

A PRUNED AND BOOTSTRAPPED AMERICAN OPTION SIMULATOR

Mark Broadie

Columbia Business School
415 Uris Hall
New York, NY 10027

Paul Glasserman

Columbia Business School
403 Uris Hall
New York, NY 10027

ABSTRACT

The pricing of American options on multiple assets or with path-dependent payoffs is an important but computationally challenging problem. In earlier work, we introduced simulation estimators for this problem which, though biased, are consistent and asymptotically unbiased. In this paper, we introduce enhancements to reduce bias. One enhancement exploits more easily computed European option prices; another uses bootstrapping for bias estimation.

1 INTRODUCTION

An option is a security granting the owner the right, but not the obligation, to buy or sell an asset at a specified price. If the option may be exercised at only a specified time, it is called *European*; if it may be exercised at any time in an interval, it is called *American*. More generally, "American" refers to any security whose cash flows can be influenced by its owner. Pricing an American option entails determining an optimal policy and is thus more difficult than pricing an otherwise equivalent European option.

In practice, American options are usually priced by applying dynamic programming to a discrete-time, discrete-space approximation to the evolution of the underlying asset or assets. The binomial method of Cox, Ross, and Rubinstein (1979) is probably the best known and most widely used such technique. See Chapter 14 of Hull (1993) for an introduction and see Broadie and Detemple (1994) for an extensive comparison of methods. Using multi-dimensional generalizations like that of Boyle et al. (1989), the binomial method and its variants are effective in pricing options involving up to three or perhaps even four assets. But because the computational requirements of these techniques typically grow exponentially in the number of state variables, higher dimensional problems quickly become intractable. The presence of

path-dependent payoffs further complicates the problem.

In Broadie and Glasserman (1995) (henceforth BG), we introduced a simulation-based method for pricing American options with finitely many exercise opportunities. The convergence rate of Monte Carlo is generally insensitive to the dimension of the problem, making it especially attractive precisely where the binomial method breaks down. The BG method produces estimates which, though biased, are consistent (converge in probability) and asymptotically unbiased. Moreover, by combining estimators that are biased high and biased low, this method produces conservative confidence intervals allowing for rigorous error control. No rigorous error control is available using the binomial method. Barraquand and Martineau (1995), Grant et al. (1994), and Tilley (1993) have proposed other simulation-based procedures for pricing American options; questions remain about the convergence of each of these.

In this paper, we investigate enhancements to our earlier estimators. Our first enhancement exploits information about European option prices to simultaneously reduce variance, bias, and execution time. Our second and more ambitious technique uses bootstrapping to estimate (and hopefully reduce) bias. We review the problem setting and our earlier estimators in the next section. The enhancements are presented in Sections 3 and 4.

2 MODEL AND ESTIMATORS

We assume that all information required to determine the payoff from exercising an option is recorded in a vector-valued Markov chain $\{S_t\}$. In practice, S_t would record all relevant information about asset prices, interest rates, exchange rates, and supplementary variables needed to eliminate path-dependence. For simplicity, we will mostly consider the case of a fixed interest rate and think of the components of S_t

as stock prices. In this case, a typical model makes the logarithms of the components of S_t random walks with normal increments correlated across components but independent over time.

2.1 The Optimal Stopping Problem

We use the following notation to determine payoffs:

- Time $t = 0, \dots, T$, means $t = t_0, \dots, t_{d-1}$, with $0 = t_0 < t_1 < \dots < t_{d-1} = T$; the t_i are the exercise opportunities.
- e^{-R_t} is the discount factor from $t - 1$ to t . We take R_t to be a component of the vector S_t and assume $R_t \geq 0$ for all t .
- $R_{0t} = \sum_{i=1}^t R_i$.
- $h_t(s)$ = payoff from exercise at time t in state s .
- $f_T(s) = g_T(s) = h_T(s)$.
- $g_t(s) = E[e^{-R_{t+1}} f_{t+1}(S_{t+1}) | S_t = s]$ = continuation value at time t in state s , $t = 0, \dots, T - 1$.
- $f_t(s) = \max\{h_t(s), g_t(s)\}$ = option value at time t in state s , $t = 0, \dots, T - 1$.

