Proceedings of the 2006 Winter Simulation Conference
L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, eds.

AMERICAN OPTIONS FROM MARS

Samuel M. T. Ehrlichman
Shane G. Henderson

School of Operations Research and Industrial Engineering
Cornell University
Ithaca, NY 14850, U.S.A.

Abstract

We develop a class of control variates for the American option pricing problem that are constructed through the use of MARS – multivariate adaptive regression splines. The splines approximate the option's value function at each time step, and the value function approximations are then used to construct a martingale that serves as the control variate. Significant variance reduction is possible even in high dimensions. The primary restriction is that we must be able to compute certain one-step conditional expectations.

1 INTRODUCTION

Simulation is now a widely used tool for pricing financial derivatives. This is true even for American options, where there is an embedded optimization problem. While simulation optimization is widely viewed as being a difficult problem, the American-option pricing problem possesses a great deal of structure that makes it more tractable than one might otherwise expect. The result is that simulation is, indeed, a viable approach to such problems, especially if one exploits the known theory of such problems to develop variance reduction techniques.

In this paper we further develop a control variate for pricing American options. The key idea, which is more precisely developed later, is as follows. We first model the price process of the underlying instruments as a Markov process. We then define a class of mean-zero martingales that are functions of the Markov process. Since the martingales have zero mean, they can be used as control variates in simulations. The trick then is to choose a good martingale from within the class of such martingales. Our martingales come from MARS (multivariate adaptive regression splines). We sketch the key ideas behind MARS later and refer the interested reader to Friedman (1991) for full details. With some care in implementation, we get very large efficiency improvements, even in some high dimensional problems.

The key idea is the use of the martingales as control variates. This idea was first developed in Henderson and Glynn (2002) for a range of estimation problems related to Markov processes. This methodology was adapted to the American option pricing problem by Bolia and Juneja (2005). The primary contribution of this paper over Bolia and Juneja (2005) is the use of MARS to select an appropriate martingale. The MARS methodology can be viewed as an automated method for selecting a linear combination of European options (not necessarily on traded assets) that mimic the payoff of the American option. In that sense, this paper is also an outgrowth of Rasmussen (2005), where a small collection of European options, evaluated at the time of exercise of the American option, are used to define a control variate in the same way that we do here. Our paper can also be viewed as an outgrowth of Laprise et al. (2006) who develop upper and lower bounds on the price of an American option (on a single asset) through a procedure that essentially involves holding a portfolio of European options.

A similar class of martingales to the one used here was also exploited in Andersen and Broadie (2004). In that paper, the martingales were used not for variance reduction but to compute an upper bound on the option price. The technique employed "simulation within simulation" in order to estimate certain conditional expectations in the martingale. Our use of MARS allows us to compute those conditional expectations without the need to resort to simulation, for the same reasons that Bolia and Juneja (2005), Rasmussen (2005) and Laprise et al. (2006) did not need to resort to simulation to compute the conditional expectations. Indeed, it is this fact that allows our martingale to be used for variance reduction as well as for computing an upper bound on the price.

As is fairly standard in American option pricing, we actually price a *Bermudan* option, that is, an option that can only be exercised at a finite number of dates. The prices we estimate can be viewed as lower bounds on the true option price, because we price the option using a specific

1-4244-0501-7/06/$20.00 ©2006 IEEE

exercise strategy that is feasible, but not necessarily optimal. Upper bounds on the true option price can be recovered using a martingale duality result established independently in Rogers (2002) and Haugh and Kogan (2004), and used in Andersen and Broadie (2004) and Bolia and Juneja (2005); we do so here as well.

To summarize, the primary contribution of this paper over and above previous work is to identify a particularly effective approach for selecting an effective martingale control variate from within an already-known class of such control variates. In work reported elsewhere, we also exploit some standard statistical tools to perform automatic reparameterizations of the problem that yield even greater variance reductions; see Ehrlichman and Henderson (2006).

The remainder of this paper is organized as follows. Section 2 briefly reviews the American option pricing problem, the method we use to determine the exercise policy, and introduces the class of martingale control variates within which we work. It also identifies the optimal martingale. Section 3 reviews the key ideas from MARS methodology, and explains how we use MARS to approximate the optimal martingale. The algorithm is then outlined in Section 4, some numerical examples are given in Section 5, and Section 6 offers some conclusions.

For further details and justification of the procedure, the reparameterizing extension mentioned above, and more extensive numerical results, see Ehrlichman and Henderson (2006).

