Primal-Dual Simulation Algorithm for Pricing Multidimensional American Options

Leif Andersen
Banc of America Securities LLC, 9 West 57th Street, New York, New York 10019, leif.andersen@bofasecurities.com

Mark Broadie
Graduate School of Business, Columbia University, 3022 Broadway,
New York, New York 10027-6902, mnb2@columbia.edu

This paper describes a practical algorithm based on Monte Carlo simulation for the pricing of multidimensional American (i.e., continuously exercisable) and Bermudan (i.e., discretely exercisable) options. The method generates both lower and upper bounds for the Bermudan option price and hence gives valid confidence intervals for the true value. Lower bounds can be generated using any number of primal algorithms. Upper bounds are generated using a new Monte Carlo algorithm based on the duality representation of the Bermudan value function suggested independently in Haugh and Kogan (2004) and Rogers (2002). Our proposed algorithm can handle virtually any type of process dynamics, factor structure, and payout specification. Computational results for a variety of multifactor equity and interest-rate options demonstrate the simplicity and efficiency of the proposed algorithm. In particular, we use the proposed method to examine and verify the tightness of frequently used exercise rules in Bermudan swaption markets.

Key words: American options; Bermudan options; Bermudan swaptions; Monte Carlo simulation; Libor market model; option pricing; multiple state variables; real options
History: Accepted by David Hsieh, finance; received April 18, 2003. This paper was with the authors $4\frac{1}{2}$ months for 1 revision.

1. Introduction

Closed-form expressions have been derived for many European options under a variety of financial models, the most notable being the Black-Scholes formula for equity options under the geometric Brownian motion model. To date, no similar expressions have been found for the prices of American options, i.e., options that can be exercised at any time up until the maturity of the option, except in (trivial) special cases. Many numerical methods for pricing American options have been proposed, and although tremendous progress has been made, the pricing of these options in multifactor models with possibly path-dependent payouts has remained a formidable challenge.

In this paper we present a simple and efficient method for pricing of American claims under general asset price or factor process dynamics. The method allows for jumps, stochastic volatility, multiple driving factors, etc., and supports virtually any type of payout specification, including path-dependent ones. The single main requirement of our method is a routine to value claims under a proposed (suboptimal) exercise strategy, a requirement that is generally easy to satisfy using existing techniques. The routine is called repeatedly to generate an upper bound

that complements the lower bound consistent with the proposed exercise strategy. The approach taken in the paper is both practical and efficient, and can be applied to a number of challenging models and option payouts of practical importance. As a particular example, we use the method to examine and verify the tightness of certain frequently applied exercise rules used for Bermudan options on swaps (i.e., swaptions) in multifactor interest-rate models.

There is a long and rich history of numerical methods for pricing American-style contingent claims. Among the earliest approaches are the explicit finite difference scheme in Brennan and Schwartz (1977) and the binomial lattice in Cox et al. (1979). The methods of Brennan and Schwartz, and Cox et al., both fall into the category of lattice-based methods, to which other finite difference methods also belong. Lattice-based methods work particularly well for American options on a single underlying asset. However, many American-style options have been introduced that depend on multiple underlying assets or state variables. Examples include spread options, outperformance options, and swaptions, to name a few. Many of these options, particularly the Bermudan interest-rate swaption, have significant economic importance. Multidimensional generalizations of the Cox et al.

binomial method were proposed in Boyle (1988), Boyle et al. (1989), He (1990), and others. A related approach involves extensions of the finite difference method to higher dimensions, as exemplified by the alternating directions implicit (ADI) method; see, e.g., Mitchell and Griffiths (2001). Adapting binomial, trinomial, or finite difference methods to higher dimensions works well for options on two or perhaps three state variables, but because their computational effort grows exponentially with the number of state variables, these methods are impractical for higher-dimensional problems. Because simulation methods do not suffer the curse of dimensionality, it is natural to consider adapting the Monte Carlo approach for this problem.

Boyle (1977) first proposed Monte Carlo simulation for the pricing of Eurpoean claims. However, it was not until much later that the possibility of using Monte Carlo simulation for pricing American-style options was suggested by Bossaerts (1989) and Tilley (1993). Broadie and Glasserman (1997) proposed a convergent algorithm based on simulated trees. Their method generates both lower and upper bounds so that valid confidence intervals on the true Bermudan[1] price can be determined.[2] The simulated tree method removes the exponential dependence of the computation (CPU) time on the problem dimension; however, the CPU time is still exponential in the number of exercise opportunities. The stochastic mesh method proposed in Broadie and Glasserman (2004) has a computation time requirement that is linear in the number of exercise opportunities and quadratic in the number of simulation paths. The stochastic mesh method also generates lower and upper bounds and converges to the true Bermudan price.

The stochastic mesh method uses a dynamic programming-style backwards recursion for approximating the price and optimal exercise policy. The weights that are used to approximate the continuation value of the option are determined by likelihood ratios. An alternate way to compute these weights based on regression was proposed in Carriere (1996), Tsitsiklis and Van Roy (1999), and Longstaff and Schwartz (2001). The computational effort in these methods is linear in the number of exercise opportunities and (nearly) linear in the number of simulation paths. Convergence results for the algorithm are given in Tsitsiklis and Van Roy (2001) and Clément et al. (2001). These methods are able to generate lower bounds by using regressions to determine approximations to continuation values, thus giving an approximation to the optimal stopping policy.

Many other simulation-based methods have been proposed, but, like the regression approach, these methods can only be used to compute lower bounds on the Bermudan option price. Andersen (2000) proposes a method that parameterizes the exercise policy and then optimizes these parameters over a set of simulated paths to determine an approximation to the optimal exercise strategy. Improvements to this method are proposed and tested in Jensen and Svenstrup (2002). Other simulation methods based on parameterizing the exercise decision include Li and Zhang (1996) and Garcia (2003). The *quantization* method recently proposed in Bally et al. (2002a, b) is another alternative that may be competitive for higher-dimensional problems. Simulation methods based on dimensionality reduction or nonparametric representations of the early exercise region include Barraquand and Martineau (1995), Carr and Yang (2001), Clewlow and Strickland (1998), and Raymar and Zwecher (1997).

Valid upper bounds based on a duality approach were recently and independently proposed in Haugh and Kogan (2004) and Rogers (2002) and are related to the earlier work of Davis and Karatzas (1994). The duality approach provides a method to compute an upper bound from the specification of some arbitrary martingale process. The tightness of the upper bound will depend critically on the choice of the martingale process.

Rogers (2002) generates a martingale process by forming a weighted average of analytically tractable martingale processes that are related to the true value function. The weights used in the average are determined by an optimization procedure conducted in a separate simulation. The choice of martingale processes is highly option and process specific and, as stated in Rogers (2002, p. 275), "appears to be more art than science." Interestingly, this approach can be used to compute an upper bound without requiring a lower bound as a starting point, and is done in a way quite different from the upper bound generated by the stochastic mesh or simulated tree algorithms. To complete a valid confidence interval with Rogers' approach also requires the determination of a lower bound, which is typically determined as the value associated with some exercise policy. The determination of a good exercise rule is also somewhat of an art, as evidenced in part by the number of papers proposing different approaches to this problem, but it appears to be a much easier problem than that of finding good martingale processes.

---

[1] Bermudan options are finitely exercisable American options, i.e., options where the holder has the right to exercise at a finite number of dates prior to the option maturity. Because of the finite nature of computer-based methods, most algorithms effectively price Bermudan options.

[2] Broadie and Detemple (1996) proposed a method for computing lower and upper bounds for American options on a single asset. Their method is based on an explicit integral representation of the American option price.

Haugh and Kogan (2004) take as their starting point an approximation to the American option value on the entire state space of the underlying stochastic process. In their paper, this approximation (which may be biased either high or low) is generated by a neural network algorithm, though regression methods could be used as well. The approximate American option value function is, in turn, used to extract a martingale process using a computationally intensive procedure.

The approach taken in this paper is similar in spirit to that of Haugh and Kogan (2004). Our numerical technique for constructing upper bounds, however, involves only straightforward Monte Carlo simulation and significantly improves the performance of the upper bound computation in several respects. Our algorithm does not build or require an approximation to the option price process throughout the state space. Instead, it uses only the information from the approximation to the optimal exercise strategy, which substantially reduces computation time and approximation error.

