Convergence of the stochastic mesh estimator for pricing Bermudan options

Athanassios N. Avramidis
Département d' Informatique et de Recherche Opérationnelle, Université de Montréal, C.P. 6128, Succ. Centre-Ville, Montréal, H3C 3J7, Canada.

Heinrich Matzinger
Faculty of Mathematics, University of Bielefeld, Postfach 100131, 33501 Bielefeld, Germany.

Broadie and Glasserman (2004) proposed a Monte Carlo algorithm they named "stochastic mesh" for pricing high-dimensional Bermudan options. Based on simulated states of the assets underlying the option at each exercise opportunity, the method produces an estimator of the option value at each sampled state. We derive an asymptotic upper bound on the probability of error of the mesh estimator under the mild assumption of the finiteness of certain moments. Both the error size and the probability bound are functions that vanish with increasing sample size. Moreover, we report the mesh method's empirical performance on test problems taken from the recent literature. We find that the mesh estimator has large positive bias that decays slowly with the sample size.

1 Introduction

In the financial markets, sophisticated, complex products are continuously offered and traded. There are many financial products whose values depend on more than one underlying asset. Examples include basket options (options on the average of several underlying assets), out-performance options (options on the maximum of several assets), spread options (options on the difference between two assets), and quantos (options whose payoff is adjusted by some stochastic variable, typically an exchange rate). Even when there is a single underlying asset, there is trend towards models with multiple stochastic factors (sources of uncertainty), eg, single-asset model with stochastic volatility. In addition, multi-factor models are gaining more acceptance and use for modeling interest rates, where models with two to four factors are common and models with up to ten factors are being tested (Broadie and Glasserman, 1997a). As computing power is steadily increasing, multifactor option-pricing models are likely to become more prevalent.

In this paper, we are concerned with the numerical pricing of Bermudan options. Bermudan options have a close relationship with American options, as the latter can be seen as the limit of Bermudan options with increasing frequency of exercise. The literature on numerical methods for option pricing often blurs the difference between Bermudan and American options. Thus, although the work we refer to addresses directly the Bermudan problem, the work titles often use the term "American". In this paper, we use the term "Bermudan" throughout.

The mathematical problem involved in pricing Bermudan options is an optimal stopping problem, where the stopping can occur on finitely many time points (stopping corresponds to exercising the option). The solution is characterized by a dynamic programming recursion. Analytical pricing formulas are available for some simple option payoffs, typically by approximation of the Bermudan by the American (continuous-time exercise) case, but practitioners design payoff structures that generally do not have known closed-form prices. The computation of prices in higher dimensions is generally a difficult task. Deterministic numerical schemes (based on partial differential equation methods) require work that grows geometrically with the number of factors (problem dimension). This work requirement renders these methods ineffective in dimensions higher than three or four.

For the higher-dimensional Bermudan option-pricing problems, Monte Carlo simulation techniques are attractive, being conceptually simple, yet able to address problems of great complexity, whether the complexity arises from the stochastic process driving the assets or from the structure of the payoff. The common theme of these techniques is to estimate the conditional expectations and the optimal stopping policy in the dynamic programming recursion via Monte Carlo sampling. We briefly review recent work in this area and its limitations.

Barraquand and Martineau (1995) partition the state space into a manageable number of cells and estimate, via Monte Carlo, a stopping policy that is constant over each cell and hence only approximately optimal. A serious drawback of this method is that it does not yield consistent estimates (in the probabilistic sense) of the optimal policy and option price. Broadie and Glasserman (1997b) use a simulated tree of the state variables and obtain convergent estimates for Bermudan option prices. The main drawback of their method is that the work is exponential in the number of exercise opportunities. Broadie and Glasserman (2004) developed a Monte Carlo algorithm they termed "stochastic mesh" for valuing Bermudan options and derived certain properties, including convergence results (a full description and comments is deferred until Section 2.2). More recently, Longstaff and Schwartz (2001) and Tsitsiklis and Van Roy (2001) proposed algorithms using Monte Carlo sampling and regression that involve two distinct types of error in approximating the option price. First, the conditional expectations in the dynamic programming recursion are approximated by their projections on a finite set of functions. Second, the optimal stopping policy corresponding to the projections (instead of the true value functions) is estimated via Monte Carlo sampling and yields and estimate of the option price. This algorithm was found accurate in extensive experiments in Longstaff and Schwartz (2001) (outperforming the mesh algorithm) and is becoming the method of choice among practitioners. Moreover, the convergence of regression-type algorithms has been established and characterized by Clément, Lamberton and Protter (2002). In practice, the regression algorithms' weakness is that care should be exercised in selecting a "good" set of functions on which to project.

In this paper we establish the convergence of the Broadie–Glasserman mesh estimator. Our main result is an asymptotic upper bound on the probability of error of the mesh estimator with respect to mesh size b, where both the error size and the upper bound on the probability of error are functions of b that vanish as $b \rightarrow \infty$. Our assumptions are different from those in the analysis in Broadie and Glasserman (2004) and are fairly mild – we require the finiteness of certain moments. Moreover, we report the mesh method's performance for test problems taken from Broadie and Glasserman (2004) and Longstaff and Schwartz (2001).

We wish to emphasize the difference in theoretical support for the mesh and regression algorithms. The convergence of the regression algorithms requires *both* the cardinality of the set of functions on which to project and the Monte Carlo sample size to go to infinity. On the other hand, the mesh algorithm does not involve projection error, and its convergence is with respect to just the Monte Carlo sample size. This suggests that mesh-type algorithms should still be of interest to practitioners. In addition, our convergence result lends support and motivates further study for enhanced mesh-type algorithms that were found empirically viable (Broadie and Glasserman, 2004; Avramidis and Hyden, 1999; Avramidis *et al.*, 2000; Boyle, Kolkiewicz and Tan, 2002).

This paper is organized as follows. Section 2 contains brief background on the problem of pricing Bermudan options and a description of the stochastic mesh method. Section 3 contains the convergence results and the supporting analysis. In Section 4, we present computational results on a subset of the test problems in Broadie and Glasserman (2004) and Longstaff and Schwartz (2001), and in Section 5 we offer conclusions. An earlier version of this work appeared in Avramidis and Matzinger (2002).

