

Tilburg University

An Irregular Grid Approach for Pricing High-Dimensional American Options

Berridge, S.J.; Schumacher, J.M.

Publication date:
2004

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):

Berridge, S. J., & Schumacher, J. M. (2004). *An Irregular Grid Approach for Pricing High-Dimensional American Options*. (CentER Discussion Paper; Vol. 2004-18). Finance.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Center



Discussion Paper

No. 2004–18

AN IRREGULAR GRID APPROACH FOR PRICING HIGH-DIMENSIONAL AMERICAN OPTIONS

By S.J. Berridge, J.M. Schumacher

February 2004

ISSN 0924-7815

An irregular grid approach for pricing high-dimensional American options

S.J. Berridge* and J.M. Schumacher
Department of Econometrics and Operations Research
and Center for Economic Research (CentER),
Tilburg University, PO Box 90153, 5000 LE Tilburg, The Netherlands.
{s.j.berridge, j.m.schumacher}@uvt.nl

February 6, 2004

Abstract

We propose and test a new method for pricing American options in a high-dimensional setting. The method is centred around the approximation of the associated complementarity problem on an irregular grid. We approximate the partial differential operator on this grid by appealing to the SDE representation of the underlying process and computing the root of the transition probability matrix of an approximating Markov chain. Experimental results in five dimensions are presented for four different payoff functions.

Keywords: American options, high-dimensional problems, free boundary problems, optimal stopping, variational inequalities, numerical methods, unstructured mesh, Markov chain approximation

MSC 2000: 35R35, 60G40, 65D15, 90C33

JEL Codes: C15, C61, C63

*Research support by the Dutch National Science Foundation (NWO).

1 Introduction

The pricing of American options has always required numerical solution methods; in high-dimensional cases even the most sophisticated methods have difficulty in providing accurate solutions. Given the practical importance of such cases, it is of considerable interest to develop solution methods which are reliable and which provide accompanying exercise and hedging strategies.

Barraquand and Martineau [1] are perhaps the first to consider pricing high-dimensional American options specifically, proposing an algorithm based on the aggregation of paths with respect to the intrinsic value. The method is difficult to analyse and has a possible lack of convergence; Boyle et al. [2] demonstrate this and propose a modification of the algorithm which leads to a low-biased estimator.

Broadie and Glasserman [5] use a stochastic tree algorithm to give both a low-biased and a high-biased estimator of the price, both asymptotically unbiased. They also argue that there exists no nontrivial unbiased estimator for the price. Their method requires an exponentially increasing amount of work in the number of exercise opportunities. In subsequent work [6] they present a related method based on a stochastic mesh which does not suffer from this problem, although this method has been found to be slow by several authors and to have a large finite-sample bias (see e.g. Fu et al. [11]).

The least squares Monte Carlo (LSM) method presented by Longstaff and Schwartz [21] attempts to approximate the price of an American option using cross-sectional information from simulated paths. The optimal exercise strategy is successively approximated backwards in time on the paths by comparing the intrinsic values to the continuation values projected onto a number of basis functions over the states. Experimental success is reported for the LSM method, although in high dimensions the basis functions must be chosen carefully. Recently Clément et al. [7] and Stentoft [26] independently provide proofs of convergence for the LSM method, showing that the convergence rate is $n^{-1/2}$ in the number of paths used. The convergence behaviour in the number of basis functions however has not been determined. Stentoft [26] and Glasserman and Yu [13] establish relationships between the paths and number of basis functions which are necessary for convergence; Stentoft finds that the number of paths should be greater than cubic in the number of basis functions to achieve convergence in probability, while Glasserman and Yu find the relationship should be exponential in the square for convergence on a worst case basis. Stentoft [25] and Moreno and Navas [22] test the LSM algorithm numerically. Stentoft suggests that basis functions up to order three are sufficient in five dimensions for arithmetic and geometric average options, but not for minimum or maximum options. Moreno and Navas find that the method is sensitive to the choice of basis functions in five dimensions.

