Convergence and Biases of Monte Carlo Estimates of American Option Prices using a Parametric Exercise Rule

Diego García*

Amos Tuck School of Business Administration, Dartmouth College†

First draft: February 1998

September 28, 2001

*I would like to thank Steve Evans for all his advice. A previous version of this paper circulated as "A Monte Carlo Method for Pricing American Options." This paper has benefited from comments by Leif Andersen, Jonathan Berk, Mark Broadie, Carl Luft, Sanjiv Das, Mark Rubinstein, and Nicholas Wonder; as well as participants in RISK Computational and Quantitative Finance (New York, 1999), the Seventh Global Finance Conference (Chicago, 2000), the First International Research Conference in Financial Risk Management (London, 2000), the Workshop on Financal Mathematics (Stanford University, 2000), and the 4th International Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing conference (Hong-Kong, 2000). The latest version of this paper can be downloaded from `http://diego-garcia.dartmouth.edu/research.html`.

†Hanover, NH 03755-9000. Email: `Diego.Garcia@Dartmouth.edu`. Phone: 1-603-646-3615.

Abstract

This paper presents an algorithm for pricing American options using Monte Carlo simulation. The method is based on using a parametric representation of the exercise boundary. It is shown that, as long as this parametric representation subsumes all relevant stopping-times, error bounds can be constructed using two different estimates, one which is biased low and one which is biased high. Both are consistent and asymptotically unbiased estimators of the true option value. Results for high-dimensional American options confirm the viability of the numerical procedure. The convergence results of the paper shed light into the biases present in other algorithms proposed in the literature.

JEL classification: C15, C63, G13.

Keywords: Monte Carlo simulation, American options, multiple state variables.

1 Introduction

The issue of pricing American securities has received a tremendous amount of attention in the last 25 years, ever since the Black-Scholes model was introduced in 1973. Most of the derivative securities traded are American style securities, so the need for a method that can generalize Black-Scholes analysis to allow for early exercise opportunities has been the subject of hundreds of research projects. The developments in the theoretical front have succeeded in characterizing the price of the American assets in the same way as European securities (see Duffie (1992), Karatzas (1988)). The main difficulty when dealing with American securities is numerical. Cox, Ross, and Rubinstein (1979) show how to use a binomial tree to deal with the early exercise feature of standard options written on a stock that follows a log-normal process. Their algorithm is a special case of what are known as "lattice methods."[1] These methods base the approximation on an appropiately chosen finite state Markov chain which converges in distribution to the continuous-time process of interest. The approximation to the optimal control is built by working backwards through the lattice formed by the state variables. These methods have the drawback that computing time grows exponentially in the number of state variables.

The focus of this paper is in the use of Monte Carlo methods to price American options. When the dimensionality of the problem is roughly higher than three, simulation has proved to be a successful technique in the derivatives pricing literature, as well as a method of solution to partial differential equations using the Feynman trick.[2] Simulation methods, introduced by Boyle (1979) into the option pricing literature, have become increasingly popular for pricing complex derivative securities (see Geweke (1996) for an overview of

[1]See Kushner and Dupuis (1992) for a complete treatment of the use of lattice methods to solve general optimal control problems. He (1990) presents convergence results for European derivatives that depend on n underlying assets under the assumption of complete markets. Amin and Khanna (1994) extend He's results to the case of American assets.

[2]The use of Monte Carlo methods for pricing derivative securities is crucial not only for high-dimensional options. Consider the case of barrier or lookback options. Standard contracts specify that the maximum value of the underlying asset, which influences the payoff of these derivative securities, should be calculated at the end of each day. Therefore, continuous time analytical approximations (see Gao, Huang, and Subrahmanyam (1998)) are bound to have substantial biases, since the distributional properties of the maximum of Brownian motion over an interval $[0, T]$ are very different than those over a finite set of points $\mathcal{T}$ (see Nahum (1999) for a thoroughout discussion of these biases). Since lattice methods can not handle well the discontinuities of maximum values of the underlying assets, it seems like Monte Carlo methods are the best suited candidate for the pricing of this type of options.

Monte Carlo methods in Economics). Their main advantage is that computational requirements grow linearly with the number of state variables: these algorithms are the most popular way to approximate high dimensional probability measures. For some time, it was argued that the Monte Carlo algorithm could not be extended to price American options due to its "forward nature." In the last few years we have seen that this is not the case.

Starting with Bossaerts (1989), different authors have modified the Monte-Carlo algorithm for European options to deal with the early exercise feature of American securities: see Boyle, Broadie, and Glasserman (1997) for a survey of the different methods; Fu et al. (1999) for a numerical study of the relative advantages of the different algorithms; and Pedersen (1999) for a comparative analysis of the competing approaches for the case of swaptions, a specially difficult American option pricing problem. Broadie and Glasserman (1997a, 1997b) propose two algorithms that derive the optimal exercise strategy recursively (see also Boyle, Kolkiewicz, and Tan (2000)). In these papers the authors suggest using two different estimators, one with a high-bias and one with a low-bias in finite samples, thereby producing a *bias-free* method for pricing American options. Carriere (1996), and Longstaff and Schwartz (2001) solve the problem estimating the continuation value of the American option through regression techniques.[3] The papers by Tilley (1993), Barraquand and Martineau (1995), and Raymar and Zwecher (1997) incorporate different aspects of the usual backwards induction algorithm by stratifying the state space and finding the optimal exercise decision in each of the subsets of the state variables.

This research project studies in detail and refines a "parametric approach"[4] to the American option pricing problem using simulation (Bossaerts (1989), Li and Zhang (1996), Grant, Vora, and Weeks (1996)). In essence, these algorithms start by representing the early decision rule by a finite number of parameters, and then find the American option price by maximizing, over the parameter space, an approximation to the price of the security (where the approximation is obtained using Monte-Carlo methods). This produces both an estimate of the optimal exercise decision and of the American option price.

An important observation is that, as long as the parameterization of the early exercise decision is complete,[5] the estimates of the American option price generated using the above algorithm will be biased high in finite samples (Theorem 1). I introduce a new estimator

[3]See also Tsitsiklis and Roy (2001).

[4]For an application of this method to swaptions see Andersen (1999) and Pedersen (1999). For an algorithm close in spirit to the one in this paper see also Ibanez and Zapatero (1999).

[5]The precise requirement is that all relevant stopping times are considered with the parametric representation of the exercise strategy, see section 2.2.

which is also shown to be consistent and has a low-bias in finite samples (Theorem 2). Putting these two theorems together generates a bias-free range for the American option price. This result has the same flavor as the results in Broadie and Glasserman (1997a, 1997b), who use completely different numerical schemes. Therefore this paper gives the precise conditions under which a bias-free algorithm can be constructed using "parametric" simulation techniques, pointing out that there may exist significant biases in the prices reported in the literature.

The numerical method presented in the paper consists of two parts: (1) an *optimization stage*, in which the optimal early exercise decision is estimated; and (2) and *valuation stage*, in which the actual option price is calculated. In the optimization stage, the algorithm searches for the optimal early exercise rule among a large space of candidate strategies for exercising the option before maturity, fixing a set of simulated paths. This generates a first estimate of the price of the American option. In the valuation stage, I calculate the value of the derivative using a standard Monte Carlo algorithm with a new set of random paths, using as the exercise strategy the one estimated in the previous step. If the same random paths are used in (1) and (2), the resulting estimated price will have a high bias. Otherwise it will be biased low. This simple characterization of the biases seems to be common to all algorithms proposed in the literature.

In the actual implementation of the method, an approximation to the exercise boundary is required in many cases of interest (in this paper I used a full parameterization only for one-dimensional problems). This causes a low-bias in the estimates, which has the same flavor as the one caused by the bucketing of Barraquand and Martineau (1995), and the approximation used in the estimation of the continuation value function in Longstaff and Schwartz (2001).[6] The convergence results suggest that the bias in these simulation methods can be decomposed into: (a) an approximation bias, which causes the estimated prices to be lower than the true value; (b) a Monte Carlo bias, which depends on whether the same paths are used on the optimization run and the evaluation stage. Figure 1 presents a graphical display of the biases of the estimators proposed in this paper.

The numerical section of the paper shows that even very parsimonious representations of the early exercise strategy will generate substantial high-biases. In section 3, I price three types of American options: standard options, maximum options on two stocks, and

[6]Of course, if we do not have an exact parametric representation of the early exercise rule, then the high-bias estimator may not have a high-bias. Namely, if we search over a restricted set of all possible exercise strategies the bias will be of unknown sign. If we insist on knowing the bias of the estimated prices, then we will want to rely on the low-bias estimator when it is not possible to represent the possible stopping times parametrically.

maximum options on five stocks; with different numbers of early exercise dates (ranging from 3 to 50 dates). In all cases I restrict the parametric representation of the relevant stopping times to have at most four parameters. This is (extremely) suboptimal, but the purpose is to illustrate the presence of substantial high biases even when we are far away from considering all possible exercise rules. It is shown that the approximation bias is less than 0.50% of the option value for all cases considered.

As long as the approximation bias is small, high biases are likely to be present, or could be generated, in existing techniques (e.g. with the methods in Barraquand and Martineau (1995), Raymar and Zwecher (1997), Longstaff and Schwartz (2001), Ibanez and Zapatero (1999)). Given the results of this paper and the sample sizes typical in previous studies, these biases could be up to 5% of the option value (a much larger fraction of the early exercise premium). This observation brings forward a criterion for comparison of existing algorithms in the literature: the size of this approximation bias. The numerical results present evidence that the approximation bias for the algorithm in this paper is negligible.

The methodology of this paper requires the existence of a parametric representation of the relevant stopping times. This may be counterfactual in some cases, since we may not have a good understanding what the early exercise rule should depend on. In some cases[7] it is possible to use properties of the payoff function (e.g. monotonicity) to obtain information about how to represent the exercise strategy. An alternative is to let the exercise rules be dependent on some abstract state variables.[8] Although the convergence properties for this case are unknown, the implementation of parametric Monte Carlo for the case of swaptions (e.g. Andersen (1999)) seems to provide practical support for this type of approach.

Section 2 presents the algorithm and the convergence results. Section 3 considers several numerical examples. Readers less interested in the technical details may want to start with the illustrative example in section 2.1 and skip sections 2.2 and 2.3.

[7]In all the examples considered in this paper it is known what the exercise rule should depend on (i.e. how to partition the state space). Other important cases for which this partitioning is easily obtained are several types of exotic options (barrier, lookback, Asian). For standard options in multi-dimensional problems it also seems feasible to partition the relevant state space in a tractable manner (see Ibanez and Zapatero (1999) for a model with stochastic interest rates).

[8]This is the approach that Longstaff and Schwartz (2001) use when arguing the flexibility and generality of their method.

2 A parametric Monte Carlo method for American options

2.1 An illustrative example

Table 1 presents a ficticious example taken from Longstaff and Schwartz (2001). The goal is to price an American put with a strike price of 1.1 and three different early exercise dates. The initial stock price is 1, and the discount factor per period is 0.9718. Longstaff and Schwartz (2001) generate 8 different random stock price paths in order to apply Monte-Carlo to price this put option. Figure 2 plots the stock price paths.