2 Background

2.1 Bermudan option pricing

Let $t = 0, 1, 2, \ldots, T$ be the set of times when the Bermudan option is exercisable, also referred to as *exercise opportunities* or simply *stages*. Except for Section 4, all references to "time" t correspond to the tth exercise opportunity, and not to time measured in the usual continuous sense. Let S_t denote the state of the stochastic factors at time (exercise opportunity) t, for $t = 0, 1, 2, \ldots, T$. The random variable S_t takes values in $\mathbb{R}^d$ with d a positive integer, allowing for the general multifactor setting. We assume that S_t is Markovian, so that the conditional distributions of future states depend only on the current state. Let $h(t,x)$ denote the payoff to the option holder from exercise at time t in state x, discounted to time zero with the possibly stochastic discount factor contained in x (the view of $h(t,x)$ as the discounted-to-time-zero payoff is adopted to simplify the notation).

The Bermudan option-pricing problem is to compute

$$q(0,x_0) = \max_{\tau} \mathrm{E}[h(\tau,S_\tau) \mid S_0 = x_0]$$

where τ is a stopping time taking values in the finite set $\{0,1,\ldots,T\}$. It is well known from arbitrage pricing theory that the arbitrage-free price of the option is obtained when expectations are taken with respect to the risk-neutral measure; see, for example, Duffie (1996) and Harrison and Pliska (1981). By the dynamic programming principle, the option value satisfies the following recursion:

$$q(t,x) = \begin{cases} h(t,x) & t=T, \text{ all } x \\ \max\{h(t,x), c(t,x)\} & 0 \le t \le T-1, \text{ all } x \end{cases}$$

where

$$c(t,x) = \mathrm{E}[q(t+1,S_{t+1}) \mid S_t = x] \tag{1}$$

is called the *continuation value* at (t,x).

2.2 The stochastic mesh method

In reviewing the method, we follow Broadie and Glasserman (2004). To begin, we generate randomly a set of *mesh points* (states) $\{S_t^j\}$, $j = 1,2,\ldots,b$ for each future exercise opportunity $t = 1,\ldots,T$. We emphasize that throughout the paper, superscripts on points S_t denote the generation index (ranging in $\{1,2,\ldots,b\}$) and not powers. For notational convenience, we define b non-random mesh points at stage zero, $S_0^j = x_0$, $j = 1,2,\ldots,b$. For $t = 1,2,\ldots,T$, let $g_t(\cdot)$ denote the probability density from which the points $\{S_t^j\}_{j=1}^b$ are sampled (to be specified later), and let $f_t(x,\cdot)$ denote the conditional risk-neutral density of S_{t+1} given $S_t = x$. We assume throughout the paper the existence of these risk-neutral densities. For notational convenience, we let $\mathcal{E} = \{0,1,\ldots,T-1\}$ denote the index set of *early-exercise opportunities* and we let $I = \{1,2,\ldots,b\}$ denote the index set of sampled points per stage. The Broadie–Glasserman mesh estimator is calculated as a backwards recursion over the set of early exercise opportunities:

$$\hat{q}(t,S_t^j) = \begin{cases} h(t,S_t^j) & j \in I,\ t=T \\ \max\{h(t,S_t^j), \hat{c}(t,S_t^j)\} & j \in I,\ t = T-1, T-2,\ldots,1,0 \end{cases}$$

where the estimate of the continuation value function, $\hat{c}(t,x)$, is

$$\hat{c}(t,x) := \sum_{j=1}^{b} \frac{\hat{q}(t+1,S_{t+1}^j) f_t(x,S_{t+1}^j)}{g_{t+1}(S_{t+1}^j)} \tag{2}$$

Note that each forward-stage value $\hat{q}(t+1,S_{t+1}^j)$ is weighed by the *likelihood ratio* (mathematically, Radon–Nikodym derivative) $f_t(x,S_{t+1}^j)/g_{t+1}(S_{t+1}^j)$.

FIGURE 1 Stratified mesh implementation illustrated for $T = 4$ and $b = 3$. An arrow points from a parent point to a child point in each simulated path. A line (including arrows) between two points at stages t and $t + 1$ indicates that a computation of the one-step transition density, f_t, between the points is required by the mesh algorithm.

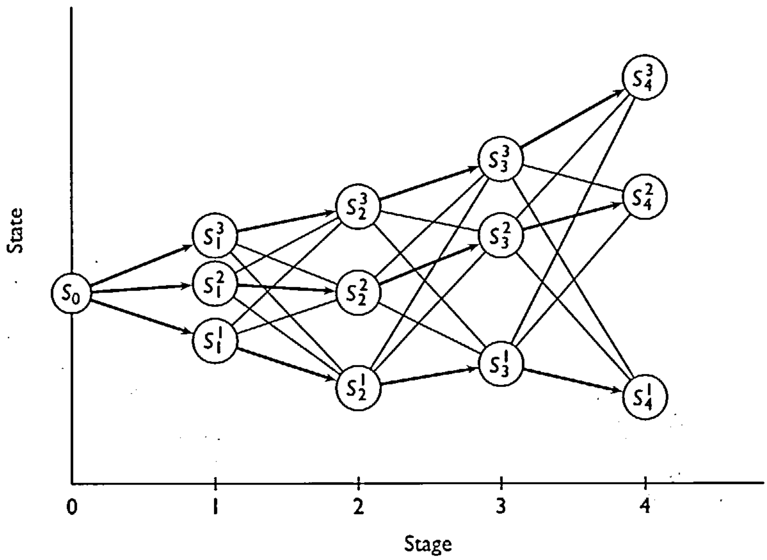

Broadie and Glasserman (2004) argued that the choice of sampling densities, $g_{t+1}(\cdot)$, is crucial to the success of the method, and they recommended sampling as follows. We generate independently b paths of S_t starting from x_0 at time 0 and let S_t^j denote the state of the jth path at time t; and then we "forget" the path to which each point belongs. This is called by the authors the *stratified implementation*. For any t, j, we call the ordered pair (S_t^j, S_{t+1}^j) a *parent* and *child*, respectively. A visual illustration of the mesh is given in Figure 1.

The authors also propose and analyze a *path estimator*. A single replication of the path estimator simulates one path of the stochastic factors underlying the option and applies a stopping policy implied by a previously estimated entire mesh; see Broadie and Glasserman (2004) for the details. Clearly, an analysis of the mesh estimator is a prerequisite for analyzing the path estimator. In this paper, our analysis is focused on the mesh estimator.

3 Convergence results

This section contains our results on the convergence of the stratified implementation of the mesh estimator. Under an assumption on the finiteness of certain moments, we will show that $\hat{q}(0, x_0)$ converges in probability to $q(0, x_0)$ as $b \to \infty$. More precisely, we derive an asymptotic upper bound on the probability of error of the mesh estimator with respect to b, where both the error size and the upper bound on the probability of error are functions of b that vanish as $b \to \infty$. We maintain all the notation established in Section 2.

