Pricing American-style securities using simulation

Mark Broadie*, Paul Glasserman

Graduate School of Business, Columbia University, New York, NY 10027, USA

Abstract

We develop a simulation algorithm for estimating the prices of American-style securities, i.e., securities with opportunities for early exercise. Our algorithm provides both point estimates and error bounds for the true security price. It generates two estimates, one biased high and one biased low, both asymptotically unbiased and converging to the true price. Combining the two estimators yields a confidence interval for the true price. The proposed algorithm is especially attractive (compared with lattice and finite-difference methods) when there are multiple state variables and a small number of exercise opportunities. Preliminary computational evidence is given.

Keywords: Monte Carlo simulation; American option pricing; Path-dependent claims; Multiple state variables; Real options

1. Introduction

There are an increasing number of important security pricing models where analytical solutions are not available. In this paper we propose a general algorithm, based on Monte Carlo simulation, for the estimation of security prices.

The algorithm can be applied to models with multiple state variables, with possible path dependencies in the state variables, and is specially designed to handle American-style securities, i.e., securities with opportunities for early exercise. [1]

Despite its widespread significance, the valuation of early-exercise features remains a difficult problem in many important settings, particularly for multifactor models. Important pricing problems frequently arise in the theory of economic investment, see, e.g., Dixit and Pindyck (1994). Examples include the development of natural resources (Brennan and Schwartz, 1985), land-use decisions (Geltner et al., 1995), the adoption of technological innovation (Grenadier and Weiss, 1994), manufacturing flexibility (Triantis and Hodder, 1990), and, more fundamentally, the option to initiate or abandon a project (Majd and Pindyck, 1987; McDonald and Siegal, 1986; Pindyck, 1988). In many of these applications, analytical results are unavailable, even for relatively simple models. Indeed, there is no analytical solution for the price of an American option on a single dividend-paying asset in the standard Black–Scholes framework.

Binomial, lattice, and finite-difference methods can be used to generate numerical solutions to pricing problems with one or two sources of uncertainty. However, realistic models of many securities require three or more state variables. Geltner et al. (1995) successfully study the value of vacant land as a call option on the maximum of two possible land uses; but they also note that in practice there may be multiple uses – single-family houses, apartment buildings, commercial or industrial uses, for example. The general model of Triantis and Hodder (1990) considers a factory making an arbitrary number of distinct products; understandably, their method leads to numerical results only in the case of two products. Other examples include quality options (e.g., multiple deliverable assets in Treasury futures contracts) and other options on multiple assets (e.g., options on the maximum of two or more asset prices and options on the difference between two asset prices). Also, realistic models of foreign currency options, swaptions, and differential swaps include stochastic exchange rates and stochastic domestic and foreign term structures; see Amin and Jarrow (1991) and Turnbull (1993) for examples. Cortazar and Schwartz (1994) use multifactor models to price commodity contingent claims. Models with stochastic volatility, interest rates, default risk, convenience yields, and asset prices are becoming increasingly common. Assets with path-dependent payoffs (e.g., American Asian options) also give rise to multi-state models, because path dependence can usually be eliminated through the inclusion of additional state variables. For models with

---

[1] A *European-style* security is one in which the owner cannot influence the cashflows of the security. This definition includes standard European options as well as, e.g., barrier options with automatic exercise when a barrier is hit. The holder of an *American-style* security can make decisions which affect its cashflows. This definition includes standard American options which allow for continuous exercise, Bermudan options which allow exercise at a finite number of dates, shout options (which allow the owner to lock-in the current intrinsic value of the option), and other securities.

multiple state variables there are few, if any, analytical results for American-style securities and many instances where formulas for European-style securities are not available. In these cases, numerical methods are the only means for obtaining pricing and hedging results.

Given the importance of valuing early-exercise features in problems with multiple state variables, the dearth of studies that address these problems must be explained, in part, by a need for effective valuation procedures. This paper addresses that need by presenting a general method, based on Monte Carlo simulation, for the valuation of assets with early-exercise features. Because the method uses random sampling, rather than the enumeration implicit in lattice and finite-difference methods, it can be applied easily to models with multiple state variables and possible path dependencies.

Most pricing problems in the literature for which closed-form solutions are not available are solved numerically by lattice or finite-difference methods. The binomial lattice method proposed by Cox et al. (1979) has been very successful for pricing claims contingent on a single state variable which follows a geometric Brownian motion process. Their method has been extended in many directions. Nelson and Ramaswamy (1990) developed binomial methods for more general diffusion processes. Hull and White (1993a) present trinomial methods for valuing interest rate-sensitive securities. Boyle et al. (1989) give a multinomial method for valuing claims on several assets whose prices follow a multivariate geometric Brownian motion process. Trigeorgis (1991) proposes a log-transformed binomial method and applies it to the pricing of real options. Zenios and Shtilman (1993) provide a procedure for sampling paths from a lattice to provide a prespecified level of accuracy. [2] Although lattice methods have many desirable features, their computational cost grows exponentially with the number of state variables. Rigorous proofs of convergence are often difficult for lattice methods. Convergence of the Cox et al. (1979) binomial method for the case of American options was only recently proved in Amin and Khanna (1994).

Simulation methods for asset pricing were introduced to finance in Boyle (1977). Since that time simulation has been successfully applied to a wide range of pricing problems. For an overview of recent developments in the use of simulation for security pricing, see Boyle et al. (1997). Simulation is an attractive method for asset pricing because of the generality in the types of assets it can handle and the ease with which it handles multiple state variables, and path dependencies. Its two major drawbacks are speed of computation and apparent inability to deal with the *free-boundary* aspect of American options. Regarding the speed of computation, note that the convergence rate of Monte Carlo methods is typically independent of the number of state variables, whereas the convergence

[2] Other extensions of the binomial approach include jump processes in Cox and Ross (1976) and trinomial and multinomial processes in Boyle (1988), Amin (1991), and Kamrad and Ritchken (1991).

rate of lattice methods is exponential in the number of state variables. Hence, simulation methods should be increasingly attractive compared to lattice methods as the number of state variables grows.

The major difficulty in valuing early-exercise features is the need to estimate optimal exercise policies as well. Standard simulation procedures are 'forward' algorithms, i.e., paths of state variables are simulated forward in time. Given a state trajectory and a pre-specified exercise policy, a path price is determined. An average over independent samples of path prices gives an unbiased estimate of the security price. By contrast, pricing procedures for assets with early-exercise features are generally 'backward' algorithms. That is, the optimal exercise strategy at the maturity of the contract is easily determined. Proceeding backwards in time, the optimal exercise strategy and corresponding price are determined via dynamic programming. The problem of using simulation to price American options stems from the difficulty of applying an inherently forward-based procedure to a problem that requires a backward procedure to solve.

The first serious attempt to apply simulation to the pricing of American options is Tilley (1993).[3] In a single state variable setting, Tilley proposes a 'bundling' algorithm for security pricing. At each time period, simulated paths are ordered by asset price and bundled into groups. An optimal exercise decision is estimated for each group. Tilley's algorithm has at least three main difficulties. First, no proof of convergence is provided. Indeed, he provides evidence that if his algorithm converges, it does not converge to the correct value. Second, his algorithm requires that all simulated paths must be stored at one time. This is necessitated by the sorting and bundling required at each time period. The storage requirement poses a significant computational problem when the number of simulated paths is large. Third, there is no stated or obvious way to generalize the algorithm to additional state variables.[4] Analysis and extensions of Tilley's method are given in Carriere (1996). More recently, Bararaquand and Martineau (1995) propose a stratification method for pricing high-dimensional American securities. Further discussion of the methods of Tilley and Barraquand and Martineau is given in Boyle et al. (1997).

In this paper we argue that there can be no general method for producing unbiased simulation estimator of American option values. We circumvent this

---

[3] Dammon and Spatt (1992) use Monte Carlo simulation to value dividend reinvestment and voluntary purchase plans which involve optimal exercise decisions. However, their algorithm is fairly specific to these contracts. They observed the potential for bias when applying Monte Carlo techniques to valuing contracts with optimal exercise provisions (see Dammon and Spatt, 1992, p. 340, footnote 13).

[4] Specifically, it is not clear how to define bundles when there are multiple state variables. No matter how the bundling procedure is defined, it is likely that most bundles will contain very few paths. Hence, the bias problem is likely to be severe and increase with the number of state variables. In the single state variable case, Tilley suggests a refined procedure for determining a 'sharp' boundary. This refined procedure does not generalize to multiple state variables.

difficulty by generating *two* estimates of the asset price based on random samples of future state trajectories and increasingly refined approximations to optimal exercise decisions. One estimate is biased high and one is biased low; both estimates are asymptotically unbiased and converge to the true price. We combine the two estimates to give a valid, conservative confidence interval for the asset price. The method has minimal storage requirements and applies directly to problems with multiple state variables. In its simplest form, the method is limited to pricing securities with a finite number of exercise opportunities. Indeed, the computational requirements of the method grow exponentially in the number of exercise opportunities. Through extrapolation we may, however, extend the method (at least approximately) to continuous exercise or an arbitrary, finite number of exercise opportunities.

The paper is organized as follows. Simulation estimators are motivated and described in the next section. Theoretical analysis of the estimators is given in Section 3. Numerical results are given in Section 4. Concluding remarks, possible extensions, and areas for future research are given in Section 5. Appendix A contains proofs and Appendix B shows that, under reasonable restrictions, there can be no unbiased estimate of American option prices. Appendix C provides detailed information on the implementation of the algorithm.