The main idea of the algorithm presented in the paper is to represent the early exercise strategy by a finite number of parameters. In this particular case the early exercise strategy can be represented by a set of stock price levels θ_t such that if the stock price at date t is below θ_t, then the put is exercised at that date ($t = 0, 1, 2, 3$), and otherwise it is held for one more period.

Table 1 shows how to calculate the value of the put following a given early exercise policy: exercise at $t = j$ if the stock price is below θ_j, where $\theta_0 = 0.95$, $\theta_1 = 0.95$, $\theta_2 = 1$, $\theta_3 = 1.1$. The value of the put under this early exercise strategy is 0.0551. Under this strategy, early exercise at date 1 occurs for four of the stock price paths. There is no early exercise at date 2.

Note that the above exercise strategy is far from optimal. If one uses the values $\theta_1 = 0.75$ and $\theta_2 = 0.8$ in the above example, the value of the put (for the same 8 paths) would be 0.1118. Early exercise at date 2 occurs only for one of the paths. The pricing algorithm in the paper calls for searching over all possible exercise strategies until one finds the one that yields the highest value for the American put. It is straightforward to check that the strategy with $\theta_0 = 0.95$, $\theta_1 = 0.75$, $\theta_2 = 0.8$ and $\theta_3 = 1.1$ is an optimal early exercise strategy: no other would achieve a higher value for the put given the 8 stock price paths considered. It is also worth noticing that this strategy is not unique: since the value function has discountinuities as θ_i is changed, many other exercise strategies would maximize the option value.

The above exercise gives us a first estimator of the value of the option. As shown in Theorem 1, this estimator has a high-bias. The implementation of the method would be complete by generating a new set of random paths, and using the exercise strategy $\theta_1 = 0.75$ and $\theta_2 = 0.8$ in order to evaluate the value of the American option for this new set of simulated trayectories. As shown in Theorem 2, the new estimated option price will have a low bias.

2.2 The American option pricing problem

Let $X_t \in \mathbb{R}^n$ be a Markov process with initial state X_0, defined at times $t = 0, \cdots, T$. I will assume that this set of random variables contains all information that is payoff relevant for the derivative security under consideration. There are $T+1$ dates in this economy: $\mathcal{T} = \{0, \ldots, T\}$. The filtration $\mathbb{F} = \{\mathcal{F}_t\}$ represents the information available to investors at each date. There exists a riskless security that yields a constant rate of return of r per unit of time[9].

An American security is defined by a payoff function $h(X_t, t) : \mathbb{R}^n \times \mathcal{T} \to \mathbb{R}$ and a expiration date T[10]. The American option pricing problem is to find

$$V_0 = \max_{\tau \in \mathcal{T}(0)} \mathbb{E}\left[e^{-r\tau} h(X_\tau, \tau)\right];$$

where $\mathcal{T}(t)$ denotes the set of stopping times[11] taking values in $\{t, t+1, \ldots, T\}$; and the expectation is taken with respect to the unique equivalent martingale measure $\mathbb{Q}$ defined using the riskless security as the numeraire.[12]

In order to characterize the optimal exercise decision, it is customary to introduce the Snell envelope of the discounted payoff

$$W_t = \max_{\tau \in \mathcal{T}(t)} \mathbb{E}\left[e^{-r(\tau - t)} h(X_\tau, \tau) | \mathcal{F}_t\right].$$

The optimal stopping time is given by (see Duffie (1992)): $\tau^* = \min\{t : W_t \leq h(X_t, t)\}$; i.e. the option should be kept alive while the value of holding it, W_t, is higher than the value of immediate exercise, $h(X_t, t)$; otherwise it is optimal to exercise early. The possible stopping times can be described without loss of generality to be of the form $\tau = \inf\{t : X_t \in \mathcal{E}_t\}$, where $\mathcal{E}_t \subset \mathbb{R}^n$ is a subset that defines the values of the state variables for which

[9] It is straightforward to generalize the algorithm to allow for stochastic interest rates.

[10] Some derivative securities pay some dividend process during their life. It is straightforward to extend the method to deal with this more general case.

[11] A stopping time is a random variable $\tau : \Omega \to \mathcal{T} \cup \{\infty\}$ such that for any t the event $\{\omega : \tau(\omega) \leq t\}$ is measurable with respect to $\mathcal{F}_t$.

[12] This is equivalent to assuming that the American option that we price is redundant given the set of traded assets. This is the standard approach in most of the asset pricing literature that deals with derivatives. The algorithm introduced in this paper could easily be adapted to finding upper and lower bounds on the derivative asset in an incomplete markets economy, under the assumption that the new asset does not influence the prices of existing securities.

early exercise occurs at date t. Let $\mathcal{E}_t^*$ denote the set of state variables for which it is optimal to exercise at date t.

The Markovian assumption allows us to write $W_t = g(X_t, t)$ for some function $g(\cdot)$. Therefore the optimal early exercise region is given by

$$\mathcal{E}_t^* = \{x \in \mathbb{R}^n : g(x,t) \leq h(x,t)\}. \tag{1}$$

The difficulty lies in the fact that we rarely know the exact functional form of the function $g(\cdot)$.[13] The method introduced in this paper is based on observing that these sets have simple structures in most cases of interest. I will make the following assumption in order to prove the convergence of the estimators.

Assumption 1 *The sets $\mathcal{E}_t$ can be written as $\mathcal{E}_t(\theta)$, where $\theta \in \Theta$ are the parameters that represent the exercise decision, and $\Theta \subset \mathbb{R}^K$ is a compact set. This parameterization can be constructed by setting $g(x,t) = G(x,t;\theta)$ for some known function $G(\cdot)$. The set of points $\{x : G(x,t;\theta) = h(x,t)\}$ has $\mathbb{P}$-measure zero for all $\theta \in \Theta$.*

The main issue is whether we can represent the sets $\mathcal{E}_t(\theta)$ by a finite dimensional set Θ.[14] The next example shows that this is possible for the case of an ordinary call option on a stock that follows a log-normal process. The following examples illustrate how it is possible to represent these sets approximately for other important cases.

Example 1 (Ordinary options) *Consider the standard Black-Scholes model, in which the price of the underlying asset follows a log-normal distribution. A call option has payoff function $h(S_t,t) = (S_t - K)^+$. With a finite set of $T+1$ exercise opportunities, it is sufficient to have the parameter space to be of dimension $T+1$ to have a perfect representation of the possible stopping times (see Kim (1990)). To see this, note that the function $g(S_t)$ in this particular case is increasing in S_t, so if equation (1) is satisfied for some $S_t = \theta_t$, then it will be satisfied for all $S_t > \theta_t$. The early exercise boundary can therefore we described without loss of generality by the $T+1$-tuple $(\theta_0, \ldots, \theta_T)$, which gives for each date the*

[13] The only exception may be Bermudan standard options, for which Geske and Johnson (1984) give analytical formulas based on the cumulative distributions of Gaussian random variables.

[14] The rest of assumptions are to assure that the exercise boundary is well behaved, namely that the probability that X_t is in the boundary of the sets $\mathcal{E}_t(\theta)$ is zero. If X_t were a discrete random variable, all the results in the paper would follow through, since Θ could be taken to be a set with a finite number of elements. This would assure uniform convergence of the estimator $V_B(\theta)$ to $V(\theta)$, which is all we need in view of Theorem 4.1.1 of Amemiya (1985).

lowest value of the stock price for which it is optimal to exercise. This example shows that the above assumption is not as restrictive as it may seem. □

Example 2 (Asian options) *The payoff function for a standard Asian call option is* $h(X_t, t) = (\bar{S}_t - K)^+$*, where* $\bar{S}_t$ *is the average stock price at date* t *(where averaging starts at some date* $t_0 < t$*). The exercise decision depends on the two state variables* $X_t = (S_t, \bar{S}_t)$*. The early exercise region is described by equation (1), i.e.* $\{x : g(x,t) \leq (x_2 - K)^+\}$*. In a Black-Scholes world, the function* $g(x,t)$ *is increasing in* x*, so the exercise region at date* t*, a subset of* $\mathbb{R}^2$*, can be described by a function* $f : \mathbb{R} \times \mathcal{T} \rightarrow \mathbb{R}$*, which gives for each value of* S_t*, the values of* $\bar{S}_t = f(S_t, t)$ *for which it is optimal to exercise the option. We can think of approximating the functions* $f(x,t)$*, by some appropiate* $F(x, t; \theta)$*, where* F *belongs to a particular space of well-behaved functions (say cubic splines or Hermite polynomials).* □

Example 3 (Maximum options) *An option on the maximum of two stocks has payoff function* $h(S_t, t) = \max(\max(S_t^1, S_t^2) - K, 0)$*. To describe the optimal early exercise region at date* t *we can break* $\mathbb{R}^2$ *into three different sets: one for which holding the option is optimal, the other two containing the values of the stock prices for which early exercise at time* t *is the decision rule. Broadie and Detemple (1997) Broadie and Detemple (1997) discuss these sets at length. We can use this characterization of the optimal stopping times in order to use the algorithm of this paper. Namely, we can introduce two functions,* $f_1(x, t; \theta)$ *and* $f_2(x, t; \theta)$*, that divide* $\mathbb{R}^2 \times \mathcal{T}$ *in the sets that define different exercise decisions.* □

Define

$$\tau(\theta) = \inf\{t : X_t \in \mathcal{E}_t(\theta)\};$$
$$V(\theta) = \mathbb{E}\left[e^{-r\tau(\theta)} h(X_{\tau(\theta)}, \tau(\theta))\right];$$

where $\tau(\theta)$ is the stopping time rule for the parameters θ, and $V(\theta)$ is the value of the option under the exercise policy dictated by θ. Denote by θ_0 the value that satisfies $V(\theta_0) = V_0$, i.e. the parameter that represents the optimal exercise policy. I will make the following assumptions.

Assumption 2 *The function* $V(\theta) : \Theta \rightarrow \mathbb{R}$ *achieves a unique maximum at* θ_0 *and is continuous and bounded.*

Assumption 3 *The American option under consideration satisfies* $h(X_{\tau(\theta)}, \tau(\theta)) < Y$ *a.s., for some random variable* Y *with* $\mathbb{E}[Y] < \infty$ *which does not depend on* θ*.*

The continuity of the value function with respect to the parameters θ seems to be a mild condition that can easily be checked for each specific application. The last assumption is purely technical. This condition is met for all the options mentioned above, since the underlying asset price provides a bound on the option's payoff, and it is an integrable random variable.

2.3 Convergence and bias of Monte Carlo estimates

The rationale of the Monte Carlo method is simple. If we simulate the underlying economy by choosing, following the law $\mathbb{P}$, B elements from Ω, then, by the Law of Large Numbers, the average of the option values should be close to their expectation. Let $e^{-r\tau(\theta)}h(X^i_{\tau(\theta)}, \tau(\theta))$ be the discounted payoff of the option for a particular realization of the stochastic process X_t, and stopping time $\tau(\theta)$. Define

$$V_B(\theta) = \frac{1}{B}\sum_{i=1}^{B} e^{-r\tau(\theta)}h(X^i_{\tau(\theta)}, \tau(\theta))$$

to be the sum of the discounted payoffs of the American option for B simulations of the underlying uncertainty of the economy following the exercise strategy dictated by θ.