Working with exercise strategies rather than value approximations inherently expands the scope of the method and allows us to take advantage of situations where good exercise strategies are known independently of value approximations. For instance, Svenstrup (2002) demonstrates that excellent exercise strategies can be obtained from relatively crude low-dimensional approximations implemented in finite difference grids. Upper bounds corresponding to these strategies can easily be obtained in our approach. Further, our algorithm is applicable for options with exercise values that are not available in closed form. Because our approach relies only on Monte Carlo simulation, the resulting algorithm is also very straightforward to implement. We note that the original approach in Haugh and Kogan (2001) generated upper bounds from supermartingales, which are always more conservative than those generated by martingale processes proposed in this paper. The more recent Haugh and Kogan (2004) approach adopts the martingale approach of this paper.

In the next section, the primal pricing problem and its dual are presented. An algorithm for computing upper bounds on Bermudan prices is developed and discussed in §3. Numerical results that demonstrate the simplicity and practicality of our suggested approach are given in §4, which is followed by brief concluding remarks.

2. The Pricing Problem

In this section we set up the notation and preliminaries for the Bermudan option pricing problem. Let $S_t = (S_t^1, \ldots, S_t^n)$ be a vector-valued Markov process

on $\mathbb{R}^n$ with fixed initial state S_0. These represent the underlying assets or state variables of the model. Let B_t denote the value at time t of \$1 invested in a riskless money market account at time 0. In the special case of a constant risk-free rate r, we have $B_t = e^{rt}$. In general, B_t may depend on the current and past state variables $S_0, \ldots, S_t$. The Bermudan option has d exercise opportunities at times $t_1 < t_2 < \cdots < t_d = T$, with $t_1 \geq 0$. The problem of pricing a Bermudan option is to find

$$\text{Primal:} \qquad Q_0 = \sup_{\tau \in \mathcal{T}} E_0\left[\frac{h_\tau}{B_\tau}\right], \qquad (1)$$

where τ is a stopping time taking values in the finite set $\mathcal{T} = \{t_1, t_2, \ldots, t_d\}$, and $h_t \geq 0$ is interpreted as the payoff from exercise at time t. We will use the terms stopping time and exercise policy interchangeably in this paper. The payoff h_t will depend on the current state and may, in general, depend on the entire history of the process until time t. The quantity h_τ/B_τ is the exercise value measured in time 0 dollars. The notation $E_t[\cdot]$ is short for the expectation conditional on the information available until time t, i.e., $E_t[\cdot] = E[\cdot \mid \mathcal{F}_t]$. We assume that the financial markets are complete and that the expectation is taken under the risk-neutral measure.[3]

The Bermudan option value at some time $t_i \in \mathcal{T}$ can also be written as:

$$Q_{t_i} = \max\left(h_{t_i}, E_{t_i}\left[\frac{B_{t_i}}{B_{t_{i+1}}}Q_{t_{i+1}}\right]\right), \qquad (2)$$

i.e., Q_{t_i} represents the maximum of exercising the option or continuing. Exercising gives a value of h_{t_i}, while the expected present value of continuing (and following an optimal exercise policy thereafter), measured in time t_i dollars, is $E_{t_i}[(B_{t_i}/B_{t_{i+1}})Q_{t_{i+1}}]$. Note that Q_{t_i} represents the value of a Bermudan option *newly issued* at time t_i, and does not equal the value of a Bermudan option issued at time 0 (because the option issued at time 0 may have been exercised prior to time t_i). The process Q_t/B_t is the smallest supermartingale[4] that dominates h_t/B_t on $t \in \mathcal{T}$, i.e., is the Snell envelope of h_t/B_t (see, e.g., Lamberton and Lapeyre 1996). The terminal condition is $Q_T = h_T$, and we are interested in computing Q_0.

Equation (1) defines the primal pricing problem. Clearly, the value achieved by following some specific

---

[3] More generally, we could eliminate B_τ in Equation (1) and define h to be the payoff in units of an arbitrary numeraire asset contained in the vector of state variables with the law of the state variables adjusted accordingly.

[4] Recall that a process Z_t is a supermartingale if $Z_t \geq E_t[Z_s]$ for all $s > t$. Roughly speaking, a supermartingale is expected to drift down over time. Discounted Bermudan and American option price functions are seen to be supermartingales, due to the loss of exercise rights as time progresses.

exercise strategy is dominated by an optimal strategy, so

$$E_0\left[\frac{h_\tau}{B_\tau}\right] \le Q_0. \tag{3}$$

In other words, any algorithm that gives a stopping rule τ can be used to compute a lower bound on the Bermudan value Q_0.

In order to define a dual problem, we first find an upper bound. Let H represent the space of adapted martingales π for which $\sup_{t\in\mathcal{T}} |\pi_t| < \infty$ and $\pi_0 = 0$. For a martingale $\pi \in H$, we have:

$$Q_0 = \sup_{\tau\in\mathcal{T}} E_0\left[\frac{h_\tau}{B_\tau} + \pi_\tau - \pi_\tau\right] = \pi_0 + \sup_{\tau\in\mathcal{T}} E_0\left[\frac{h_\tau}{B_\tau} - \pi_\tau\right]$$

$$\le \pi_0 + E_0\left[\max_{t\in\mathcal{T}}\left(\frac{h_t}{B_t} - \pi_t\right)\right], \tag{4}$$

where the second equality follows from the martingale property of π and the Optional Sampling Theorem.

Because $\pi \in H$ was arbitrary, the inequality in (4) holds after taking the infimum:

$$Q_0 \le \inf_{\pi\in H}\left(\pi_0 + E_0\left[\max_{t\in\mathcal{T}}\left(\frac{h_t}{B_t} - \pi_t\right)\right]\right).$$

The right-hand side of the previous equation defines a dual problem:

$$\text{Dual:} \quad \inf_{\pi\in H}\left(\pi_0 + E_0\left[\max_{t\in\mathcal{T}}\left(\frac{h_t}{B_t} - \pi_t\right)\right]\right). \tag{5}$$

The "duality gap" will be zero if the upper bound in (4) holds with equality. To find a martingale that gives a tight upper bound, we rely on the supermartingale property of Q_t/B_t, which allows for a Doob-Meyer decomposition of the form

$$\frac{Q_t}{B_t} = M_t - A_t, \tag{6}$$

where M_t is a martingale and A_t is an increasing process with $A_0 = 0$. Now take $\pi_t = M_t$ in Equation (4) to get:

$$Q_0 \le Q_0 + E_0\left[\max_{t\in\mathcal{T}}\left(\frac{h_t}{B_t} - \frac{Q_t}{B_t} - A_t\right)\right] \le Q_0, \tag{7}$$

where the second inequality follows because $Q_t \ge h_t$ and $A_t \ge 0$. Thus, the inequality in (4) holds with equality (i.e., there is no "duality gap") when π_t is taken to be the martingale component of the discounted price process, Q_t/B_t.

To get a good lower bound, we need to find an exercise policy (stopping time) τ that is, loosely speaking, close to some optimal policy τ^*. The algorithms mentioned earlier, and others, can all be used to generate candidate exercise policies τ. To get a good upper

bound, Equation (7) suggests we take π_t to be the martingale component of a good approximation to the discounted price process Q_t/B_t. Our approach will be to find an exercise policy τ that defines a lower bound price process L_t and then take π_t to be the martingale component of L_t/B_t. Computational issues are important because the upper bound requires computing the martingale π, which in turn depends on the lower bound L_t function. In many algorithms, however, the lower bound function is not available at every point in the state space and must itself be computed, e.g., by simulation or some interpolation scheme. Because noise or simulation error can propagate through this process, choices made at each stage will affect the final result. The next section addresses these issues.

3. Computing the Upper Bound
This section describes an algorithm for computing an upper bound for the Bermudan option price. For any given exercise strategy, we can always define an adapted exercise indicator process l_t that equals 1 if exercise should take place at time t (given $\mathcal{F}_t$) and 0 otherwise. For all $0 \le t \le t_d$ we define t-indexed stopping times τ_t as

$$\tau_t = \inf\{u \in \mathcal{T} \cap [t, T]: l_t = 1\}.$$

Thus, τ_t denotes the first instance at time t or later at which an option that is newly issued at time t (or simply alive at time t) should be exercised according to the given strategy. Notice that we in effect associate the term "exercise strategy" with a *sequence* of stopping times. With this definition, a discounted lower bound price process L_t/B_t can be defined as

$$\frac{L_t}{B_t} = E_t\left[\frac{h_{\tau_t}}{B_{\tau_t}}\right], \tag{8}$$

i.e., L_t is the value at time t of following the chosen exercise policy from time t onward. L_t is also seen to be the value of an option that is newly issued at time t and exercised according to the exercise indicator process. Defining the sequence of stopping times, τ_t, is useful because we will need to track the evolution of lower bound process defined in terms of these newly issued options. If τ_t is close to an optimal strategy τ_t^* that solves

$$E_t\left[\frac{h_{\tau_t^*}}{B_{\tau_t^*}}\right] = \sup_{\tau_t\in\mathcal{T}\cap[t, T]} E_t\left[\frac{h_{\tau_t}}{B_{\tau_t}}\right] = \frac{Q_t}{B_t},$$

then L_t should be close to Q_t. We will use the lower bound process L_t as the basis for computing the upper bound through Equation (4). As L_t can be computed for any adapted exercise strategy, specification of an exercise strategy through the indicator process l_t (or, equivalently, through the sequence of stopping

times τ_t) in principle suffices to compute an upper bound.[5]

