}. \end{aligned} \tag{17}$$

(The full covariance matrix is irrelevant for our purpose.) This allows us to compute the conditional expectations $E_{t-1}\hat{L}_t$ easily.

Table 1 shows our computational results for Bermudan-Asian options. For all examples, we considered an option maturing in 6 months with monthly exercise/averaging dates; the annualized risk-free rate was $12r = .06$; the initial asset price was $S_0 = 100$. For the Heston examples, the *annualized* model parameters were $\kappa = 1.5, \sigma = .2, \theta = .36, \rho = -.75, V_0 = .4$. The stopping times were fit using the polynomials of degree up to 4 in S_t, A_t, and (for the Heston model) V_t, for $t = 1, \ldots, T-1$.

Table 1: Asian Option Results

Model	K	Naïve L_0	MARS L_0	MARS U_0	VR	TR
BS ($\sigma = .3$)	95	2.77 (.07)	2.73 (.00)	2.78 (.00)	210	85
BS ($\sigma = .3$)	115	15.92 (.14)	15.86 (.01)	15.95 (.01)	230	44
BS ($\sigma = .6$)	95	7.88 (.15)	7.80 (.01)	7.94 (.01)	190	71
BS ($\sigma = .6$)	115	20.57 (.23)	20.48 (.02)	20.65 (.01)	230	56
Heston	95	5.04 (.11)	4.96 (.01)	5.06 (.01)	150	61
Heston	115	17.73 (.11)	17.66 (.01)	17.78 (.01)	200	50

Parenthesized values are 95% confidence interval half -widths. VR=Variance Reduction, TR=Time Reduction, defined at the top of Section 5.

In these examples, the reduction in variance is dramatic, ranging from about 150 times to 250 times variance reduction. Similarly, for an approximate 95% confidence interval with (relative) width .001, one needs to do about 50 times more work with the naïve estimator than with the one using the MARS-based control variate. Finally, observe that the closeness of the (MARS) estimates of L_0 and U_0 suggests that the stopping time found by the LSM algorithm is quite good.

5.2 Basket Options

Next, we consider options on baskets of d assets whose prices are given by $(X_t : t = 0, \ldots, T) = (S_t(i) : t = 0, \ldots, T; i = 1, \ldots, d)$. Specifically, we test call options on the maximum and on the average of the assets, which have respective payoff functions

$$g_{\max}(t,x) = \left(\vee_{i=1}^{d} x(i) - K \right)_+ , \quad g_{\mathrm{avg}}(t,x) = \left(\frac{1}{d} \sum_{i=1}^{d} x(i) - K \right)_+ .$$

The underlying assets are assumed to follow the multidimensional Black-Scholes model, which is discretized as

$$S_t(i) = S_{t-1}(i) \exp\left(r - \delta - \frac{1}{2}\sigma_i^2 + \sigma_i W_t(i)\right), \tag{18}$$

for $i = 1, \ldots, d$, where $W_t = (W_t(1), \ldots, W_t(d))$ is a sequence of independent (in time) multivariate normal random variates with mean zero, unit variance, and a specified correlation matrix (see below). Here, δ is the dividend rate paid by each of the stocks per time step.

We take the annualized risk-free rate $12r$ to be .05, the dividend rate $12\delta = .1$, the annualized volatility to be $\sqrt{12}\sigma = .2$, the expiration to be 3 years, and the strike price to be $K = 100$. The dimension d of the problem takes the values $d = 2, 3, 5, 10$. We test several values of the initial prices $(S_0(i))$, which are taken to be identical for $i = 1, \ldots, d$. For the payoff function g_{avg}, we take the basis functions ϕ for fitting the stopping time τ to be the polynomials of degree up to two in the d asset prices. For the function g_{max}, we take the basis functions to be the polynomials of degree up to two in the *order statistics* of the asset prices, which is similar to the choice of basis functions for such options in Longstaff and Schwartz (2001).