2. Overview of the method

We begin our discussion by focusing on a standard call option on a single underlying asset which pays dividends at a continuous rate. The typical simulation approach to European option pricing is to use simulation to estimate the expectation

$$C = \mathrm{E}[e^{-rT}(S_T - K)^+]$$

under the risk-neutral measure. As usual, r denotes the riskless rate of interest, T the option maturity, K the strike price, and S_T the terminal stock price. The American option pricing problem is to find

$$C = \max_{\tau} \mathrm{E}[e^{-r\tau}(S_\tau - K)^+] \tag{1}$$

over all stopping times $\tau \leq T$. Throughout the paper, we focus on a discrete time approximation to this problem where we restrict the exercise opportunities to lie in the finite set of times $0 = t_0 < t_1 < \cdots < t_d = T$. The analogous procedure for an American option would be to simulate a path of asset prices, say, $S_0, S_1, \ldots, S_T$, at corresponding times $0 = t_0 < t_1 < \cdots < t_d = T$; then compute a discounted option value corresponding to this path, and finally average the results over many simulated paths. The main question is how to compute a discounted option value corresponding to the asset price path.
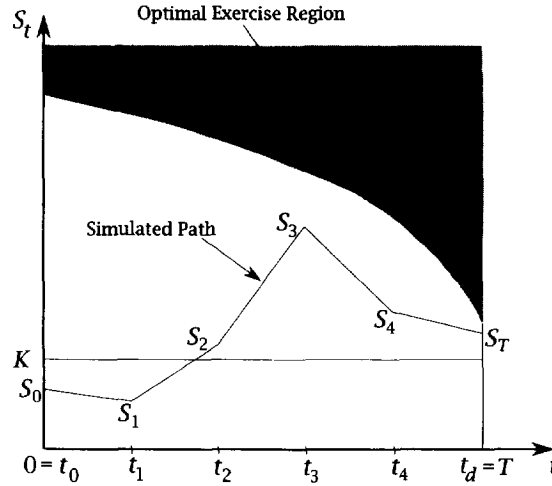
Fig. 1. Sample simulated path.

If the optimal stopping policy were known, the path estimate would be $e^{-r\tau}$ $(S_\tau - K)^+$. But the optimal stopping policy is not known and must also be determined via the simulation. A natural idea is to compute the optimal stopping time for the simulated path. This gives the path estimate

$$\max_{i=0,\ldots,d} e^{-rt_i}(S_{t_i} - K)^+.$$

However, this path estimate corresponds to the *perfect foresight solution* and hence it tends to overestimate the option value. Indeed, the overestimate follows from the inequality $\max_{i=0,\ldots,d} e^{-rt_i}(S_{t_i} - K)^+ \geq e^{-r\tau}(S_\tau - K)^+$. An illustration of this phenomenon is given in Fig. 1. Following the optimal but unknown exercise strategy, the option would be exercised at maturity since the path never entered the optimal exercise region. However, the perfect foresight strategy would achieve a higher path value by exercising at time t_3. The expected value of these path prices is not equal to the American option price given by (1) and so this procedure does not solve the pricing problem. Increasing the simulation effort by simulating many paths does not remove the bias problem.

The previous discussion illustrates the difficulties involved in applying standard simulation methodology to an American option pricing problem. Indeed, it seems unlikely that an unbiased estimator exists for this problem. Thus, in order to develop valid error bounds on the true option price, we introduce two estimators, one biased high and one biased low, but both asymptotically unbiased as the computational effort increases. These estimators are based on simulated trees. The simulated trees are parameterized by b, the number of branches per node.
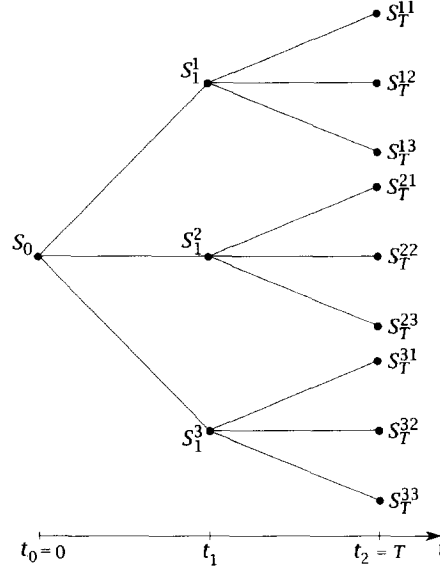
Fig. 2. Simulated tree for $b = 3$.

State variables are simulated at the finite number of possible decision points, i.e., exercise times. For purposes of exposition, we introduce the estimators informally and in a limited context. In the next section, we give a precise general formulation.

An illustration of a tree for $b = 3$ is given in Fig. 2. The connections between nodes indicates the dependence structure of the stock prices. For example, both S_T^{11} and S_T^{12} depend on S_1^1 but neither depends on S_1^2. It is important to keep in mind that in Fig. 2 the nodes at a fixed time appear according to the order in which they are generated, *not according to their node values*, as would be the case in a lattice method. For example, the node labeled S_1^1 need not correspond to a higher asset value than the node labeled S_1^2. Indeed, since a node label may, in general, record multiple state variables, there may be no natural ordering of the node values. Thus, Fig. 2 applies as well in the case of a multifactor model as a single-factor model.

2.1. The high estimator Θ

Let C denote the price of an American call option with $d + 1$ exercise opportunities at times t_i, $i = 0, \ldots, d$. We denote our first estimator by Θ. It is defined as the call value estimate obtained by a dynamic programming (DP) algorithm applied to the simulated tree. At the terminal date, the option value is known.
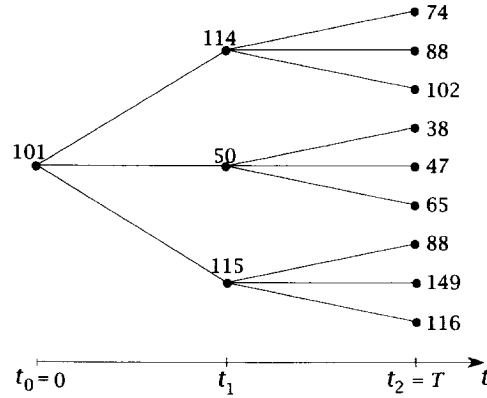
Fig. 3. Stock price tree.
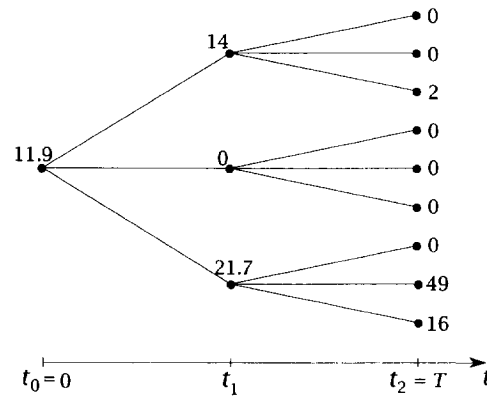
Fig. 4. Θ ('high') estimate.

At each prior date, the option value is defined to be the maximum of the immediate exercise value and the expectation of the succeeding discounted option values. Finally, Θ is the estimated option value at the initial node. A formal definition of Θ is given in the next section. A numerical illustration is given in Figs. 3 and 4. The parameters are $S_0 = 101$, $K = 100$, $T = 1$, and $r = 0$. For this particular tree, the American option price estimate is $\Theta = 11.9$.

The Θ estimator gives an estimate of the true option price which is biased upward, that is, $E[\Theta] \geq C$. We refer to this estimator as the 'high' estimator. Some intuition for the bias is given next; the precise argument is given in Appendix A. Any simulated tree will not perfectly represent the distribution of stock prices. If at some node future stock prices are too high, the dynamic programming

algorithm may choose not to exercise and receive a value higher than the 'optimal' decision to exercise. Likewise, if future stock prices are too low, the dynamic programming algorithm may choose to exercise even when the 'optimal' decision is not to exercise. In each case, the DP algorithm takes advantage of knowledge of the future to overestimate the option value.

Although biased, the high estimator is consistent, i.e., Θ converges to C as b increases. Precise statements about the type of convergence are given in the next section; proofs are given in Appendix A. The consistency of the Θ estimator distinguishes it from the estimator based on a single stock price path discussed earlier.

2.2. The low estimator θ

Next we propose an estimator that is biased low. The idea is to separate the branches at each node into two sets. The first set of branches is used to decide whether or not to exercise, and the second set is used to estimate the continuation value, if necessary. As we will see, this separation removes the upward bias in the estimator, but instead leads to an estimator with downward bias.

This idea is illustrated in Fig. 5. The numerical values are based on the stock price tree in Fig. 3. At each node the second two branches are used to determine the exercise decision and the first branch is used to determine the continuation value, if necessary. For example, at the bottom node at time t_1 the decision is made by comparing the immediate exercise value (15) with the discounted expected value of not exercising ($32.5 = 0.5 * 16 + 0.5 * 49$). Hence, the decision is made to continue, but the value assigned to this action is 0 (based on the first branch which leads to a terminal stock price of 88).

Why is this estimator biased low? To gain some intuition, consider the time just prior to expiration. The exercise decision is based here on unbiased information from the maturity date. If the correct decision is inferred from this information, the estimator would be unbiased. But with a finite sample, there is a positive probability of inferring a suboptimal decision. In this case, the value assigned to this node will be an unbiased estimate of the lower value associated with the incorrect decision. The expected node value is a weighted average of an unbiased estimate (based on the correct decision) and an estimate which is biased low (based on the incorrect decision). The net effect is an estimate which is biased low.