If the option were European, so there is no dependence on θ in the above expression, the estimator would be a consistent and unbiased estimate of the true option price. In the case of American options, the algorithm must also find an estimate for the optimal early exercise rule. The way in which we obtain the parameters that characterize the optimal stopping times may in principle introduce biases in the estimated asset prices.

The optimal exercise rule for the simulated data, $\hat{\theta}_B$, is defined as

$$\hat{V}_B \equiv V_B(\hat{\theta}_B) = \max_{\theta\in\Theta} V_B(\theta); \tag{2}$$

where it is understood that if $\hat{\theta}_B$ is not unique, we would pick any one value. Note that $V_B(\theta)$ is a discontinuous function. Nevertheless the optimization problem of (2) is well defined.[15] The optimum value will belong to a set of points which converge to the true population value. The next theorem gives some properties of the estimator $\hat{V}_B$.[16][17]

[15]The reader may have noted that the setup of the optimization problem above is very similar to that of M-estimators. The simulated data can be taken to be the observations, and the early exercise rule to be the parameters to be estimated. The convergence result is subsumed by the results in simulation-based optimization (e.g. Dupacova and Wets (1988), King and Rockafellar (1993) and Shapiro (1993)). The proof is nevertheless of interest for its simplicity (compared to these references), which makes use of the added structure of the American option pricing problem outlined in the previous section.

[16]All proofs are contained in the Appendix.

[17]Broadie and Glasserman (1997a, 1997b) have proved convergence of the estimated option prices to the true asset value under fairly general conditions on the American option pricing problem under second-moment

Theorem 1 *The estimator $\hat{V}_B$ converges almost surely and in $\mathcal{L}^1$ to $V(\theta_0)$. In particular, the estimator $\hat{V}_B$ is a consistent and asymptotically unbiased estimate of $V(\theta_0)$. For finite B, $\hat{V}_B$ is biased high. The estimate of the early exercise region, $\hat{\theta}_B$, converges almost surely to the true population value θ_0.*

The intuition for the bias in the above result is simple: for finite B the estimated option price will pick a parameter ($\hat{\theta}_B$) such that the estimate of the option value, $V_B(\hat{\theta}_B)$, is higher than with any other parameter value, even the true optimal exercise decision, $V_B(\theta_0)$. Since $\mathbb{E}\left[V_B(\theta_0)\right] = V(\theta_0)$, the estimate $\hat{V}_B$ will have a high bias in finite samples.

The existence of this bias suggests the introduction of an alternative estimator to be able to find probabilistic bounds on the true option price. An obvious candidate can be constructed by using the estimated exercise rule $\hat{\theta}_B$ on a new set of simulated data consisting of b paths. Define:

$$\hat{v}_B^b = \frac{1}{b}\sum_{j=1}^{b} e^{-r\tau(\hat{\theta}_B)} h\left(X^j_{\tau(\hat{\theta}_B)}, \tau(\hat{\theta}_B)\right) \tag{3}$$

where the simulated data X_t^j are independent of that used to estimate $\hat{\theta}_B$. It is natural, although by no means necessary, to consider the case where the number of paths used to estimate $\hat{\theta}_B$ is equal to those used to estimate $\hat{v}_B^b$. I will use the abbreviated notation $\hat{v}_B$ when this is the case.

The next theorem shows that $\hat{v}_B$ is a consistent estimate of the true option price, that it is asymptotically unbiased, and that for finite B this estimate is biased low.

Theorem 2 *The estimator $\hat{v}_B^b$ converges in probability and in $\mathcal{L}^1$ to $V(\theta_0)$. For finite B, the estimator $\hat{v}_B^b$ is biased low.*

The nature of the biases is very different. $\hat{V}_B$ has a high bias caused by the fact that the true measure induced by X_t is being misrepresented when a finite number of simulations are used and the optimization routine takes advantage of this. The bias in $\hat{v}_B^b$ is not driven by the misrepresentation of the measure, but rather by the distributional properties of $\hat{\theta}_B$. The results in the numerical section suggest that the distribution of the estimates of the early exercise parameters is usually tight with respect to its effect on the price of the derivative. As it will be shown in the numerical section, these biases are also very different

assumptions. Bossaerts (1989) and Longstaff and Schwartz (2001) prove convergence of their algorithms relying heavily on the one-dimensional nature of their problem. I provide a simple convergence proof for both estimators using the sample analog of the American option price. It should be noted that the conditions for convergence in this paper are the weakest in the literature, since only finite first moments are assumed.

in magnitude: the size of the bias of the low estimator is smaller than that of the high estimator.[18]

In the numerical results we will report the following confidence bin:[19]

$$\left[\hat{v}_B^b - z_{\alpha/2} s(\hat{v}_B^b), \hat{V}_B + z_{\alpha/2} s(\hat{V}_B)\right]; \tag{4}$$

where z_α is the $1-\alpha$ quantile of the standard normal distribution, and $s(\hat{V}_B)$ and $s(\hat{v}_B^b)$ are the standard deviations of $\hat{V}_B$ and $\hat{v}_B^b$. This error bound heuristically comes by appealing to the Central Limit Theorem.[20] Also by standard asymptotic arguments it is to be expected that

$$\sqrt{B}\left(\hat{\theta}_B - \theta_0\right) \sim N\left[0, V_{\theta\theta'}^{-1} V_\theta V_\theta^\top V_{\theta\theta'}^{-1}\right]$$

where $V_{\theta\theta'}$ denotes the matrix of second derivatives, and V_θ the vector of first derivatives. The variance-covariance matrix in the above expression can be estimated: (i) using sample analogs of each quantity (which under appropiate conditions on second moments and differentiability of $V(\cdot)$ will be consistent) or (ii) by using the sample properties of $\hat{\theta}_B$ (when the algorithm is runned multiple times).

Another interesting observation is that these first and second derivatives may be very different depending on the parameters of the option. Consider the case of an out-of-the-money option (for which its current exercise value is zero). It is likely that it would be difficult to get a low variance for the $\hat{\theta}$ with this type of option, since very few of the simulated paths will lead to early exercise, so the option price is relatively insensitive to the parameters of the exercise boundary. The algorithm would improve if we use an estimate for θ obtained through an option that is in-the-money, which we can do simply by altering the initial value of the stock (by noting that the optimal boundary does not depend on the initial state). The variance of θ will be smaller, which will result in a lower expected bias for $\hat{v}_B^b$. As a matter of fact, this line of argument leads to a criterion of optimality for the initial value of the state variables in the simulation: those which minimize the variance of the estimates of the early exercise decision. I present some numerical evidence in the next

[18]A similar result has also been reported in Raymar and Zwecher (1997), that note that their 3-stage estimator is "better behaved" than their 2-stage estimator.

[19]See Fishman (1996) for different methods to assess the errors of Monte Carlo simulations.

[20]I do not present a proof of this result, since the focus of the paper is in studying the biases of the algorithm. The setup is not identical to those of standard estimators (see Amemiya (1985)), since $V_B(\theta)$ is not a continuous function of the parameters. Nevertheless, the assumptions on the continuity of $V(\theta)$ and its derivatives, together with a finite second moment on $h(X_\tau, \tau)$ for all τ seem to be sufficient to guarantee asymptotic normality of $\hat{V}_B$, $\hat{v}_B$ and $\hat{\theta}_B$. See any of the references in footnote 15.

section that follows this idea. A rigorous analysis of this "variance-reduction" technique[21] is an interesting topic for future research.

2.4 Numerical algorithm

The following steps give a formal outline of the algorithm:

1. Parameterize the exercise boundary, using a finite number of parameters $\theta \in \Theta$. Fix some initial values for these parameters, say θ^0.
2. Generate B paths of X_t. Let X_t^j denote the realized path of the simulation j at time t.
3. Evaluate an approximation to the value of the option, $V_B(\theta^{(i)})$, using the stopping time $\tau(\theta^{(i)})$ and the simulated paths of X_t. Namely:
   (a) For each $t = 0, \ldots, T$ define $\tau_j = \min\{t : X_t^j \in \mathcal{E}_t(\theta^{(i)})\}$. Note that $\tau_j = t$ if the exercise rule that θ represents calls for early exercise at time t for the realization of the path $\{X_t^j\}$.
   (b) Calculate
   $$V_B(\theta^{(i)}) = \frac{1}{B}\sum_{j=1}^{B} e^{-r\tau_j} h(X_{\tau_j}^j, \tau_j).$$
4. Change the parameters of the exercise boundary and go back to the evaluation step. Continue until the algorithm finds a maximum for the function $V_B(\cdot)$.
5. Generate b paths of X_t, independent of those used to calculate $\hat{\theta}_B$. Evaluate
$$\hat{v}_B^b = \sum_{i=1}^{b} e^{-r\tau(\hat{\theta}_B)} h\left(X_{\tau(\hat{\theta}_B)}^i, \tau(\hat{\theta}_B)\right).$$

All the optimization routines presented in this paper use a variant of the simplex method (see Press, Teukolsky, Vetterling, and Flannery (1992)). This choice of optimization algorithm responds to the discontinuity of the value function, together with the presence of multiple local maxima.

[21] Note that this improvement to the algorithm is different from the standard variance-reduction techniques which only concern themselves with the variance of the actual option price, not the early exercise strategy. The main advantage of variance-reduction of the estimated stopping times is to reduce the bias of the option prices, not their variance.

3 Numerical results

In the following section I will consider pricing options for which the relevant underlying state variables are log-normally distributed stock prices. The simulated paths were generated using

$$S_{t+1} = S_t e^{\left(r-\delta-\sigma^2/2\right)\Delta+\sigma\sqrt{\Delta}Z_{t+1}}$$

where Z_{t+1} is a standard normal random variable, Δ is the time step used in the approximation, r is the annual interest rate, δ is the dividend yield of the underlying asset, and S_t denotes the stock price at date t.

3.1 Standard American options

In this section, ordinary call options with four different exercise dates ($t = 0, T/3, 2T/3, T$) are studied. The payoff for an ordinary call is $h(S_t, t) = \max(S_t - K, 0)$. This type of option is of particular interest, since it will be simple to find a representation of the exercise regions using a small number of parameters. No approximation bias should be present in our results. The exercise regions can be described by $\mathcal{E}_{T/3}(\theta) = [\theta_1, \infty)$ and $\mathcal{E}_{2T/3}(\theta) = [\theta_2, \infty)$. The exercise region at maturity is $S_T \in [K, \infty)$. The two parameters $\theta = (\theta_1, \theta_2)$ represent all the relevant stopping times in the American option problem, i.e. all the decisions that the option holder can make before maturity.

I use the following base case parameters, so the early exercise feature of the options has some significance: $r = 0.05$, $\delta = 0.10$, $\sigma = 0.20$. The option has time to maturity of 1 year, and a strike price of $K = 100$. These values are taken from Broadie and Glasserman (1997a). The true value of the option in this case can be found by the formulas given in Geske and Johnson (1984).