2 MATHEMATICAL FRAMEWORK

We adopt much of the notation of Andersen and Broadie (2004) in what follows.

Let $(X_t : t = 0, 1, \ldots, T)$ be a discrete-time Markov process on $\mathbb{R}^d$, for some finite d, and some fixed integer T. This process represents the price dynamics of the underlying, and potentially additional information that is used to compute the payoff of the option when exercised. For notational ease we assume that the option can be exercised at any time $t = 0, 1, 2, \ldots, T$. We assume an arbitrage-free market and therefore work with a risk-neutral pricing measure. See, e.g., Duffie (2001) or Glasserman (2004) for background.

Let $g(t, X_t) \geq 0$ represent the payoff as a function of time and state if the option is exercised at time t, $t = 0, \ldots, T$. We assume that $Eg^2(t, X_t) < \infty$. Let r be the riskless interest rate which, for simplicity we assume is constant and deterministic, so that \$1 at time 0 is worth e^{rt} dollars at time t.

Let $\mathbb{F} = (\mathscr{F}_t : t = 0, \ldots, T)$ be the natural filtration of (X_t), and let $\mathcal{T}(t)$ be the set of all stopping times (with respect to $\mathbb{F}$) taking values in $\{t, t+1, \ldots, T\}$, $t = 0, \ldots, T$. Each stopping time in $\mathcal{T}(t)$ represents an exercise strategy for an option newly released at time t. Let $E_t[\cdot] = E[\cdot \,|\, \mathscr{F}_t]$.

The Bermudan option pricing problem is to compute the option price Q_0, where

$$Q_t = \sup_{\tau \in \mathcal{T}(t)} E_t \left[e^{-r(\tau - t)} g(\tau, X_\tau) \right], \tag{1}$$

for $t = 0, \ldots, T$. The standard theory of option pricing, e.g., Duffie (2001), ensures that there exists an optimal family of stopping times $(\tau_t^* : t = 0, \ldots, T)$ such that τ_t^* attains the supremum in (1) for each t. This optimal family satisfies the recursion

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

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

Similarly, the family of option values $(Q_t : t = 0, \ldots, T)$ satisfies the recursion

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

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

2.1 Lower and Upper Bounds

We adopt the least-squares Monte Carlo (LSM) method (Longstaff and Schwartz 2001) for developing an approximation to the optimal stopping time family (τ_t^*) and thereby the option price family (Q_t). The stopping times (τ_t) we obtain are *sub-optimal*, so that the option prices (L_t) implied by the algorithm are lower bounds on the true option prices (Q_t). Here $L_t = E_t e^{-r(\tau_t - t)} g(\tau_t, X_{\tau_t})$.

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 $\mathbb{F}$. By the optional sampling theorem, for any $t \geq 0$,

$$\begin{aligned} Q_t &= e^{rt} \sup_{\tau \in \mathcal{T}(t)} E_t \left[e^{-r\tau} g(\tau, X_\tau) - \pi_\tau + \pi_\tau \right] \\ &= e^{rt} \sup_{\tau \in \mathcal{T}(t)} E_t \left[e^{-r\tau} g(\tau, X_\tau) - \pi_\tau \right] + e^{rt} \pi_t \\ &\leq e^{rt} E_t \max_{s=t,\ldots,T} \left[e^{-rs} g(s, X_s) - \pi_s \right] + e^{rt} \pi_t \\ &=: U_t. \end{aligned} \tag{2}$$

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.

2.2 Martingales

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

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

Then (π_t) is evidently a martingale, and can be used to obtain an upper bound on the option price as in (2).

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 if we can compute the stopping time (τ_t) (as is possible using the LSM method), then conditional on $\mathscr{F}_t$, we can compute L_t by averaging conditionally independent replicates of $e^{-r(\tau_t - t)} g(\tau_t, X_{\tau_t})$. But in a slight extension of an observation in Bolia and Juneja (2005), Ehrlichman and Henderson (2006) observe that if $h_t(X_t) = L_t$, then

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

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

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

If $\mathscr{F}_0$ is the trivial sigma field, so that X_0 is deterministic, then taking $t = 0$ we get a zero variance estimator of L_0, the lower bound on the option price at time 0. Moreover, Andersen and Broadie (2004) show that with this choice of martingale, the inequality in (2) is tight *provided* the suboptimal stopping times (τ_t) are equal to the true optimal stopping times (τ_t^*). They also show that in the more general setting, this "duality gap" is bounded above by $E\sum_{s=1}^{T}(Q_s - L_s)$.