We divide each test further into three cases:

1. The assets' returns are uncorrelated,
2. The correlation between $W_t(i)$ and $W_t(j)$, for $i \neq j$, is a constant ρ, and
3. We randomly generate a correlation matrix for $(W_t(i), i = 1, \ldots, d)$, $t = 1, \ldots, T$, using the method of Marsaglia and Olkin (1984).

We test both the control variate based on MARS and the control variate based on LDA-MARS as in Section 3.2. For the LDA-MARS tests, we partition the sample paths at each time step t into three groups of approximately equal size corresponding to low, medium, and high values of $g(\tau_t, X_{\tau_t})$, and take the first two eigenvalues of the matrix (10), resulting in a total of nine LDA directions. (These are included in addition to, not instead of, the canonical directions.)

We apply the MARS and LDA-MARS fitting algorithms to the logarithm of S_t. The conditional distribution of $\log S_t$ given $\mathscr{F}_{t-1}$ is multivariate Gaussian with mean $\log S_{t-1} + r - \frac{1}{2}\sigma^2$, variance σ^2, and correlation matrix given by $\mathrm{Cor}_{t-1}(\log S_t(i), \log S_t(j)) = \rho$ for $1 \leq i < j \leq d$. Therefore, for a direction $a \in \mathbb{R}^d$, $\|a\| = 1$, the conditional distribution of $a' \log S_t$ given $\mathscr{F}_{t-1}$ is Gaussian with mean and variance

$$E_{t-1} a' \log S_t = \sum_{i=1}^{d} a(i) \left(\log S_{t-1}(i) + r - \frac{1}{2}\sigma^2\right),$$

$$\mathrm{Var}_{t-1}\, a' \log S_t = \sigma^2 a' C a = \sigma^2 \left(\rho \left(\sum_{i=1}^{d} a(i)\right)^2 + (1-\rho) \sum_{i=1}^{d} a^2(i)\right),$$

where C is the $d \times d$ matrix with 1 on the diagonal and ρ off the diagonal. This allows us to compute the conditional expectations (8).

For the call on the average, the variance reduction using LDA-MARS is quite good, resulting in a speed-up factor of between about 5 and 50 for both the uncorrelated case and the randomly correlated case, and between about 25 and 110 for the positively correlated case. There is some degradation of performance as the dimension increases from 2 to 10. We also observe that the variance reduction is much greater for options at-the-money than out-of-the-money.