Table 2 presents the results of running a simple optimization for different starting values for the stock price, with $B = 40,000$.[22] Computational time for each of the rows is on the order of half a minute. The relative errors, defined as $(\hat{v}_B - P)/P$, are small compared to the confidence bins (given by equation (4) with $\alpha = 0.05$). The algorithm performs better for options that are in-the-money, i.e. those for which the current exercisable value is positive. This is due to the fact that the early exercise decision plays a bigger role when the derivative is in-the-money. For an option with an initial stock value of $S_0 = 70$, the Monte Carlo method is going to have few paths that reach 110 range, which is where the

[22] For the option with $S_0 = 120$ the continuation value was estimated as 18.90, with a confidence interval [18.78, 18.98].

optimal exercise boundary lies. This is going to result in a higher bias for V_B, since the measure is not well represented in the relevant range. Note that this also results in a high variance for $\hat{\theta}$, which will further cause $\hat{v}_B$ to have a larger low bias. Estimating $\hat{\theta}$ using an initial value that makes the early exercise decision relevant (say $S_0 = 110$) and then using this estimate to price the option with $S_0 = 70$ produced substantially better results.

The algorithm is easily generalized to a larger set of early exercise opportunities. First note that the exercise regions can be shown to be of the form $\mathcal{E}_t = [g(t), \infty)$ for some function $g(t)$ (see Kim (1990)). I will use piecewise cubic Hermite polynomials (see Press, Teukolsky, Vetterling, and Flannery (1992)) to approximate this exercise boundary $g(\cdot)$. In order to have a parameterization of the boundary that is easy to interpret, I fix some points $\{t_0, \ldots, t_K\}$, and use as parameters the function values at those points, i. e. $f(t_i) = \theta_i$. The boundary function $g(t; \theta)$ is defined to be the piecewise cubic Hermite polynomial that interpolates through the points $f(t_i)$. After some experimentation, I settle on specifying this type of Hermite polynomial with the points $x = 0, 0.7T, T$, where $f(0) = \theta_1$, $f(0.7T) = \theta_2$, and $f(T) = K$.[23] This seems to be a rich functional form for the exercise boundary in this problem.[24]

I apply the algorithm to price an American call with $K = 100$, $S = 110$, $T = 0.5$, $r = 0.03$, $\delta = 0.07$, $\sigma = 0.2$.[25] The true price can be calculated very accurately using existing techniques (see Ju (1998)). Exercise is allowed at 40 different equally spaced dates.

The method works well for standard American call options. Figure 3 presents the estimated early exercise boundaries for two values of B. The variability of the estimated boundaries is quite small for as few as 16000 simulations. In order to see the asymptotic behavior of our estimates, figure 4 presents the variability of the option prices ($\hat{V}_B$ and $\hat{v}_B$) for different simulation values. With $B = 16000$ the confidence bounds have a size on the order of 15 cents, and the bias is smaller than 4 cents for both estimators. The finite sample bias of $\hat{v}_B$ is almost negligible. Nevertheless, it is worth noticing that $\hat{v}_B$ has a bias that was not present in the previous application, since there is an extra low bias (on the order of one cent for this particular option) in the estimator due to the approximation of

[23]It can be shown (Kim (1990)) that $g(T) = K$ when $\delta \geq r$.

[24]I experimented with more general functional forms for the early exercise frontier with similar results to those reported in the paper.

[25]The choice of a short-term option poses the most difficulties for the algorithm under consideration, since the early exercise boundary will be fairly non-linear. Simulations done by the author on long-term options ($T = 3$) yield extremely good results by simply using a linear exercise boundary.

the exercise boundary by the specific functional form that I use.[26]

3.2 Maximum options on 2 stocks

In this section I consider the pricing of an American call option on the maximum of 2 stocks. This option yields the payoff $h(X_t, t) = \max(\max(S_t^1, S_t^2) - K, 0)$. This type of derivative has become popular as a test of Monte Carlo methods in higher dimensions, since accurate prices can be obtained in 2 dimensions, and the early exercise decision is non-trivial. Some general properties of the optimal exercise regions were studied in Broadie and Detemple (1997). Theoretically, the decision of whether to exercise at date t or not depends on the values of the two stock prices at that date. Motivated by their qualitative charaterizations, and taking advantage of the symmetry of the problem[27] I choose the following parsimonious representation of the sets that will determine the optimal stopping time:

$$\tau = \inf\{t : \max(\max(S_t^1, S_t^2) - K, 0) > \theta_t^1; |S_t^1 - S_t^2| > \theta_t^2\}$$

Figure 6 plots these sets together with the density of the two stock prices at each date. The first parameter, θ_t^1, measures how deep in the money the option should be in order for early exercise to be optimal. The second parameter, θ_t^2, measures the "push" from the asset with a lower value to the maximum price. If the second asset is far away from the maximum, it is more likely that early exercise is optimal. This simplification introduces a low bias in the estimates.

I consider pricing this type of call option with base parameters $K = 100$, $\sigma_1 = \sigma_2 = 0.20$, $\rho = 0.3$, $r = 0.05$, $\delta = 0.10$, $T = 1$. I will start by considering Bermudan options, for which there are 4 different early exercise opportunities, at $t = 0, T/3, 2T/3, T$ (this is the same case considered in Broadie and Glasserman (1997a)).

Figure 5 presents the results of running a number of different repetitions of the algorithm for initial values $S_0^1 = S_0^2 = 120$. The bias for $\hat{V}_B$ is substantial for small values of B. The bias for the low estimator seems to be consistently within ten cents of the true option value.

In order to generate estimates for a different range of options, I will use estimates for $\hat{v}_B^b$ obtained from a single run of $\hat{V}_B$. Using $B = 100,000$ a value $\hat{V}_B = 26.08$ was obtained

[26] We are comparing our prices to those produced by a method that considers continuous exercise, which makes the algorithm look worse than what it should: the right price to compare with would be that of an option with a fixed finite number of early exercise opportunities (equal to the number of time steps in the simulation).

[27] Symmetry is only used for comparability with those results in the literature. In principle I could parameterize these sets in a richer manner. The results that follow suggest that this may not be necessary in order to get good approximations to the American option prices.

(the true value is 25.98), with a standard error of 0.051. The estimated parameters are $\hat{\theta}_1^1 = 15.01$, $\hat{\theta}_2^1 = 12.04$, $\hat{\theta}_1^2 = 14.07$ and $\hat{\theta}_2^2 = 9.03$. The optimization algorithm took about 5 minutes to run.[28] Figure 6 presents the estimated exercise regions as well as the density function of the two stocks.

Table 3 presents the resulting option values, using $\hat{v}_B$, for different stock prices using the estimated exercise regions of the previous simulation. Each of the rows in the table took about 90 seconds to be calculated. There is a low bias in the estimates, so that some of the confidence intervals do not contain the true option value. This bias is nevertheless of small practical importance (no larger than 2 cents for any of the options), and could be corrected by specifying a more general exercise region. Using 400,000 paths, the bias present in both estimators was found to be less than one cent. The accuracy of the method is overall surprisingly good, given the crude approximation of the exercise boundary that I use.

3.3 Maximum options on 5 assets

I proceed to price the same type of maximum option, where the maximum is over the prices of 5 different assets. For this type of option, in order to represent the early exercise strategy, we would need to specify sets in $\mathbb{R}^5$, since the optimal decision at time t depends on the levels of all five of the stock prices at that date. In order to find a parsimonious representation of the possible stopping times, I reduce the problem dimensionality by considering only two state variables in the early exercise decision: the maximum price at time t, $S_t^{(1)}$, and the second ordered statistic, $S_t^{(2)}$. The stopping times that I consider are of the form

$$\tau(\theta) = \inf \left\{ t \geq 0 : \max(S_t^{(1)} - K, 0) > \theta_t^1 ; S_t^{(1)} - S_t^{(2)} > \theta_t^2 \right\}$$

The parameters have a similar interpretation as in the 2-dimensional case. θ_t^1 measures the moneyness of the option. The immediate exercise value of the option needs to be high enough for early exercise to be optimal. θ_t^2 measures the "push" by the second-order statistic to the maximum. The higher the second order statistic is, the less likely it will be to exercise the option early. There are other obvious ways to describe the early exercise decision that merit attention. The third order statistic, or the average of the second and the third order statistics, seem to be good candidates to be included as part of the description of the stopping times.

[28]It should be noted that in these simulations I did not use any variance control techniques, which have proven to be very successful for these types of problems. Computational time could be significantly reduced.

I investigate the base case scenario considered in Broadie and Glasserman (1997b). All assets have the same volatility, $\sigma = 0.2$, and dividend yield $\delta = 0.10$. The correlation among the assets is 0. The risk-free rate is $r = 0.05$, the time to maturity of the option is $T = 3$, and the strike price is $K = 100$. I pick the starting values $S_0^1 = 130$ and $S_0^2 = \cdots = S_0^5 = 90$. For these parameters the option is in-the-money, and the difference between the first order and the second order statistic will be substantial throughout the life of the security. Therefore, both types of early exercise parameters will have significant effects on the early exercise decision for most of the generated paths. The estimated option value is 30.749, with the early exercise parameters $\theta_1^1 = 26.16$, $\theta_2^1 = 18.22$, $\theta_1^2 = 20.96$, $\theta_2^2 = 13.04$.

For this option Broadie and Glasserman (1997b) report option prices for the initial values $S_1 = \cdots = S_5 = 90, 100, 110$. Using the parameters from the previous optimization run, I run the algorithm to calculate $\hat{v}_B^b$ with 100,000 paths. Each row of Table 4 took less than 2 seconds.[29] The low-biases found for in-the-money maximum options with this parameter values were all less than 0.3% of the option value.

The algorithm can be easily extended to deal with an arbitrary number of early exercise dates. Consider the following stopping times

$$\tau(\theta) = \inf\left\{t \geq 0 : \max(S_t^{(1)} - K, 0) > f_1(t, \theta^1); S_t^{(1)} - S_t^{(2)} > f_2(t, \theta^2)\right\}.$$

The functions $f_i(t, \theta^i)$ are a smooth time interpolation of the "moneyness" and "pull" parameters of the previous example. As in the case of the standard American call, I use piecewise cubic Hermite polynomials to do the time interpolation, specifying these polynomials by the levels at $t = 0, 0.7T, T$, where $f_i(0, \theta^i) = \theta_1^i$, $f_i(0.7T, \theta^i) = \theta_2^i$, $f_i(T, \theta^i) = 0$.

The only reported prices for this type of option as those in Broadie and Glasserman (1997b), who consider nine possible early exercise dates, and the same parameters as in the previous example. Figure 7 presents the estimates for different values of B. The reported biases are all below 1% of the option value. Note that this is a small number even as a proportion of the American option premium reported in Broadie and Glasserman (1997b), while we are using an extremely parsimonious set of stopping times.

[29] It should be noted that not only is the computational effort much smaller, but the results presented in their paper make heavy use of variance reduction techniques, whereas in these examples I have implemented a brute force simulation method.