We begin by observing some distributional properties of the stratified implementation. Let π be a random permutation of the integers in $\{1, 2, \ldots, b\}$ chosen with equal probability from all possible such permutations, and let $\mathcal{F}_t$ be the σ-field $\mathcal{F}_t = \sigma(S_t^1, S_t^2, \ldots, S_t^b)$. Then

$$\text{conditional on } \mathcal{F}_t, \quad S_{t+1}^{\pi(1)}, S_{t+1}^{\pi(2)}, \ldots, S_{t+1}^{\pi(b)} \overset{\text{i.d.}}{\sim} g_{t+1}(\cdot) := \frac{1}{b}\sum_{i=1}^{b} f_t(S_t^i, \cdot) \tag{3}$$

where $\overset{\text{i.d.}}{\sim}$ means "are identically distributed with density …". Note that the density $g_{t+1}(\cdot)$ is defined conditionally on $\mathcal{F}_t$. Also note that $S_{t+1}^{\pi(1)}, S_{t+1}^{\pi(2)}, \ldots, S_{t+1}^{\pi(b)}$ are conditionally dependent random vectors. On the other hand,

$$\text{conditional on } \mathcal{F}_t, \quad S_{t+1}^1, S_{t+1}^2, \ldots, S_{t+1}^b \text{ are independent} \tag{4}$$

Also note that $S_{t+1}^1, S_{t+1}^2, \ldots, S_{t+1}^b$ are conditionally *not* identically distributed; they are unconditionally independent and identically distributed.

Our upper bound depends on a constant, C, bounding certain moments as follows. Let $\{(S_t^j : t = 0, 1, \ldots, T)\}_{j=1}^3$ denote the mesh points along three paths which are independent and identically distributed (with $S_0^j = x_0$), as prescribed by the path-generation step of the stratified implementation of the mesh algorithm. We make the following assumptions:

$$\max_{t \in \mathcal{E}} \mathrm{E}\left[\frac{f_t^4(S_t^1, S_{t+1}^2)}{f_t^4(S_t^\ell, S_{t+1}^2)} \max_{t+1 \le r \le T}\{h^4(r, S_r^2)\}\right] \le \frac{C}{8} \quad \text{for } \ell = 1, 2, 3 \tag{5}$$

$$\max_{t \in \mathcal{E}} \mathrm{E}\left[\frac{f_t^4(S_t^1, S_{t+1}^1)}{f_t^4(S_t^\ell, S_{t+1}^1)} \max_{t+1 \le r \le T}\{h^4(r, S_r^1)\}\right] < \infty \quad \text{for } \ell = 1, 2, 3 \tag{6}$$

$$\max_{t \in \mathcal{E}} \mathrm{E}\left[\frac{f_t^4(S_t^1, S_{t+1}^2)}{f_t^4(S_t^\ell, S_{t+1}^2)}\right] \le \frac{C}{8} \quad \text{for } \ell = 1, 2, 3 \tag{7}$$

$$\max_{t \in \mathcal{E}} \mathrm{E}\left[\frac{f_t^4(S_t^1, S_{t+1}^1)}{f_t^4(S_t^\ell, S_{t+1}^1)}\right] < \infty \quad \text{for } \ell = 1, 2, 3 \tag{8}$$

In words, condition (5) says that the constant $C/8$ is a simultaneous upper bound (taken over all early-exercise stages, t) on the expectation of the fourth power of the random variable defined as the ratio of the one-step transition density $f_t(S_t^1, S_{t+1}^2)$ (ie, going from S_t^1 to its non-child, S_{t+1}^2) to the one-step transition density $f_t(S_t^k, S_{t+1}^2)$ ($k = 1,2,3$ yield distinct cases) times the maximum future payoff over a path that starts at S_{t+1}^2. The role of C in (7) has a similar interpretation. Our main result is as follows.

THEOREM 1 *Suppose that b mesh paths* $\{(S_t^j : t = 0, 1, \ldots, T)\}_{j=1}^b$ *are generated independently with* $S_0^j = x_0$ *for all j, where* $x_0 \in \mathbb{R}^d$ *is the known state at time 0. Under assumptions (5)–(8),*

$$P\left\{ \left| \widehat{q}(0,x_0) - q(0,x_0) \right| > \left(1+\frac{\delta}{b^\gamma}\right)^T - 1 \right\}$$

$$\leq \frac{6CT}{\delta^4 b^{1-4\gamma}} + O(b^{-2+4\gamma}) \quad \text{for any } \delta > 0,\ 0 < \gamma < \frac{1}{4}$$

PROOF We start with a few definitions. The time index $t \in \mathcal{E}$, unless explicitly stated otherwise. Let

$$\bar{c}(t,x) := \frac{1}{b}\sum_{j=1}^{b} \frac{q(t+1, S_{t+1}^j) f(x, S_{t+1}^j)}{g_{t+1}(S_{t+1}^j)} \tag{9}$$

That is, if the function $q(t+1, \cdot)$ were known (which of course is not the case), then $\bar{c}(t,x)$ would be a natural estimate of $c(t,x)$. Fix $\delta > 0$ and $0 < \gamma < 1/4$, and define the events

$$A_1(t) := \left\{ \omega : \left| \bar{c}(t, S_t^i)(\omega) - c(t, S_t^i)(\omega) \right| \leq \frac{\delta}{b^\gamma}, \quad \text{for all } i \in I \right\} \tag{10}$$

and

$$A_2(t) := \left\{ \omega : \left| \frac{1}{b}\sum_{j=1}^{b} \frac{f_t(S_t^i, S_{t+1}^j)}{g_{t+1}(S_{t+1}^j)}(\omega) - 1 \right| \leq \frac{\delta}{b^\gamma}, \quad \text{for all } i \in I \right\} \tag{11}$$

where ω denotes a generic point in the sample space. Let A_1 be the event that $A_1(t)$ holds for each $t \in \mathcal{E}$, ie,

$$A_1 := \cap_{t \in \mathcal{E}} A_1(t)$$

Similarly, define $A_2 := \cap_{t \in \mathcal{E}} A_2(t)$. Finally, define the event of direct interest

$$A := \left\{ \omega : \widehat{q}(0,x_0)(\omega) - q(0,x_0) \right| \leq \left(1 + \frac{\delta}{b^\gamma}\right)^T - 1 \right\}$$

For notational simplicity, we will suppress the dependence of all random variables on ω for the remainder of the paper.