Rather than using the simple low estimator just described, we propose a modification. At each node, we use branch 1 to estimate the continuation value and the other $b-1$ branches to estimate the exercise decision. This process is repeated $b - 1$ times, using branch 2 to estimate the continuation value, then branch 3, etc. The b values obtained are averaged to determine the option value estimate at the node. A formal definition of this estimator is given in the next section, but first a numerical example is given in Fig. 6. Consider the bottom node at
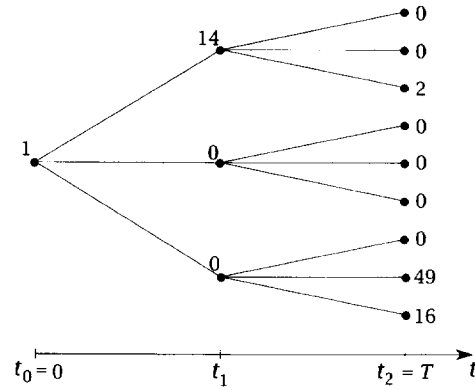
Fig. 5. Simple low estimate.
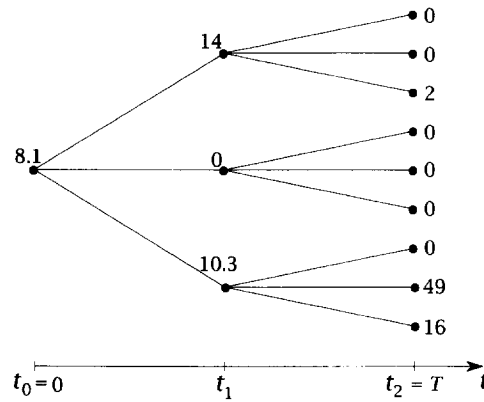
Fig. 6. θ ('low') estimate.

time t_1. As before, when branch 1 is used to determine the continuation value (and branches 2 and 3 are used to determine the decision 'continue'), the estimate is 0. When branch 2 is used to determine the continuation value (and branches 1 and 3 are used to determine the decision 'continue'), the estimate is 16. When branch 3 is used to determine the continuation value, branches 1 and 2 are used to determine the decision 'exercise', so the estimate is 15. These values are averaged to give an estimate for the node of $10.3(=(0 + 16 + 15)/3)$.

The resulting estimate at the initial node of the tree using this revised procedure is denoted θ. The θ estimator is biased downward but is also consistent. That is, $E[\theta] \leq C$ and θ converges to C. We refer to θ as the 'low' estimator.

3. Analysis of the estimators

Our objective in this section is to give a precise specification of our estimators, supplementing the description in Section 2, and to state their theoretical properties. Proofs are deferred to Appendices A and B.

We use the following notation:

- Time is indexed by $t = 0, 1, \ldots, T$. This is a slight abuse of notation, for we really mean that $t = t_0, t_1, \ldots, t_d$, with $0 = t_0 < t_1 < \cdots < t_d = T$, where t_i is the time of the ith exercise opportunity.
- $\{S_t : t = 0, 1, \ldots, T\}$ is a (possibly vector-valued) risk-neutralized Markov chain recording all state variables.
- 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.
- $h_t(s)$ is the payoff from exercise at time t in state s.
- $f_T(s) = h_T(s)$, i.e., at expiration the option is worth the payoff from immediate exercise.
- $g_t(s) = E[e^{-R_{t+1}} f_{t+1}(S_{t+1}) \mid S_t = s]$ is the continuation value at time t in state s.
- $f_t(s) = \max\{h_t(s), g_t(s)\}$ is the option value at time t in state s.

This framework is sufficiently general to encompass most pricing models which allow early exercise opportunities. The Markov assumption is not an essential restriction since we allow multiple state variables. Indeed, some of the variables could incorporate information about the past (for example, whether or not a barrier has been crossed, or the maximum stock price to date), recorded to eliminate path dependence. The framework is general enough to allow for stochastic interest rates, stochastic volatilities, and similar features.

A random tree with b branches per node is represented by the array

$$\{S_t^{i_1 \cdots i_t} : t = 0, 1, \ldots, T; \ i_j = 1, \ldots, b; \ j = 1, \ldots, t\}.$$

See Fig. 2 for an illustration. The joint distribution of the elements of this array is specified as follows: S_0 is the fixed initial state; $S_{t+1}^{i_1 \cdots i_t j}$, $j = 1, \ldots, b$, are conditionally independent of each other and of all $S_u^{i'_1 \cdots i'_u}$ with $u < t$ or $i'_t \neq i_t$, given $S_t^{i_1 \cdots i_t}$, and given $S_t^{i_1 \cdots i_t}$, each $S_{t+1}^{i_1 \cdots i_t j}$ has the distribution of $[S_t \mid S_{t-1} = S_{t-1}^{i_1 \cdots i_{t-1}}]$. Thus, each sequence

$$S_0, S_1^{i_1}, S_2^{i_1 i_2}, \ldots, S_T^{i_1 \cdots i_T}$$

is a realization of the Markov chain $\{S_t : t = 0, 1, \ldots, T\}$, and two such sequences evolve independently of each other once they differ in some i_t. To lighten notation, we omit superscripts when doing so does not introduce ambiguity.

3.1. The high estimator Θ

The high estimator Θ is defined recursively by

$$\Theta_T^{i_1\cdots i_T} = f_T(S_T^{i_1\cdots i_T}) \tag{2}$$

and

$$\Theta_t^{i_1\cdots i_t} = \max\left\{ h_t(S_t^{i_1\cdots i_t}), \frac{1}{b} \sum_{j=1}^{b} e^{-R_{t+1}^{i_1\cdots i_t j}} \Theta_{t+1}^{i_1\cdots i_t j} \right\}, \tag{3}$$

for $t = 0, \ldots, T - 1$. At each node, this estimator chooses the maximum of the early exercise payoff and the continuation value estimated from all successor nodes. (Here and throughout, a quantity subscripted by $t = 0$ is understood to have no superscript; thus, $\Theta_t^{i_1\cdots i_t}$ becomes simply Θ_0 at $t = 0$.) This estimator depends on the branching parameter and we sometimes include b as an explicit argument. Let $\overline{\Theta}_0$ denote the sample mean of n independent replications of Θ_0. Define the p-norm of the random variable X by $(E|X|^p)^{1/p}$.

Theorem 1 (High-estimator consistency). *Suppose* $E[|h_t(S_t)|^{p'}] < \infty$ *for all* t, *for some* $p' > 1$. *Then* $\overline{\Theta}_0(b)$ *converges to* $f_0(S_0)$ *in* p-norm, *for any* $0 < p < p'$, *as* $b \to \infty$ *with* n *arbitrary* (n *may or may not increase to infinity*). *In particular,* $\overline{\Theta}_0(b)$ *converges to* $f_0(S_0)$ *in probability and is thus a consistent estimator of the option value.*

A consequence of this result is that

$$E[\Theta_0(b)] \to f_0(S_0)$$

as $b \to \infty$, so the estimator is asymptotically unbiased. For finite b, it is useful to know that the bias is always positive:

Theorem 2 (High-estimator bias). *The high estimator is indeed biased high, i.e.,*

$$E[\Theta_0(b)] \geq f_0(S_0)$$

for all b.

3.2. The low estimator θ

The low estimator θ is defined recursively as follows. First let

$$\theta_T^{i_1\cdots i_T} = f_T(S_T^{i_1\cdots i_T}). \tag{4}$$

Next define

$$
\eta_t^{i_1\cdots i_t,j} =
\begin{cases}
h_t(S_t^{i_1\cdots i_t}) & \text{if } h_t(S_t^{i_1\cdots i_t}) \geq \dfrac{1}{b-1}\displaystyle\sum_{\substack{i=1 \\ i\neq j}}^{b} e^{-R_{t+1}^{i_1\cdots i_t i}}\,\theta_{t+1}^{i_1\cdots i_t,i}, \\[4mm]
e^{-R_{t+1}^{j}}\,\theta_{t+1}^{i_1\cdots i_t,j} & \text{if } h_t(S_t^{i_1\cdots i_t}) < \dfrac{1}{b-1}\displaystyle\sum_{\substack{i=1 \\ i\neq j}}^{b} e^{-R_{t+1}^{i_1\cdots i_t i}}\,\theta_{t+1}^{i_1\cdots i_t,i}
\end{cases}
\tag{5}
$$

for $j = 1,\ldots,b$. Then let

$$
\theta_t^{i_1\cdots i_t} = \frac{1}{b}\sum_{j=1}^{b} \eta_t^{i_1\cdots i_t,j},
\tag{6}
$$

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

This estimator is also consistent.

Theorem 3 (Low-estimator consistency). *Suppose* $P(h_t(S_t) \neq g_t(S_t)) = 1$ *for all* t. *Then Theorem 1 holds for the low estimator as well.*

The additional condition imposed here, ensuring that the optimal exercise policy is almost surely never indifferent between continuation and immediate exercise, greatly simplifies the analysis of the estimator and seems harmless in practice.

As in the case of the high estimator, Theorem 3 implies that $\theta_0(b)$ is asymptotically unbiased. It is convenient that for finite b the biases of the two estimators have opposite signs, as the following result ensures:

Theorem 4 (Low-estimator bias). *The bias of the low estimator is negative, i.e.,*

$$
E[\theta_0(b)] \leq f_0(S_0)
$$

for all b.

It is easy to define generalizations of the low estimator. For example, at each node, use b_1 branches to determine the exercise decision and b_2 branches to evaluate the resulting payoff, with $b_1 + b_2 = b$ and both $b_i \to \infty$ as $b \to \infty$. These alternative estimators are consistent and biased low.

3.3. Comparison of the estimators

Our last result on the estimators orders their values:

Theorem 5. On every realization of the array $\{S_t^{i_1\cdots i_t}: t = 0,1,\ldots,T;\ i_j = 1,\ldots,b;\ j = 1,\ldots,t\}$, *the low estimator is less than or equal to the high*

estimator. In short,

$$\theta_t^{i_1\cdots i_t} \le \Theta_t^{i_1\cdots i_t}$$

with probability one, for all $i_1,\ldots,i_t$ and all $t = 0, 1, \ldots, T$.