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

$$h_t(X_t) \approx L_t$$

for each t, in some sense. (We do so in the mean-square sense.) We then use the induced martingale (evaluated at time τ_t) as a control variate in estimating L_0. To compute this martingale, we need to be able to compute the conditional expectation $E_{s-1}h_s(X_s)$ efficiently. Andersen and Broadie (2004) do so using simulation. We instead 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). However, our class of functions, which come from MARS, differs from that in Bolia and Juneja (2005), and is closer to that used in Rasmussen (2005). We now explore MARS.

3 MARS

Multivariate adaptive regression splines (Friedman 1991), or MARS, is essentially a nonparametric regression technique. MARS has been used in many contexts. We use MARS to approximate the value function L_t. Interestingly, MARS has been used to approximate value functions of stochastic dynamic programs before (Chen et al. 1999), although in a very different setting.

We use a restricted version of MARS that can be described as follows. Given predictors $x_1, \dots, x_N \in \mathbb{R}^d$ and responses $y_1, \dots, y_N$, MARS fits a model of the form

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

Here $x(i)$ denotes the ith component of x, the J_is are integers, α_0 and the $\alpha_{i,j}$s are constants, and $q_{i,j} \in \{-1, 1\}$. The "knots" $k_{i,j}$ are selected from the ith coordinates of the predictors, i.e., $k_{i,j} \in \{x_n(i) : 1 \le n \le N\}$ for each j.

We now give the essential idea behind how MARS fits this model. For full details see Friedman (1991), or for a summary, see Hastie et al. (2001). MARS proceeds in a stepwise manner. For each $n = 1, \dots, N$ and $i = 1, \dots, d$, MARS considers adding the 2 basis functions

$$(x(i) - x_n(i))_+, \quad (-[x(i) - x_n(i)])_+. \tag{3}$$

A basis function is added if the improvement in fit exceeds a given threshold, up to a maximum number of basis functions. After this step is completed, MARS then removes some of the basis functions, so long as a certain quantity that balances fit and number of basis functions is not reduced. Based on an argument in Friedman (1991), Ehrlichman and Henderson (2006) states that the fitting time is bounded by a term that is of the order $d^5 N$. Our experiments below indicate that this bound may be quite pessimistic.

We apply MARS to select the martingale as follows. Consider time step t, and suppose that we have N samples $X_t^{(1)}, \dots, X_t^{(N)}$ of X_t. Suppose further that corresponding to these samples, we have N *noisy* samples $Y_t^{(1)}, \dots, Y_t^{(N)}$ of L_t. (The noisy sample $Y_t^{(i)}$ is simply the payoff attained by the LSM method on the post-t portion of the path $X_0^{(i)}, X_1^{(i)}, \dots, X_T^{(i)}$. Since the post-$t$ path is a single realization of the future evolution of (X_s) beyond time t, $Y_t^{(i)}$ can be viewed as a noisy sample of L_t.) We then use MARS to fit an approximation $\hat{L}_t$ (previously denoted h_t)

of L_t as

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

After we fit such an approximation for each $t = T, \ldots, 1$, we then perform a "production run" in which the martingale and corresponding estimators are computed.

In computing the resulting martingale $\hat{\pi} = (\hat{\pi}_t : t = 0, \ldots, T)$, we need to be able to compute the conditional expectations

$$E_{t-1}(q_{i,j}[X_t(i) - k_{i,j}])_+,$$

for each t. But this is simply the value of a European option on a single asset which, depending on the complexity of the Markov process, can be computed very easily. For certain pricing problems the calculation can be difficult. Therefore, in many cases instead of applying the MARS algorithm directly to (X_t), we instead apply MARS to a transformation of (X_t).

4 ALGORITHM SUMMARY

We give a simple summary of the algorithm here. For full details see Ehrlichman and Henderson (2006).