CLAIM 1 $A \supset A_1 \cap A_2$

PROOF OF CLAIM 1 We assume that events A_1 and A_2 hold and show by a recursive argument going backwards in time that event A must hold. We start by showing that an error bound that holds uniformly over all estimates at time $t + 1$ can be iterated backwards in time. Fix $\varepsilon > 0$ and suppose that for some t $(0 < t \le T - 1)$ the error of the estimates at the forward points satisfies

$$\left|\widehat{q}(t+1, S^j_{t+1}) - q(t+1, S^j_{t+1})\right| \le \varepsilon \qquad \text{for all } j \in I \tag{12}$$

Then

$$\begin{aligned}
&|\widehat{c}(t,x) - c(t,x)| \\
&= \frac{1}{b}\left|\sum_{j=1}^{b} \frac{\widehat{q}(t+1, S^j_{t+1}) f_t(x, S^j_{t+1})}{g_{t+1}(S^j_{t+1})} - \sum_{j=1}^{b} \frac{q(t+1, S^j_{t+1}) f_t(x, S^j_{t+1})}{g_{t+1}(S^j_{t+1})}\right| \\
&= \frac{1}{b}\left|\sum_{j=1}^{b}\left(\widehat{q}(t+1, S^j_{t+1}) - q(t+1, S^j_{t+1})\right) \frac{f_t(x, S^j_{t+1})}{g_{t+1}(S^j_{t+1})}\right| \\
&\le \frac{\varepsilon}{b}\sum_{j=1}^{b} \frac{f_t(x, S^j_{t+1})}{g_{t+1}(S^j_{t+1})} \\
&\le \varepsilon\left(1 + \frac{\delta}{b^\gamma}\right) \qquad \text{for all } x \in \{S^1_t, S^2_t, \ldots, S^b_t\}
\end{aligned} \tag{13}$$

where the last inequality follows since A_2 holds. So if (12) holds, then the error of $\widehat{q}$ at stage t $(0 \le t \le T - 1)$ is bounded uniformly on j as follows:

$$\begin{aligned}
&|\widehat{q}(t, S^j_t) - q(t, S^j_t)| \\
&= \left|\max\left\{h(t, S^j_t), \widehat{c}(t, S^j_t)\right\} - \max\left\{h(t, S^j_t), c(t, S^j_t)\right\}\right| \\
&\le \left|\widehat{c}(t, S^j_t) - c(t, S^j_t)\right| \\
&\le \left|\widehat{c}(t, S^j_t) - c(t, S^j_t)\right| + \left|c(t, S^j_t) - c(t, S^j_t)\right| \\
&\le \varepsilon\left(1 + \frac{\delta}{b^\gamma}\right) + \frac{\delta}{b^\gamma} \qquad \text{for all } j \in I
\end{aligned} \tag{14}$$

where in the last inequality we used (13) and that event A_1 holds.

Now the recursive bounding is as follows. We start the error bounding with the special case $t = T - 1$, where we observe that $\widehat{c}(T-1, S^j_{T-1}) - \bar{c}(T-1, S^j_{T-1}) = 0$

for all j, and so the definition of the event $A_1(T-1)$ implies that (14) holds for $t = T-1$ with $\varepsilon = 0$. Iterating the bounding argument in (14) with $t = T-2, T-3, \ldots, 0$, we get

$$|\widehat{q}(0, x_0) - q(0, x_0)| \leq \frac{\delta}{b^\gamma} \sum_{j=0}^{T-1} \left(1 + \frac{\delta}{b^\gamma}\right)^j$$

$$= \left(1 + \frac{\delta}{b^\gamma}\right)^T - 1$$

which completes the proof of Claim 1. □

Letting A^c denote the complement of the event A, we have $P(A^c) \leq P(A_1^c) + P(A_2^c)$. To complete the proof, we will show that

$$P(A_1^c) \leq \frac{3CT}{\delta^4 b^{1-4\gamma}} + O(b^{-2+4\gamma}) \tag{15}$$

and

$$P(A_2^c) \leq \frac{3CT}{\delta^4 b^{1-4\gamma}} + O(b^{-2+4\gamma}) \tag{16}$$

We first obtain the upper bound for $P(A_1^c)$. Define the event

$$A_1(t, i) = \left\{\omega : \left|\bar{c}(t, S_t^i)(\omega) - c(t, S_t^i)(\omega)\right| \leq \frac{\delta}{b^\gamma}\right\}$$

Recall that $A_1 = \cap_{t=0}^{T-1} A_1(t) = \cap_{t=0}^{T-1} \cap_{i=1}^{b} A_1(t, i)$, so

$$P(A_1^c) \leq \sum_{t=0}^{T-1} \sum_{i=1}^{b} P\big(A_1^c(t, i)\big) = b \sum_{t=0}^{T-1} P\big(A_1^c(t, 1)\big) \tag{17}$$

the equality holding because $(S_t^i, \{S_t^j\}_{j=1}^b)$, $i = 1, \ldots, b$, the paths are unconditionally identically distributed. We will show that

$$P\big(A_1^c(t, 1)\big) \leq \frac{3C}{\delta^4 b^{2-4\gamma}} + O(b^{-3-4\gamma}) \quad \text{for all } t \in \mathcal{E} \tag{18}$$

which, in view of (17), proves (15).

The key for proving (18) is that $\bar{c}(t, S_t^1) - c(t, S_t^1)$ can be written as the sum of b random variables which, conditional on $\mathcal{F}_t$, have mean zero and are independent.

CLAIM 2 $\bar{c}(t,S_t^1) - c(t,S_t^1) = (1/b)\sum_{j=1}^{b} Z^j(t)$, where

$$Z^j(t) := \frac{q(t+1,S_{t+1}^j)f(S_t^1,S_{t+1}^j)}{g_{t+1}(S_{t+1}^j)} - \mathrm{E}\left[\frac{q(t+1,S_{t+1}^j)f(S_t^1,S_{t+1}^j)}{g_{t+1}(S_{t+1}^j)}\middle|\mathcal{F}_t\right],\quad j\in I \tag{19}$$

where we recall that $\mathcal{F}_t$ is the σ-field $\mathcal{F}_t = \sigma(S_t^1, S_t^2, \ldots, S_t^b)$.