Our goal is to compute

$$f_0(S_0) = \max_{\tau} E[e^{-R_{0\tau}} h_{\tau}(S_{\tau})], \quad (1)$$

where the maximum is over all stopping times τ taking values in $\{0, 1, \dots, T\}$. The optimal policy stops at

$$\tau = \inf\{0 \leq t \leq T : h_t(S_t) \geq g_t(S_t)\};$$

i.e., the first time the immediate exercise value is at least as great as the continuation value. That the market price of an option can be represented in this way is a consequence of the general theory of the pricing of contingent claims. For an entry into the connection between simulation and this theory see Boyle et al. (1995); for textbook treatments see Duffie (1992) and Hull (1993).

Broadie and Glasserman (1995) argue that, in general, there is no unbiased estimator of (1). As an alternative, they introduce two estimators, one biased high and one biased low, both consistent and asymptotically unbiased. We discuss these next.

2.2 The Estimators

Our method simulates random trees determined by the evolution of S_t , rather than just sample paths. Given a value of the *branching parameter* b , the evolution of the tree can be described recursively as follows. From the (fixed) initial state S_0 , we generate b




Figure 1: A tree with $b = 3$

independent samples S_1^1, \dots, S_1^b of the state at time $t = 1$; from each node value $S_t^{i_1 \dots i_t}$ we generate b samples $S_{t+1}^{i_1 \dots i_t j}$, $j = 1, \dots, b$, conditionally independent of each other given $S_t^{i_1 \dots i_t}$ and each having the distribution of S_{t+1} given $S_t = S_t^{i_1 \dots i_t}$. Thus, each sequence $S_0, S_1^{i_1}, S_2^{i_1 i_2}, \dots, S_T^{i_1 \dots i_T}$ is a realization of the Markov chain S_t . Figure 1 illustrates a tree with $b = 3$.

Our *high estimator* is simply the result of applying dynamic programming to the random tree. More precisely, working backwards through the tree using the recursions

$$\Theta_T^{i_1 \dots i_T} = h_T(S_T^{i_1 \dots i_T})$$

and

$$\Theta_t^{i_1 \dots i_t} = \max \left\{ h_t(S_t^{i_1 \dots i_t}), \frac{1}{b} \sum_{j=1}^b e^{-R_{t+1}^{i_1 \dots i_t j}} \Theta_{t+1}^{i_1 \dots i_t j} \right\}, \quad (2)$$

we compute the high estimator $\Theta = \Theta_0$. That this estimator is biased high is a consequence of Jensen's inequality; that it is also consistent and asymptotically unbiased as $b \rightarrow \infty$ is proved in BG (1995).

The high estimator uses all branches emanating from a node to approximate both the optimal action (stop or continue) and the value of this decision. Our *low estimator* differs in that it separates the branches

used to determine the action from those used to determine the payoff. Here is a verbal description:

1. At each node in the tree, reserve one successor node. Average the discounted low estimator values at the other $b - 1$ successor nodes.
2. If the average obtained is less than the immediate exercise value, set the node value equal to the immediate exercise value; otherwise, set the node value equal to the discounted value from the reserved node.
3. Average the resulting node value over all b ways of selecting the reserved successor node.
4. Repeat these steps backwards through the tree.

Now we give a more precise formulation. First let

$$\theta_T^{i_1 \dots i_T} = h_T(S_T^{i_1 \dots i_T}).$$

Next, set

$$\eta_t^{i_1 \dots i_j} = h_t(S_t^{i_1 \dots i_j})$$

if

$$h_t(S_t^{i_1 \dots i_t}) \geq \frac{1}{b-1} \sum_{\substack{i=1 \\ i \neq j}}^b e^{-R_{t+1}^{i_1 \dots i_t}} \theta_{t+1}^{i_1 \dots i_t}, \quad (3)$$

and

$$\eta_t^{i_1 \dots i_j} = e^{-R_{t+1}^j} \theta_{t+1}^{i_1 \dots i_j}$$

otherwise. Then let

$$\theta_t^{i_1 \dots i_t} = \frac{1}{b} \sum_{j=1}^b \eta_t^{i_1 \dots i_j}, \quad (4)$$

for $t = 0, \dots, T - 1$. The low estimator is $\theta = \theta_0$. BG (1995) show that this estimator is biased low but also consistent and asymptotically unbiased as $b \rightarrow \infty$.