Table 2: Basket Option Results: Call on Average

d	S_0	Naïve L_0	MARS L_0	LMARS L_0	LMARS U_0	MVR	MTR	LMVR	LMTR
Uncorrelated asset prices.									
2	90	1.98 (.07)	2.00 (.04)	1.99 (.01)	2.08 (.01)	3.2	2.8	51.0	43.0
2	100	4.90 (.10)	4.94 (.05)	4.93 (.01)	5.06 (.01)	4.4	3.9	73.0	59.0
3	90	1.08 (.05)	1.09 (.03)	1.10 (.01)	1.26 (.01)	2.2	1.8	18.0	15.0
3	100	3.61 (.07)	3.62 (.04)	3.63 (.01)	3.85 (.01)	3.2	2.9	32.0	29.0
5	90	0.39 (.02)	0.41 (.02)	0.42 (.01)	0.58 (.01)	1.4	1.1	6.4	4.3
5	100	2.32 (.05)	2.36 (.03)	2.37 (.01)	2.59 (.01)	2.4	1.9	19.0	15.0
10	90	0.05 (.01)	0.05 (.01)	0.05 (.00)	0.15 (.00)	1.0	0.5	2.3	0.2
10	100	1.18 (.03)	1.21 (.02)	1.25 (.01)	1.42 (.01)	1.9	1.3	13.0	8.3
Correlated asset prices ($\rho = .45$ for all asset pairs).									
2	90	3.08 (.09)	3.11 (.03)	3.09 (.01)	3.15 (.01)	7.3	6.8	83.0	78.0
2	100	6.39 (.13)	6.43 (.04)	6.38 (.01)	6.48 (.01)	11.0	9.7	120.0	110.0
3	90	2.61 (.08)	2.64 (.03)	2.62 (.01)	2.70 (.01)	6.0	5.1	67.0	56.0
3	100	5.77 (.12)	5.82 (.04)	5.80 (.01)	5.91 (.01)	8.6	7.4	95.0	83.0
5	90	2.15 (.07)	2.25 (.03)	2.24 (.01)	2.33 (.01)	4.6	3.4	60.0	48.0
5	100	5.27 (.10)	5.32 (.04)	5.30 (.01)	5.41 (.01)	6.9	5.8	91.0	75.0
10	90	1.77 (.06)	1.92 (.03)	1.96 (.01)	2.08 (.01)	4.0	2.8	35.0	23.0
10	100	4.71 (.10)	4.87 (.04)	4.90 (.01)	5.02 (.01)	7.0	5.7	80.0	64.0
Correlated asset prices (random correlation matrix).									
5	90	0.07 (.01)	0.08 (.01)	0.08 (.01)	0.15 (.00)	2.5	0.9	3.7	0.5
5	100	0.85 (.02)	0.88 (.01)	0.89 (.01)	1.03 (.01)	3.2	2.5	11.0	8.1
10	90	0.09 (.01)	0.09 (.01)	0.10 (.01)	0.21 (.01)	1.2	0.6	2.9	0.3
10	100	1.38 (.03)	1.40 (.02)	1.42 (.01)	1.62 (.01)	2.1	2.0	11.0	9.3

Parenthesized values are 95% confidence interval half-widths. MVR/LMVR = MARS/LMARS variance reduction. MTR/LMTR = MARS/LMARS time reduction, defined at the top of Section 5.

It is natural to expect that LDA-MARS should perform significantly better than MARS for an option on the average of stocks, as there is one linear direction (namely, $a = (1, \ldots, 1)$) that is likely to capture much of the variation in the value function. It is also plausible that the effect of the control variate is stronger when the assets are positively correlated, and that the degradation with dimension is smaller in that case as well, since under this correlation structure much of the variance of the assets' returns is driven by a single factor. Both of these observations are borne out in the results.