PROOF OF CLAIM 2

$$\frac{1}{b}\sum_{j=1}^{b} Z^j(t) = \frac{1}{b}\sum_{j=1}^{b}\left(\frac{q(t+1,S_{t+1}^j)f(S_t^1,S_{t+1}^j)}{g_{t+1}(S_{t+1}^j)} - \mathrm{E}\left[\frac{q(t+1,S_{t+1}^j)f(S_t^1,S_{t+1}^j)}{g_{t+1}(S_{t+1}^j)}\middle|\mathcal{F}_t\right]\right)$$

$$= \bar{c}(t,S_t^1) - \mathrm{E}\left[\frac{1}{b}\sum_{j=1}^{b}\frac{q(t+1,S_{t+1}^j)f(S_t^1,S_{t+1}^j)}{g_{t+1}(S_{t+1}^j)}\middle|\mathcal{F}_t\right]$$

$$= \bar{c}(t,S_t^1) - \mathrm{E}\left[\frac{1}{b}\sum_{j=1}^{b}\frac{q(t+1,S_{t+1}^{\pi(j)})f(S_t^1,S_{t+1}^{\pi(j)})}{g_{t+1}(S_{t+1}^{\pi(j)})}\middle|\mathcal{F}_t\right]$$

$$= \bar{c}(t,S_t^1) - \mathrm{E}\left[\frac{q(t+1,X)f(S_t^1,X)}{g_{t+1}(X)}\middle|\mathcal{F}_t\right]$$

where X represents a random variable obtained by choosing one of the points $S_{t+1}^1, S_{t+1}^2, \ldots, S_{t+1}^b$ at random with equal probability. The key behind the third step is the invariance of the sum inside the expectation with respect to permutations of the $\{S_{t+1}^j\}_{j=1}^b$. The conditional distribution of X when conditioned under $\mathcal{F}_t$ has the density $g_{t+1}(\cdot)$ in (3), so

$$\mathrm{E}\left[\frac{q(t+1,X)f(S_t^1,X)}{g_{t+1}(X)}\middle|\mathcal{F}_t\right] = \mathrm{E}\left[q(t+1,S_{t+1}^1)\middle|\mathcal{F}_t\right]$$

$$= c(t,S_t^1)$$

which completes the proof of Claim 2. □

Our upper bound for the probability $P(A_1^c(t,1))$ will use Markov's inequality with the fourth moment of the deviation $\bar{c}(t,S_t^1) - c(t,S_t^1)$. We will show that this fourth moment goes to zero sufficiently fast with b. First, we need two lemmas.

LEMMA 1 *Suppose that Y is a non-negative random variable with* $\mathrm{E}[Y^4] < \infty$. *Then* $\mathrm{E}[(Y - \mathrm{E}[Y|\mathcal{F}])^4] \le 8\mathrm{E}[Y^4]$, *where $\mathcal{F}$ is an arbitrary σ-field.*

PROOF

$$\begin{aligned}
&\mathrm{E}\big[(Y-\mathrm{E}[Y|\mathcal{F}])^4\big] \\
&= \mathrm{E}\big(Y^4 - 4Y^3\mathrm{E}[Y|\mathcal{F}] + 6Y^2\mathrm{E}^2[Y|\mathcal{F}] - 4Y\mathrm{E}^3[Y|\mathcal{F}] + \mathrm{E}^4[Y|\mathcal{F}]\big) \\
&\le \mathrm{E}[Y^4] + 6\mathrm{E}\big(Y^2\mathrm{E}^2[Y|\mathcal{F}]\big) + \mathrm{E}\big(\mathrm{E}^4[Y|\mathcal{F}]\big) \\
&\le \mathrm{E}[Y^4] + 6\sqrt{\mathrm{E}[Y^4]}\sqrt{\mathrm{E}\big(\mathrm{E}^4[Y|\mathcal{F}]\big)} + \mathrm{E}\big(\mathrm{E}[Y^4|\mathcal{F}]\big) \\
&\le 2\mathrm{E}[Y^4] + 6\sqrt{\mathrm{E}[Y^4]}\sqrt{\mathrm{E}[Y^4]} \\
&= 8\mathrm{E}[Y^4]
\end{aligned}$$

In the second step, we dropped non-positive random variables from the expectation. In the third step, we used the Cauchy–Schwartz inequality for the second term and Jensen's inequality for the third term, and in the fourth step we used again Jensen's inequality inside the second square root. □

LEMMA 2 *Let $\mathcal{F}$ denote an arbitrary σ-field, and let $Z_1, Z_2, \ldots, Z_b$ be random variables which, conditional on $\mathcal{F}$, have mean zero, are conditionally independent of each other, and such that* $\mathrm{E}[Z_1^4] < \infty$ *and* $\mathrm{E}[Z_j^4] \le C$ *for each* $j \ne 1$, *where the expectations are unconditional and C is a constant. Then*

$$\mathrm{E}\left[\left(\frac{1}{b}\sum_{j=1}^{b} Z_j\right)^4\right] \le \frac{3C}{b^2} + O(b^{-3})$$

PROOF $\mathrm{E}[(\sum_{j=1}^{b} Z_j)^4] = \sum \mathrm{E}[\mathrm{E}[Z_{j_1}Z_{j_2}Z_{j_3}Z_{j_4} \,|\, \mathcal{F}]]$, where the four indices are ranging independently from 1 to b. Since $\mathrm{E}[Z_{j_1} \,|\, \mathcal{F}] = 0$, the conditional independence of the $Z's$ implies that the summand vanishes if there is one index different from the three others. This leaves terms of the form $\mathrm{E}[\mathrm{E}[Z_{j_1}^4 \,|\, \mathcal{F}]]$, of which there are b, and terms of the form $\mathrm{E}[\mathrm{E}[Z_{j_1}^2 Z_{j_2}^2 \,|\, \mathcal{F}]]$ for $j_1 \ne j_2$, of which there are $3b(b-1)$. For each of the two different forms, the number of terms with any index equal to 1 is $O(b^{-1})$ of the total number of such terms, and so the finiteness of $\mathrm{E}[Z_1^4]$ implies that the relative contribution of these terms to the total is $O(b^{-1})$. Now focusing on terms where all indices are different than 1, we have

$$\mathrm{E}\big[\mathrm{E}[Z_{j_1}^4 \,|\, \mathcal{F}]\big] = \mathrm{E}\big[Z_{j_1}^4\big] \le C$$

and

$$\mathrm{E}\big[\mathrm{E}[Z_{j_1}^2 Z_{j_2}^2 \,|\, \mathcal{F}]\big] = \mathrm{E}\big[Z_{j_1}^2 Z_{j_2}^2\big] \le \sqrt{\mathrm{E}\big[Z_{j_1}^4\big]}\,\sqrt{\mathrm{E}\big[Z_{j_2}^4\big]} \le C$$

Hence

$$\mathrm{E}\left[\left(\sum_{j=1}^{b} Z_j\right)^4\right] \le bC\left(1+O(b^{-1})\right) + 3b(b-1)C\left(1+O(b^{-1})\right)$$

which completes the proof of Lemma 2. □

Claim 3 For each $t \in \mathcal{E}$, the conditions of Lemma 2 are valid for $Z_j = Z^j(t)$ and $\mathcal{F} = \mathcal{F}_t$.