For a fixed branching parameter b , we may compute multiple replications of the high and low estimators and estimate the standard deviation of each. By taking the upper confidence limit of the high estimator and the lower confidence limit of the low estimator we obtain a conservative confidence interval for the true price. If we let d be the number of exercise opportunities and n the number of replications, then the work required to carry this out grows like nb^{d-1} . In particular, increasing the branching parameter is typically far more costly than increasing the number of replications. However, increasing b is essential to reducing bias and thus reducing the width of the confidence interval. This difficulty motivates our investigation into techniques for reducing or estimating bias.

3 PRUNING WITH EUROPEAN VALUES

Pricing a European option is generally easier than pricing the American counterpart: there is no optimization involved in the European price. It is therefore natural to try to exploit information obtained from the European case in pricing the American option. In BG (1995), we showed through examples that the European price provides a highly effective control variate. Here we illustrate two other uses, both of which reduce the number of nodes in the tree.

1. Pruning at the end. Let $T - 1$ denote the penultimate exercise opportunity. The optimal action at time $T - 1$ depends on which is greater, the immediate exercise value $h_{T-1}(S_{T-1})$ or the continuation value $g_{T-1}(S_{T-1})$. The estimators Θ and θ implicitly estimate the continuation value at each node. But at time $T - 1$ the continuation value is just the value of a European option initiated at time $T - 1$ and maturing at time T . Computing this value directly and efficiently eliminates the need to branch at the penultimate node. This reduces the work required to $O(nb^{d-2})$. Intuitively, it seems likely to reduce both bias and variance as well. Bias and convergence results for the original BG estimators continue to apply: we have merely reduced the number of time steps and replaced the terminal payoff $h_T(\cdot)$ with a new terminal payoff $E[e^{-R_T} h_T(S_T) | S_{T-1} = \cdot]$.

2. Intermediate pruning. The sole reason for branching (as opposed to simulating sample paths in the usual way) is to allow for consistent estimation of the optimal action at a node. Suppose that at time t there is a node corresponding to state s . If we knew that $h_t(s) < g_t(s)$, we would know that the optimal action is to continue and there would be no need to branch; it would suffice to generate just one successor node. Of course, in general we do not know $g_t(s)$ since g_t is itself the value function of an optimal stopping problem. But if we can find an easily computed lower bound $\ell(s) \leq$ (respectively $<$) $g_t(s)$, we can check if $h_t(s) <$ (respectively \leq) $\ell(s)$. If this holds, stopping is guaranteed to be suboptimal so there is no need to branch. If there is no branching out of node $i_1 \dots i_t$, then (2) gets replaced by $\Theta_t^{i_1 \dots i_t} = \exp(-R_{t+1}^{i_1 \dots i_t}) \Theta_{t+1}^{i_1 \dots i_t}$ and (4) gets replaced with $\theta_t^{i_1 \dots i_t} = \exp(-R_{t+1}^{i_1 \dots i_t}) \theta_{t+1}^{i_1 \dots i_t}$. This does not alter the convergence of the estimators, provided the number of replications (equivalently, the number of branches out of S_0) increases to infinity. It is generally not possible to determine in advance the reduction in work per run resulting from pruning because the amount of pruning is random.

In virtually all practical examples, the value of an

option remains strictly positive throughout its existence. Thus, a simple choice of bound is simply $\ell(s) = 0$: at any node at which the immediate exercise value is zero, there is no need to branch. This test is free because this choice of ℓ requires no computational effort.

In the case $h_t(s) > 0$, we may decide to compare the immediate exercise value with a more refined bound. Natural choices are $\ell_{t+1}(s), \dots, \ell_T(s)$, where $\ell_k(s)$ is the value of a European option initiated in state s and maturing at time k . Each of these corresponds to a particular (suboptimal) exercise policy for the American option and thus provides a lower bound on $g_t(s)$. If $h_t(s) < \ell_k(s)$ for any $k = t + 1, \dots, T$, there is no need to branch.