Table 3: Basket Option Results: Call on Max

d	S_0	Naïve L_0	MARS L_0	LMARS L_0	LMARS U_0	MVR	MTR	LMVR	LMTR
Uncorrelated asset prices.									
2	90	7.92 (.16)	8.05 (.03)	8.08 (.03)	8.40 (.02)	23	18	37	30
2	100	13.77 (.20)	13.88 (.05)	13.90 (.03)	14.46 (.03)	18	17	40	37
2	110	21.27 (.23)	21.33 (.06)	21.33 (.04)	22.09 (.04)	14	11	39	33
3	90	11.15 (.18)	11.20 (.05)	11.17(.04)	11.92 (.04)	14	10	17	13
3	100	18.58 (.23)	18.57 (.07)	18.60 (.06)	20.01 (.05)	11	8.9	16	12
3	110	27.42 (.26)	27.42 (.09)	27.43 (.07)	29.22 (.06)	8.7	6.9	16	12
5	90	16.27 (.21)	16.46 (.08)	16.46 (.08)	18.18 (.07)	6.4	4.8	7.5	5.1
5	100	25.83 (.25)	25.98 (.11)	25.97 (.10)	28.79 (.08)	5.5	4.3	6.8	5.0
5	110	36.46 (.28)	36.58 (.13)	36.53 (.11)	40.20 (.10)	4.8	3.7	6.5	4.7
10	90	25.70 (.24)	25.85 (.12)	25.87 (.12)	29.20 (.11)	3.9	2.9	4.3	3.1
10	100	37.73 (.27)	39.74 (.14)	37.94 (.13)	42.14 (.12)	3.6	2.7	4.2	3.1
10	110	50.19 (.30)	50.47 (.16)	50.45 (.15)	55.25 (.13)	3.6	2.7	4.1	3.0
Correlated asset prices ($\rho = .45$ for all asset pairs).									
2	90	7.11 (.16)	7.18 (.04)	7.19 (.03)	7.68 (.03)	16	13	28	22
2	100	12.20 (.20)	12.28 (.05)	12.35 (.04)	13.08 (.03)	15	14	31	28
2	110	19.02 (.24)	19.00 (.06)	19.03 (.04)	20.00 (.04)	15	12	35	25
3	90	9.26 (.19)	9.27 (.06)	9.30 (.05)	10.47 (.04)	11	8.1	15	11
3	100	15.33 (.23)	15.35 (.07)	15.38 (.06)	17.04 (.05)	11	8.4	17	13
3	110	22.97 (.27)	22.95 (.08)	22.97 (.07)	24.98 (.06)	11	8.3	17	13
5	90	12.33 (.21)	12.46 (.08)	12.49(.08)	14.53 (.07)	6.9	5.3	8.1	5.8
5	100	19.65 (.26)	19.80 (.10)	19.84 (.09)	22.94 (.09)	6.7	5.1	8.0	5.8
5	110	28.36 (.30)	28.52 (.12)	28.55 (.10)	32.42 (.10)	6.8	5.1	9.0	6.4
10	90	17.46 (.26)	17.64 (.11)	17.64(.11)	20.86(.12)	5.5	3.8	6.0	3.7
10	100	26.43 (.31)	26.58 (.13)	26.67 (.12)	31.20 (.13)	5.4	3.6	6.2	4.8
10	110	36.50 (.35)	36.74 (.15)	36.69 (.14)	42.21 (.15)	5.5	4.1	6.5	4.7
Correlated asset prices (random correlation matrix).									
5	90	15.78 (.20)	15.89 (.06)	15.88 (.06)	17.16 (.06)	9.7	7.4	10	7.6
5	100	25.45 (.23)	25.61 (.08)	25.63 (.07)	27.70 (.07)	8.8	6.8	9.4	6.8
5	110	36.35 (.25)	36.52 (.09)	36.53 (.09)	39.25 (.07)	7.7	5.9	8.5	6.0
10	90	23.55 (.24)	23.79 (.12)	23.79 (.11)	27.02 (.11)	4.3	2.9	4.9	3.1
10	100	35.10 (.28)	35.35 (.14)	35.41 (.13)	39.69 (.12)	4.0	3.4	4.7	3.6
10	110	47.23 (.30)	47.56 (.15)	47.53 (.14)	52.56 (.13)	3.9	3.0	4.5	3.2

Parenthesized values are 95% confidence interval half -widths. MVR/LMVR = MARS/LMARS variance reduction. MTR/LMTR = MARS/LMARS time reduction, defined at the top of Section 5.

The results are somewhat less dramatic for the case of the option on the maximum. This is most likely due to the fact that the payoff function g_{max} is highly non-separable, so the fitted functions $\hat{L}$ are poor approximations for the true value functions L. In fact, not only is $g_{\max}$ non-separable, but it cannot even be represented exactly in the form (9). Thus, even when LDA directions are used, and even in the correlated assets case, the performance degrades quickly to a variance reduction factor of only around 2 or 3 as the dimension increases. Still, the method seems to be able to provide about a threefold decrease in computation time even in this case. We also observe that the effects of correlation are much less noticable for the call on the max than for the call on the average.

The first nine rows of the first panel of Table 3 may be compared with Table 2 of Andersen and Broadie (2004). Our results (using LDA-MARS) include confidence intervals that are approximately twice the width of the ones reported in Andersen and Broadie (2004), although we use 20,000 simulation trials to their 2,000,000 trials. In order to get confidence intervals of the same order, we would need to use approximately 80,000 trials – still quite a bit fewer than 2,000,000. On the other hand, the "duality gaps" between the upper and lower bounds are much tighter in Andersen and Broadie (2004) than in our study. This stands to reason; our upper bounds are wholly reliant on the approximation $\hat{\pi}$ for π; in contrast, they compute π explicitly by running additional simulation trials.