To apply the upper bound in Equation (4), we now define a martingale π_t by $\pi_0 = L_0$, $\pi_1 = L_1/B_1$, and for $2 \le k \le d$,

$$\pi_k = \pi_{k-1} + \frac{L_k}{B_k} - \frac{L_{k-1}}{B_{k-1}} - l_{t_{k-1}} E_{k-1}\left[\frac{L_k}{B_k} - \frac{L_{k-1}}{B_{k-1}}\right], \quad (9)$$

where we use the simplified notation L_k for L_{t_k}. With this definition π_k is easily seen to be a martingale. If continuation is indicated at t_{k-1}, i.e., if $l_{t_{k-1}} = 0$, then

$$\frac{L_{k-1}}{B_{k-1}} \equiv E_{k-1}\left[\frac{h_{\tau_{k-1}}}{B_{\tau_{k-1}}}\right] = E_{k-1}\left[E_k\left[\frac{h_{\tau_{k-1}}}{B_{\tau_{k-1}}}\right]\right] = E_{k-1}\left[\frac{L_k}{B_k}\right]. \quad (10)$$

Therefore, the discounted lower bound process is a martingale in the continuation region. Thus, $E_{k-1}(\pi_k) = \pi_{k-1}$ and π is a martingale in the continuation region as well. If exercise is indicated at t_{k-1}, i.e., if $l_{t_{k-1}} = 1$, then

$$E_{k-1}(\pi_k) = \pi_{k-1} + E_{k-1}\left[\frac{L_k}{B_k} - \frac{L_{k-1}}{B_{k-1}}\right] - E_{k-1}\left[\frac{L_k}{B_k} - \frac{L_{k-1}}{B_{k-1}}\right]$$

$$= \pi_{k-1},$$

and π is a martingale in the exercise region.

Because π is a martingale, Equation (4) from the previous section gives an upper bound on the price of a Bermudan option:

$$Q_0 \le L_0 + E_0\left[\max_{t \in \mathcal{T}}\left(\frac{h_t}{B_t} - \pi_t\right)\right].$$

This equation shows that an upper bound is given by the lower bound plus a term that is the value at time 0 of a lookback option that pays $\max_{t \in \mathcal{T}}(h_t/B_t - \pi_t)$. The "perfect foresight" nature of lookbacks can make these options quite expensive, and hence the upper bound could be quite loose.

To investigate the tightness of this bound, rewrite Equation (9) as

$$\pi_k = L_0 + \sum_{i=0}^{k-1}\left[\frac{L_{i+1}}{B_{i+1}} - E_i\left(\frac{L_{i+1}}{B_{i+1}}\right)\right], \qquad k = 1, 2, \ldots, d, \quad (11)$$

where we have used the fact that $E_{k-1}[L_k/B_k - L_{k-1}/B_{k-1}]$ is zero whenever $l_{t_{k-1}}$ is zero. Define the difference process $e_k = (Q_k - L_k)/B_k$, where Q_k is the true Bermudan price at time t_k, and note that $e_k \ge 0$. Substituting in Equation (11) gives

$$\pi_k = M_k - e_0 - \sum_{i=0}^{k-1}\left[e_{i+1} - E_i(e_{i+1})\right], \qquad k = 1, 2, \ldots, d, \quad (12)$$

where M_k is the optimal martingale defined from Q_k in Equation (6). Let $\mathcal{K} = \{1, \ldots, d\}$. Then the upper bound can be written as:

$$U_0 = L_0 + E_0\left[\max_{k \in \mathcal{K}}\left(\frac{h_k}{B_k} - \pi_k\right)\right]$$

$$= L_0 + E_0\left[\max_{k \in \mathcal{K}}\left(\frac{h_k}{B_k} - M_k + e_0 + \sum_{i=0}^{k-1}[e_{i+1} - E_i(e_{i+1})]\right)\right]$$

$$\le Q_0 + E_0\left[\max_{k \in \mathcal{K}}\left(\sum_{i=0}^{k-1}[e_{i+1} - E_i(e_{i+1})]\right)\right],$$

where the inequality follows because $h_k/B_k \le Q_k/B_k \le M_k$. It appears difficult to bound this expression in general, but because $e_{i+1} \ge 0$ and $E_i(e_{i+1}) \ge 0$, we have

$$U_0 \le Q_0 + E_0\left[\max_{k \in \mathcal{K}}\left(\sum_{i=0}^{k-1} e_{i+1}\right)\right] \le Q_0 + E_0\left[\sum_{i=1}^{d} e_i\right]. \quad (13)$$

Equation (13) shows that if the lower bound is uniformly close to the true price, e.g., if $e_k = (Q_k - L_k)/B_k \le \epsilon$, then the upper bound will differ by at most $d\epsilon$ from the true value. Thus, the upper bound deteriorates at most linearly with the number of exercise dates, though in practice the upper bound appears to be much better. This is illustrated next for the case of a single asset.

Lower and Upper Bound: Numerical Results for a Single Asset

To illustrate and compare the lower and upper bounds, we use a finely spaced binomial lattice to compute the bounds for Bermudan and American call options on a single asset in the standard Black-Scholes model. Suboptimal exercise policies are specified as fixed multiples of the optimal exercise policy. For example, suppose at time t the optimal policy is to exercise when $S_t \ge S_t^*$. For a fixed fraction f, define the suboptimal exercise policy by exercising when $S_t \ge f S_t^*$, for all $t \in \mathcal{T}$ and $t < T$ (where at time T the option is exercised if it is in the money). Exercising too early corresponds to taking $f < 1$ and exercising too late corresponds to $f > 1$. Clearly, the suboptimal policy approaches the optimal policy as f approaches 1. The dependence of the lower and upper bounds on f is given in Table 1 below.

Table 1 shows that the lower and upper bounds improve as f approaches 1 (i.e., as τ approaches τ^*). Even though the upper bound deteriorates as the number of exercise opportunities increases, the lower bound deteriorates as well. So, for most parameter values, the lower and upper bounds are comparable in magnitude. In the cases with the largest differences, the upper bound is tighter than the lower bound. Most importantly, Table 1 indicates that if the lower bound is reasonably tight, we can expect that

Table 1   Lower and Upper Bounds on a Single Asset

	$d = 2$		$d = 10$		American	
f	Lower	Upper	Lower	Upper	Lower	Upper
0.86	−0.93	0.11	−2.19	0.66	−2.60	2.45
0.88	−0.56	0.08	−1.35	0.49	−1.57	1.77
0.91	−0.29	0.05	−0.71	0.30	−0.83	1.11
0.94	−0.12	0.03	−0.29	0.13	−0.35	0.54
0.97	−0.03	0.01	−0.07	0.03	−0.08	0.13
1.00	0.00	0.00	0.00	0.00	0.00	0.00
1.03	−0.03	0.01	−0.05	0.03	−0.05	0.04
1.06	−0.10	0.05	−0.17	0.11	−0.19	0.13
1.10	−0.19	0.10	−0.34	0.25	−0.37	0.29
1.13	−0.30	0.15	−0.53	0.41	−0.57	0.50
1.17	−0.42	0.21	−0.73	0.59	−0.78	0.74
True	7.18		7.98		8.17	

Notes. Parameters: The initial stock and option strike prices are $S = K = 100$, the interest rate is $r = 5\%$, the dividend rate is $\delta = 10\%$, the option matures in $T = 3$ years, and the stock volatility parameter is $\sigma = 20\%$. In the Bermudan cases, there are exercise opportunities at times $t_i = iT/d$, $i = 0, 1, \ldots, d$. Continuous exercise is allowed in the American case. The call option payoff is $h(S_t) = \max(S_t - K, 0)$ when exercised at time t when the stock price is S_t. In the first column f gives the ratio of the critical exercise price under the suboptimal policy to the optimal critical exercise price. For comparison, the corresponding European option value is 6.02. The table shows the differences between the bounds and the true value. For example, for $f = 0.86$ and $d = 2$, the lower bound is $7.18 - 0.93$ and the upper bound is $7.18 + 0.11$. The numbers in this table were computed using a binomial lattice with 2,000 time steps. Monte Carlo simulations were used to compute the upper bounds.

the upper bound will be reasonably tight as well. Also, the upper bound deteriorates quite slowly as the number of exercise dates increases. This example illustrates the practical usefulness of the upper bound and indicates the potential benefit from a computationally efficient method for computing the upper bound.