4 Conclusion

The numerical solution to the American option pricing problem concerns the weak convergence of probability measures. When the number of state variables is large, a researcher usually has no choice but to turn to Monte-Carlo methods. This paper suggests how to construct estimates for the American option values by introducing two estimators, one biased high and one biased low, showing that both are asymptotically unbiased. The numerical results of presented in this paper, together with those in the literature, suggest that very simple representations of the exercise regions can produce very accurate option prices. For all the derivative assets that I have considered, the relevant sets can be expressed by a few number of parameters: they are fairly flat regions.[30] Moreover, the impact on the option value of a slightly wrong representation of the exercise boundary is small. It should be noted that this fact has been used in the development of several highly efficient numerical methods for the pricing of low-dimensional derivatives, by approximating the exercise boundary by constants or exponential functions (Broadie and Detemple (1996), Ju (1998), Gao, Huang, and Subrahmanyam (1998)).

The method presented in the paper is obviously more general than its treatment in optimal stopping time problems. Different optimal control problems can be treated in the same way, and estimates for the optimal controls and value functions can be obtained in a similar fashion (see Krusell and Smith (1996) and Smith (1991) for a Monte Carlo algorithm in a standard stochastic dynamic programming problem in Macroeconomics). If the controls influence the dynamic evolution of the state variables, one will need to use a weak-approximation scheme to generate the state variables (see Kloeden and Platen (1992)), which increases the computational requirements of the algorithm. Preliminary work by the author on the standard optimal investment problem under uncertainty looks very promising.

[30] The algorithm of this paper allows the researcher to investigate the magnitude of the approximation bias by simply increasing the dimension of Θ_K. It is difficult to think of an option for which more than 15-20 parameters would not suffice for a "rich" representation of the exercise boundary. The optimization problem would not be easy to solve (due to the discontinuity of the objective function), but if one looks at the numerical achievements of the non-parametric Statistics literature, one should not see this issue as unsolvable.

Appendix

Proof of Theorem 1.

Consider

$$\hat{V}_B^* = \sup_{\theta \in \Theta^*} \frac{1}{B} \sum_{i=1}^{B} e^{-r\tau(\theta)} h(X^i_{\tau(\theta)}, \tau(\theta))$$

where Θ^* is a finite subset of Θ. For each θ, the function $V_B(\theta)$ is the sum of i.i.d. random variables. By the Strong Law of Large Numbers $V_B(\theta)$ converges almost surely to $V(\theta)$, for each $\theta \in \Theta^*$, and since Θ^* is finite

$$\lim_{B \to \infty} \hat{V}_B^* = \sup_{\theta \in \Theta^*} V(\theta).$$

This argument yields $\liminf_{B\to\infty} \hat{V}_B \geq \sup_{\theta \in \Theta^*} V(\theta)$; but since this is true for any finite set Θ^* we have

$$\liminf_{B \to \infty} \hat{V}_B \geq \sup_{\theta \in \Theta} V(\theta) = V(\theta_0).$$

In order to get an upper bound, first consider $H(X, \theta) \equiv \limsup_{\theta' \to \theta} e^{-r\tau(\theta')} h(X_{\tau(\theta')}, \tau(\theta'))$. The function $H(X, \cdot)$ is upper semicontinuous (see exercise 49 in chapter 2 of Royden (1988)). Upper semicontinuity is equivalent to $H(X, \theta) = \inf_N H_N(X, \theta)$ for some collection of functions such that $H_1(X, \theta) \geq H_2(X, \theta) \geq \ldots$ and $H_N(X, \cdot)$ is continuous (see Lemma 6-9 IV of Taylor (1965)).

I first show that the expectation of the limiting random variable $H(X, \theta)$ equals $V(\theta)$ for all $\theta \in \Theta$. Fox a fixed realization $\{X_t\}$, the function $e^{-r\tau(\theta')} h(X_{\tau(\theta')}, \tau(\theta'))$ has a finite number of discontinuities in θ, caused by the change in the early exercise decision. These discontinuities occur at those points θ where X_t falls at the boundary of at least one of the sets $\mathcal{E}_t(\theta)$. Therefore, for a fixed realization of $\{X_t\}$, $H(\theta, \theta) \neq e^{-r\tau(\theta)} h(X_{\tau(\theta)}, \tau(\theta))$ if and only if X_t happens to fall on one of the boundary points of the sets $\mathcal{E}_t(\theta)$. Since this event has probability zero it follows that $\mathbb{E}[H(X, \theta)] = V(\theta)$.

The random variable $\frac{1}{B} \sum_{i=1}^{B} H_N(X^i, \cdot)$ takes values on the separable Banach space of continuous functions on Θ. By the Strong Law of Large Numbers on this space (see Laha and Rohatgi (1979) Theorem 7.2.1), $\frac{1}{B} \sum_{i=1}^{B} H_N(X^i, \cdot)$ converges almost surely to $\mathbb{E}[H_N(X, \cdot)]$. Hence $\sup_{\theta \in \Theta} \frac{1}{B} \sum_{i=1}^{B} H_N(X^i, \cdot)$ converges almost surely to $\sup_{\theta \in \Theta} \mathbb{E}[H_N(X, \theta)]$, and

$$\limsup_{B \to \infty} \hat{V}_B \leq \lim_{B \to \infty} \sup_{\theta \in \Theta} \frac{1}{B} \sum_{i=1}^{B} H_N(X^i, \theta) = \sup_{\theta \in \Theta} \mathbb{E}[H_N(X, \theta)].$$

Now note that by assumption there exists an integrable random variable Y that is a bound on $H_N(X,\theta)$,[31] so $\mathbb{E}[H_N(X,\theta)]$ converges (as $N \to \infty$) to $\mathbb{E}[H(X,\theta)]$ pointwise. Moreover, as functions of θ, the expectations $\mathbb{E}[H(X,\theta)]$ and $\mathbb{E}[H_N(X,\theta)]$ are continuous. By Dini's Theorem (see Proposition 11 in chapter 9 of Royden (1988)) the convergence of $\mathbb{E}[H_N(X,\theta)]$ to $\mathbb{E}[H(X,\theta)]$ is uniform in θ. This fact, together with the above equation gives

$$\limsup_{B\to\infty} \hat{V}_B \leq \sup_{\theta\in\Theta} \mathbb{E}[H(X,\theta)] = V(\theta_0).$$

In order to show convergence in the mean, note that

$$\mathbb{E}\left[|\hat{V}_B - V(\theta_0)|\right] \leq \mathbb{E}\left[\frac{1}{B}\sum_{i=1}^{B} Y^i\right] + V(\theta_0)$$

since we assumed the existence of random variables Y_i that dominate $e^{-r\tau(\theta)}h(X_{\tau(\theta)},\tau(\theta))$ for all θ. By the Strong Law of Large Numbers $\frac{1}{B}\sum_{i=1}^{B} Y^i$ converges almost surely and in $\mathcal{L}^1$ to $\mathbb{E}[Y]$. Therefore the conditions of the generalized dominated convergence theorem hold (see Proposition 18 from chapter 11 in Royden (1988)), and

$$\lim_{B\to\infty} \mathbb{E}\left[|\hat{V}_B - V(\theta_0)|\right] = \mathbb{E}\left[\lim_{B\to\infty} |\hat{V}_B - V(\theta_0)|\right] = 0,$$

i. e. $\hat{V}_B$ converges to $V(\theta_0)$ in $\mathcal{L}^1$.

Finally, we see that

$$\begin{aligned} \mathbb{E}\left[\hat{V}_B\right] &= \mathbb{E}\left[\sup_{\theta\in\Theta} \frac{1}{B}\sum_{i=1}^{B} e^{-r\tau(\theta)}h(X^i_{\tau(\theta)},\tau(\theta))\right] \\ &\geq \mathbb{E}\left[\frac{1}{B}\sum_{i=1}^{B} e^{-r\tau(\theta_0)}h(X^i_{\tau(\theta_0)},\tau(\theta_0))\right] = V(\theta_0) \end{aligned}$$

so that $\hat{V}_B$ is biased high.

To complete the proof, suppose that $\hat{\theta}_B$ does not converges to θ_0. Then, by the compactness of Θ, there exists $\tilde{\theta} \neq \theta_0$, and a subsequence $\{B_k\}$, such that $\hat{\theta}_{B_k}$ converges to $\tilde{\theta}$. Note that

$$\hat{V}_{B_k} \leq \frac{1}{B}\sum_{i=1}^{B} H_N(X^i, \hat{\theta}_{B_k})$$

for all N.

The previous argument shows that the right-hand side of the above equation converges almost surely to the function $\mathbb{E}[H_N(X,\theta)]$ evaluated at $\theta = \tilde{\theta}$. Letting $N \to \infty$ we get that

$$\limsup_{k\to\infty} \hat{V}_{B_k} \leq \mathbb{E}[H(X,\theta)]\big|_{\theta=\tilde{\theta}} = V(\tilde{\theta}) < V(\theta_0).$$

[31] To see this note that we can always replace H_N by H_N^* in the following way $H_N^*(X,\theta) = \min(H_N(X,\theta), Y)$. Since we have assumed that $\inf_N H_N(X,\theta)$ is dominated by some integrable random variable Y, the sequence H_N^* is bounded by Y.

Since θ_0 was assumed to achieve the unique maximum we have reached a contradiction, so $\hat{\theta}_B$ converges almost surely to θ_0. $\square$

Proof of Theorem 2.

First note that

$$\mathbb{E}\left[\left|\hat{v}_B^b - V(\theta_0)\right|\right] \leq \mathbb{E}\left[\left|\hat{v}_B^b - V(\hat{\theta}_B)\right|\right] + \mathbb{E}\left[\left|V(\hat{\theta}_B) - V(\theta_0)\right|\right].$$

¿From Theorem 1 the second term in the expression above goes to zero.

The Strong Law of Large Numbers implies that

$$\lim_{b\to\infty} \mathbb{E}\left[\left|\hat{v}_B^b - V(\hat{\theta}_B)\right| |\hat{\theta}_B\right] = 0$$

for all values of $\hat{\theta}_B$. Moreover, the above conditional expectation is bounded above by $\frac{1}{b}\sum_{i=1}^{b} Y^i + V(\theta_0)$, so by the generalized dominated convergence theorem

$$\lim_{B,b\to\infty} \mathbb{E}\left[\left|\hat{v}_B^b - V(\hat{\theta}_B)\right|\right] = \lim_{B,b\to\infty} \mathbb{E}\left[\mathbb{E}\left[\left|\hat{v}_B^b - V(\hat{\theta}_B)\right| |\hat{\theta}_B\right]\right] = \mathbb{E}\left[\lim_{B,b\to\infty} \mathbb{E}\left[\left|\hat{v}_B^b - V(\hat{\theta}_B)\right| |\hat{\theta}_B\right]\right] = 0.$$

Convergence in probability follows from $\mathcal{L}^1$ convergence.