5.3 Barrier Options

Finally we test our method on a variety of barrier options: the up-and-out call, the up-and-out put, and the down-and-out put, all on a single asset. Unlike a vanilla Bermudan call, a Bermudan up-and-out call on an asset that does not pay dividends may have an optimal exercise policy other than the trivial one $\tau_0 = T$. Again, we test both the Black-Scholes and the Heston models.

Although it would be possible to accommodate the path dependence of barrier options by expanding the state space, we adopt a different approach. Let $\mathscr{B} \subset \mathbb{R}^d$ denote the region in which the option is knocked out. Assume the payoff function g satisfies $g(\cdot, x) \equiv 0$ for all $x \in \mathscr{B}$. For each $t = 0, \ldots, T$, let ν_t be the first hitting time of $\mathscr{B}$ between t and T, or $T + 1$ if there is no such hitting time, i.e.,

$$\nu_t = \inf\{s = t, \ldots, T + 1 : (s, X_s) \in \{t, \ldots, T\} \times \mathscr{B} \cup \{T + 1\} \times \mathbb{R}^d\}.$$

We now redefine our value function to be

$$Q_t = \sup_{\tau \in \mathcal{T}(t)} E_t \left[e^{-r(\tau \wedge \nu_t - t)} g\left(\tau \wedge \nu_t, X_{\tau \wedge \nu_t}\right) \right]. \tag{19}$$

The stopping times τ_t^* solving (19) satisfy

$$\tau_T^* \wedge \nu_T = T,$$

$$\tau_t^* \wedge \nu_t = \begin{cases} t & \text{if } \nu_t = t \text{ or if } g(t, X_t) \geq e^{-r} E_t Q_{t+1}, \\ \tau_{t+1}^* & \text{otherwise,} \end{cases}$$

for $t = 0, \ldots, T - 1$. The suboptimal stopping times $(\tau_t : t = 0, \ldots, T)$ are defined analogously to those in Section 2.1. In this setting the analogous martingale π satisfies

$$\pi_{\tau_t \wedge \nu_t} - \pi_t = e^{-r\tau_t \wedge \nu_t} g\left(\tau_t \wedge \nu_t, X_{\tau_t \wedge \nu_t}\right) - e^{-rt} L_t,$$

similar to Proposition 2.1, and we have

$$Q_0 \leq E \max_{t=0,\ldots,T} \left[g\left(t \wedge \nu_0, X_{t \wedge \nu_0}\right) - \pi_{t \wedge \nu_0} \right].$$

In other words, the martingale π evaluated only as far as the hitting time of the knock-out region, both for computing the control variate and the upper bound. This leads to Algorithms 3 and 4, which are modifications of Algorithms 1 and 2, respectively. The only difference between Algorithms 1 and 3 occurs on