Computing the Upper Bound via Simulation

Let l_t be the exercise indicator process of an arbitrary exercise policy and let L_t denote the option value under this policy, as defined in Equation (8). Given L_t, define the martingale π as in Equation (9). Then, an upper bound is given by

$$L_0 + E_0 \left[\max_{t \in \mathcal{T}} \left(\frac{h_t}{B_t} - \pi_t \right) \right] = L_0 + \Delta_0. \tag{14}$$

The quantity Δ_0 can be estimated by the following procedure.

Summary of the Simulation Procedure to Estimate the Upper Bound.

1. Simulate a path of state variables: $S_0 = (S_0^1, \ldots, S_0^n)$, $S_1 = (S_1^1, \ldots, S_1^n)$, $\ldots$, $S_d = (S_d^1, \ldots S_d^n)$ and the associated $B_1, \ldots, B_d$ factors used for discounting.

2. Along the path generated in Step 1, check the exercise policy at each time $k = 1, \ldots, d$. If continuation is recommended at time k, i.e., $l_{t_k} = 0$, follow the procedure in Step 2a, otherwise use 2b.

(a) ($l_{t_k} = 0$). Launch a "simulation within a simulation" to estimate:

$$\frac{L_k}{B_k} = E_k \left[\frac{h_{\tau_k}}{B_{\tau_k}} \right].$$

In particular, use N_2 subpaths starting from S_k which are stopped according to τ_k, and average $h_{\tau_k}(S_{\tau_k})/B_{\tau_k}$ over these subpaths. Also, note in this case that in defining π in Equation (9) the term

$$l_{t_k} E_k \left[\frac{L_{k+1}}{B_{k+1}} - \frac{L_k}{B_k} \right]$$

is zero.

(b) ($l_{t_k} = 1$). Set L_k/B_k equal to the discounted exercise value h_k/B_k. If $k < d$, launch a "simulation within a simulation" to estimate:

$$E_k \left[\frac{L_{k+1}}{B_{k+1}} \right] = E_k \left[\frac{h_{\tau_{k+1}}}{B_{\tau_{k+1}}} \right].$$

In particular, use N_3 subpaths starting from S_k, which are stopped at the first time $t \geq t_{k+1}$ such that $l_t = 1$ and average $h_{\tau_{k+1}}(S_{\tau_{k+1}})/B_{\tau_{k+1}}$ over these subpaths. These quantities will be used to estimate the term

$$l_{t_k} E_k \left[\frac{L_{k+1}}{B_{k+1}} - \frac{L_k}{B_k} \right] = E_k \left[\frac{L_{k+1}}{B_{k+1}} \right] - \frac{h_k}{B_k}$$

in defining π in Equation (9).

3. Set $\pi_1, \pi_2, \ldots, \pi_d$ as in Equation (9) and compute

$$\max_{k \in \mathcal{K}} \left(\frac{h_k}{B_k} - \pi_k \right).$$

Repeat Steps 1–3 for N_1 simulation trials, then average the results in Step 3 to estimate the quantity Δ_0 in the upper bound in Equation (14).

Here are some comments on the procedure. In Step 1, a path of the state variables is simulated. The simulation should include at least the potential exercise times $t_1, t_2, \ldots, t_d = T$, but might include additional times. The intermediate times might be necessary to record information about path-dependent payoffs or because a simulation discretization scheme (e.g., Euler or Milstein) requires smaller time steps between exercise dates.

Step 2 is used to estimate π_k from the lower bound process L_k. Suppose at time k that τ_k specifies continuation. In this case,

$$\frac{L_k}{B_k} = E_k \left(\frac{L_{k+1}}{B_{k+1}} \right) = E_k \left(\frac{h_{\tau_k}}{B_{\tau_k}} \right)$$

(see Equation (10)) and simulation is used in Step 2a to estimate the quantity L_k/B_k by estimating the expectation on the right, $E_k(h_{\tau_k}/B_{\tau_k})$. Suppose at time k that τ_k specifies exercise. Then L_k/B_k is simply h_k/B_k. However, the quantity $E_k(L_{k+1}/B_{k+1})$ now

represents the value (in time 0 dollars) of the strategy that continues at time k (even though τ_k specifies exercise) and then exercises according to the τ_{k+1} exercise policy. In short, it represents the *continuation value* at time k, given the information available at time k. In Step 2b this quantity is estimated through

$$E_k\left(\frac{L_{k+1}}{B_{k+1}}\right) = E_k\left[\frac{h_{\tau_{k+1}}}{B_{\tau_{k+1}}}\right].$$

Because this simulation is started at S_k in the exercise region, it is likely that l_t will recommend exercise at time t_{k+1}, so this "inner simulation" is likely to be fast. Thus, Steps 2a and 2b use the same "inner" simulation procedure, with only a slight difference in interpretation. Also note that the inner simulations in Step 2 are *not* recursive—further simulations are not run on each subpath.

The quantity $E_k(L_{k+1}/B_{k+1})$ from Step 2b could be estimated with one-step simulation (or integration procedure) using an estimate of $L_{k+1}(S_{k+1})/B_{k+1}$ at states S_{k+1} at time $k + 1$. The estimate could be obtained through a functional approximation (e.g., regression, spline, or neural network). A functional approximation approach, though, requires having good estimates of $L_{k+1}(S_{k+1})/B_{k+1}$ over a range of states S_{k+1}. Instead, we estimate the quantity $E_k(L_{k+1}/B_{k+1})$ directly from the exercise indicator process l_t. This approach does not require any function approximation, and the resulting accuracy of the estimate is more easily controlled.

Similarly, the quantity $L_k(S_k)/B_k$ at a continuation point required in Step 2a could be taken from a functional approximation. Instead, we estimate this quantity directly from the exercise indicator process l_t.

Finally, note that Steps 2a and 2b are used to estimate conditional expectations of the forms L_k/B_k and $E_k(L_{k+1}/B_{k+1})$. Monte Carlo simulation is only one convenient and general method for approximating these expectations. Given any particular pricing problem, it might be that these quantities can be computed analytically or by some other numerical procedure that is faster and/or more accurate than simulation. If such other numerical methods are available, then, of course, they should be employed.

Figure 1 illustrates the simulation procedure. At time t_1, the state S_1 is in the continuation region specified by the exercise policy. In order to estimate L_1/B_1, we begin a simulation from the state S_1 and simulate subpaths $S_2^i, S_3^i, \ldots, S_{\tau_1}^i$ for $i = 1, 2, \ldots, N_2$. Each subpath is stopped when specified by the exercise policy τ_1. The discounted payoffs $h_{\tau_1}(S_{\tau_1}^i)/B_{\tau_1}$ for $i = 1, \ldots, N_2$ are averaged to give a simulation estimate of L_1/B_1. At time t_2, the state S_2 is in the continuation region, and L_2/B_2 is estimated using subpaths $S_3^i, S_4^i, \ldots, S_{\tau_2}^i$ for $i = 1, 2, \ldots, N_2$. The discounted payoffs $h_{\tau_2}(S_{\tau_2}^i)/B_{\tau_2}$ for $i = 1, \ldots, N_2$ are averaged to give a

Figure 1   Simulation in a Simulation
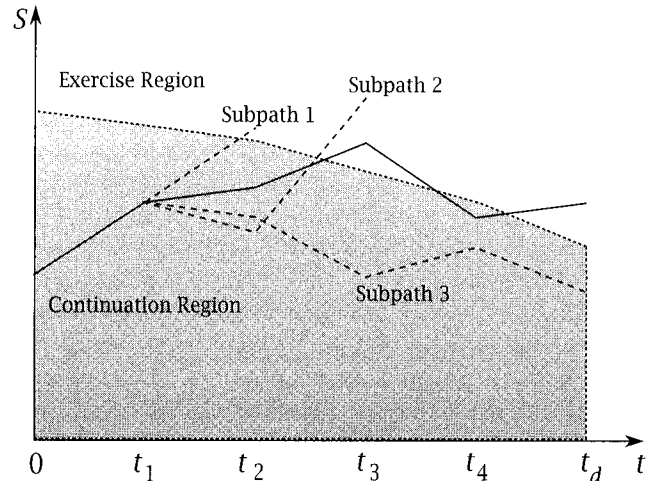
Notes. This figure illustrates the computation of L_1/B_1 using three subpaths that start from S_1 at time t_1. The exercise policy is defined by continuing (i.e., not exercising) in the lower shaded area and exercising in the upper region. Subpath 1 is exercised at time t_2, subpath 2 at time t_3, and subpath 3 is not exercised. Additional subpaths would be used to compute L_2/B_2 starting from S_2. Because S_3 lies in the exercise region, the term L_3/B_3 is set to h_3/B_3. Additional subpaths starting from S_3 are used to compute $E_3[L_4/B_4]$, which represents the value of continuing at time t_3.

simulation estimate of L_2/B_2. At time t_3, the state S_3 is in the exercise region. Therefore, L_3/B_3 is set to h_3/B_3. In order to estimate $E_3(L_4/B_4)$, we begin a simulation from the state S_3 and simulate subpaths $S_4^i, \ldots, S_{\tau_4}^i$ for $i = 1, 2, \ldots, N_3$. Each subpath is stopped when specified by the exercise policy τ_4, and the resulting h_{τ_4}/B_{τ_4} are averaged to give a simulation estimate of $E_3(L_4/B_4)$.