In order to see that the estimate is biased low, simply note that

$$\mathbb{E}\left[\hat{v}_B^b\right] = \mathbb{E}\left[\mathbb{E}\left[\hat{v}_B^b|\hat{\theta}_B\right]\right] = \mathbb{E}\left[V(\hat{\theta}_B)\right] \leq V(\theta_0)$$

since $V(\cdot)$ is assumed to achieve a unique maximum at θ_0. $\square$

References

Amemiya, T., 1985, *Advanced Econometrics.* Harvard University Press, Cambridge.

Amin, K., and A. Khanna, 1994, "Convergence of American option values from discrete-to continuous-time," *Mathematical Finance*, 4(4), 289–304.

Andersen, L., 1999, "A Simple Approach to the Pricing of Bermudan Swaptions in the Multi-Factor LIBOR Market Model," working paper, General Re Financial Products.

Barraquand, J., and D. Martineau, 1995, "Numerical Valuation of High Dimensional Multivariate American Securities," *Journal of Financial & Quantitative Analysis*, 30(3), 383–405.

Bossaerts, P., 1989, "Simulation Estimators of Optimal Early Exercise," working paper, Carnegie Mellon University.

Boyle, P., 1979, "Options: a Monte Carlo Approach," *Journal of Financial Economics*, 4, 323–338.

Boyle, P., M. Broadie, and P. Glasserman, 1997, "Monte Carlo Methods for Security Pricing," *Journal of Economic Dynamics & Control*, 21, 1267–1321.

Boyle, P. P., A. W. Kolkiewicz, and K. S. Tan, 2000, "Pricing American Style Options Using Low Discrepancy Mesh Methods," working paper, University of Waterloo.

Broadie, M., and J. Detemple, 1996, "American Option Valuation: New Bounds, Approximations, and a Comparison of Existing Methods," *Review of Financial Studies*, 9(4), 1211–1250.

——— , 1997, "The Valuation of American options on Multiple Assets," *Mathematical Finance*, 7(3), 241–286.

Broadie, M., and P. Glasserman, 1997a, "Pricing American-style Securities using Simulation," *Journal of Economic Dynamics & Control*, 21, 1323–1352.

——— , 1997b, "A Stochastic Mesh Method for Pricing High-Dimensional American Options," working paper, Columbia University.

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

Cox, J. C., S. A. Ross, and M. Rubinstein, 1979, "Option Pricing: A Simplified Approach," *Journal of Financial Economics*, 7(3), 229–263.

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

Dupacova, J., and R. Wets, 1988, "Asymptotic Behavior of Statistical Estimators and of Optimal Solutions of Stochastic Optimization Problems," *The Annals of Statistics*, 16(4), 1517–1549.

Fishman, G., 1996, *Monte Carlo : Concepts, Algorithms, and Applications.* Springer-Verlag, New York.

Gao, B., J. Huang, and M. Subrahmanyam, 1998, "The Valuation of American Barrier Options Using the Decomposition Technique," working paper, New York University.

Geske, R., and H. Johnson, 1984, "The American Put Option Valued Analytically," *Journal of Finance*, 39(5), 1511–1524.

Geweke, J., 1996, "Monte-Carlo Simulation and Numerical Integration," in *Handbook of Computational Economics*, ed. by H. Amman, D. Kendrick, and J. Rust. Elsevier, New York.

Grant, D., G. Vora, and D. Weeks, 1996, "Simulation and the Early-Exercise Option Problem," *Journal of Financial Engineering*, 5(3), 211–227.

He, H., 1990, "Convergence from Discrete to Continuous Time Contingent Claims Prices," *Review of Financial Studies*, 3, 523–546.

Ibanez, A., and F. Zapatero, 1999, "Monte Carlo Valuation of American Options Through Computation of the Optimal Exercise Frontier," working paper, Marshall School of Business, USC.

Ju, N., 1998, "Pricing by American Option by Approximating its Early Exercise Boundary as a Multipiece Exponential Function," *Review of Financial Studies*, 11(3), 627–646.

Karatzas, I., 1988, "On the Pricing of American Options," *Applied Mathematics and Optimization*, 17, 37–60.

Kim, I. J., 1990, "The Analytic Valuation of American Options," *Review of Financial Studies*, 3(4), 547–572.

King, A. J., and R. T. Rockafellar, 1993, "Asymptotic Theory for Solutions in Statistical Estimation and Stochastic Programming," *Mathematics of Operations Research*, 18(1), 148–162.

Kloeden, P. E., and E. Platen, 1992, *Numerical Solution of Stochastic Differential Equations.* Springer-Verlag, New York.

Krusell, P., and A. Smith, 1996, "Rules of Thumb in Macroeconomic Equilibrium: A Quantitative Analysis," *Journal of Economic Dynamics & Control*, 20(4), 527–558.

Kushner, H., and P. Dupuis, 1992, *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, New York.

Li, B., and G. Zhang, 1996, "Hurdles Removed," working paper, Merrill Lynch.

Longstaff, F., and E. Schwartz, 2001, "Valuing American Options by Simulation: a Simple Least-squares Approach," *Review of Financial Studies*, 14(1), 113–147.

Nahum, E., 1999, "On the Pricing of Lookback Options," working paper, Dept. of Statistics, UC Berkeley.

Pedersen, M., 1999, "Bermudan Swaptions in the LIBOR Market Model," working paper, SimCorp A/S.

Press, W., S. Teukolsky, W. Vetterling, and B. Flannery, 1992, *Numerical Recipes in C : the Art of Scientific Computing*. Cambridge University Press.

Raymar, S., and M. Zwecher, 1997, "Monte Carlo Estimation of American Call Options on the Maximum of Several Stocks," *Journal of Derivatives*, 5(1), 7–23.

Shapiro, A., 1993, "Asymptotic Behavior of Optimal Solutions in Stochastic Programming," *Mathematics of Operations Research*, 18(4), 829–845.

Smith, A., 1991, "Solving Stochastic Dynamic Programming Problems using Rules of Thumb," working paper 816, Queen's University.

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

Tsitsiklis, J. N., and B. V. Roy, 2001, "Regression Methods for Pricing Complex American-Style Options," *IEEE Transactions on Neural Networks*, 12(4), 694–703.

	Stock price paths			Exercise values			
Path	t=1	t=2	t=3	t=3	t=2	t=1	t=0
1	1.09	1.08	1.34	-	0.02	0.01	0.01
2	1.16	1.26	1.54	-	-	-	-
3	1.22	1.07	1.03	0.07	0.03	-	-
4	0.93	0.97	0.92	0.18	0.17	0.17	0.17
5	1.11	1.56	1.52	-	-	-	-
6	0.76	0.77	0.90	0.20	0.19	0.19	0.18
7	0.92	0.84	1.01	0.09	0.09	0.08	0.08
8	0.88	1.22	1.34	-	-	-	-

Table 1: Numerical example of the pricing algorithm for an American put with a strike of $K = 1.1$, with 3 possible early exercise dates. The exercise value columns record the value upon exercise following the early exercise strategy: exercise at $t = j$ if the stock price is abouve θ_j, where $\theta_0 = 0.95$, $\theta_1 = 0.95$, $\theta_2 = 1$, $\theta_3 = 1.1$. The value of the put under this strategy is 0.0551. The example is taken from Longstaff and Schwartz (2001).

S_0	$\hat{V}_B$	$s(\hat{V}_B)$	$\hat{v}_B$	$s(\hat{v}_B)$	Confidence bin	Point	True value	Error
70	0.124	0.002	0.124	0.002	[0.121, 0.127]	0.124	0.121	2.49%
80	0.682	0.007	0.673	0.007	[0.661, 0.694]	0.677	0.670	1.80%
90	2.318	0.019	2.314	0.019	[2.283, 2.349]	2.316	2.303	0.66%
100	5.745	0.037	5.725	0.037	[5.664, 5.807]	5.735	5.731	0.25%
110	11.340	0.057	11.317	0.066	[11.209, 11.434]	11.329	11.341	-0.01%
120	20	-	20	-	[20, 20]	20	20	0%

Table 2: American call option prices. Parameter values: $r = 0.05$, $\delta = 0.10$, $\sigma = 0.20$, $T = 1$ year, $K = 100$.

S	$\hat{v}_B$	$\sigma(\hat{v}_B)$	Confidence bin	True value	Relative error
70	0.2361	0.0009	[0.2343, 0.2379]	0.237	0.37%
80	1.2582	0.0022	[1.2540, 1.2625]	1.259	0.06%
90	4.066	0.004	[4.0582, 4.0738]	4.077	0.27%
100	9.3647	0.0052	[9.3544, 9.3749]	9.361	0.04%
110	16.9206	0.007	[16.9069, 16.934]	16.924	0.02%
120	25.9614	0.007	[25.9477, 25.975]	25.980	0.07%
130	35.7506	0.0092	[35.7325, 35.768]	35.763	0.03%

Table 3: Maximum options on 2 stocks. Parameter values $\sigma_1 = \sigma_2 = 0.20$, $\rho = 0.3$, $r = 0.05$, $\delta = 0.10$, $T = 1$, $K = 100$. The value for S is the initial value of both assets (assumed to be the same).

S	$\hat{v}_B$	$\sigma(\hat{v}_B)$	Confidence bin	B&G point estimate	Relative error
90	16.008	0.055	[15.900, 16.115]	16.006	0.01%
100	25.234	0.068	[25.101, 25.366]	25.284	0.20%
110	35.537	0.086	[35.367, 35.706]	35.695	0.44%

Table 4: Maximum options on five underlying assets. Parameter values: $\sigma_1 = \cdots = \sigma_5 = 0.20$, $\rho = 0$, $r = 0.05$, $\delta = 0.10$, $T = 1$, $K = 100$.

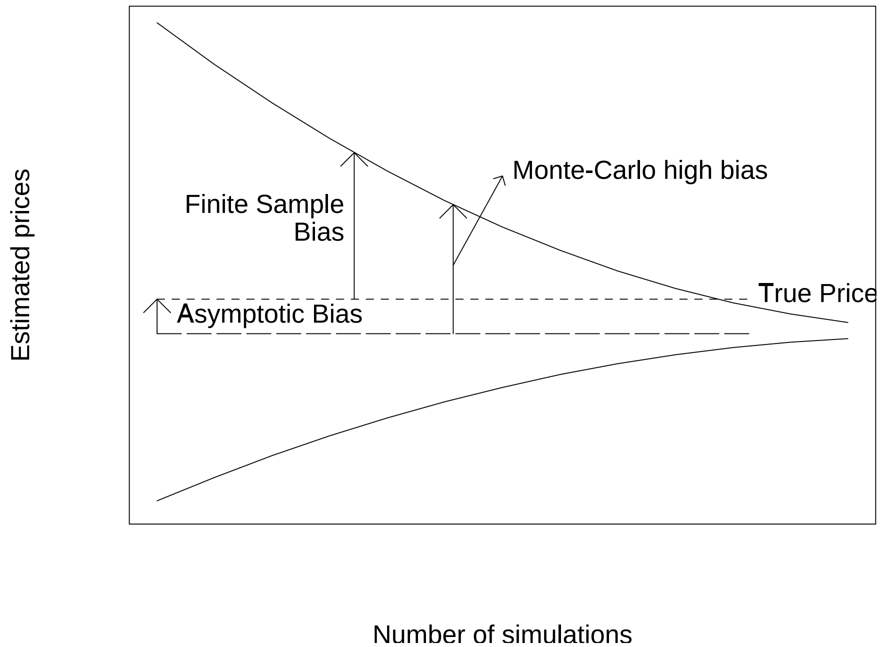

Figure 1: The straight lines represent the true option price and the asymptotic value to which both estimators converge. The curves give an indication of the expected values of both estimators. When the parametric representation subsumes all possible stopping times "asymptotic bias" is zero and the two estimators converge to the true American option price. When this is not the case the high-biases estimator converges to an option value that is lower than the theoretical one, but it may still be biased high in finite samples. The low-biased estimator always has an expected value less than the true American option price.