We have examined the effectiveness of these two pruning techniques in pricing an American option on the maximum of two assets. In this example, the state is two-dimensional with components $S_t^{(1)}$ and $S_t^{(2)}$. These asset prices evolve according to the rule

$$S_{t_j+1}^{(i)} = S_{t_j}^{(i)} \exp\left\{\left(r - \delta_i - \frac{1}{2}\sigma_i^2\right)(t_{j+1} - t_j) + \sqrt{t_{j+1} - t_j} W_j^{(i)}\right\}, \quad i = 1, 2,$$

where r is the (constant) interest rate, δ_1, δ_2 are the dividend yields, σ_1, σ_2 are the asset volatilities, and $(W_j^{(1)}, W_j^{(2)})$ are mean-zero normal random variates with standard deviations σ_1, σ_2 and correlation ρ . For simplicity, we take $S_0^{(1)} = S_0^{(2)} \equiv s_0$ and let this common initial asset price vary. The immediate exercise value is given by $h(s_1, s_2) = \max\{\max\{s_1, s_2\} - K, 0\}$, where K is the *strike price*. This corresponds to an option to buy either asset at price K .

Numerical results are illustrated in Table 1 for the following parameters: the annualized interest rate $r = .05$; both assets have dividend yields $\delta_i = 0.10$ and volatilities $\sigma_i = 0.20$; their correlation is $\rho = 0.30$; the strike price K is 100; the time to expiration T is 1 year; there are four exercise opportunities. We implemented the European value at the penultimate step and intermediate pruning based on first checking if $h(s_1, s_2) > 0$ and, if so, then comparing with a single European option maturing at time T . (The European price is computed from the formula of Johnson 1987.) At nodes passing both these tests we generate $b = 50$ branches. For each value of s_0 , the two rows in the table show results without and with pruning, respectively, based on approximately equal CPU times. The values labeled "True" were obtained using the binomial method of Boyle et al. (1989) with 800 time steps. The values labeled "Est" are the averages of the corresponding high and low estimators. Taking the midpoint is a fairly arbitrary way of compromis-

ing between the two; the "% Error" was computed from this estimate. Confidence intervals were computed as explained at the end of Section 2. Because these intervals are conservative, the actual coverage is usually much greater than the nominal coverage. Indeed, we have found that a nominal level of 90% often yields 99% actual coverage.

The results of Table 1 indicate substantial gains from pruning, as evidenced by the percent errors and confidence interval widths. The potential for savings is even greater than that indicated by the tables. If, for example, $\delta_1 = \delta_2 = 0$, then the optimal policy holds the option to expiration; with pruning, the algorithm will in fact never branch.

4 BOOTSTRAPPING

Our second approach to the issue of bias uses a bootstrapped estimate of bias based on Efron and Tibshirani (1993). We first briefly review their technique, then discuss how we apply it in our setting.

4.1 Background

Efron and Tibshirani consider the problem of estimating the bias of an estimator defined on random samples from a distribution on the real line. Let F be the distribution and let $t(F)$ be a scalar statistic. Let X_1, \dots, X_n be i.i.d. samples from F and write X for (X_1, \dots, X_n) . Let $\gamma(X)$ be an estimator of $t(F)$. By definition, its bias is

$$\text{bias}_F(\gamma) = E_F[\gamma(X)] - t(F),$$

the subscript F indicating the distribution from which the observations are drawn.

Let \hat{F} denote the empirical distribution of the observations (X_1, \dots, X_n) . The bootstrap estimate of bias is

$$\text{bias}_{\hat{F}}(\gamma) = E_{\hat{F}}[\gamma(X)] - t(\hat{F}).$$

In this expression, $t(\hat{F})$ is the value of the statistic of interest for the empirical distribution and $E_{\hat{F}}[\gamma(X)]$ is the expectation of the estimator γ when the observations are drawn from \hat{F} rather than F . Typically, $t(\hat{F})$ can be evaluated directly, but $E_{\hat{F}}[\gamma(X)]$ requires Monte Carlo: B random samples of size n are drawn from \hat{F} and the sample mean of their γ values is used to approximate $E_{\hat{F}}[\gamma(X)]$.

Efron and Tibshirani recommend a slight variant of this estimator. To each of the B bootstrap random samples there corresponds a vector of frequencies (P_1, \dots, P_n) , in which P_i records the fraction of observations in the sample equal to the i^{th} original observation X_i . Let \bar{F} be the distribution assigning