To briefly demonstrate that the simulation algorithm will produce an upper bound, notice that the Monte Carlo simulations embedded in the algorithm above introduces noise in the estimates for the martingale π. In particular, the terms L_k/B_k in (9) are effectively replaced by $L_k/B_k + \epsilon_k$, where ϵ_k is a pure noise term with mean zero and standard deviation proportional to $1/\sqrt{N_2}$ (if S_k is a point of continuation; otherwise, if S_k is a point of exercise the standard deviation is zero). Similarly, the terms $E_k(L_{k+1}/B_{k+1})$ will be replaced by $E_k(L_{k+1}/B_{k+1}) + \epsilon_k'$, where ϵ_k' is a pure noise term with mean zero and standard deviation proportional to $1/\sqrt{N_3}$. Compared with (9), we get

$$\hat{\pi}_k = \hat{\pi}_{k-1} + \frac{L_k}{B_k} + \epsilon_k - \frac{L_{k-1}}{B_{k-1}} - \epsilon_{k-1}$$
$$- l_{t_{k-1}}\left(E_{k-1}\left[\frac{L_k}{B_k} - \frac{L_{k-1}}{B_{k-1}}\right] + \epsilon_k' - \epsilon_{k-1}\right). \quad (15)$$

By induction, we can write

$$\hat{\pi}_k = \pi_k + \tilde{\epsilon}_k,$$

where $\tilde{\epsilon}_k$ is a sum of mean zero noise terms. Rather than computing $\max_{k \in \mathfrak{R}}(h_k/B_k - \pi_k)$ exactly along

each path, our algorithm instead generates the noisy estimate:

$$\max_{k \in \mathcal{K}}\left(\frac{h_k}{B_k} - \pi_k - \tilde{\epsilon}_k\right).$$

Let m denote the random index in the exercise set at which $h_k/B_k - \pi_k$ takes its maximum. Then

$$E_0\left[\max_{k \in \mathcal{K}}\left(\frac{h_k}{B_k} - \pi_k - \tilde{\epsilon}_k\right)\right] \geq E_0\left[\frac{h_m}{B_m} - \pi_m - \tilde{\epsilon}_m\right]$$

$$= E_0\left[\frac{h_m}{B_m} - \pi_m\right]$$

$$= E_0\left[\max_{k \in \mathcal{K}}\left(\frac{h_k}{B_k} - \pi_k\right)\right],$$

where the first equality follows from the zero mean of $\tilde{\epsilon}_m$. Thus, our algorithm's estimate of Δ_0, and thereby of the price upper bound, will be biased high for finite samples N_2 and N_3, but still yielding a valid upper bound. Of course, the higher we set N_2 and N_3, the lower this bias will be.

Before proceeding to concrete numerical examples, we discuss how to use the lower and upper bound results to construct confidence intervals for Bermudan option prices. Suppose that the Monte Carlo estimate of the lower bound is $\hat{L}_0$ with a sample standard deviation $\hat{s}_L$ based on N independent simulation trials. Also, let the simulation estimate of Δ_0, determined from the algorithm above, be denoted $\hat{\Delta}_0$ with sample standard deviation $\hat{s}_\Delta$ based on N_1 trials. With z_x denoting the xth percentile of a standard Gaussian distribution, asymptotically a $100(1-\alpha)\%$-probability confidence interval for the Bermudan price Q_0 must be tighter[6] than

$$\left[\hat{L}_0 - z_{1-\alpha/2}\frac{\hat{s}_L}{\sqrt{N}}, \quad \hat{L}_0 + \hat{\Delta}_0 + z_{1-\alpha/2}\sqrt{\frac{\hat{s}_L^2}{N} + \frac{\hat{s}_\Delta^2}{N_1}}\right]. \quad (16)$$

The confidence interval in (16) is conservative because of the low bias in $\hat{L}_0$ (i.e., $E_0[\hat{L}_0] \leq Q_0$) and the high bias in $\hat{L}_0 + \hat{\Delta}_0$, which comes from the nature of the upper bound and the additional high bias described in the discussion after Equation (15). The standard error for the upper bound is based on the assumption that the lower bound estimate $(\hat{L}_0)$ and the upper bound increment $(\hat{\Delta}_0)$ are computed using independent simulation trials. We choose to separate the lower and

upper bound computations in this fashion because the time required to compute the lower bound estimate is typically less than required to estimate the increment $\hat{\Delta}_0$. The computational effort and precision associated with these two quantities can be set separately through the choices of N and N_1. This freedom appears to be well worth the cost of both standard errors appearing in the upper endpoint of the confidence limit in (16). The lower and upper bounds can be combined in many ways to give a point estimate of the price. Based on the limited results in Table 1, the obvious point estimate

$$\hat{Q}_0 = \hat{L}_0 + \tfrac{1}{2}\hat{\Delta}_0 \quad (17)$$

appears to give better price estimates than either the lower or upper bound alone.

The computation time required to approximate Δ_0 is, in the worst case, proportional to:

$$n \times N_1 \times \max(N_2, N_3) \times d^2. \quad (18)$$

Equation (18) says the worst-case CPU time is linear in the problem dimension, n. Each of N_1 outer simulation trials involves simulating paths with exercise opportunities at steps $1, 2, \ldots, d$. In addition, each of the N_2 or N_3 inner simulation trials involves simulating additional paths of up to d steps each. Putting this together gives the result in Equation (18). In practice, the inner simulation trials are often stopped very quickly, and so the actual running time of the algorithm appears to be closer to linear in d.

4. Computational Results

In this section we test the method on two classes of problems: the pricing of multiasset equity options, and the pricing of interest-rate derivatives in single- and multifactor term structure models. In particular, we price max-call equity options, a problem which has become a standard test case in the literature. We also price Bermudan swaptions in the Libor market model, which is a problem of significant practical interest. We focus on the determination of upper bounds through our method, and we illustrate the procedure using two different methods for determining lower bounds. For the equity options, we use a regression method for estimating continuation values and determining an exercise strategy. For the interest-rate derivatives, we use a method that parameterizes the exercise region and then optimizes over these parameters to determine an exercise policy.

Equity Max-Options
Our first test of the method is to price equity options in a multiasset Black-Scholes framework. In particular, we suppose that the risk-neutral dynamics of n

---

[6] In addition to the random Monte Carlo error, simulation of some models will also involve a systematic error stemming from the time discretization of the process for S_t and/or B_t. An example of this is the simulation of the Libor market model, see, e.g., Andersen and Andreasen (2000). Such discretization errors are not accounted for in (16), nor in any previous arguments. For the case of Libor market models, Andersen and Andreasen (2000) demonstrate that the discretization errors associated with typical schemes are often very small relative to the random Monte Carlo error.

assets follow correlated geometric Brownian motion processes, i.e.,

$$\frac{dS_t^i}{S_t^i} = (r - \delta_i)\,dt + \sigma_i\,dW_t^i, \qquad (19)$$

where W_t^i, $i = 1, \dots, n$, are standard Brownian motion processes and the instantaneous correlation of W^i and W^j is ρ_{ij}. For simplicity, in our numerical results we take $\delta_i = \delta$ and $\rho_{ij} = \rho$ for all $i, j = 1, \dots, n$ and $i \neq j$. The interest rate r is assumed to be constant, so the value of the money market account at time t is $B_t = e^{rt}$. Exercise opportunities are equally spaced at times $t_i = iT/d$, $i = 0, 1, \dots, d$. The option that we price is the max-call option, which has a payoff upon exercise at time t of[7]

$$h_t(S_t) = (\max(S_t^1, \dots, S_t^n) - K)^+.$$

Properties of the exercise region for this option are studied in Broadie and Detemple (1997). Numerical results for this option using the stochastic mesh method are given in Broadie and Glasserman (2004).