To construct the naïve estimator we run the LSM method in a first phase, thereby obtaining the stopping times (τ_t). Then, in a second phase we independently compute the sample average of a number of i.i.d. replicates of

$$e^{-r\tau_0} g(\tau_0, X_{\tau_0}).$$

To construct the MARS-based estimator, we again proceed in 2 phases. In Phase 1, we apply the LSM method to compute the stopping times (τ_t), while simultaneously fitting the approximation functions $\hat{L}_t$ using MARS. In Phase 2 we generate mutually independent paths of $(X_t : t = 0, \ldots, T)$ that are independent of the paths used in Phase 1. The Phase 2 paths are then used to estimate the lower bound via a sample average of terms of the form

$$e^{-r\tau_0} g(\tau_0, X_{\tau_0}) - \hat{\pi}_{\tau_0}.$$

In our implementation, we include the usual control variate multiplicative constant β in front of the control $\hat{\pi}_{\tau_0}$, and estimate its optimal value using the usual methodology; see, e.g., Law and Kelton (2000) for details. We also use the Phase 2 paths to estimate the upper bound via a sample average of terms of the form

$$\max_{t=0,1,\ldots,T} \left[e^{-rt} g(t, X_t) - \hat{\pi}_t \right].$$

The LSM method in Phase 1 involves a regression that is performed separately from the MARS fit of the lower bounds. One might consider using MARS for the LSM regressions as well, perhaps due to programming convenience, but in our experiments we have not seen significant gains in the performance of the policy represented by the stopping times computed by such a method. Moreover, it would be slower to use MARS than to use simple regression.

5 NUMERICAL EXAMPLES

We now apply the methodology above to price a range of derivatives. Our calculations were performed using R (R Development Core Team 2005). In order to avoid stating results that are strongly platform-dependent (like computation times), we instead report *ratios* of computation times. To be able to compare "apples with apples" we proceed as follows.

In all experiments we fix the runlengths for Phase 1 and Phase 2 to 10,000 and 20,000 respectively. 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 runlengths ($\tilde{n}$ and n for the naïve and MARS-based estimators respectively) required to achieve a confidence interval halfwidth 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 runlengths, 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

$$\text{TR} = \tilde{R}/R \tag{4}$$

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

$$\text{VR} = \tilde{s}^2/s^2 \tag{5}$$

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.

5.1 Example: Asian Options

We first price Bermudan-Asian options under the Black-Scholes model. More precisely, we have that $X_t = (S_t, A_t)$, where S_0 is deterministic, and $S_1, \ldots, S_T$ are generated via

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

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

$$A_t = \frac{1}{t}\sum_{s=1}^{t} S_s.$$

(Note that A_0 is undefined.) The averaging dates are assumed to coincide with the possible exercise dates, which now exclude the date $t = 0$.

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 $t \geq 1$.

As noted at the end of Section 3, the MARS fitting is applied to a transformation of (X_t) rather than to (X_t) itself. Indeed, we replace S_t by its logarithm, and the arithmetic average A_t by the logarithm of the geometric average

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

This yields an approximation

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

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

$$E_{t-1} \log S_t = \log S_{t-1} + r - \frac{\sigma^2}{2},$$

for $t = 1, \ldots, T$, and

$$\begin{aligned} E_0 \log \tilde{A}_1 &= E_0 \log S_1, \\ E_{t-1} \log \tilde{A}_t &= \frac{1}{t}\Bigg((t-1)\log \tilde{A}_{t-1} + \log S_{t-1} \\ &\qquad + r - \frac{\sigma^2}{2}\Bigg), \end{aligned}$$

for $t = 2, \ldots, T$, and conditional marginal variances

$$\text{Var}_{t-1} \log \begin{bmatrix} S_t \\ \tilde{A}_t \end{bmatrix} = \begin{bmatrix} \sigma^2 \\ (\sigma/t)^2 \end{bmatrix},$$

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

Table 1 gives the results for Bermudan-Asian options. We considered an option maturing in 6 months with monthly exercise/averaging dates, so that $T = 6$. The annualized risk-free rate was $12r = .06$, and the initial asset price was $S_0 = 100$. The stopping times were fit in the LSM method using polynomials of degree up to 4 in S_t and A_t for each $t = 1, \ldots, T-1$.

In Table 1 the column "Naïve L_0" gives the naïve estimate together with a 95% confidence interval halfwidth in parentheses. The corresponding information for the MARS-based estimate is given in the column "MARS L_0," and the MARS-based estimate of the upper bound appears under "MARS U_0." The columns headed "VR" and "TR" respectively give the variance reduction ratio and efficiency improvement ratios given by (5) and (4), to two significant figures.

The reduction in variance is dramatic, approximately equal to a factor of 200 in the examples we tried. Of course, one must take into account computational time. When we do so, the efficiency improvement factor ranges from about 40 to 80. Finally, the estimated bounds on the option price are very close, suggesting that the stopping times found by the LSM method are quite good.