s_0	θ	$se(\theta)$	Θ	$se(\Theta)$	[90% CI]	Est.	True	% Err.
70	0.224	0.006	0.227	0.007	[0.214, 0.238]	0.226	0.236	4.25
70	0.239	0.001	0.239	0.001	[0.237, 0.241]	0.239	0.236	1.27
80	1.211	0.022	1.231	0.023	[1.175, 1.268]	1.221	1.260	3.12
80	1.255	0.004	1.257	0.004	[1.249, 1.263]	1.256	1.260	0.35
90	3.933	0.051	4.036	0.052	[3.850, 4.122]	3.985	4.076	2.23
90	4.077	0.013	4.092	0.013	[4.056, 4.113]	4.085	4.076	0.22
100	9.090	0.089	9.361	0.091	[8.944, 9.510]	9.225	9.356	1.40
100	9.288	0.032	9.359	0.033	[9.234, 9.412]	9.323	9.356	0.35
110	16.558	0.128	17.005	0.130	[16.347, 17.219]	16.782	16.925	0.85
110	16.775	0.061	16.942	0.062	[16.674, 17.043]	16.858	16.925	0.40
120	25.477	0.163	26.111	0.162	[25.209, 26.377]	25.794	25.979	0.71
120	25.838	0.084	26.107	0.084	[25.700, 26.246]	25.973	25.979	0.02
130	35.203	0.187	35.924	0.185	[34.896, 36.230]	35.564	35.763	0.56
130	35.606	0.100	35.948	0.100	[35.442, 36.113]	35.777	35.763	0.04

Table 1: Improvement from pruning: for each s_0 , first row is without pruning, second row is with pruning

mass \bar{P}_i to X_i , where \bar{P}_i is the average of the P_i over the B bootstrap samples. The recommended estimator of bias is then

$$\frac{1}{B} \sum_{j=1}^B \gamma(X^{(j)}) - t(\bar{F}), \quad (5)$$

where $X^{(j)}$ denotes the j^{th} bootstrap sample.

4.2 Bootstrapping the High and Low Estimators

Applying the bootstrap method to Θ and θ is not altogether straightforward because these estimators are defined on random trees, rather than on sets of scalars. We implement it as follows. We view a random tree as the state space of a Markov chain. The transition structure of the chain is determined by the branches of the tree, and the chain selects any of the b branches emanating from a node with equal probability. (For simplicity, we explain the method without the pruning of Section 3.) The law of this Markov chain plays the role of \hat{F} in the scalar setting: it is an empirical version of the true law of S_t .

The next step is to generate bootstrap samples from the empirical law. In the scalar setting, this merely requires sampling with replacement n times from the original observations (X_1, \dots, X_n) . In the BG setting, generating a bootstrap sample means generating b^{d-1} paths of the empirical Markov chain. (There are b^{d-1} paths through the original tree, so this is the appropriate number of paths to generate for a bootstrap sample.) Each bootstrap path has probability $1/b$ of choosing any of the b branches emanating from a node; however, the actual proportion of paths in a bootstrap sample choosing a branch is likely to

differ somewhat from $1/b$. For each node and for each bootstrap sample we thus obtain a vector of frequencies (P_1, \dots, P_b) for the b branches that emanate from that node. In fact, it is only these frequencies that matter, so it suffices to generate them directly, without actually simulating paths through the tree. Once we have a vector of such resampled frequencies at a node, we compute bootstrapped estimator values $\Theta^{(j)}$ and $\theta^{(j)}$ by replacing the uniform weights in (2), (3), and (4) with the weights (P_1, \dots, P_b) . This yields bootstrap values $\Theta^{(j)}$ and $\theta^{(j)}$, $j = 1, \dots, B$.

Paralleling the scalar case, at each node we also record the vector $(\bar{P}_1, \dots, \bar{P}_b)$ of average frequencies over all B bootstrap samples and compute estimators $\Theta_{\bar{B}}$ and $\theta_{\bar{B}}$ from these weights. Paralleling (5), the resulting estimators of bias are

$$\widehat{\text{bias}}(\Theta) = \frac{1}{B} \sum_{j=1}^B \Theta^{(j)} - \Theta_{\bar{B}}$$

and

$$\widehat{\text{bias}}(\theta) = \frac{1}{B} \sum_{j=1}^B \theta^{(j)} - \theta_{\bar{B}}.$$

The bias-adjusted high and low estimators are

$$\Theta - \widehat{\text{bias}}(\Theta) \quad \text{and} \quad \theta - \widehat{\text{bias}}(\theta).$$