Function approximation methods for determining lower bounds are described in Carriere (1996), Tsitsiklis and Van Roy (1999), Longstaff and Schwartz (2001), and others. The basic idea is to use a functional approximation scheme (e.g., splines, linear regression, neural network, or a similar method) to estimate the continuation value of the option at each exercise time. Improvements based on the control variate technique are investigated in Rasmussen (2002). We follow the Longstaff and Schwartz approach and use linear regression. The method, though, is not well specified until the precise set of regression basis functions is chosen. In our tests, we use a slightly different set of basis functions than Longstaff and Schwartz. In particular, we use a set of 12 functions, consisting of the largest and the second largest asset prices, three polynomials of degree two (e.g., the squares of each asset price and the product of the two), four polynomials of degree three, the value of a European max-call option on the largest two assets, and the square and the cube of this value (and a constant term is also included). In particular, the use of the European max-call value and its powers is new. Different choices of basis functions may lead to better lower bounds; however, this choice of 13 basis functions is sufficient for purposes of illustration, because we mainly wish to focus on the determination of upper bounds given approximate exercise policies (and hence lower bounds). For recent work investigating other basis functions and the trade-off between the number of simulated paths and the number of basis functions; see Glasserman and Yu (2003).

[7] The notation x^+ means $\max(x, 0)$.

Table 2 Bermudan Max-Call Equity Options in Multiasset Black-Scholes Models

S_0	Lower bound	Upper bound	95% CI	Binomial
$n = 2$ assets:				
90	8.065 (0.006)	8.069 (0.007)	[8.053, 8.082]	8.075
100	13.907 (0.008)	13.915 (0.010)	[13.892, 13.934]	13.902
110	21.333 (0.009)	21.340 (0.010)	[21.316, 21.359]	21.345
$n = 3$ assets:				
90	11.279 (0.007)	11.290 (0.009)	[11.265, 11.308]	11.29
100	18.678 (0.009)	18.703 (0.013)	[18.661, 18.728]	18.69
110	27.531 (0.010)	27.627 (0.019)	[27.512, 27.663]	27.58
$n = 5$ assets:				
90	16.618 (0.008)	16.634 (0.010)	[16.602, 16.655]	
100	26.128 (0.010)	26.253 (0.020)	[26.109, 26.292]	
110	36.725 (0.011)	36.798 (0.017)	[36.704, 36.832]	

Notes. The payoff of the max-call option is:

$$(\max(S_t^1, \dots, S_t^n) - K)^+.$$

The parameters are: $K = 100$, $r = 5\%$, $\delta = 10\%$, $\rho = 0$, $T = 3$, and $\sigma = 20\%$. The initial vector is $S_0 = (S, \dots, S)$, with $S = 90$, 100, or 110 as indicated in the table. Exercise opportunities are equally spaced at times $t_i = iT/d$, $i = 0, 1, \dots, d$, with $d = 9$. Values in parentheses are the standard errors. $N_0 = 200{,}000$ trials were used to estimate regression coefficients that determine the exercise policy; $N = 2{,}000{,}000$ trials were used to estimate the lower bound; $N_1 = 1{,}500$ and $N_2 = N_3 = 10{,}000$ were the parameter settings used for the upper bound. The regression estimates of the continuation value used the largest two assets in a regression with 12 independent variables. Binomial values were determined from the multidimensional BEG routine of Boyle et al. (1989). For $n = 2$ assets, the BEG results were from 1,350 time steps with an approximate error of 0.003. For $n = 3$ assets, the BEG results were from 270 time steps with an approximate error of 0.015.

Results for $n = 2, 3$, and 5 assets are given in Table 2. When the true value is unknown, the accuracy of the method can be estimated in a number of ways. Denote the confidence interval by $[a, b]$ and take the half-width of the interval, $(b - a)/2$, and divide by the midpoint of the interval, $(a + b)/2$. The result, $(b - a)/(b + a)$, is a conservative error estimate, which ranges from 0.10% to 0.35% for the results in Table 2. In general, the errors appear to be smaller for the lower-dimensional options (which is to be expected given the choice of basis functions) and for the out-of-the-money options, i.e., when the early exercise premium is lower. The confidence intervals are consistent with and somewhat tighter than the stochastic mesh results reported in Broadie and Glasserman (2004). For $n = 2$ and $n = 3$ it is feasible to compare the lower and upper bound results with the multidimensional binomial method of Boyle et al. (1989). In every case, the binomial point estimate of the Bermudan option value lies within the 95% confidence intervals determined through simulation. A less conservative error estimate can be obtained using the binomial value, b_0, as an estimate of the true value Q_0. In this case, the absolute relative error can be approximated as: $|\hat{L}_0 + \frac{1}{2}\hat{\Delta}_0 - b_0|/b_0$.

For $n = 2$ and 3 when a binomial estimate is available, the absolute relative error ranges from 0.00% to 0.10%.

The results reported in Table 2 are fairly remarkable given the simplicity of the method, the relatively limited effort in determining the lower bounds, and the absence of any variance reduction techniques. Clearly, improvements in the lower bound will lead to tighter upper bounds. Introducing appropriate variance reduction techniques will reduce the standard errors of the lower and upper bounds and further narrow the confidence intervals.

Although little effort was made to optimize the computer code, a rough breakdown of computational effort is instructive. Determining the regression coefficients took approximately 1%–10% of the total CPU time. Determining the lower bound given the regression coefficients took about 5%–30% of the CPU time, and then determining the upper bound took about 60%–95% of the CPU time. Therefore, the determination of the upper bound took between 2 and 15 times the effort of determining the lower bound. This is very reasonable given the use of nested simulations to determine the upper bound. Of course, these results are highly dependent on the number of exercise opportunities. With larger values of d, the time spent in determining the upper bound relative to the lower bound will grow.

Bermudan Swaptions in the Libor Market Model

Next we examine upper bounds for Bermudan swaptions in the Libor market model framework of Brace et al. (1997), Jamshidian (1997), and Miltersen et al. (1997). To make our discussion precise, we first introduce some new notation. Let $P(t, T)$ denote the time t price of \$1 received with certainty at time T. Using the dates in our exercise set as a tenor reference, we can define Libor-style discrete forwards as

$$F_i(t) = \frac{P(t, t_i)/P(t, t_{i+1})}{t_{i+1} - t_i}.$$

Typically, the accrual periods $t_{i+1} - t_i$ are either three or six months. Following Andersen and Andreasen (2000), the dynamics of forward rates are assumed to satisfy

$$dF_i(t) = \mu_i(t)\,dt + \varphi(F_i(t))\lambda_i^T(t)\,dW(t), \qquad (20)$$

where W is an l-dimensional Brownian motion, $\lambda_i(t)$ is an l-dimensional bounded deterministic function of time, and φ is a "skew" function satisfying certain regularity conditions. The drift in Equation (20) can be determined by arbitrage restrictions, and depends on the numeraire asset chosen. Typically, the drift will be a function of multiple forwards; see Andersen and Andreasen (2000) for details.

Consider now a regular fixed-for-floating interest-rate swap exchanging a fixed coupon θ for discretely compounded floating rates, with the first payment exchange[8] at time t_2, and the last exchange at time t_d. As seen by the fixed payer, at time $t < t_1$ the value of the swap is

$$s(t) = \sum_{i=1}^{d-1} P(t, t_{i+1})[F_i(t) - \theta](t_{i+1} - t_i). \qquad (21)$$

For $t > t_1$, we adjust the formula above to account for cash flows that were made in the past and should no longer be counted. Specifically, if $t_k < t \le t_{k+1}$ for some value of k, we simply modify the sum in Equation (21) to start at $i = k$ rather than $i = 1$.

A *European swaption* $c_i(t)$ with maturity t_i entitles its holder to exercise into the swap s at time t_i. Assuming that the option is held by the fixed side payer (a *payer swaption*), the time t_i payout is thus $c_i(t_i) \equiv s(t_i)^+$. A *Bermudan swaption* is the right to exercise into the above-described swap at any one[9] of the dates $t_1, t_2, \ldots, t_d$. The date t_1 is often known as the lock-out date of the Bermudan swaption. The Libor market model (20) is jointly Markov in the entire set of Libor forward rates, often totalling more than 30 or 40 variables. As such, lattice-based methods are unsuited for numerical work, and pricing generally requires Monte Carlo simulation. We here wish to test our upper-bound algorithm on the pricing of Bermudan swaptions.