4. Numerical results

In this section we provide some numerical results to illustrate the simulation method. Implementation details are given in Appendix C. We begin by pricing a standard American call option on a single asset which pays continuous dividends and whose price is governed by a geometric Brownian motion process. In particular, we assume that the risk neutralized price of the underlying asset, S_t, satisfies the stochastic differential equation

$$\mathrm{d}S_t = S_t[(r - \delta)\,\mathrm{d}t + \sigma\,\mathrm{d}z_t], \tag{7}$$

where z_t is a standard Brownian motion process. In Eq. (7), r is the riskless interest rate, δ is the dividend rate, and $\sigma > 0$ is the volatility parameter. Under the risk neutral measure, $\ln(S_i/S_{i-1})$ is normally distributed with mean $(r - \delta - \sigma^2/2)(t_i - t_{i-1})$ and variance $\sigma^2(t_i - t_{i-1})$. Given S_{i-1}, S_i can be simulated using

$$S_i = S_{i-1}\,\mathrm{e}^{(r-\delta-\sigma^2/2)(t_i-t_{i-1})+\sigma\sqrt{t_i-t_{i-1}}Z}, \tag{8}$$

where Z is a standard normal random variable.[5] The parameters were chosen so the early exercise opportunity would have significant value.[6] The results are given in Tables 1 and 2.

In the tables, the confidence intervals are given by

$$[\max\{(S_0 - K)^+, \theta - z_{\alpha/2}s(\theta)/\sqrt{n}\}, \Theta + z_{\alpha/2}s(\Theta)/\sqrt{n}], \tag{9}$$

where $z_{\alpha/2}$ is the $1 - \alpha/2$ quantile of the standard normal distribution, and $s(\theta)$ and $s(\Theta)$ are the sample standard deviations of θ and Θ, respectively. We take the upper confidence limit from the high estimator. The lower confidence limit is taken from the low estimator, except that the lower limit is truncated at the

---

[5] If dividends are discrete, one approach is to replace Eq. (8) by

$$S_i = S_{i-1}\,\mathrm{e}^{(r-\sigma^2/2)(t_i-t_{i-1})+\sigma\sqrt{t_i-t_{i-1}}Z} - D_i,$$

where D_i is the dividend paid at time t_i.

[6] McDonald and Schroder (1990) show that in this setting the value of an American call option is equal to the value of an American put with the following change of parameters: $S_t \to K$, $K \to S_t$, $r \to \delta$, and $\delta \to r$. So the call option results in Tables 1 and 2 can also be viewed as put option results with $r = 10\%$ and $\delta = 5\%$.

Table 1
American call option on a single asset

S	Low est θ	Std err of θ	High est Θ	Std err of Θ	90% confidence bounds	Point est	True value	Rel error
70	0.115	0.004	0.117	0.004	[0.108, 0.124]	0.116	0.121	4.36%
80	0.649	0.016	0.662	0.016	[0.624, 0.688]	0.656	0.670	2.15%
90	2.251	0.039	2.316	0.040	[2.187, 2.382]	2.283	2.303	0.87%
100	5.628	0.076	5.824	0.078	[5.502, 5.952]	5.726	5.731	0.09%
110	10.988	0.156	11.603	0.113	[10.732, 11.789]	11.296	11.341	0.40%
120	19.743	0.139	20.329	0.069	[20.000, 20.442]	20.164	20.000	0.82%
130	29.763	0.124	30.154	0.049	[30.000, 30.235]	30.077	30.000	0.26%

Option parameters: $K = 100$, $r = 0.05$, $\delta = 0.10$, $T = 1.0$, $\sigma = 0.2$, and four exercise opportunities at times $0, T/3, 2T/3$, and T. The 'true value' corresponds to the same four exercise opportunities. Simulation parameters: $n = 100$, $b = 50$.

Table 2
American call option on a single asset – with control variate

S	Low est θ	Std err of θ	High est Θ	Std err of Θ	90% Confidence bounds	Point est	True value	Rel error
70	0.121	0.000	0.122	0.000	[0.120, 0.122]	0.121	0.121	0.08%
80	0.663	0.001	0.676	0.001	[0.662, 0.677]	0.670	0.670	0.07%
90	2.268	0.005	2.334	0.002	[2.260, 2.337]	2.301	2.303	0.11%
100	5.631	0.013	5.828	0.007	[5.611, 5.840]	5.730	5.731	0.02%
110	10.957	0.085	11.576	0.017	[10.816, 11.605]	11.266	11.341	0.66%
120	19.742	0.139	20.306	0.043	[20.000, 20.376]	20.153	20.000	0.76%
130	29.773	0.122	30.138	0.038	[30.000, 30.200]	30.069	30.000	0.23%

Option parameters: $K = 100$, $r = 0.05$, $\delta = 0.10$, $T = 1.0$, $\sigma = 0.2$, and four exercise opportunities at times $0, T/3, 2T/3$, and T. The 'true value' corresponds to the same four exercise opportunities. Simulation parameters: $n = 100$, $b = 50$. The European option value is used as a control.

immediate exercise value of the option, which is a trivial lower bound on the true value. The point estimate is given by the simple average

$$0.5 \max\{(S_0 - K)^+, \theta\} + 0.5\Theta.$$

With a finite number of exercise opportunities, the true value of the call option can be obtained from the formula in Geske and Johnson (1984). The numbers in the 'Rel error' column are defined by

$$\left| \frac{C - \hat{C}}{C} \right|,$$

where C is the true option value and $\hat{C}$ is the simulation point estimate. (The reported relative errors are based on more significant digits than are shown in the tables.)

The results in Table 1 are consistent with the theoretical developments in the previous section. Indeed, the θ estimator is biased low, the Θ estimator is biased high, and $\theta \leq \Theta$. For these parameters, the relative error ranges from 0.09% to 4.36%. Each row in Table 1 can be computed in about 1 min on a PC with a 133 MHz Pentium processor. The results in Table 2 are obtained using the control variate technique. See, e.g., Section 11.4 of Law and Kelton (1991) for a discussion of this technique. The European option value is used as the control variate in Table 2. Table 2 shows significant reductions in the standard errors of the estimates compared to the corresponding entries in Table 1, except for deep in-the-money options.[7] In all cases the relative error is less than 1% with the control variate technique.

Because of the bias of the estimators and the definition in (9), the reported confidence intervals are *conservative*. That is, the true value will fall in the confidence interval more times, on average, than suggested by the confidence level. For example, for the option in Table 1 with $S = K = 100$ we ran 1000 simulation trials with different random number seeds. We then computed confidence intervals with $z_{\alpha/2} = 0.5$, corresponding to a 38% confidence level, and found that over 90% of the intervals contained the true value. Similarly, with $z_{\alpha/2} = 1.0$, corresponding to a 68% confidence level, 96% of the intervals contained the true value. With $z_{\alpha/2} = 1.645$, corresponding to a 90% confidence level, 99% of the intervals contained the true value. Results for other securities will depend on the model and simulation parameters, which influence the bias in the estimators. The reported confidence intervals will be more conservative when the biases are larger.

4.1. Higher-dimensional results

Next we price American call options on the maximum of k assets. The payoff upon exercise of this option is $(\max_{i=1,...,k} S^i - K)^+$. Under the risk neutral measure asset prices are assumed to follow correlated geometric Brownian motion processes, i.e.,

$$dS_t^i = S_t^i[(r - \delta_i)\,dt + \sigma_i\,dz_t^i], \tag{10}$$

where z_t^i is a standard Brownian motion process and the instantaneous correlation of z^i and z^j is ρ_{ij}. For simplicity, in our numerical results we take $\delta_i = \delta$ and $\rho_{ij} = \rho$ for all $i,j = 1,...,k$ and $i \neq j$.

---

[7] A control variate is most effective when it is highly correlated with the quantity being estimated. The European option and American option payoffs are less correlated when the option is deep in-the-money.

Table 3
American max-option on two assets

S	Low est Θ	Std err of Θ	High est Θ	Std err of Θ	90% confidence bounds	Point est	True value	Rel error
70	0.247	0.008	0.250	0.008	[0.234, 0.263]	0.249	0.237	4.90%
80	1.225	0.021	1.246	0.021	[1.191, 1.281]	1.235	1.259	−1.85%
90	4.019	0.049	4.116	0.051	[3.938, 4.200]	4.067	4.077	−0.24%
100	9.228	0.093	9.487	0.095	[9.075, 9.644]	9.358	9.361	−0.03%
110	16.775	0.132	17.241	0.134	[16.558, 17.461]	17.008	16.924	0.50%
120	25.747	0.141	26.369	0.140	[25.515, 26.599]	26.058	25.980	0.30%
130	35.541	0.194	36.254	0.200	[35.221, 36.583]	35.898	35.763	0.38%

Option parameters: $S^1 = S^2 = S$ as indicated in the table. Also $K = 100$, $r = 0.05$, $\delta = 0.10$, $T = 1.0$, $\sigma = 0.2$, $\rho = 0.3$, and four exercise opportunities at times $0, T/3, 2T/3$, and T. The 'true value' corresponds to the same four exercise opportunities. Simulation parameters: $n = 100$, $b = 50$.

Tables 3 and 4 give results for two assets and Tables 5 and 6 give results for five assets. For the two asset case, the true value of the option can be approximated using, for example, the multivariate algorithm of Boyle et al. (1989) or Kamrad and Ritchken (1991).[8] The European max-option is used as a control variate in Table 4.[9] The European max-option value is used as the control variate in Table 4. The confidence intervals are considerably narrower with the control and all relative errors are less than one percent in Table 4. Each row of Table 3 and 4 can be computed in about two minutes on a PC with a 133 MHz Pentium processor.