Proof of Claim 3 Fix $t \in \mathcal{E}$. The $Z^j(t)$ have conditional mean zero by definition (19). Moreover, conditional on $\mathcal{F}_t$, each of the variables $Z^1(t), Z^2(t), \ldots, Z^b(t)$ is a function of the single random variable $S^1_{t+1}, S^2_{t+1}, \ldots, S^b_{t+1}$, respectively. Thus the $Z^j(t)$ are conditionally independent, in view of (4). To bound the unconditional fourth moments, we apply Lemma 1 with $Y = q(t+1, S^j_{t+1}) f_t(S^1_t, S^j_{t+1}) / g_{t+1}(S^j_{t+1})$ and $\mathcal{F} = \mathcal{F}_t$:

$$\begin{aligned}\mathrm{E}\left[(Z^j(t))^4\right] &\le 8\mathrm{E}\left[\frac{q^4(t+1, S^j_{t+1}) f_t^4(S^1_t, S^j_{t+1})}{g^4_{t+1}(S^j_{t+1})}\right] \\ &\le 8\mathrm{E}\left[\frac{\max_{t+1\le r\le T}\{h^4(r, S^j_r)\} f_t^4(S^1_t, S^j_{t+1})}{g^4_{t+1}(S^j_{t+1})}\right] \quad \text{for all } j \in I \end{aligned} \tag{20}$$

First, we will prove the finite upper bound for $j \ne 1$. The $\{Z^j(t)\}_{j=2}^b$ are unconditionally identically distributed and, continuing from (20), we have

$$\begin{aligned}\mathrm{E}\left[(Z^2(t))^4\right] &\le 8\mathrm{E}\left[\max_{t+1\le r\le T}\{h^4(r, S^2_r)\} f_t^4(S^1_t, S^2_{t+1}) \frac{1}{b}\sum_{k=1}^{b} \frac{1}{f_t^4\left(S^k_t, S^2_{t+1}\right)}\right] \\ &\le C\end{aligned}$$

To obtain the first inequality, we upper-bounded the factor $g^{-4}_{t+1}(S^2_{t+1})$ by expanding $g_{t+1}(\cdot)$ according to (3) and using the fact (Jensen's inequality) that for any $x_1, x_2, \ldots, x_b > 0$

$$\left(\frac{1}{b}\sum_{i=1}^{b} x_i\right)^{-4} \le \frac{1}{b}\sum_{i=1}^{b} x_i^{-4}$$

The second inequality then follows from assumption (5). An analogous argument combined with assumption (6) shows that $\mathrm{E}[(Z^1(t))^4] < \infty$. □

Now we have

$$P(A_1^c(t,1)) = P\left(\left|\bar{c}(t,S_t^1) - c(t,S_t^1)\right| \geq \frac{\delta}{b^\gamma}\right)$$

$$= P\left(\frac{1}{b}\left|\sum_{j=1}^{b} Z^j(t)\right| \geq \frac{\delta}{b^\gamma}\right)$$

$$\leq \frac{\mathrm{E}\left[\left(\frac{1}{b}\sum_{j=1}^{b} Z^j(t)\right)^4\right] b^{4\gamma}}{\delta^4} \tag{21}$$

$$\leq \frac{3C}{\delta^4 b^{2-4\gamma}} + O(b^{-3+4\gamma}) \quad \text{for all } t \in \mathcal{E} \tag{22}$$

In step three, we used Markov's inequality with power 4, and in step four we applied Lemma 2 for $Z_j = Z^j(t)$ and $\mathcal{F} = \mathcal{F}_t$ (which is valid, as proved in Claim 3). This is precisely what was required in (18), and completes the proof of (15).

To prove the probability bound $P(A_2^c) \leq 3CT/(\delta^4 b^{1-4\gamma}) + O(b^{-2+4\gamma})$, we simply observe that the event A_2 in (11) can be written in the form of event A_1 in (10) by choosing $q(\cdot,\cdot) = 1$ in (9) (implying $c(\cdot,\cdot) = 1$). Then we argue as we did to obtain the bound for the event A_1^c, using assumptions (7) and (8) in place of (5) and (6), respectively. This completes the proof of Theorem 1. □

The following result shows that the rate of convergence may be sharpened using moments analogous to (5)–(8) but of higher order.

THEOREM 2 *Suppose that the mesh paths $\{S_t^j\}_{j=1}^b$ are generated independently with $S_0^j = x_0$ for all j, where $x_0 \in \mathbb{R}^d$ is the known state at time 0. Under assumptions (5)–(8) where we replace the power 4 by the power 8 and let C_1 be the corresponding constant,*

$$P\left\{\left|\widehat{q}(0,x_0) - q(0,x_0)\right| > \left(1+\frac{\delta}{b^\gamma}\right)^T - 1\right\}$$

$$\leq \frac{2520\,C_1 T}{\delta^8 b^{5-8\gamma}} + O(b^{-6+8\gamma}) \quad \text{for any } \delta > 0,\ 0 < \gamma < \frac{5}{8}$$

SKETCH OF PROOF One can show that $P(A_1^c(t,1)) \leq 1260 C_1/(\delta^8 b^{5-8\gamma}) + O(b^{-6+8\gamma})$ by arguing analogously to (21)–(22), using Markov's inequality with power 8 instead of power 4 and a result analogous to Lemma 2 for the eighth moment. The other steps in the proof parallel the proof of Theorem 1. □