Algorithm 3 Phase One (Barrier Option): Fit Stopping Times and Control Variate

simulate $(X_0(n), \ldots, X_T(n) : n = 1, \ldots, N_1)$
$Y(n) \leftarrow g(T, X_T(n))$, for $n = 1, \ldots, N_1$
$cv(n) \leftarrow 0$
for $t = T-1, \ldots, 0$ **do**
  $\boldsymbol{\theta}_t \leftarrow \operatorname{argmin}_{\boldsymbol{\theta}} \sum_{n=1}^{N_1} \left(Y(n) - \hat{h}(X_{t+1}(n); \boldsymbol{\theta}) \right)^2$
  $cv(n) \leftarrow cv(n) + \hat{h}(X_{t+1}(n); \boldsymbol{\theta}_t) - E\left[\hat{h}(X_{t+1}; \boldsymbol{\theta}_t) \mid X_t(n) \right]$, for $n = 1, \ldots, N_1$
  $\boldsymbol{\alpha}_t \leftarrow \operatorname{argmin}_{\boldsymbol{\alpha}} \sum_{n=1}^{N_1} 1_{[g(t, X_t(n)) > 0]} e^{-r} \left(\boldsymbol{\alpha}' \boldsymbol{\phi}(X_t(n)) - (Y(n) - cv(n)) \right)^2$
  for $n = 1, \ldots, N_1$ **do**
    if $g(t, X_t(n)) > [\boldsymbol{\alpha}_t' \boldsymbol{\phi}(X_t(n))]_+$ **or** $X_t(n) \in \mathscr{B}$ **then**
      $Y(n) \leftarrow g(t, X_t(n)); cv(n) \leftarrow 0$
    else
      $Y(n) \leftarrow e^{-r} Y(n); cv(n) \leftarrow e^{-r} cv(n)$
    end if
  end for
end for

Algorithm 4 Phase Two (Barrier Option): Compute Option Price Lower and Upper Bounds

simulate *new* paths $({}_0(n), \ldots, X_T(n) : n = 1, \ldots, N_2)$
$\hat{\pi}_t(n) \leftarrow \sum_{s=1}^{t} e^{-rs} \left(\hat{h}(X_s(n); \boldsymbol{\theta}_s) - E\left[\hat{h}(X_s; \boldsymbol{\theta}_s) \mid X_{s-1}(n) \right] \right)$, for $n = 1, \ldots, N_2$, $t = 1, \ldots, T$
$\tau_0(n) \leftarrow T \wedge \min\{t = 0, \ldots, T-1 : g(t, X_t(n)) > \boldsymbol{\alpha}_t' \boldsymbol{\phi}(X_t(n))\}$, for $n = 1, \ldots, N_2$
$\nu_0(n) \leftarrow (T+1) \wedge \min\{t = 0, \ldots, T : X_t(n) \in \mathscr{B}\}$
$L_0 \leftarrow \frac{1}{N_2} \sum_{n=1}^{N_2} e^{-r(\tau_0(n) \wedge \nu_0(n))} g((\tau_0(n) \wedge \nu_0(n)), X_{(\tau_0(n) \wedge \nu_0(n))}(n)) - \hat{\pi}_{(\tau_0(n) \wedge \nu_0(n))}(n)$
$U_0 \leftarrow \frac{1}{N_2} \sum_{n=1}^{N_2} \max_{t=0,\ldots,\nu_0(n)} \left(g(t, X_t(n)) - \hat{\pi}_t(n) \right)$

line 9, which in the barrier option case says to zero out the cash flow and the control variate upon exercise *or knockout*. Algorithm 4 differs from Algorithm 2 in that the exercise time τ_0 is replaced with the minimum of the exercise time and knock out time $\tau_0 \wedge \nu_0$.

In our numerical experiments on Bermudan barrier options, the underlying dynamics of $(X_t : t = 0, \ldots, T) = (S_t : t = 0, \ldots, T)$ or $(X_t : t = 0, \ldots, T) = ((S_t, V_t) : t = 0, \ldots, T)$ follow (18) or (16) accordingly, and the appropriate parameters of the one-step conditional distributions are still given by (17) (without the A_t term). The model parameters are as described in Section 5.1. Tables 4 and 5 report the computational results for barrier options in the Black-Scholes and Heston models, respectively.