Results for five assets are reported in Tables 5 and 6. Relative errors are not reported because the true value is unknown. A multinomial lattice with k assets and n time steps has on the order of n^k terminal nodes. With $k = 5$ the computations are prohibitive for n as small as 50. And even if the computations could be done, the resulting value would not be very accurate.[10] However, as shown in Tables 5 and 6, the simulation method is able to produce valid confidence intervals in a reasonable amount of computing time. As before, the confidence interval widths are considerably narrowed with the use of the control variate. The confidence interval halfwidths are within 1% of the midpoint of the interval.

[8] In Tables 3 and 4 we approximated the true value using the algorithm of Kamrad and Ritchken (1991). Let x denote the lattice value with 600 time steps and y the value with 1200 time steps. Our final estimate, z, is based on these two values and two-point Richardson extrapolation, i.e., $z = 2y - x$. Additional comparisons with the known European max-option value indicated that this procedure resulted in an error of about 0.001.

[9] A formula for this option is given in Johnson (1987). For the k-asset case, cumulative k-variate normal probabilities need to be evaluated. We used the algorithm in Schervish (1984, 1985) for these computations.

[10] Our computational experiments using lattice methods with $k = 3$ and $n = 50$ indicate errors on the order of 0.10.

Table 4
American max-option on two assets – with control variate

S	Low est θ	Std err of θ	High est Θ	Std err of Θ	90% confidence bounds	Point est	True value	Rel error
70	0.235	0.000	0.238	0.000	[0.235, 0.239]	0.237	0.237	−0.03%
80	1.250	0.001	1.271	0.001	[1.248, 1.273]	1.261	1.259	0.15%
90	4.035	0.005	4.133	0.004	[4.027, 4.139]	4.084	4.077	0.16%
100	9.245	0.011	9.505	0.008	[9.227, 9.518]	9.375	9.361	0.16%
110	16.718	0.020	17.183	0.015	[16.685, 17.207]	16.951	16.924	0.16%
120	25.682	0.026	26.304	0.019	[25.639, 26.336]	25.993	25.980	0.05%
130	35.452	0.031	36.162	0.021	[35.400, 36.197]	35.807	35.763	0.12%

Option parameters: $S^1 = S^2 = S$ as indicated in the table. Also $K = 100$, $r = 0.05$, $\delta = 0.10$, $T = 1.0$, $\sigma = 0.2$, $\rho = 0.3$, and four exercise opportunities at times $0, T/3, 2T/3$, and T. The 'true value' corresponds to the same four exercise opportunities. Simulation parameters: $n = 100$, $b = 50$. The European max-option value is used as a control.

Table 5
American max-option on five assets

S	Low est θ	Std err of θ	High est Θ	Std err of Θ	90% confidence bounds	Point est
70	0.556	0.012	0.561	0.012	[0.536, 0.581]	0.558
80	2.641	0.038	2.681	0.039	[2.578, 2.746]	2.661
90	7.799	0.076	7.941	0.078	[7.674, 8.069]	7.870
100	15.832	0.120	16.118	0.122	[15.634, 16.319]	15.975
110	25.607	0.151	26.022	0.154	[25.359, 26.276]	25.815
120	36.379	0.157	36.847	0.158	[36.121, 37.107]	36.613
130	47.053	0.163	47.618	0.164	[46.785, 47.888]	47.335

Option parameters: $S^i = S$, $i = 1, \ldots, 5$ as indicated in the table. Also $K = 100$, $r = 0.05$, $\delta = 0.10$, $T = 1.0$, $\sigma = 0.2$, $\rho = 0.3$, and four exercise opportunities at times $0, T/3, 2T/3$, and T. Simulation parameters: $n = 100$, $b = 50$.

The importance of variance reduction is clearly demonstrated in these results. Broadie and Glasserman (1995) investigate other enhancements of the basic method described in this paper. In particular, they investigate pruning techniques to reduce the number of nodes in the tree. For example, if a European value is easily computed, then that value can be used at the penultimate time step. At an intermediate stage in the tree, branching can be eliminated if the optimal decision is known or easily computed (for example, if an option is out-of-the-money then exercise is suboptimal). They also investigate bootstrapping to estimate and reduce the bias in the low and high estimators. Numerical results indicate substantial benefits from the pruning techniques. Although bootstrapping succeeds in reducing bias, the increased computational effort largely offsets the gain from the improved estimates.

Table 6
American max-option on five assets – with control variate

S	Low est θ	Std err of θ	High est Θ	Std err of Θ	90% confidence bounds	Point est
70	0.552	0.000	0.557	0.000	[0.551, 0.557]	0.554
80	2.690	0.002	2.731	0.001	[2.687, 2.733]	2.711
90	7.752	0.005	7.893	0.004	[7.744, 7.899]	7.823
100	15.761	0.010	16.046	0.007	[15.745, 16.058]	15.904
110	25.600	0.013	26.015	0.009	[25.579, 26.030]	25.808
120	36.267	0.017	36.734	0.012	[36.238, 36.753]	36.500
130	47.122	0.019	47.689	0.013	[47.091, 47.710]	47.405

Option parameters: $S^i = S$, $i = 1, \ldots, 5$ as indicated in the table. Also $K = 100$, $r = 0.05$, $\delta = 0.10$, $T = 1.0$, $\sigma = 0.2$, $\rho = 0.3$, and four exercise opportunities at times $0, T/3, 2T/3$, and T. Simulation parameters: $n = 100$, $b = 50$. The European max-option value is used as a control.

Additional variance reduction techniques are investigated in Broadie, Glasserman, and Jain (1996). In particular, they find that Latin hypercube sampling works particularly well in this setting. For example, in Table 6 the 90% confidence interval in the row $S = 100$ is [15.745, 16.058]. Using Latin hypercube sampling in addition to the control variate gives a 90% confidence interval of [15.865, 15.907]. The halfwidth of the interval is 0.13% of the midpoint of the interval.

4.2. Discrete versus continuous exercise

Because of the discrete nature of the simulation method our numerical results are limited to a finite number of exercise opportunities. In addition, the computational cost of our method is exponential in the number of exercise opportunities. With current technology, this limits the number of exercise opportunities that we can consider.[11] In practice, there are many securities with a limited number of exercise opportunities. For example, in the case of a call option on a stock which pays discrete dividends, it is well known that exercise is only optimal just prior to ex-dividend dates. For stock options whose maturity is 1 yr or less, this means that there will typically be at most four opportunities for optimal early exercise. Another example is over-the-counter options which often have 'structured exercise opportunities', i.e., where exercise is allowed only at a finite number of prespecified dates. As a final example, Jamshidian (1996, p. 20) writes 'In the Libor and swap markets there are hardly any options with continuous exercise.'

In many cases, however, the security under consideration allows for continuous exercise. To estimate this value under our method, we need to use an extrapolation procedure. Geske and Johnson (1984) were the first to investigate Richardson

---

[11] Conveniently, the method is very well suited to parallel computing.

extrapolation in a finance context. Let $C(d)$ denote the price of an option with $d + 1$ exercise opportunities at times $t_0, \ldots, t_d$ and let C denote the price with continuous exercise. For small d, $C(d)$ may be a poor estimate of C. However, an extrapolated estimate of the form $\sum_{i=1}^{d} a_i C(i)$ may provide an excellent estimate of C. Richardson exptrapolation offers a particular recipe for choosing (in advance) the coefficients a_i. Richardson exptrapolation can also be used to extrapolate to an arbitrary, finite number of exercise opportunities.

Broadie et al. (1996) investigate this technique and find that it substantially improves the estimate of the continuous exercise price. For example, consider the case of two assets with $S^1 = S^2 = 100$. The true value with continuous exercise is $C = 9.637$ (which is greater than the value of 9.361 with only four exercise opportunities). The European value is $C(1) = 8.932$. Simulation estimates give $C(2) = 9.248$ and $C(3) = 9.357$. The relative error of the $C(3)$ estimate is -2.9%. The estimate obtained by extrapolating $C(1)$, $C(2)$, and $C(3)$ is 9.581 which has a relative error of 0.6%. Further, similar results are reported in Broadie et al. (1996).

5. Conclusions

In this paper we have shown that there is essentially no unbiased simulation estimator of the value of early exercise. To circumvent this difficulty, we developed a method which generates two estimates, one biased high and one biased low. Both estimates are asymptotically unbiased and converge to the true price. They can be combined to give a valid confidence interval on the true price. Because of estimator bias, the confidence intervals obtained are conservative: actual coverage tends to exceed the nominal stated coverage.

The method is most promising for pricing American-style securities with multiple state variables. Although our estimators were developed for options with two decisions, exercise or not, they are easily extended to a finite number of decisions. For example, the management of a firm may have the opportunity to choose between initiating a project, expanding or contracting to one of several levels, or abandoning a project. Hence, this method could be particularly useful for the valuation of complex real options.

Preliminary computational evidence given in the paper is quite encouraging. For a five asset problem, relative errors of less than one percent are obtained with modest computational effort. Our proposed algorithm can be naturally parallelized in several ways and this could lead to large improvements in computation time on parallel machines, e.g., on a network of workstations.

This work can be extended in several directions. Variations of the low estimator, e.g., using b_1 branches to determine the exercise decision and b_2 branches to evaluate the resulting payoff, remain to be explored. The number of branches per node does not need to be constant throughout the tree. The convergence rate

of the algorithm and the effect of the choice of n and b on the error remain to be explored. Alternative variance reduction techniques, including other control variates, could be tested. Quasi-Monte Carlo methods, also termed low discrepancy methods, are another promising avenue of exploration (see Birge, 1994; Joy et al., 1996; Owen, 1994, 1995; Paskov and Traub, 1995). Additional computational testing on other American-style securities remains to be done.