Tsitsiklis and Van Roy [27] propose a method similar to LSM where approximate value functions are projected onto an orthogonal set of basis functions, the orthogonality being with respect to a suitably chosen inner product which in general changes between time periods. They provide a proof of convergence but no empirical results. The method differs from LSM in that the projection is used to determine an approximate value function rather than an exercise rule.

Boyle et al. [4] recently extended the stochastic mesh method of Broadie and Glasserman [6] with their low discrepancy mesh (LDM) method. This involves generating a set of low discrepancy interconnected paths and using a dynamic programming approach to find prices on the mesh.

An interesting alternative approach is proposed independently by Rogers [24] and Haugh and Kogan [16] and later developed by Jamshidian [18] and Kolodko and Schoenmakers [19]. They use a dual formulation of the problem in which a minimisation is performed over martingales. The method is sensitive to the choice of basis martingales chosen to perform the minimisation, and so requires the basis to be well-chosen in order to give an accurate solution. The method gives a high-biased estimator.

In this paper, we propose a new approach to solving the American option pricing problem inspired by the success of numerical integration in high dimensions and related to the method of lines for solving partial differential equations (PDEs).

We first perform a discretisation of the state space using quasi-Monte Carlo (QMC) points, the points being taken with respect to an importance sampling distribution related to the transition density of the process at expiry. We then propose an approximation to the partial differential operator on this grid by taking the logarithm of a transition probability matrix $P^{(T-t)}$ which approximates the joint density of the underlyings at the expiry of the option, $T - t$. This approximation is then used to formulate linear complementarity problems (LCPs) at successive time points, working back from the option expiry.

We propose an implementation of this method in which the matrix logarithm of $P^{(T-t)}$ does not need to be calculated explicitly, but instead a root of the matrix can be calculated. The root operation is cheaper than the logarithm, although the logarithm allows variation of the time step without recalculation. The computational elements of the method are thus the QMC trials, the generation of the matrix $P^{(T-t)}$, the matrix root and solving an LCP at each time step. For approximating the European option price this method amounts to performing a numerical integration with importance sampling, which is known to be an efficient method in high dimensions as long as the importance sampling distribution is chosen appropriately.

The remainder of the paper is organised as follows. In Section 2 we present a mathematical formulation of the problems to be solved numerically and in Section 3 we show how an irregular grid method can be used to solve the problem. We then

present the experimental setup in Section 4, results in Section 5 and concluding remarks in Section 6.

2 Formulation

We consider a complete and arbitrage-free market described by state variable $X(s) \in \mathbb{R}^d$ for $s \in [t, T]$ which follows a Markov diffusion process

$$dX(s) = \mu(X(s), s)ds + \sigma(X(s), s)dW(s) \quad (2.1)$$

with initial condition $X(t) = x_t$, and a derivative product on $X(s)$ with intrinsic value $\psi(X(s), s)$ at time s and value $V(s) = v(X(s), s)$ for some pricing function $v(x, s)$. The process $V(s)$ satisfies

$$dV(s) = \mu_V(X(s), s)ds + \sigma_V(X(s), s)dW(s) \quad (2.2)$$

where μ_V and σ_V can be expressed in terms of μ and σ by means of Itô's lemma. The terminal value is given by $v(\cdot, T) = \psi(\cdot, T)$.

In such a market there exists a unique equivalent martingale measure under which all price processes are martingales. The risk-neutral process in this case is given by

$$dX(s) = \mu_{RN}(X(s), s)ds + \sigma(X(s), s)dW(s) \quad (2.3)$$

where μ_{RN} is the risk-neutral drift.

Our objective is to provide approximations for the current value $v(x_t, t)$ of the derivative product and the corresponding optimal exercise and hedging strategies τ and H :

$$\tau : \mathbb{R}^d \times [t, T] \rightarrow \{0, 1\} \quad (2.4)$$

$$H : \mathbb{R}^d \times [t, T] \rightarrow \mathbb{R}^d. \quad (2.5)$$

In the following, we appeal to the complementarity formulation of the American option price which is presented for example in Jaillet et al. [17]. Let \mathcal{L} be the related diffusion operator

$$\mathcal{L} = \frac{1}{2}\text{tr}\sigma\sigma' \frac{\partial^2}{\partial x^2} + \mu_{RN} \frac{\partial}{\partial x} - r \quad (2.6)$$