4 Empirical performance

We report empirical results on the performance of the mesh estimator. The first set of test problems is a subset of those in Broadie and Glasserman (2004). Under the risk-neutral measure, the d assets are independent, and each follows a geometric Brownian motion process:

$$dS_\tau(k) = S_\tau(k)\left[(r-\delta)d\tau + \sigma dW_\tau(k)\right], \qquad k = 1,\ldots,d \tag{23}$$

where $W_\tau(k)$, $k = 1,\ldots,d$ are independent Brownian motions, r is the riskless interest rate, δ is the dividend rate, and σ is the constant volatility. Exercise opportunities occur at the set of calendar times $\tau_t = t\mathcal{T}/T$, $t = 0,1,\ldots,T$, where $\mathcal{T}$ is the calendar option expiration time. Under the risk-neutral measure, the random variables $\log(S_{\tau_t}(k)/S_{\tau_{t-1}}(k))$ for $k = 1,\ldots,d$ are independent and normally distributed with mean $(r-\delta-\sigma^2/2)(\tau_t - \tau_{t-1})$ and variance $\sigma^2(\tau_t - \tau_{t-1})$. We consider two types of option payoff. *A maximum call* option is a call option on the maximum of the assets with payoff equal to

$$h\big(t,(S(k),k=1,\ldots,d)\big) = e^{-r\tau_t}\left(\max_{1\le k\le d} S(k) - K\right)^+$$

where $(x)^+ := \max(x,0)$. *A geometric average call* option is a call option on the geometric average of the assets with payoff equal to

$$h\big(t,(S(k),k=1,\ldots,d)\big) = e^{-r\tau_t}\left(\left(\prod_{k=1}^{d} S(k)\right)^{\frac{1}{d}} - K\right)^+$$

The second set of test problems is a subset of those in Longstaff and Schwartz (2001) for a put option on a single asset. The dynamics of the asset are as in (23).

Tables 1, 2 and 3 contain results for the maximum call, geometric average call, and put option, respectively. For each choice of parameter values, we consider both out-of-the-money and in-the-money cases. Within each table, each of several different panels (separated by horizontal lines) shows the estimator accuracy as the mesh size b takes increasing values in the set $\{200, 400, 800, 1600\}$. Our performance measures are the relative bias (RB), relative standard error (RSE), and relative root mean square error (RRMSE) of $\hat{q}$, defined as the bias, standard error, and root mean square error divided by the true option value, respectively. We approximated the true option values using the results in Broadie and Glasserman (2004) as follows. For the max option, we used the most accurate estimates in that paper, which have a relative error less than 0.35% with 99% confidence. For the geometric average option, the values are calculated from a single-asset binomial tree, presumably with negligible error. For the put option, we used the values from Longstaff and Schwartz (2001), which were calculated with negligible error with a finite-difference algorithm. The estimates $\widehat{RB}$, $\widehat{RSE}$ and $\widehat{RRMSE}$ in these tables are based on 64 independent replications of $\hat{q}$. For brevity, we omit the standard errors of these estimates; relative to the magnitude of the estimated quantities our own estimation error is negligible.

TABLE 1 Maximum call option on five assets. The short-term interest rate is $r = 0.05$, the dividend rate is $\delta = 0.1$, the volatility is $\sigma = 0.2$, the strike price is $K = 100$, and the time to expiration is $\mathcal{T} = 3$ years. The number of exercise opportunities is $T + 1$. The value of all assets at time zero is x_0. "CPU" refers to CPU time in seconds per replication of the mesh estimator $\hat{q}$ on a SUN Ultra 5 workstation. The columns $\widehat{RB}$, $\widehat{RSE}$, and $\widehat{RRMSE}$ refer to estimates of the relative bias, relative standard error, and relative root mean square error of $\hat{q}$. The option values for the six cases (in order of appearance) are 16.006, 35.695, 16.474, 36.497, 16.659 and 36.782.

T	x_0	Number of mesh points	CPU	$\widehat{RB}$	$\widehat{RSE}$	$\widehat{RRMSE}$
3	90	200	3.3	0.175	0.093	0.198
		400	8.4	0.127	0.052	0.137
		800	24.1	0.089	0.038	0.097
		1600	78.1	0.064	0.023	0.068
	110	200	3.3	0.149	0.044	0.155
		400	8.4	0.115	0.036	0.121
		800	24.3	0.074	0.021	0.077
		1600	78.0	0.054	0.015	0.056
6	90	200	6.6	0.402	0.098	0.414
		400	17.0	0.337	0.066	0.343
		800	49.0	0.288	0.043	0.291
		1600	158.5	0.231	0.029	0.233
	110	200	6.6	0.370	0.066	0.376
		400	16.9	0.331	0.038	0.333
		800	48.7	0.256	0.023	0.257
		1600	158.5	0.203	0.018	0.204
9	90	200	9.9	0.557	0.096	0.566
		400	25.6	0.521	0.064	0.525
		800	73.2	0.466	0.042	0.468
		1600	238.4	0.402	0.032	0.403
	110	200	9.8	0.556	0.061	0.559
		400	25.5	0.503	0.040	0.505
		800	73.2	0.445	0.026	0.446
		1600	239.4	0.368	0.021	0.368

TABLE 2 Geometric average call option on d assets. The short-term interest rate is $r = 0.03$, the dividend rate is $\delta = 0.05$, the volatility is $\sigma = 0.4$, the strike price is $K = 100$, and the time to expiration is $\mathcal{T} = 1$ year. There are 11 exercise opportunities. The value of all assets at time zero is x_0. "CPU" refers to CPU time in seconds per replication of the mesh estimator $\hat{q}$ on a SUN Ultra 5 workstation. The columns $\widehat{\text{RB}}$, $\widehat{\text{RSE}}$, and $\widehat{\text{RRMSE}}$ refer to estimates of the relative bias, relative standard error, and relative root mean square error of $\hat{q}$. The option values for the four cases (in order of appearance) are 1.362, 10.211, 0.761 and 10.

d	x_0	Number of mesh points	CPU	$\widehat{\text{RB}}$	$\widehat{\text{RSE}}$	$\widehat{\text{RRMSE}}$
5	90	200	10.9	0.621	0.320	0.699
		400	28.4	0.610	0.218	0.647
		800	80.7	0.584	0.139	0.601
		1600	260.3	0.493	0.090	0.502
	110	200	11.0	0.533	0.101	0.542
		400	28.6	0.460	0.061	0.464
		800	81.7	0.367	0.042	0.370
		1600	260.4	0.277	0.032	0.279
7	90	200	15.4	0.628	0.336	0.712
		400	39.5	0.635	0.269	0.690
		800	112.9	0.605	0.198	0.636
		1600	362.9	0.610	0.141	0.626
	110	200	15.4	0.477	0.100	0.488
		400	39.3	0.455	0.061	0.459
		800	112.6	0.396	0.041	0.398
		1600	365.3	0.338	0.029	0.340

It is obvious that the mesh estimator is highly positively biased, with bias being the dominant factor in the estimator's overall error, as measured by root mean square error. The most fatal flow of the algorithm illustrated by these results is that the bias decays slowly with b, and this appears to be the general pattern over further experiments that are not reported here. In view of the quadratic growth of the algorithm work with b, extrapolation from these tables suggests that considerable bias will persist for most feasible sample sizes.