Appendix A

In this appendix, we prove Theorems 1–5. First, we introduce some additional notation. With X a random variable, we write $\|X\|$ for the p-norm $(\mathrm{E}|X|^p)^{1/p}$ of X, the value of p depending on the context. The notation $\|X\|_{S_t}$ indicates the *conditional* norm $(\mathrm{E}[|X|^p|S_t])^{1/p}$. This coincides with the unconditional norm when $t = 0$, because S_0 is deterministic. A quantity like Θ_t appearing without a superscript indicates a generic copy of $\Theta_t^{i_1 \cdots i_t}$, and Θ_{t+1}^j similarly indicates a generic copy of $\Theta_{t+1}^{i_1 \cdots i_t j}$. A statement like '$\Theta_{t+1}^j$, $j = 1, \ldots, b$, are conditionally independent given S_t' is short for '$\Theta_{t+1}^{i_1 \cdots i_t j}$, $j = 1, \ldots, b$, are conditionally independent given $S_t^{i_1 \cdots i_t}$, for all $i_1, \ldots, i_t$'.

We need the following preliminary result:

Lemma A.1.   If $\|h_t(S_t)\| < \infty$ for all t, for some $p \geq 1$, then the following are also finite for all $0 \leq t_1 \leq t_2 \leq T$:

(i) $\|f_{t_2}(S_{t_2})\|_{S_{t_1}}$;

(ii) $\sup_b \|\Theta_{t_2}(b)\|_{S_{t_1}}$;

(iii) $\sup_b \|\theta_{t_2}(b)\|_{S_{t_1}}$.

Proof. If every $h_t(S_t)$ has finite pth moment, then each $\|h_{t_2}(S_{t_2})\|_{S_{t_1}}$ is finite and so are $\|f_{t_2}(S_{t_2})\|_{S_{t_1}}$ and $\|g_{t_2}(S_{t_2})\|_{S_{t_1}}$ because the max, discounting, and conditional expectation operations preserve finiteness of moments. For $\|\Theta_{t_2}(b)\|_{S_{t_1}}$, fix t_1 and proceed backwards by induction on t_2 from T to t_1. At $t_2 = T$, (ii) follows from (i). At $t_2 < T$, we have

$$\sup_b \|\Theta_{t_2}(b)\|_{S_{t_1}} \leq \|h_{t_2}(S_{t_2})\|_{S_{t_1}} + \sup_b \left\| \frac{1}{b} \sum_{j=1}^{b} e^{-R_{t_2+1}^j} \Theta_{t_2+1}^j(b) \right\|_{S_{t_1}}$$

$$\leq \|h_{t_2}(S_{t_2})\|_{S_{t_1}} + \sup_b \|\Theta_{t_2+1}(b)\|_{S_{t_1}},$$

which is then also finite. The argument for (iii) is essentially the same. $\square$

Proof of Theorem 1. Take $n = 1$. We will prove that convergence holds for each t (not just $t = 0$), proceeding backwards by induction. We claim that if

$\|\Theta_{t+1}(b) - f(S_{t+1})\|_{S_{t+1}} \to 0$, then $\|\Theta_t(b) - f(S_t)\|_{S_t} \to 0$. Since $\Theta_T \equiv f(S_T)$, the theorem will be proved once we verify this claim.

We have

$$\|\Theta_t(b) - f(S_t)\|_{S_t}$$

$$= \left\| \max\left\{ h_t(S_t), \frac{1}{b}\sum_{i=1}^{b} e^{-R_{t+1}^i} \Theta_{t+1}^i(b)\right\} - \max\{h_t(S_t), g_t(S_t)\} \right\|_{S_t}$$

$$\leq \left\| \frac{1}{b}\sum_{i=1}^{b} e^{-R_{t+1}^i} \Theta_{t+1}^i(b) - g_t(S_t) \right\|_{S_t}$$

$$\leq \left\| \frac{1}{b}\sum_{i=1}^{b} e^{-R_{t+1}^i} \{\Theta_{t+1}^i(b) - f_{t+1}(S_{t+1}^i)\} \right\|_{S_t}$$

$$+ \left\| \frac{1}{b}\sum_{i=1}^{b} e^{-R_{t+1}^i} f_{t+1}(S_{t+1}^i) - g_t(S_t) \right\|_{S_t}$$

$$\equiv B_1 + B_2, \quad \text{say.}$$

Given S_t, the $e^{-R_{t+1}^i} f(S_{t+1}^i)$, $i = 1, \ldots, b$, are i.i.d. with mean $g(S_t)$ and finite p-norm. By Theorem I.4.1 of Gut (1988), we thus have $B_2 \to 0$. For B_1 we have

$$B_1 \leq \|\Theta_{t+1}(b) - f_{t+1}(S_{t+1})\|_{S_t}.$$

The induction hypothesis implies that

$$\|\Theta_{t+1}(b) - f_{t+1}(S_{t+1})\|_{S_{t+1}} \to 0.$$

By a standard condition for uniform integrability (see Gut, 1988, p. 178), we also have

$$\|\Theta_{t+1}(b) - f_{t+1}(S_{t+1})\|_{S_t} \to 0 \tag{A.1}$$

if

$$\sup_b E[|\Theta_{t+1}(b) - f_{t+1}(S_{t+1})|^{p+\varepsilon}|S_t] < \infty \tag{A.2}$$

for some $\varepsilon > 0$. From Lemma A.1, we know that in fact both

$$\sup_b E[|\Theta_{t+1}(b)|^{p+\varepsilon}|S_t] \quad \text{and} \quad E[|f_{t+1}(S_{t+1})|^{p+\varepsilon}|S_t]$$

are finite, so (A.2) holds and therefore (A.1) holds.

The same argument applies for the average of n independent replications of $\Theta_0(b)$, no matter how n changes with b. $\square$

Proof of Theorem 2. We prove more generally that $E[\Theta_t|S_t] \geq f_t(S_t)$ for $t = 0$, $1, \ldots, T$. The proof proceeds by backwards induction. By definition, $\Theta_T = f_T(S_T)$ so trivially $E[\Theta_T|S_T] \geq f_T(S_T)$. Our induction hypothesis is $E[\Theta_{t+1}|S_{t+1}] \geq f_{t+1}(S_{t+1})$. By Jensen's inequality and the definition of Θ_t, we have

$$E[\Theta_t|S_t] \geq \max\{h_t(S_t), E[e^{-R_{t+1}}\Theta_{t+1}|S_t]\}$$

$$= \max\{h_t(S_t), E[e^{-R_{t+1}}E[\Theta_{t+1}|S_{t+1}]|S_t]\}$$

$$\geq \max\{h_t(S_t), E[e^{-R_{t+1}}f_{t+1}(S_{t+1})|S_t]\}$$

$$= \max\{h_t(S_t), g_t(S_t)\} = f_t(S_t). \qquad \square$$

Proof of Theorem 3. The proof is similar to that of Theorem 1 and proceeds backwards by induction on t.

Convergence is automatic at $t = T$. Suppose now that $\|\theta_{t+1}(b) - f(S_{t+1})\|_{S_{t+1}} \to 0$. Define

$$Y_t^j(b) = \frac{1}{b-1} \sum_{\substack{i=1 \\ i \neq j}}^{b} e^{-R_{t+1}^i} \theta_{t+1}^i, \qquad (A.3)$$

and recall that $g_t(S_t) \neq h_t(S_t)$, with probability one. We claim the following hold:

(a) $\|(1/b)\sum_{j=1}^{b} e^{-R_{t+1}^j} \theta_{t+1}^j(b) - g(S_t)\|_{S_t} \to 0$;

(b) $\|Y_t^j(b) - g(S_t)\|_{S_t} \to 0$.

(c) $\|\mathbf{1}_{\{h_t(S_t) \geq Y_t^j(b)\}} - \mathbf{1}_{\{h_t(S_t) > g_t(S_t)\}}\|_{S_t} \to 0$.

The proof of (a) is just the same as the proof of a corresponding step in the proof of Theorem 1. The same argument applies to (b), because the estimators in (a) and (b) differ only in the omission of one term in $Y_t^j(b)$. For (c), suppose that $h_t(S_t) < g_t(S_t)$. Then

$$\|\mathbf{1}_{\{h_t(S_t) \geq Y_t^j(b)\}} - \mathbf{1}_{\{h_t(S_t) > g_t(S_t)\}}\|_{S_t} = \|\mathbf{1}_{\{h_t(S_t) \geq Y_t^j(b)\}}\|_{S_t}$$

$$= P(Y_t^j(b) \leq h_t(S_t)|S_t)^{1/p} \to 0,$$

because (b) holds and convergence in p-norm implies convergence in probability. The same argument applies if $h_t(S_t) > g_t(S_t)$.