The performance of the control variate for the barrier option examples seems to have huge variability, with variance reduction factors ranging from 8.5 to 350 just within the Black-Scholes cases. Why is there such a discrepancy in the quality of the algorithm between these two examples? The cases for which the control variate is very successful are the "up-and-out" puts, which knock out when the option is deep out-of-the-money. The less successful cases are the "up-and-out" call and the "down-and-out" put, which knock out in-the-money and as such have discontinuous payoff functions. Indeed, this discontinuity is notorious for causing headaches among traders, especially in currency markets, who must hedge these options. Our problem here is that MARS has difficulty fitting these functions as well as it fits the smoother "up-and-out" put value functions. We believe that with a little manual tweaking we could get the performance for the discontinuous cases to improve significantly. However, in the spirit of having a fully automated procedure we have not pursued this line of inquiry. An interesting future research project would be to develop a version of MARS that is more robust to discontinuities in the target function.

Table 4: Barrier Option Results (Black-Scholes, $\sigma = .3$)

Stgy.	K	B	Naïve L_0	MARS L_0	MARS U_0	VR	TR
Call	95	130	13.38 (.16)	13.36 (.04)	14.11 (.04)	18.0	15.0
Call	115	130	3.19 (.06)	3.19 (.02)	3.44 (.02)	8.5	7.3
Put	95	70	6.61 (.11)	6.7 (.03)	7.25 (.03)	14.0	13.0
Put	115	70	18.34 (.19)	18.45 (.02)	19.30 (.04)	79.0	53.0
Put	95	130	6.98(.14)	7.09 (.01)	7.13 (.01)	180.0	140.0
Put	115	130	17.72 (.20)	17.86 (.01)	17.92 (.01)	350.0	140.0

Parenthesized values are 95% confidence interval half -widths. VR=Variance Reduction, TR=Time Reduction, defined at the top of Section 5.

Table 5: Barrier Option Results (Heston)

Stgy.	K	B	Naïve L_0	MARS L_0	MARS U_0	VR	TR
Call	95	130	14.06 (.15)	13.98 (.05)	14.82 (.05)	9.9	8.5
Call	115	130	3.36 (.06)	3.33 (.03)	3.66 (.02)	5.0	4.3
Put	95	70	8.95 (.12)	8.99 (.05)	9.82 (.04)	5.2	4.5
Put	115	70	22.12 (.20)	22.20 (.06)	23.72 (.06)	9.6	7.9
Put	95	130	12.93(.24)	13.10(.02)	13.16 (.02)	170.0	120.0
Put	115	130	22.04 (.33)	22.34 (.03)	22.51 (.02)	170.0	110.0

Parenthesized values are 95% confidence interval half -widths. VR=Variance Reduction, TR=Time Reduction, defined at the top of Section 5.

6 Conclusion

We have presented a new, automated procedure for finding control variates for the American option pricing problem. The key advantages of our method are its degree of applicability to many option types and stochastic processes, without requiring much additional implementation overhead, and its use of off-the-shelf software. Our method works extremely well for problems of moderate dimension (up to about 5), and for problems where much of the variability of the underlying processes can be explained with a moderate number of parameters. Moreover, the method can "discover" such structure automatically as a result of using an adaptive fitting procedure.

A possible area of future research would be to apply this technique in conjunction with quasi-Monte Carlo methodology. This would likely result in even greater variance reductions, although that remains to be seen. The good news is that the overall procedure would not change in any substantive way. Finally, this paper suggests that there is promise in applying techniques from the (vast) statistical data mining literature to the American option pricing problem. This is a direction we hope to continue to pursue.

Acknowledgments

Samuel Ehrlichman has been supported by an NDSEG Fellowship from the U.S. Department of Defense and the ASEE. This work was supported in part by NSF Grant DMI-0400287.

References

Andersen, L. and M. Broadie (2004). Primal-dual simulation algorithm for pricing multidimensional American options. *Management Sci. 50*(9), 1222–1234.

Bolia, N. and S. Juneja (2005). Function-approximation-based perfect control variates for pricing American options. In M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Jones (Eds.), *Proceedings of the Winter Simulation Conference*, Piscataway, New Jersey. Institute of Electrical and Electronics Engineers, Inc.

Broadie, M. and P. Glasserman (2004). A stochastic mesh method for pricing high-dimensional American options. *Journal of Computational Finance 7*(4), 35–72.