Table 1: Asian Option Results

$\sigma\sqrt{12}$	K	Naïve L_0	MARS L_0	MARS U_0	VR	TR
.3	95	2.77 (.07)	2.73 (.00)	2.78 (.00)	210	85
.3	115	15.92 (.14)	15.86 (.01)	15.95 (.01)	230	44
.6	95	7.88 (.15)	7.80 (.01)	7.94 (.01)	190	71
.6	115	20.57 (.23)	20.48 (.02)	20.65 (.01)	230	56

5.2 Example: Basket Options

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

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

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

$$S_t(i) = S_{t-1}(i) \exp\left(r - \frac{1}{2}\sigma_i^2 + \sigma_i W_t(i)\right), \qquad (6)$$

for $i = 1, \dots, d$, where $W_t = (W_t(1), \dots, W_t(d))$ is a sequence of independent (in time) multivariate $\mathcal{N}(0,1)$ random variates with a specified correlation matrix. To simplify things, we assume that there is a single constant $\rho \in (-1,1)$ such that $\text{Cor}(W_t(i), W_t(j)) = \rho$ for all $1 \le i < j \le d$, and we take $\sigma_1 = \cdots = \sigma_d = \sigma$. We take the annualized risk-free rate $12r$ to be .06 and the initial asset prices to be $S_0(i) = 100$, for $i = 1, \dots, d$. The dimension d of the problem takes the values $d = 2, 5, 10, 20$. For the payoff function g_{avg}, we take the basis functions for fitting the stopping times in the LSM method 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 in the spirit of a suggestion in Longstaff and Schwartz (2001).

As in the Asian case, we apply the MARS-fitting algorithm to the logarithm of S_t. The conditional distribution of $\log S_t$ given $\mathscr{F}_{t-1}$ is multivariate Gaussian with mean $\log S_{t-1} + r - \sigma^2/2$, and variance σ^2. Again, this information is sufficient to compute the required conditional expectations. The results are given in Tables 2 and 3. Again, columns VR and TR are to two significant figures.

Table 2 presents the results for the put on the average. The efficiency improvements seen here are not as strong as those for the Asian option. Nevertheless, the factors, ranging from about 2 to 3 for the uncorrelated case to about 7 to 9 for the correlated case, represent useful improvements. Note that the upper bound estimates are quite poor; we have been able to obtain far better results for both variance reduction and the upper bound using the extension of this work described in Ehrlichman and Henderson (2006), and the results are reported there.

There is some degradation in performance as the dimension increases from 2 to 20 in the uncorrelated case. The performance is better in the correlated case. It is plausible that the efficiency improvements are stronger when the assets are positively correlated, and that the degradation with dimension is smaller in that case as well, because a large fraction of the assets' returns is driven by a single factor that is well represented in the individual (marginal) prices.

The results for the put on the maximum (Table 3) are less encouraging. In the uncorrelated case, the use of MARS actually led to a *reduction* in efficiency for dimensions 5 and higher. (We did not report the results for $d = 10, 20$.) The results for the correlated case are stronger, but not as strong as we might like. 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.

6 DISCUSSION AND CONCLUSION

We have presented a technique for (almost) automatically determining an effective martingale-based control variate for pricing American options. The method employs separable MARS approximations of the value functions. The key advantages of this approach are that one can use off-the-shelf software for fitting the approximation, the procedure itself is highly automated, and the one-step conditional expectations that help define the martingale can usually be very easily computed.

The method works extremely well when the separable approximations are accurate, providing substantial efficiency improvements and excellent upper bounds on the option price. However, for problems where the separable approximation is not as accurate, the results are correspondingly poorer. An extension of these methods described in