We now claim that from (a) and (c) it follows that

$$\left\| \frac{1}{b} \sum_{j=1}^{b} e^{-R_{t+1}^j} \theta_{t+1}^j(b) \mathbf{1}_{\{h_t(S_t) < Y_t^j(b)\}} - g_t(S_t)\mathbf{1}_{\{h_t(S_t) < g_t(S_t)\}} \right\|_{S_t} \to 0, \qquad (A.4)$$

but we postpone a proof of this claim until the end. An immediate consequence of (c) is that

$$\left\| h_t(S_t)\mathbf{1}_{\{h_t(S_t) \geq Y_t^j(b)\}} - h_t(S_t)\mathbf{1}_{\{h_t(S_t) > g_t(S_t)\}} \right\|_{S_t} \to 0. \tag{A.5}$$

Now observe that

$$\left\| \theta_t(b) - f(S_t) \right\|_{S_t}$$

$$= \left\| \frac{1}{b}\sum_{j=1}^{b} \eta_t^j(b) - f(S_t) \right\|_{S_t}$$

$$= \left\| \frac{1}{b}\sum_{j=1}^{b}\left(h_t(S_t)\mathbf{1}_{\{h_t(S_t) \geq Y_t^j(b)\}} + e^{-R_{t+1}^j}\theta_{t+1}^j(b)\mathbf{1}_{\{h_t(S_t) < Y_t^j(b)\}} \right) - f_t(S_t) \right\|_{S_t}$$

$$\leq \left\| \frac{1}{b}\sum_{j=1}^{b} e^{-R_{t+1}^j}\theta_{t+1}^j(b)\mathbf{1}_{\{h_t(S_t) < Y_t^j(b)\}} - g_t(S_t)\mathbf{1}_{\{h_t(S_t) < g_t(S_t)\}} \right\|_{S_t}$$

$$+ \left\| h_t(S_t)\mathbf{1}_{\{h_t(S_t) \geq Y_t^j(b)\}} - h_t(S_t)\mathbf{1}_{\{h_t(S_t) > g_t(S_t)\}} \right\|_{S_t} \to 0$$

by (A.4) and (A.5).

It remains to prove (A.4). Using a self-explanatory simplification of the notation, what we need to show becomes

$$\left\| \frac{1}{b}\sum_{j=1}^{b} u_j(b)v_j(b) - uv \right\| \to 0.$$

Now,

$$\left\| \frac{1}{b}\sum_{j=1}^{b} u_j(b)v_j(b) - uv \right\| \leq \left\| \frac{1}{b}\sum_{j=1}^{b} [u_j(b)v_j(b) - u_j(b)v] \right\|$$

$$+ \left\| \frac{1}{b}\sum_{j=1}^{b} u_j(b)v - uv \right\|$$

$$\leq \|u_1(b)\| \cdot \|v_1(b) - v\| + \|v\| \cdot \left\| \frac{1}{b}\sum_{j=1}^{b} u_j(b) - u \right\|.$$

But $\|v_1(b) - v\| \to 0$ by (c), and $\|(1/b)\sum_{j=1}^{b} u_j(b) - u\| \to 0$ by (a). □

Proof of Theorem 4. We use backwards induction. By definition, $\theta_T = f_T(S_T)$ so trivially $\mathrm{E}[\theta_T|S_T] \leq f_T(S_T)$. Our induction hypothesis is $\mathrm{E}[\theta_{t+1}|S_{t+1}] \leq f_{t+1}(S_{t+1})$.

Let Y_t^j be as in (A.3). From the definition of θ_t, notice that $\mathrm{E}[\theta_t|S_t] = \mathrm{E}[\eta_t^j|S_t]$ for any $j = 1,\ldots,b$. Since Y_t^j is conditionally independent of θ_{t+1}^j given S_t, we have

$$\mathrm{E}[\eta_t^j|S_t] = \mathrm{E}[h_t(S_t)\mathbf{1}_{\{h_t(S_t) \geq Y_{t+1}^j\}}|S_t] + \mathrm{E}[e^{-R_{t+1}^j}\theta_{t+1}^j\mathbf{1}_{\{h_t(S_t) < Y_{t-1}^j\}}|S_t]$$

$$= h_t(S_t)P(h_t(S_t) \geq Y_{t+1}^j|S_t) + \mathrm{E}[e^{-R_{t-1}^j}\theta_{t+1}^j|S_t]P(h_t(S_t) < Y_{t+1}^j|S_t)$$

$$= h_t(S_t)p + \mathrm{E}[e^{-R_{t+1}^j}\theta_{t+1}^j|S_t](1 - p) \quad \text{say.}$$

But then

$$\mathrm{E}[\theta_t|S_t] = h_t(S_t)p + \mathrm{E}[e^{-R_{t+1}^j}\mathrm{E}[\theta_{t+1}^j|S_{t+1}]|S_t](1 - p)$$

$$\leq h_t(S_t)p + \mathrm{E}[e^{-R_{t+1}^j}f(S_{t+1})|S_t](1 - p)$$

$$= h_t(S_t)p + g_t(S_t)(1 - p)$$

$$\leq \max\{h_t(S_t), g_t(S_t)\} = f_t(S_t). \qquad \square$$

Proof of Theorem 5. The proof is by backwards induction. The induction starts since $\theta_T = \Theta_T = f_T(S_T)$. The induction hypothesis is $\theta_{t+1}^j \leq \Theta_{t+1}^j$, $j = 1,\ldots,b$.

Let Y_t^j be as in (A.3). If $Y_t^1,\ldots,Y_t^b$ are all less than or equal to $h_t(S_t)$, then $\eta_t^j = h_t(S_t)$, $j = 1,\ldots,b$, so $\theta_t = h_t(S_t) \leq \Theta_t$. Suppose, now, that at least one Y_t^j is greater than $h_t(S_t)$. Then

$$\frac{1}{b}\sum_{j=1}^{b}\eta_t^j = \frac{1}{b}\sum_{j=1}^{b}(h_t(S_t)\mathbf{1}_{\{h_t(S_t) \geq Y_t^j\}} + e^{-R_{t+1}^j}\theta_{t+1}^j\mathbf{1}_{\{h_t(S_t) < Y_t^j\}})$$

$$= \left(\frac{1}{b}\sum_{j=1}^{b}\mathbf{1}_{\{h_t(S_t) \geq Y_t^j\}}\right)h_t(S_t) + \left(\frac{1}{b}\sum_{j=1}^{b}\mathbf{1}_{\{h_t(S_t) < Y_t^j\}}\right)$$

$$\times \frac{\sum_{j=1}^{b}e^{-R_{t+1}^j}\theta_{t+1}^j\mathbf{1}_{\{h_t(S_t) < Y_t^j\}}}{\sum_{j=1}^{b}\mathbf{1}_{\{h_t(S_t) < Y_t^j\}}}$$

$$= ph_t(S_t) + (1 - p)\frac{\sum_{j=1}^{b}e^{-R_{t+1}^j}\theta_{t+1}^j\mathbf{1}_{\{h_t(S_t) < Y_t^j\}}}{\sum_{j=1}^{b}\mathbf{1}_{\{h_t(S_t) < Y_t^j\}}} \quad \text{say.} \qquad \text{(A.6)}$$

Without loss of generality, suppose that $Y_t^1, \ldots, Y_t^k$ are greater than $h_t(S_t)$ and $Y_t^{k+1}, \ldots, Y_t^b$ are less than or equal to $h_t(S_t)$. Then the ratio appearing in (A.6) equals $k^{-1} \sum_{j=1}^{k} e^{-R_{t+1}^j} \theta_{t+1}^j$. For any $i \leq k < j \leq b$, we have $Y_t^i > Y_t^j$ which implies $e^{-R_{t+1}^i} \theta_{t+1}^i \leq e^{-R_{t+1}^j} \theta_{t+1}^j$. Thus,

$$\max\{e^{-R_{t+1}^1}\theta_{t+1}^1, \ldots, e^{-R_{t+1}^k}\theta_{t+1}^k\} \leq \min\{e^{-R_{t+1}^{k+1}}\theta_{t+1}^{k+1}, \ldots, e^{-R_{t+1}^{k+1}}\theta_{t+1}^{k+1}\}.$$

But then

$$\frac{1}{k} \sum_{i=1}^{k} e^{-R_{t+1}^i}\theta_{t+1}^i \leq \frac{1}{b} \sum_{i=1}^{b} e^{-R_{t+1}^i}\theta_{t+1}^i.$$

From (A.6), we now get

$$\frac{1}{b} \sum_{j=1}^{b} \eta_t^j \leq p h_t(S_t) + (1-p)\frac{1}{b} \sum_{i=1}^{b} e^{-R_{t+1}^i}\theta_{t+1}^i$$

$$\leq p h_t(S_t) + (1-p)\frac{1}{b} \sum_{i=1}^{b} e^{-R_{t+1}^i}\Theta_{t+1}^i$$

$$\leq \max\left\{ h_t(S_t), \frac{1}{b} \sum_{i=1}^{b} e^{-R_{t+1}^i}\Theta_{t+1}^i \right\}$$

$$= \Theta_t. \quad \square$$

Appendix B

In this appendix we argue that, under some restrictions, there can be no general unbiased estimator of the price of American options. The key step is to reduce the problem to the unbiased estimation of

$$\alpha \triangleq \max\{a, E[X]\}$$

for an arbitrary constant a, with the distribution of X varying over a sufficiently rich class $\mathscr{H}$. To be concrete, we take $\mathscr{H}$ to include all constants and at least one random variable Y for which $P(Y < a)$ and $P(Y > a)$ are both strictly positive. Notice that α is the value of the option to choose between a sure payment of a or a risky payment of X. Any reasonably general method for producing unbiased estimates of American option prices should be able to generate unbiased estimates of α for all $X \in \mathscr{H}$.

We first show that there is no function $g : \mathbb{R} \to \mathbb{R}$ such that $E[g(X)] = \max\{a, E[X]\}$ for all $X \in \mathscr{H}$. For this equation to hold for all constant X we must have $g(x) = \max\{a, x\}$. But then $E[g(Y)] = E[\max\{a, Y\}] > \max\{a, E[Y]\}$, by the conditions imposed on Y.