Carriere, J. F. (1996). Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics 19*(1), 19–30.

Chen, V. C. P., D. Ruppert, and C. A. Shoemaker (1999). Applying experimental design and regression splines to high-dimensional continuous-state stochastic dynamic programming. *Oper. Res. 47*(1), 38–53.

Clément, E., D. Lamberton, and P. Protter (2002). An analysis of a least squares regression method for American option pricing. *Finance Stoch. 6*(4), 449–471.

Cox, J. C., S. A. Ross, and M. Rubinstein (1979). Option pricing: A simplified approach. *Journal of Financial Economics 7*, 229–263.

Duffie, D. (2001). *Dynamic Asset Pricing Theory* (Third ed.). Princeton, New Jersey: Princeton University Press.

Fisher, R. (1936). The use of multiple measurements in taxonomic problems. *Annals of Eugenics 7*(2), 179–188.

Friedman, J. H. (1991). Multivariate adaptive regression splines. *Ann. Statist. 19*(1), 1–141. With discussion and a rejoinder by the author.

Friedman, J. H., E. Grosse, and W. Stuetzle (1983). Multidimensional additive spline approximation. *SIAM J. Sci. Statist. Comput. 4*(2), 291–301.

Friedman, J. H. and W. Stuetzle (1981). Projection pursuit regression. *J. Amer. Statist. Assoc. 76*(376), 817–823.

Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*, Volume 53 of *Applications of Mathematics*. New York: Springer-Verlag.

Glasserman, P. and B. Yu (2004). Simulation for American options: regression now or regression later? In *Monte Carlo and quasi-Monte Carlo methods 2002*, pp. 213–226. Berlin: Springer.

Hastie, T., R. Tibshirani, and J. H. Friedman (2001). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. New York: Springer Verlag.

Haugh, M. B. and L. Kogan (2004). Pricing American options: a duality approach. *Oper. Res. 52*(2), 258–270.

Henderson, S. G. and P. W. Glynn (2002). Approximating martingales for variance reduction in Markov process simulation. *Math. Oper. Res. 27*(2), 253–271.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies 6*(2), 327–43. available at http://ideas.repec.org/a/oup/rfinst/v6y1993i2p327-43.html.

Jones, L. K. (1987). On a conjecture of Huber concerning the convergence of projection pursuit regression. *Ann. Statist. 15*(2), 880–882.

Laprise, S. B., M. C. Fu, S. I. Marcus, A. E. B. Lim, and H. Zhang (2006, January). Pricing American-style derivatives with European call options. *Management Scuence 52*(1), 95–110.

Leisch, F., K. Hornik, and B. D. Ripley (2005). *mda: Mixture and flexible discriminant analysis*. R Foundation for Statistical Computing. R package version 0.3-1.

Longstaff, F. A. and E. S. Schwartz (2001). Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies 14*(1), 113–47.

Marsaglia, G. and I. Olkin (1984). Generating Correlation Matrices. *SIAM Journal on Scientific and Statistical Computing 5*, 470.

R Development Core Team (2005). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. ISBN 3-900051-07-0.

Rasmussen, N. S. (2005). Control variates for Monte Carlo valuation of American options. *J. Comp. Finance 9*(1), 84–102.

Rogers, L. C. G. (2002). Monte Carlo valuation of American options. *Math. Finance 12*(3), 271–286.

Rudin, W. (1987). *Real and Complex Analysis* (Third ed.). McGraw-Hill series in higher mathematics. Boston, MA: McGraw-Hill.

Tsitsiklis, J. N. and B. Van Roy (2001, July). Regression methods for pricing complex American-style options. *IEEE-NN 12*, 694–703.

Zhang, H., C.-Y. Yu, H. Zhu, and J. Shi (2003). Identification of linear directions in multivariate adaptive spline models. *J. Amer. Statist. Assoc. 98*(462), 369–376.