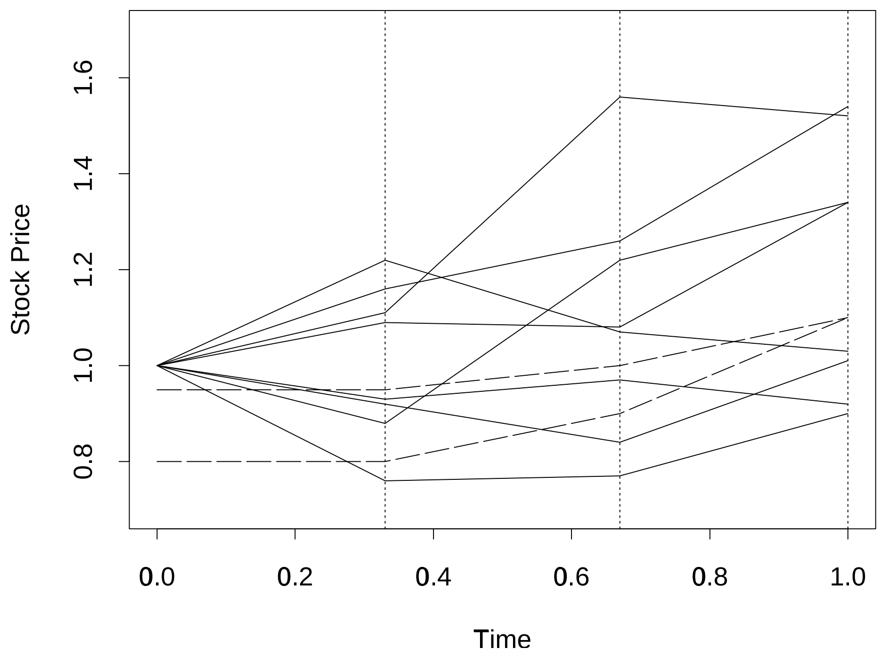

Figure 2: The solid lines represent the eight stock price paths of the numerical example presented in table 1. The dotted lines are two possible exercise rules. The first exercise rule is the one considered in table 1. The second is the optimal one (note that it is not unique).

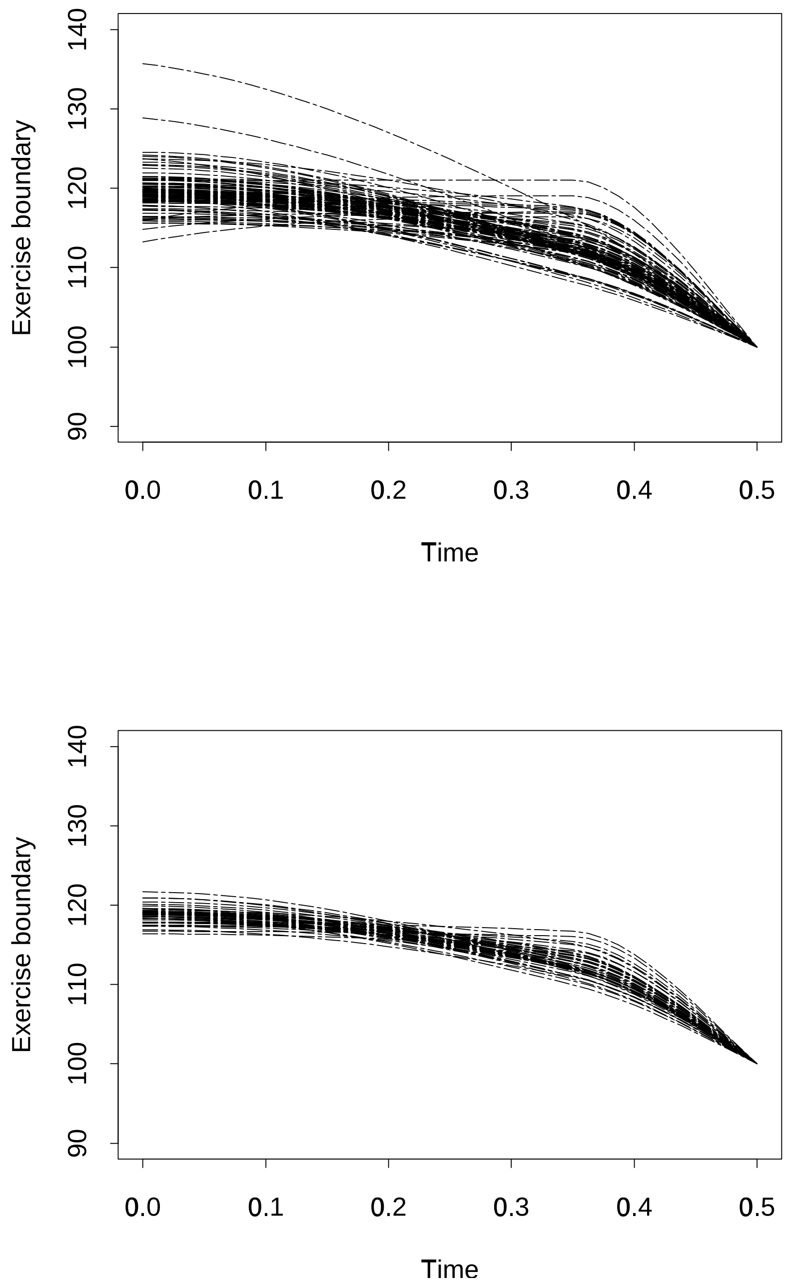

Figure 3: Graphs of estimated exercise boundaries for an American call option with parameters $\sigma = 0.20$, $r = 0.03$, $\delta = 0.07$, $T = 0.5$, $S = 110$, $K = 100$, $N = 40$. The top panel presents the results for $B = 1{,}000$, and the lower panel for $B = 16{,}000$.

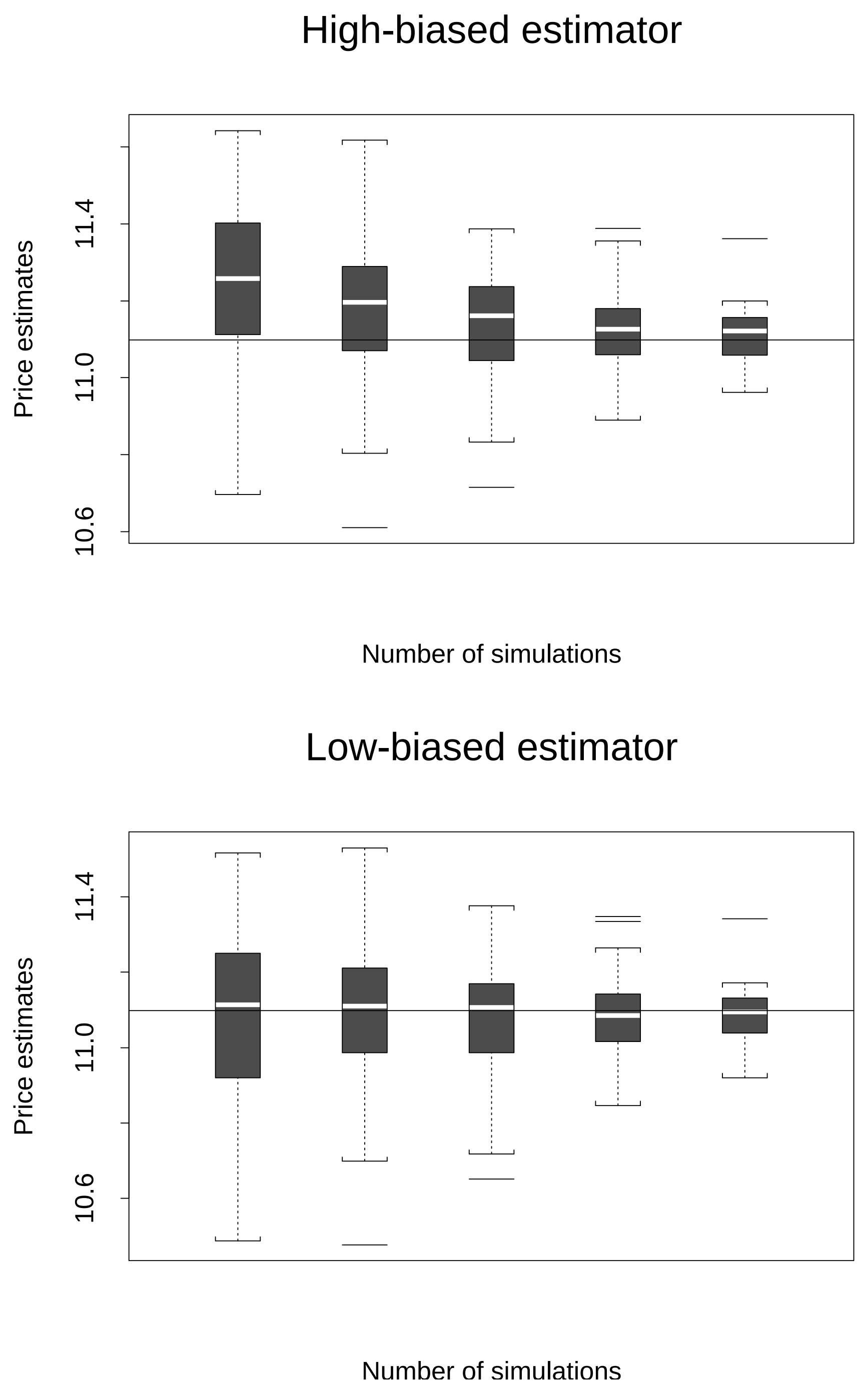

Figure 4: Boxplot of the estimates of the option values for the standard American call with 40 equally spaced exercise dates. Parameter values $\sigma = 0.20$, $r = 0.03$, $\delta = 0.07$, $T = 0.5$, $S = 110$, $K = 100$. From left to right, $B =$ 1000, 2,000, 4,000, 8,000, 16,000. The solid line marks the true value of the option (11.10). The top panel presents the estimates using $\hat{V}_B$, while the bottom panel are the values of $\hat{v}_B$.