Table 1 suggests that the bias and overall error increase rapidly with the number of exercise opportunities, $T + 1$. This is expected in view of Theorem 1, where the upper bound on the estimator error grows geometrically with the number of exercise opportunities. Contrasting Table 3 to Tables 1 and 2 suggests that the estimation error generally increases with the problem dimension, d. Our results shed some light on the effect of problem dimension via the moments that appear in (5) and (7).

TABLE 3 Put option on single asset. The short-term interest rate is $r = 0.06$, the dividend rate is $\delta = 0$ (no dividends), the volatility is $\sigma = 0.4$, the strike price is $K = 40$, and the time to expiration is $\mathcal{T}$ years. There are 50 exercise opportunities per year. The value of the asset at time zero is x_0. "CPU" refers to CPU time in seconds per replication of the mesh estimator $\hat{q}$ on a SUN Ultra 5 workstation. The columns RB, RSE, and RRMSE refer to estimates of the relative bias, relative standard error, and relative root mean square error of $\hat{q}$. The option values for the four cases (in order of appearance) are 7.101, 3.948, 8.508 and 5.647.

$\mathcal{T}$	x_0	Number of mesh points	CPU	$\widehat{\text{RB}}$	$\widehat{\text{RSE}}$	$\widehat{\text{RRMSE}}$
1	36	200	3.8	0.315	0.069	0.322
		400	10.8	0.207	0.047	0.212
		800	35.2	0.129	0.032	0.133
		1600	138.3	0.085	0.021	0.088
	44	200	3.8	0.354	0.098	0.367
		400	10.8	0.241	0.067	0.250
		800	35.3	0.164	0.049	0.171
		1600	138.2	0.100	0.030	0.105
2	36	200	7.7	0.467	0.082	0.475
		400	21.5	0.306	0.050	0.310
		800	70.8	0.187	0.033	0.189
		1600	276.8	0.122	0.020	0.123
	44	200	7.7	0.519	0.096	0.528
		400	21.6	0.353	0.067	0.359
		800	71.2	0.232	0.049	0.237
		1600	287.3	0.154	0.031	0.157

5 Conclusion

We have established the convergence of the original version of the stochastic mesh estimator proposed by Broadie and Glasserman (2004) for pricing Bermudan options. In particular, we have derived an asymptotic upper bound on the probability of error with respect to the mesh size, b. Both the error size and the upper bound on the probability of error are functions of b that vanish as $b \to \infty$. Our results hold under mild finiteness-of-moment assumptions and thus provide support for the use of the algorithm in most situations of practical interest. We note, however, that the method's applicability may be limited by the requirement that risk-neutral transition densities between exercise opportunities should be easily computable.

Our computational experience with the original estimator (with no bias reduction and variance reduction enhancements) shows very poor behavior – specifically, large positive bias. The bias is present even for a small number of exercise opportunities, and decays slowly with the mesh size. The poor empirical performance of the original mesh estimator reported here and elsewhere (Avramidis and Hyden, 1999; Boyle, Kolkiewicz and Tan, 2002) suggests that the constant C appearing in our probability bound may be very large in typical applications. This observation is consistent with the experience of many researchers that likelihood ratios are often highly variable random variables. Inspection of the relevant inequalities suggests that C may grow fast with the problem dimension, underlying the inherent difficulties of the computational problem in higher dimensions.

On a positive note, several computational enhancements of the mesh algorithm have been proposed and studied experimentally. Broadie and Glasserman (2004) reported successful implementations that yield acceptable error via careful adaptation of effective variance reduction techniques. Avramidis and Hyden (1999) and Avramidis *et al.* (2000) reduced the bias and overall error to reasonable levels by devising a biased-low estimator and averaging it with the biased-high mesh estimator studied here and, further, by employing importance sampling. Boyle, Kolkiewicz and Tan (2002) developed further the bias-reduction technique, used low-discrepancy sequences, and observed considerable error reduction. The analysis presented in this paper can probably be adapted to establish convergence of some of these enhanced mesh-type algorithms. Hence, there is reason to believe that enhanced mesh-type algorithms that are convergent and have good small-sample performance can be devised, thus adding a useful tool to the arsenal of computational techniques for pricing Bermudan options.

REFERENCES

Avramidis, A. N., and Hyden, P. (1999). Efficiency improvements for pricing American options with a stochastic mesh. In P. A. Farrington, H. B. Nembhard, D. T. Sturrock and G. W. Evans (eds), *Proceedings of the 1999 Winter Simulation Conference*, pp. 344–50. IEEE Press, Piscataway, New Jersey.

Avramidis, A. N., and Matzinger, H. (2002). Convergence of the stochastic mesh estimator for pricing American options. In E. Yücesan, C. H. Chen, J. L. Snowdon and J. M. Charnes (eds), *Proceedings of the 2002 Winter Simulation Conference*, pp. 1560–7. IEEE Press, Piscataway, New Jersey.

Avramidis, A. N., Zinchenko, Y., Coleman, T., and Verma, A. (2000). Efficiency improvements for pricing American options with a stochastic mesh: parallel implementation. *Financial Engineering News* **19**, 1–7.

Barraquand, J., and Martineau, D. (1995). Numerical valuation of high dimensional multivariate American Securities. *Journal of Financial and Quantitative Analysis* **30**(2), 383–405.

Boyle, P. P., and Kolkiewicz, A. W., and Tan, K. S. (2002). Pricing American derivatives using simulation: a biased low approach. In K.-T. Fang, F. J. Hickernell and H. Niederreiter (eds), *Monte Carlo and quasi-Monte Carlo methods 2000*, pp. 181–200. Springer-Verlag, Berlin.

Broadie, M., and Glasserman, P. (1997a). Monte Carlo methods for pricing high-dimensional American options: an overview. *Electronic Journal of Financial Risk* **3** (December), 15–37.

Broadie, M., and Glasserman, P. (1997b). Pricing American-style securities using simulation. *Journal of Economic Dynamics and Control* **21**(8–9), 1323–52.

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

Clément, E., and Lamberton, D., and Protter, P. (2002). An analysis of a least squares regression algorithm for American option pricing. *Finance and Stochastics* **17**, 448–71.

Duffie, D. (1996). *Dynamic asset pricing theory*, Second edition. Princeton University Press, Princeton, New Jersey.

Harrison, M., and Pliska, S. (1981). Martingales and stochastic integrals in the theory of continuous trading. *Stochastic Processes and Their Applications* **11**, 215–60.

Longstaff, F. A., and Schwartz, E. S. (2001). Valuing American options by simulation: a simple least squares approach. *Review of Financial Studies* **14**(1), 113–47.

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