Table 2: Basket Option Results: Put on Average

d	$\sigma\sqrt{12}$	Naïve L_0	MARS L_0	MARS U_0	VR	TR
Uncorrelated asset prices ($\rho = 0$).						
2	.3	6.23 (.11)	6.25 (.05)	8.15 (.04)	4.1	3.6
2	.6	14.68 (.22)	14.67 (.10)	17.90 (.07)	4.7	4.1
5	.3	3.38 (.06)	3.38 (.04)	5.07 (.03)	2.7	2.2
5	.6	8.84 (.14)	8.80 (.08)	12.00 (.07)	3.2	2.6
10	.3	2.05 (.04)	2.04 (.03)	3.27 (.03)	2.4	1.8
10	.6	5.83 (.10)	5.81 (.06)	8.43 (.06)	2.7	2.1
20	.3	1.07 (.02)	1.12 (.02)	1.97 (.02)	2.0	1.6
20	.6	3.50 (.06)	3.59 (.04)	5.61 (.04)	2.3	1.9
Correlated asset prices ($\rho = .45$).						
2	.3	7.75 (.13)	7.80 (.04)	8.79 (.02)	10.0	8.7
2	.6	17.50 (.25)	17.55 (.08)	19.38 (.05)	10.0	8.9
5	.3	6.58 (.11)	6.65 (.04)	7.73 (.03)	8.4	6.8
5	.6	15.12 (.23)	15.19 (.08)	17.29 (.05)	8.5	7.0
10	.3	6.13 (.10)	6.22 (.03)	7.16 (.02)	9.1	7.4
10	.6	14.15 (.21)	14.25 (.07)	16.17 (.05)	9.4	7.5
20	.3	5.69 (.10)	5.96 (.03)	6.79 (.02)	11.0	8.7
20	.6	13.42 (.20)	13.74 (.06)	15.41 (.05)	11.0	8.3

Table 3: Basket Option Results: Put on Max

d	$\sigma\sqrt{12}$	Naïve L_0	MARS L_0	MARS U_0	VR	TR
Uncorrelated asset prices ($\rho = 0$).						
2	.3	3.83 (.08)	3.83 (.06)	6.00 (.05)	2.2	1.7
2	.6	9.83 (.19)	9.83 (.11)	13.92 (.09)	2.6	2.0
5	.3	0.39 (.02)	0.39 (.02)	0.95 (.02)	1.2	0.87
5	.6	1.54 (.06)	1.54 (.05)	3.32 (.06)	1.3	0.94
Correlated asset prices ($\rho = .45$).						
2	.3	5.32 (.11)	5.38 (.06)	7.19 (.04)	3.5	2.8
2	.6	12.75 (.22)	12.84 (.11)	16.13 (.07)	4.1	3.3
5	.3	2.29 (.06)	2.30 (.04)	3.76 (.04)	2.2	1.7
5	.6	6.07 (.15)	6.17 (.09)	9.25 (.07)	2.4	2.0
10	.3	1.05 (.04)	1.10 (.03)	2.06 (.03)	1.8	1.4
10	.6	3.13 (.10)	3.23 (.07)	5.52 (.06)	2.0	1.6
20	.3	0.44 (.03)	0.45 (.02)	1.10 (.02)	1.6	1.3
20	.6	1.46 (.07)	1.50 (.05)	3.05 (.04)	1.8	1.5

Ehrlichman and Henderson (2006) gives tremendous improvement in many cases where MARS alone does not do so well.

One could apply quasi-Monte Carlo methodology (including the randomized variants) in conjunction with the methodology described here. 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.

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 *Proceedings of the Winter Simulation Conference*, ed. M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Jones. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

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.

Duffie, D. 2001. *Dynamic asset pricing theory*. Third ed. Princeton, New Jersey: Princeton University Press.

Ehrlichman, S., and S. G. Henderson. 2006. Adaptive contol variates for pricing multi-dimensional American options. Working paper, Cornell University Dept. of ORIE.

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

Glasserman, P. 2004. *Monte Carlo methods in financial engineering*, Volume 53 of *Applications of Mathematics*. New York: Springer-Verlag.

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.

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.

Law, A. M., and W. D. Kelton. 2000. *Simulation modeling and analysis*. 3rd ed. New York: McGraw-Hill.

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.

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.

AUTHOR BIOGRAPHIES

SAMUEL M. T. EHRLICHMAN is a Ph.D. student in the School of Operations Research and Industrial Engineering at Cornell University. He received a John McMullen Graduate Fellowship from Cornell and is currently an NDSEG Fellow. His e-mail address is <se52@cornell.edu>.

SHANE G. HENDERSON is an associate professor in the School of Operations Research and Industrial Engineering at Cornell University. He has previously held positions in the Department of Industrial and Operations Engineering at the University of Michigan and the Department of Engineering Science at the University of Auckland. He is the simulation area editor at *Operations Research*, and an associate editor for the *ACM Transactions on Modeling and Computer Simulation* and *Operations Research Letters*. He likes cats but is allergic to them. His research interests include discrete-event simulation and simulation optimization.