Table 2 shows results with and without the bootstrap bias adjustment. The option parameters are the same as those used for Table 1. For each value of s_0 , the first row gives results without the bootstrap adjustment, the second row gives results with the adjustment. The number of replications for the unadjusted estimates was increased to make the CPU times for all rows approximately equal. All results use

s_0	θ	$se(\theta)$	$\widehat{bias}(\theta)$	Θ	$se(\Theta)$	$\widehat{bias}(\Theta)$	[90% CI]	Est.	True	% Err.
70	0.244	0.003	—	0.244	0.003	—	[0.239, 0.250]	0.244	0.236	3.58
70	0.244	0.006	0.000	0.244	0.006	0.000	[0.234, 0.253]	0.244	0.236	3.38
80	1.265	0.010	—	1.267	0.010	—	[1.249, 1.283]	1.266	1.260	0.45
80	1.249	0.015	-0.001	1.249	0.015	0.001	[1.223, 1.274]	1.249	1.260	0.92
90	4.069	0.019	—	4.105	0.020	—	[4.037, 4.137]	4.087	4.076	0.28
90	4.063	0.029	-0.004	4.080	0.030	0.013	[4.015, 4.128]	4.071	4.076	0.10
100	9.298	0.030	—	9.476	0.030	—	[9.249, 9.526]	9.387	9.356	0.33
100	9.319	0.048	-0.007	9.421	0.048	0.065	[9.240, 9.500]	9.370	9.356	0.15
110	16.769	0.038	—	17.200	0.039	—	[16.705, 17.264]	16.984	16.925	0.35
110	16.708	0.055	-0.008	16.972	0.055	0.161	[16.617, 17.063]	16.840	16.925	0.50
120	25.793	0.045	—	26.493	0.045	—	[25.720, 26.568]	26.143	25.979	0.63
120	25.929	0.075	-0.026	26.344	0.075	0.263	[25.805, 26.467]	26.136	25.979	0.61
130	35.484	0.049	—	36.408	0.050	—	[35.403, 36.490]	35.946	35.763	0.51
130	35.594	0.081	-0.019	36.155	0.080	0.346	[35.460, 36.287]	35.875	35.763	0.31

Table 2: Bootstrapping results, two-asset case: for each s_0 , first row is without bootstrapping, second row is with bootstrapping

the pruning techniques discussed in Section 3. The branching parameter b is 10 in all cases; the adjusted estimates use $B = 100$ bootstrap replications.

The results are mixed. At lower values of s_0 , the optimal exercise policy almost always holds the option to expiration, and there is very little branching after pruning. As a result, there is very little bias (the high and low estimators give nearly identical results) and no benefit from bootstrapping. In particular, the bias reduction is not sufficient to offset the increase in standard error that results from bootstrapping. At high values of s_0 , there is more branching; the bias adjustments become more significant and result in narrower interval estimates. The relative error of the point estimate appears to decrease slightly. However, recall that our point estimate is (somewhat arbitrarily) the average of the high and low estimates. Because the bias adjustment appears to be more significant for the high than the low estimate, the midpoint may not make optimal use of these adjustments.

Table 3 shows results for an American option on the maximum of five assets. The parameters of this option are the same as those of the two-asset example of Table 2; in particular, correlations between all pairs of assets are equal. We have no way of evaluating the true price in the five-asset case so we report only interval estimates. At each value of s_0 , the three rows implement the following techniques, respectively: pruning; pruning plus a control variate; pruning, control variate, and bootstrapping. Because computing even the European price in the five-asset case is burdensome, we implemented only intermediate pruning, not pruning at the end. For the intermediate pruning we used the European price of an

option on the maximum of the two assets with the highest value at a node; this is easier to compute than a five-asset option and still provides an effective lower bound. For the control variate, we were able to use the five-asset European price (computed from Johnson 1987) because using the control variate technique requires evaluating this price just once. All rows are based on branching parameter $b = 50$. The bootstrapped results use $B = 100$ bootstrap replications. The number of replications in each case was chosen to make the CPU times approximately equal.

The first conclusion to be drawn from the interval estimates in Table 3 is that the method is viable for American options on five assets. In particular, the halfwidths of the interval estimates using the control variate are within 1% of the interval midpoint. The results also suggest that the use of the control variate may be the single most effective way to increase efficiency. Bootstrapping for bias estimation appears to provide some further narrowing of the interval estimates. As always, caution is necessary in extrapolating from limited computational results, especially since the example and the method both involve multiple parameters.

ACKNOWLEDGMENT

We thank Gautam Jain for computational assistance.