Now suppose that there is a function $g : \mathbb{R}^n \to \mathbb{R}$ for $n > 1$ such that $E[g(X_1, X_2, \ldots, X_n)] = \max\{a, E[X]\}$ for all $X \in \mathscr{H}$, where the X_i are i.i.d. with the distribution X. For this case, we impose the additional requirement that $g \geq a$, so that the estimated value is never less than the obvious lower bound a. We assume that $\mathscr{H}$ includes at least one random variable Y with $E[Y] < a$, one random variable Z with $E[Z] > a$, and both having densities with the same support. Now to have $E[g(Y_1, \ldots, Y_n)] = a$, we must have $g \equiv a$ on the support of Y. But then $E[g(Z_1, \ldots, Z_n)] = a$, a contradiction. [12]

Appendix C

In this appendix we provide detailed information for the implementation of the method. Some care must be take because of the potentially large number of terminal nodes in a tree. Using a *depth-first* procedure, the storage requirements for the algorithm are minimal. In particular, the memory required is only $O(bd)$.

The pseudo-code in Fig. 7 indicates a procedure for processing a single tree and computing a tree estimate using a depth-first procedure. By tree estimate, we mean the high estimate Θ, the low estimate θ, or the estimate of a control variate. The final tree estimate is the average of tree estimates over n independent replications of trees.

In the pseudo-code, 'node value' refers to the computation of Θ, θ, or the control at a given node in the tree. For example, for Θ, the 'node value' is computed using Eq. (2) at a terminal node (i.e., case 1 or case 2 with $j = d$ in Fig. 7) or Eq. (3) at an intermediate node (i.e., case 3 or case 4 with $j < d$ in Fig. 7). For θ, the 'node value' is computed using Eq. (4) at a terminal node or Eq. (6) at an intermediate node. For a control value, the 'node value' is the usual present value operator.

In the pseudo-code, 'state variable' refers to the computation of the state variable $S_t^{i_1 \cdots i_j}$. For example, in the case of an option on a single asset whose price follows a lognormal distribution, the 'state variable' is computed using Eq. (8). Other details of the simulation depend on the particular pricing problem and are left unspecified in the pseudo-code.

---

[12] We thank Chris Rogers for suggesting this argument.

```
/* allocate storage */
integer vector w(j), for j = 1 to d by 1;
real matrix v(i, j), for i = 1 to b by 1, j = 1 to d by 1;

/* initialize parameters */
v(1, 1) = S; w(1) = 1;
for j = 2 to d by 1;
    v(1, j) = 'state variable'; w(j) = 1;
    end for

/* process tree */
j = d;
while (j > 0);
        case 1:  (j = d and w(j) < b)
                 v(w(j), j) = 'node value';
                 v(w(j) + 1, j) = 'state variable'; w(j) = w(j) + 1;
                 end case 1
        case 2:  (j = d and w(j) = b)
                 v(w(j), j) = 'node value';
                 w(j) = 0; j = j - 1;
                 end case 2
        case 3:  (j < d and w(j) < b)
                 v(w(j), j) = 'node value';
                 if (j > 1);
                     v(w(j) + 1, j) = 'state variable'; w(j) = w(j) + 1;
                     for i = j + 1 to d by 1;
                         v(1, i) = 'state variable'; w(i) = 1;
                         end for
                     j = d;
                 else j = 0; end if
                 end case 3
        case 4:  (j < d and w(j) = b)
                 v(w(j), j) = 'node value';
                 w(j) = 0; j = j - 1;
                 end case 4
        end while;
/* return tree estimate */
'tree estimate' = v(1, 1);
```

Fig. 7. Simulation algorithm Pseudo-code.

References

Amin, K.I., 1991, On the computation of continuous time option prices using discrete approximations, Journal of Financial and Quantitative Analysis 26, 477–495.

Amin, K.I. and J.N. Bodurtha, Jr., 1995, Discrete-time valuation of American options with stochastic interest rates, Review of Financial Studies 8, 193–234.

Amin, K.I. and R.A. Jarrow, 1991, Pricing foreign currency options under stochastic interest rates, Journal of International Money and Finance 10, 310–329.

Amin, K.I. and A. Khanna, 1994, Convergence of American option values from discrete- to continuous-time financial models, Mathematical Finance 4, 289–304.

Barraquand, J. and D. Martineau, 1995, Numerical valuation of high dimensional multivariate American securities, Journal of Financial and Quantitative Analysis 30, 383–405.

Birge, J.R., 1994, Quasi-Monte Carlo approaches to option pricing, Technical report 94-19 (Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI 48109).

Boyle, P.P., 1977, Options: A Monte Carlo approach, Journal of Financial Economics 4, 323–338.

Boyle, P.P., 1988, A lattice framework for option pricing with two state variables, Journal of Financial and Quantitative Analysis 23, 1–12.

Boyle, P.P., M. Broadie and P. Glasserman, 1997, Monte Carlo methods for security pricing, Journal of Economic Dynamics and Control, this issue.

Boyle, P.P., J. Evnine and S. Gibbs, 1989, Numerical evaluation of multivariate contingent claims, Review of Financial Studies 2, 241–250.

Brennan, M.J. and E.S. Schwartz, 1985, Evaluating natural resource investments, Journal of Business 58, 135–157.

Broadie, M. and P. Glasserman, 1995, A pruned and bootstrapped American option simulator, in: Alexopoulos, Kang, Lilegdon and Goldsman, eds., Proceedings of the 1995 Winter Simulation Conference (The Society for Computer Simulation, San Diego, CA) 212–219.

Broadie, M., P. Glasserman and G. Jain, 1996, Enhanced Monte Carlo estimates for American option prices, Working paper (Columbia University) Journal of Derivatives, to appear.

Carriere, J.F., 1996, Valuation of the early-exercise price for derivative securities using simulations and splines, Insurance: Mathematics and Economics 19, 19–30.

Cortazar, G. and E. Schwartz, 1994, The valuation of commodity contingent claims, Journal of Derivatives 1, 27–39.

Cox, J.C. and S.A. Ross, 1976, The valuation of options for alternative stochastic processes, Journal of Financial Economics 3, 145–166.

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

Dammon, R.M. and C.S. Spatt, 1992, An option-theoretic approach to the valuation of dividend reinvestment and voluntary purchase plans, Journal of Finance 47, 331–348.

Dixit, A. and R. Pindyck, 1994, Investment under uncertainty (Princeton University Press, Princeton).

Geltner, D., T. Riddiough and S. Stojanovic, 1995, Insights on the effect of land use choice: The perpetual option on the best of two underlying assets, Working paper (MIT); Journal of Urban Economics, to appear.

Geske, R. and H.E. Johnson, 1984, The American put options valued analytically, Journal of Finance 39, 1511–1524.

Grenadier, S.R. and A.M. Weiss, 1994, Optimal migration strategies for firms facing technological innovations: An option pricing approach, Working paper (Stanford University Graduate School of Business).

Gut, A., 1988, Stopped random walks (Springer, New York).

Hull, J. and A. White, 1993a, One-factor interest rate models and the valuation of interest-rate derivative securities, Journal of Financial and Quantitative Analysis 28, 235–254.

Hull, J. and A. White, 1993b, Efficient procedures for valuing European and American path-dependent options, Journal of Derivatives 1, 21–31.

Jamshidian, F., 1996, Libor and swap market models and measures, Working paper (Sakura Global Capital, London).

Johnson, H., 1987, Options on the maximum or the minimum of several assets, Journal of Financial and Quantitative Analysis 22, 227–283.

Joy C., P.P. Boyle and K.S. Tan, 1996, Quasi-Monte Carlo methods in numerical finance, Management Science 42, 926–938.

Kamrad, B. and P. Ritchken, 1991, Multinomial approximating models for options with k state variables, Management Science 37, 1640–1652.

Law, A.M. and W.D. Kelton, 1991, Simulation modeling and analysis, 2nd ed. (McGraw-Hill, New York).

Majd, S. and R.S. Pindyck, 1987, Time to build, option value, and investment decisions, Journal of Financial Economics 18, 7–27.

McDonald, R. and M. Schroder, 1990, A parity result for American options, Working paper (Northwestern University).

McDonald, R. and D. Siegel, 1985, Investment and the valuation of firms when there is an option to shut down, International Economic Review 26, 331–349.

Nelson, D. and K. Ramaswamy, 1990, Simple binomial processes as dffusion approximation in financial models, Review of Financial Studies 3, 393–430.

Owen, A., 1994, Randomly permuted (t,m,s)-nets and (t,s)-sequences, Working paper (Department of Statistics, Stanford University).

Owen, A., 1995, Monte Carlo variance of scrambled equidistribution quadrature, in: H. Niederreiter and P.J.S. Shiue, eds., Monte Carlo and quasi-Monte Carlo methods in scientific computing (Springer, Berlin).

Paskov, S. and J. Traub, 1995, Faster valuation of financial derivatives, Journal of Portfolio Management 22, 113–120.

Pindyck, R.S., 1988, Irreversible investment, capacity choice, and the value of the firm, American Economic Review 78, 969–985.

Schervish, M., 1984, Algorithm AS195: Multivariate normal probabilities with error bound, Applied Statistics 33, 81–94.

Schervish, M., 1985, Correction to 'multivariate normal probabilities with error bound', Applied Statistics 34, 103–104.

Tilley, J.A., 1993, Valuing American options in a path simulation model, Transactions of the Society of Actuaries 45, 83–104.

Triantis, A.J. and J.E. Hodder, 1990, Valuing flexibility as a complex option, Journal of Finance 45, 549–565.

Trigeorgis, L., 1991, A log-transformed binomial numerical analysis method for valuing complex multi-option investments, Journal of Financial and Quantitative Analysis 26, 309–326.

Turnbull, S.M., 1993, Pricing and hedging diff swaps, Journal of Financial Engineering 2, 297–333.

Zenios, S.A. and M.S. Shtilman, 1993, Constructing optimal samples from a binomial lattice, Journal of Information and Optimization Sciences 14, 125–147.