As a first step, we need a reasonable exercise strategy. Andersen (2000) proposes a variety of such strategies, the simplest and fastest of which is the following:

Exercise strategy 1:

$$l_{t_i} = 1 \quad \text{if and only if } s(t_i) > H(t_i). \qquad (22)$$

In other words, we simply exercise if the proceeds from doing so are bigger than some deterministic function of time H. An optimization algorithm for determining H is discussed in Andersen (2000).

To focus on a specific example, assume that the forward curve is flat at 10%, the accrual periods are three months, and the skew function is $\varphi(x) = x$. We will consider two volatility scenarios: a one-factor scenario with $\lambda_i(t) = 0.2$ for all i and t, and a two-factor scenario with

$$\lambda_i(t) = [0.15, 0.15 - \sqrt{0.009(t_i - t)}]^T, \qquad t \le t_i.$$

---

[8] As is common practice, we assume that interest-rate payments are made in arrears. That is, the payment associated with the floating forward rate observed at time t_1 will be paid at time t_2, and so on.

[9] The right to exercise at time t_d is worthless as the swap is expired on this date. We include it to ensure notational consistency with the rest of the paper.

Table 3 Bermudan Payer Swaptions in a One-Factor Libor Market Model

t_1	t_d	θ	Lower bound Strategy 1	Lower bound Strategy 2	$\hat{\Delta}_0$ Strategy 1		95% CI	Lower-upper average
0.25	1.25	8%	184.6 (0.1)	184.6 (0.1)	0.02	(0.01)	[184.5, 184.7]	184.6
0.25	1.25	10%	49.1 (0.1)	48.9 (0.1)	0.02	(0.003)	[48.8, 49.4]	49.1
0.25	1.25	12%	8.9 (0.1)	8.7 (0.1)	0.004	(0.001)	[8.7, 9.1]	8.9
1	3	8%	355.6 (0.4)	355.1 (0.4)	0.07	(0.01)	[354.9, 356.4]	355.6
1	3	10%	157.8 (0.5)	156.8 (0.5)	0.2	(0.02)	[156.9, 158.9]	157.9
1	3	12%	61.8 (0.4)	61.0 (0.3)	0.04	(0.01)	[61.1, 62.5]	61.8
1	6	8%	807.2 (0.9)	808.0 (0.9)	0.23	(0.03)	[805.4, 809.3]	807.3
1	6	10%	417.8 (0.9)	416.9 (0.9)	0.63	(0.06)	[415.9, 420.3]	418.1
1	6	12%	212.7 (0.9)	212.6 (0.9)	0.33	(0.04)	[210.9, 214.8]	212.9
1	11	8%	1,381.6 (1.6)	1,380.2 (1.6)	1.3	(0.1)	[1,378.4, 1,386.2]	1,382.3
1	11	10%	812.9 (1.4)	813.2 (1.4)	1.3	(0.1)	[810.0, 817.0]	813.5
1	11	12%	495.8 (1.5)	496.7 (1.4)	0.7	(0.1)	[492.7, 499.6]	496.2
3	6	8%	493.2 (0.8)	493.3 (0.8)	0.08	(0.01)	[491.6, 494.9]	493.2
3	6	10%	293.6 (0.9)	293.0 (0.9)	0.65	(0.07)	[291.8, 296.1]	293.9
3	6	12%	170.3 (0.8)	169.9 (0.8)	0.53	(0.06)	[168.6, 172.5]	170.6

Notes. The numbers in the table were generated in an Euler-discretized, log-normal Libor market model with one factor, $\varphi(x) = x$, and $\lambda_i(t) = 0.2$ for all i and t. The accrual periods are three months and the initial forward curve is flat at 10%. All numbers are in basis points, with numbers in parentheses denoting sample standard deviations. The first three columns denote the lock-out date, the final maturity, and the coupon of the Bermudan payer swaption, respectively. The lower bounds in the fourth and fifth columns were generated using $N = 50{,}000$ paths with antithetic sampling; they are identical to numbers reported in Table 6a in Andersen (2000). The estimates $\hat{\Delta}_0$ reported in the sixth column were generated with $N_1 = 750$ and $N_2 = N_3 = 300$. The seventh column reports the 95% confidence interval for the true price as determined by Equation (16). The lower-upper average in the last column were computed as the sum of the fourth column and one-half of the sixth column as in Equation (17).

Results for a variety of Bermudan swaptions are shown in Tables 3 and 4 below. In the tables, we have also included lower bound estimates from an alternative exercise strategy (Strategy 2), of the following form:

Exercise Strategy 2:

$l_{t_i} = 1$ if and only if

$$s(t_i) > \max(c_{i+1}(t_i), \ldots, c_{d-1}(t_i)) + H_2(t_i). \quad (23)$$

That is, we exercise if the proceeds from doing so exceed the maximum price of the European options underlying the Bermudan structure, plus some deterministic spread to be found by optimization. The rationale behind this strategy is discussed in Andersen (2000).

For the one-factor scenario in Table 3, the spread between the lower and upper bounds generated from Strategy 1 are very low, never more than one or more basis points, leading us to conclude that Strategy 1 very accurately captures the correct exercise decision for the data in Table 3. In the two-factor scenario in Table 4, the spreads between upper and lower estimates are, not surprisingly, wider than for the one-factor case, although still relatively small for most of the contracts examined. Reasonably significant spreads,[10] in the order of 15 to 20 basis points,

can be observed for the 11-year contract with one-year lockout. The suboptimality of Exercise Strategy 1 for this particular case is also reflected in the fact that the more complicated Exercise Strategy 2 here picks up significant additional value relative to Strategy 1. In fact, Strategy 2 produces prices that lie very close to the average of the upper and lower bound, suggesting that this strategy is close to optimal. Using Strategy 2 to form an upper bound confirms this: The spread between the upper and lower bounds for the 11-year contract with one-year lockout is reduced to 7.3, 6.3, and 3.5 basis points for coupons of 8%, 10%, and 12%, respectively.

As a general rule, we typically find that Strategy 1 works quite well for most Bermudan swaption contracts, even in a multifactor framework. Indeed, computing upper and lower bounds based on this strategy for all the different contracts and models in Andersen (2000) (covering a variety of factor, volatility, and skew scenarios), the only set of data that produced estimates for Δ_0 in excess of six basis points was that reported in our Table 4 above. As we have seen, Strategy 1 must still be approached with some care in a multifactor setting, primarily for contracts with short lock-out periods and long maturities. In such cases, Strategy 2 will often pick up the loss in value. In any case, the existence of an upper bound will always allow one to estimate the error involved in a particular strategy and, if necessary, improve it.

[10] Bid-offer spreads in Bermudan swaptions markets are generally quite high, often around 5%–10% of the contract value.

Table 4 **Bermudan Payer Swaptions in a Two-Factor Libor Market Model**

t_1	t_d	θ	Lower bound Strategy 1	Lower bound Strategy 2	$\hat{\Delta}_0$ Strategy 1	95% CI	Lower-upper average
0.25	1.25	8%	184.0 (0.0)	184.0 (0.0)	0.05 (0.01)	[183.9, 184.1]	184.0
0.25	1.25	10%	43.3 (0.1)	43.2 (0.1)	0.06 (0.01)	[43.1, 43.6]	43.3
0.25	1.25	12%	5.6 (0.1)	5.6 (0.1)	0.01 (0.003)	[5.5, 5.7]	5.6
1	3	8%	339.7 (0.2)	339.4 (0.2)	0.4 (0.1)	[339.2, 340.6]	339.9
1	3	10%	125.8 (0.3)	125.7 (0.3)	0.7 (0.1)	[125.1, 127.2]	126.2
1	3	12%	36.9 (0.2)	36.6 (0.2)	0.2 (0.0)	[36.4, 37.6]	37.0
1	6	8%	750.2 (0.6)	751.6 (0.6)	3.7 (0.3)	[749.0, 755.2]	752.1
1	6	10%	317.0 (0.7)	319.4 (0.7)	5.0 (0.3)	[315.6, 323.5]	319.5
1	6	12%	127.7 (0.6)	129.2 (0.6)	2.6 (0.2)	[126.5, 131.6]	129.0
1	11	8%	1,247.3 (1.2)	1,253.7 (1.3)	18.1 (1.4)	[1,245.1, 1,269.0]	1,256.3
1	11	10%	620.8 (1.1)	633.2 (1.3)	20.8 (1.2)	[618.4, 645.0]	631.2
1	11	12%	327.1 (1.2)	337.0 (1.2)	14.8 (1.0)	[324.7, 345.0]	334.5
3	6	8%	444.7 (0.6)	445.2 (0.6)	0.8 (0.1)	[443.6, 446.6]	445.1
3	6	10%	226.9 (0.7)	227.5 (0.7)	1.2 (0.1)	[225.5, 229.5]	227.5
3	6	12%	107.1 (0.6)	107.6 (0.6)	0.8 (0.1)	[105.9, 109.0]	107.5

Notes. The numbers above were generated in an Euler-discretized, log-normal Libor market model with two factors, $\varphi(x) = x$, and $\lambda_i(t) = [0.15, 0.15 - \sqrt{0.009(t_i - t)}]^T$, $t \le t_i$. The accrual periods are three months and the initial forward curve is flat at 10%. All numbers are in basis points, with numbers in parentheses denoting sample standard deviations. The first three columns denote the lock-out date, the final maturity, and the coupon of the Bermudan payer swaption, respectively. The lower bounds in the fourth and fifth columns were generated using $N = 50,000$ paths with antithetic sampling; they are identical to numbers reported in Table 6b in Andersen (2000). The estimates $\hat{\Delta}_0$ reported in the sixth column were generated with $N_1 = 750$ and $N_2 = N_3 = 300$. The seventh column reports the 95% confidence interval for the true price as determined by Equation (16). The lower-upper average in the last column were computed as the sum of the fourth column and one-half of the sixth column as in Equation (17).