REFERENCES

- Barraquand, J., and D. Martineau. 1995. Numerical Valuation of High Dimensional Multivariate American Securities, working paper, Salomon Brothers

s_0	θ	$se(\theta)$	Θ	$se(\Theta)$	[90% CI]
70	0.551	0.012	0.556	0.012	[0.531, 0.576]
70	0.552	0.000	0.555	0.000	[0.552, 0.555]
70	0.552	0.000	0.554	0.000	[0.552, 0.554]
80	2.687	0.036	2.730	0.037	[2.628, 2.791]
80	2.699	0.001	2.721	0.001	[2.698, 2.722]
80	2.700	0.002	2.709	0.001	[2.697, 2.711]
90	7.760	0.071	7.906	0.073	[7.643, 8.026]
90	7.787	0.002	7.873	0.002	[7.784, 7.875]
90	7.790	0.010	7.829	0.005	[7.773, 7.838]
100	15.787	0.105	16.070	0.107	[15.614, 16.246]
100	15.799	0.006	16.006	0.004	[15.790, 16.013]
100	15.852	0.028	15.925	0.015	[15.806, 15.950]
110	25.655	0.132	26.060	0.132	[25.437, 26.277]
110	25.649	0.009	25.970	0.007	[25.634, 25.982]
110	25.701	0.047	25.822	0.024	[25.624, 25.862]
120	36.310	0.150	36.806	0.150	[36.063, 37.052]
120	36.340	0.013	36.739	0.010	[36.318, 36.754]
120	36.308	0.070	36.566	0.034	[36.192, 36.622]
130	47.245	0.164	47.785	0.164	[46.975, 48.055]
130	47.227	0.015	47.699	0.011	[47.202, 47.717]
130	47.275	0.079	47.510	0.039	[47.145, 47.574]

Table 3: Five-asset case: for each s_0 , first row uses pruning, second row uses pruning and control variate, third row uses pruning, control variate, and bootstrapping

International Ltd., to appear in *Journal of Financial and Quantitative Analysis*.

Boyle, P., M. Broadie, and P. Glasserman. 1995. Recent Advances in Simulation for Security Pricing, in this volume.

Boyle, P., J. Evnine, and S. Gibbs. 1989. Numerical Evaluation of Multivariate Contingent Claims, *Review of Financial Studies*, Vol. 2, 241–250.

Broadie, M., and J. Detemple. 1994. American Option Valuation: New Bounds, Approximations, and a Comparison of Existing Methods, working paper, Columbia University (to appear in the *Review of Financial Studies*).

Broadie, M., and P. Glasserman. 1995. Pricing American-Style Securities Using Simulation, working paper, Columbia University.

Cox, J.C., S.C. Ross, and M. Rubinstein. 1979. Option Pricing: A Simplified Approach, *Journal of Financial Economics*, Vol. 7, 229–263.

Duffie, D. 1992. *Dynamic Asset Pricing Theory*, Princeton University Press, Princeton, New Jersey.

Efron, B., and R.J. Tibshirani. 1993. *An Introduction to the Bootstrap*, Chapman & Hall, New York.

Grant, D., G. Vora, and D. Weeks. 1994. Path-Dependent Options: Extending the Monte Carlo Simulation Approach, working paper, Anderson School of Management, University of New Mexico,

Albuquerque, New Mexico 87131–6471.

Hull, J. 1993. *Options, Futures, and Other Derivative Securities*, 2nd edition, Prentice-Hall, Englewood Cliffs, New Jersey.

Johnson, H. 1987. Options on the Maximum or the Minimum of Several Assets, *Journal of Financial and Quantitative Analysis*, Vol. 22, 227–283.

Tilley, J.A. 1993. Valuing American Options in a Path Simulation Model, *Transactions of the Society of Actuaries*, Vol. 45, 83–104.

AUTHOR BIOGRAPHIES

MARK BROADIE joined the faculty of the Columbia University Graduate School of Business in 1983. His current research interests include portfolio optimization and the pricing of derivative securities. He holds a Ph.D. from Stanford University and a B.S. from Cornell University.

PAUL GLASSERMAN joined the faculty of the Columbia University Graduate School of Business in 1991. Prior to that, he was a Member of Technical Staff in the Operations Research department of AT&T Bell Laboratories in Holmdel, NJ. He holds a Ph.D. from Harvard University and an A.B. from Princeton University.