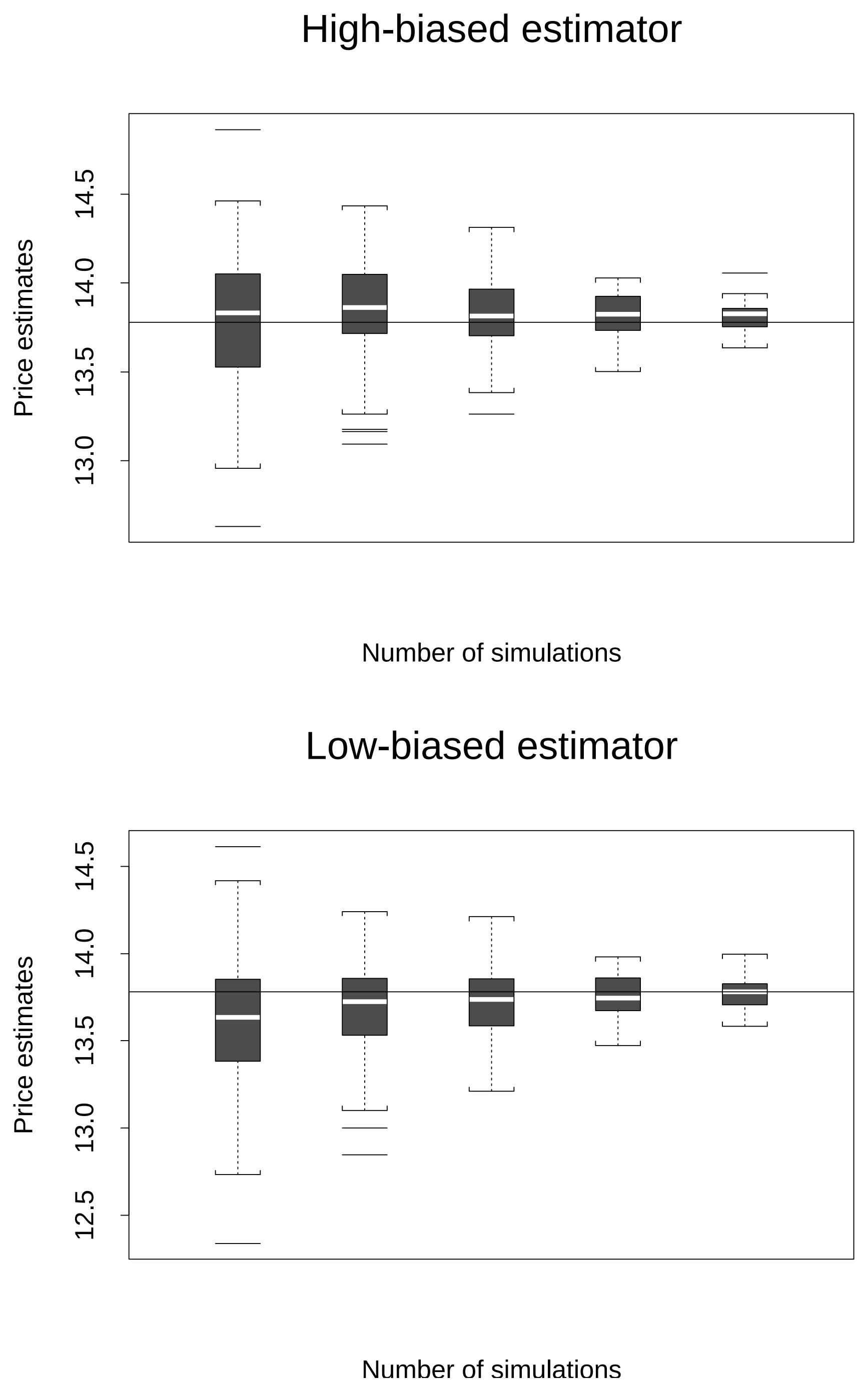

Figure 5: Estimated option values for the 2 asset maximum option. Parameter values $\sigma = 0.20$, $r = 0.05$, $\delta = 0.10$, $T = 1$, $S_0^1 = S_0^2 = 120$, $K = 100$. The true option value is 25.98. ¿From left to right, $B =$ 2,000, 4,000, 8,000, 16,000, 32,000.

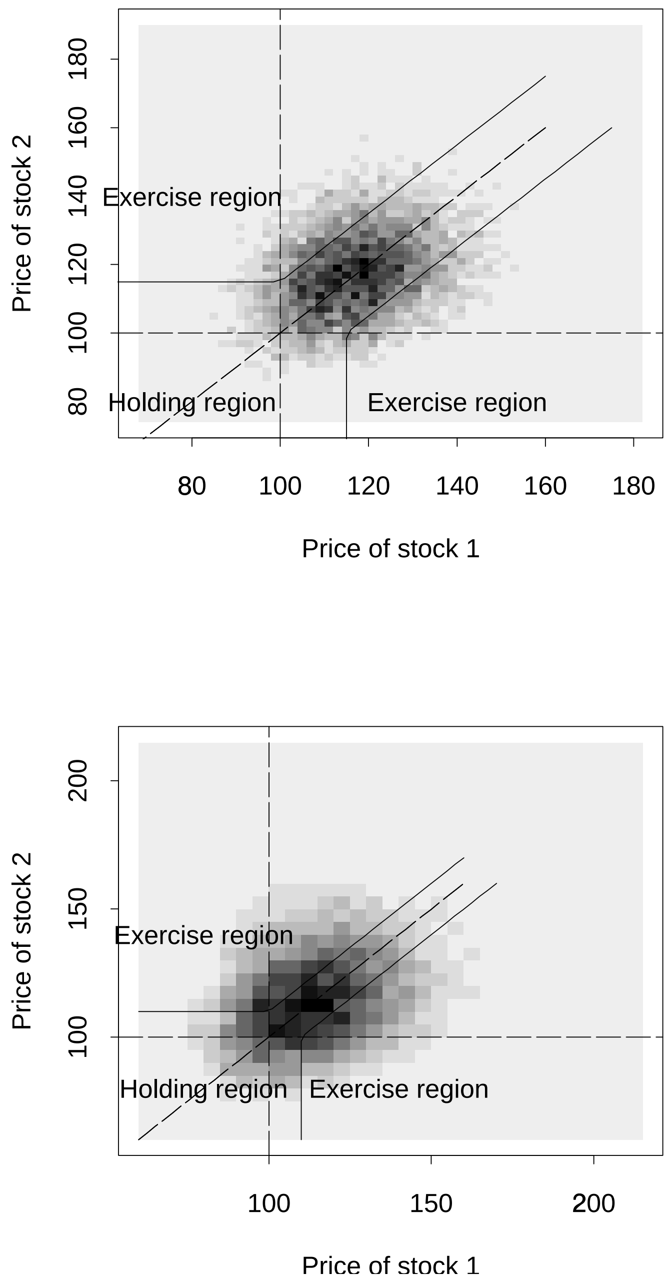

Figure 6: The top panel presents the estimated exercise regions, as well as the density function of the two stock prices at $t = T/3$. The bottom figure presents the results for $t = 2T/3$.

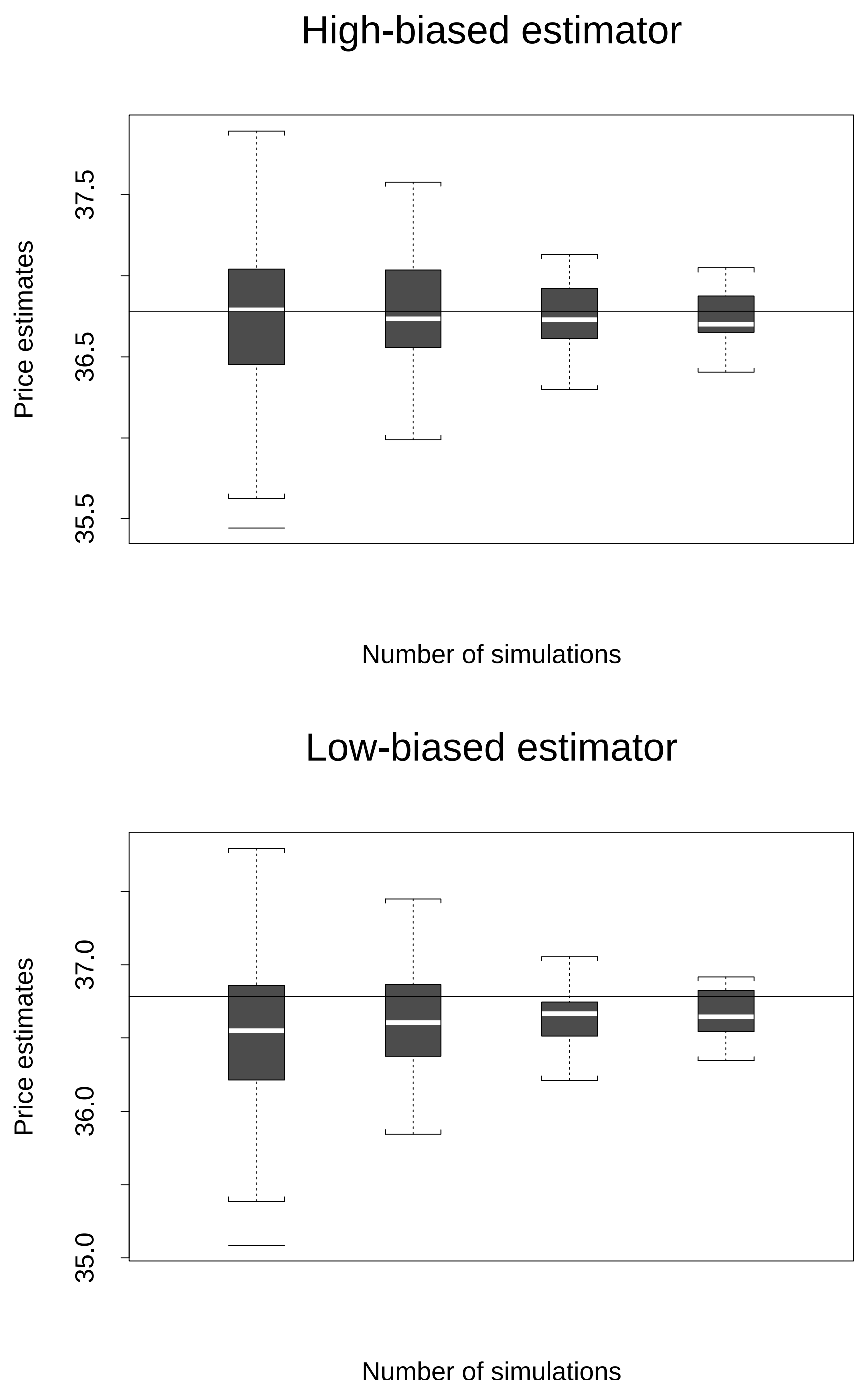

Figure 7: Boxplots for the 5 asset American maximum option with 9 equally spaced early exercise opportunities. The parameter values are set to $\sigma = 0.20$, $r = 0.05$, $\delta = 0.10$, $T = 3$, $S_0^1 = \dots S_0^5 = 120$, $K = 100$. From left to right, $B = 2{,}000$, $4{,}000$, $8{,}000$, $16{,}000$. The solid line represents the point estimate of Broadie and Glasserman (1997b).