This section has focused primarily on establishing and verifying the basic Monte Carlo algorithm for the upper bound computation, with little emphasis put on maximizing numerical efficiency. The application of standard variance reduction methods would clearly improve the running time and error bounds for the method. Specialized variance reduction methods for computing lower and upper bounds with this simulation approach are developed and investigated in Broadie and Cao (2003). Extensions of this method for the computation of lower and upper bounds on option Greeks are given in Kaniel et al. (2003).

5. Conclusions

American-style contingent claims continue to be an extremely important component of the market for financial derivatives. Financial models with multiple driving factors (including, for example, stochastic volatility and jump components) are growing in importance as empirical evidence mounts. However, empirical and theoretical work with American-style derivatives in these more realistic, multifactor models has been significantly hampered by computational issues. Our paper contributes to this challenging and important area by proposing a simple, efficient, and general method for generating valid price intervals for American-style options. For higher-dimensional pricing problems, the algorithm is most naturally implemented as a Monte Carlo simulation that is used

to estimate the value of a dual problem associated with the primal option pricing problem. When coupled with any of a number of algorithms for estimating lower bounds on Bermudan prices, the upper bound can be used both to compute a better price estimate and to determine whether more effort is required to improve the lower bound. In many cases of practical importance, lower bounds can be determined very quickly, and then the associated upper bound can be used to demonstrate the tightness of the lower bound.

References

Andersen, L. 2000. A simple approach to the pricing of Bermudan swaptions in the multi-factor Libor market model. *J. Comput. Finance* 3(1) 5–32.

Andersen, L., J. Andreasen. 2000. Volatility skews and extensions of the Libor market model. *Appl. Math. Finance* 7(1) 1–32.

Bally, V., G. Pages, J. Printems. 2002a. A quantization method for pricing and hedging multi-dimensional American style options. Working paper, INRIA (The French National Institute for Research in Computer Science and Control).

Bally, V., G. Pages, J. Printems. 2002b. First order schemes in the numerical quantization method. *Math. Finance* **13**(1) 1–16.

Barraquand, J., D. Martineau. 1995. Numerical valuation of high dimensional multivariate American securities. *J. Financial Quant. Anal.* **30**(3) 383–405.

Bossaerts, P. 1989. Simulation estimators of optimal early exercise. Working paper, Carnegie Mellon University, Pittsburgh, PA.

Boyle, P. 1977. Options: A Monte Carlo approach. *J. Financial Econom.* **4** 323–338.

Boyle, P. 1988. A lattice framework for option pricing with two state variables. *J. Financial Quant. Anal.* **23** 1–12.

Boyle, P., J. Evnine, S. Gibbs. 1989. Numerical evaluation of multivariate contingent claims. *Rev. Financial Stud.* **2** 241–250.

Brace, A., D. Gatarek, M. Musiela. 1997. The market model of interest rate dynamics. *Math. Finance* **7** 127–155.

Brennan, M., E. Schwartz. 1977. The valuation of American put options. *J. Finance* **32** 449–462.

Broadie, M., M. Cao. 2003. Improved bounds for pricing American options using Monte Carlo simulation. Working paper, Columbia University, New York.

Broadie, M., J. Detemple. 1996. American option valuation: New bounds, approximations, and a comparison of existing methods. *Rev. Financial Stud.* **9**(4) 1211–1250.

Broadie, M., J. Detemple. 1997. The valuation of American options on multiple assets. *Math. Finance* **7** 241–286.

Broadie, M., P. Glasserman. 1997. Pricing American-style securities using simulation. *J. Econom. Dynam. Control* **21**(8–9) 1323–1352.

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

Carr, P., G. Yang. 2001. Simulating Bermudan interest rate derivatives. *Quantitative Analysis in Financial Markets*, Vol. II. M. Avellaneda, ed. World Scientific, 295–316.

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

Clément, E., D. Lamberton, P. Protter. 2001 An analysis of a least-squares regression algorithm for American option pricing. *Finance Stochastics* **6** 449–471.

Clewlow, L., C. Strickland. 1998. Pricing interest rate exotics by Monte Carlo simulation. Chapter 17 in: *Monte Carlo: Methodologies and Applications for Pricing and Risk Management*. Risk Books, London, U.K.

Cox, J., S. Ross, M. Rubinstein. 1979. Option pricing: A simplified approach. *J. Financial Econom.* **7** 229–263.

Davis, M., I. Karatzas. 1994. A deterministic approach to optimal stopping, with applications. P. Kelly, ed. *Probability, Statistics and Optimization: A Tribute to Peter Whittle.* John Wiley & Sons, New York and Chichester, 455–466.

Garcia, D. 2003. Convergence and biases of Monte Carlo estimates of American option prices using a parametric exercise rule. *J. Econom. Dynamics Control* **27**(10) 1855–1879.

Glasserman, P., B. Yu. 2003. Number of paths versus number of basis functions in American option pricing. *Ann. Appl. Probab.* Forthcoming.

Haugh, M., L. Kogan. 2001. Approximating pricing and exercising of high-dimensional American options: A duality approach. Working paper, MIT, Boston, MA.

Haugh, M., L. Kogan. 2004. Approximating pricing and exercising of high-dimensional American options: A duality approach. *Oper. Res.* **52**(2) 258–270.

He, H. 1990. Convergence from discrete- to continuous-time contingent claim prices. *Rev. Financial Stud.* **3** 523–546.

Jamshidian, F. 1997. LIBOR and swap market models and measures. *Finance Stochastics* **1** 293–330.

Jensen, M., M. Svenstrup. 2002. Efficient control variates and strategies for Bermudan swaptions in a Libor market model. Working paper, Aarhus School of Business.

Kaniel, R., S. Tompaidis, A. Zemlianov. 2003. Efficient computation of hedging parameters for discretely exercisable options. Working paper, University of Texas at Austin, Austin, TX.

Lamberton, D., B. Lapeyre. 1996. *Introduction to Stochastic Calculus Applied to Finance.* Chapman & Hall/CRC, Boca Raton, FL.

Li, B., G. Zhang. 1996. Hurdles removed. Working paper, Merrill Lynch, New York.

Longstaff, F., E. Schwartz. 2001. Valuing American options by simulation: A simple least-squares approach. *Rev. Financial Stud.* **14** 113–148.

Miltersen, K. R., K. Sandmann, D. Sandermann. 1997. Closed-form solutions for term structure derivatives with lognormal interest rates. *J. Finance* **52** 409–430.

Mitchell, A., D. F. Griffiths. 2001. *Finite Difference and Related Methods for Differential Equations.* John Wiley & Sons, New York.

Rasmussen, N. 2002. Efficient control variates for Monte-Carlo valuation of American options. Working paper, Aarhus School of Business.

Raymar, S., M. Zwecher. 1997. A Monte Carlo valuation of American call options on the maximum of several stocks. *J. Derivatives* **5**(1) 7–23.

Rogers, C. 2002. Monte Carlo valuation of American options. *Math. Finance* **12** 271–286.

Svenstrup, M. 2002. On the suboptimality of single-factor exercise strategies for Bermudan swaptions. Working paper, University of Aarhus.

Tilley, J. 1993. Valuing American options in a path simulation model. *Trans. Soc. Actuaries* **45** 83–104.

Tsitsiklis, J., B. Van Roy. 1999. Optimal stopping of Markov processes: Hilbert space theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives. *IEEE Trans. Automatic Control* **44** 1840–1851.

Tsitsiklis, J. N., B. Van Roy. 2001. Regression methods for pricing complex American-style options. *IEEE Trans. Neural Networks* **12**(4) 694–703.