where r is the risk-free rate. Then the option value is found by solving the complementarity problem

$$\left\{ \begin{array}{l} \frac{\partial v}{\partial t} + \mathcal{L}v \leq 0 \\ v - \psi \geq 0 \\ (\frac{\partial v}{\partial t} + \mathcal{L}v)(v - \psi) = 0 \end{array} \right. \quad (2.7)$$

for $(x, s) \in \mathbb{R}^d \times [t, T]$ with the terminal condition $v(\cdot, T) \equiv \psi(\cdot, T)$.

3 Methodology

To solve the complementarity problem (2.7) we first form a semidiscrete complementarity problem by discretising the state space but leaving time continuous. This involves sampling the state space using QMC trials and finding a suitable approximation of \mathcal{L} . We then use standard time stepping techniques to form a system of fully discrete LCPs. There are many methods for solving LCPs; examples include projected successive overrelaxation (PSOR) and linear programming.

We first present and motivate each step of the algorithm separately, and then summarise by providing a concise statement of the algorithm.

3.1 State space discretisation

We first consider a semidiscrete approximation to the complementarity problem (2.7) in which the state space is discretised and time left continuous. This is often called the method of lines. In the pricing problem this amounts to approximating (2.7) by a system of ordinary differential equations with complementarity conditions.

The choice of a constant grid in the state space has the advantage that Crank-Nicolson and implicit solutions can be easily considered. This seems advantageous since, in the case of solving PDEs without complementarity conditions, the Crank-Nicolson method is known to have a convergence rate of δt^2 rather than δt for other first order schemes. Furthermore when solving discretised complementarity problems, the implicit scheme is the only time stepping method known to be unconditionally stable (see Glowinski et al. [14]).

The choice of grid begs importance sampling considerations. That is, in order to obtain a more accurate approximation, more grid points should be placed at states which are more likely to be visited by the process, and at locations where the value function has greater magnitude.

We denote the grid by $\mathcal{X} = \{x_1, \dots, x_n\} \subseteq \mathbb{R}^d$, and the corresponding operator approximation by A . The construction of A will be considered in Section 3.2.

Assuming that \mathcal{X} and A are given, we form the corresponding semidiscrete complementarity problem

$$\begin{cases} \frac{dv}{dt}(s) + Av(s) \leq 0 \\ v(s) - \psi \geq 0 \\ \left(\frac{dv}{dt}(s) + Av(s)\right)' (v(s) - \psi) = 0 \end{cases} \quad (3.1)$$

for $s \in [t, T]$ with terminal condition $v_i(T) = \psi(x_i)$ for each $i = 1, \dots, n$. Note that $v(s)$ is now a time-dependent vector in \mathbb{R}^n where n is the number of grid

points.

It is also instructive to view the semi-discrete setting as a Markov chain approximation to the optimal stopping problem. That is, the process $X(s)$ is approximated by a process restricted to the grid \mathcal{X} . The operator A gives transition intensities on this grid.

3.2 Approximating the partial differential operator

We now propose a method for specifying A in (3.1) for a given grid \mathcal{X} . The method is inspired by numerical integration, and in the European case the resulting method will reduce to numerical integration with importance sampling. This property is emphasised in Glasserman [12] as a favourable property of the stochastic mesh method presented by Broadie and Glasserman [6].

We assume that the grid \mathcal{X} has been generated using random or QMC draws with respect to a certain density $g(x)$. We also assume that the joint density $f_{x,T-t}$ of the stochastic process is available for arbitrary initial points x and time horizons $T - t$, although in principle one could adapt the following construction to the case where the density is not known explicitly, but for example the process can be simulated.

Denote by $P^{(T-t)}$ the matrix with entries

$$p_{ij}^{(T-t)} = \frac{1}{\sum_{k=1}^n \tilde{f}_{x_i, T-t}(x_k)} \cdot \tilde{f}_{x_i, T-t}(x_j) \quad (3.2)$$

where the weights are given by

$$\tilde{f}_{x_i, T-t}(x) = \frac{f_{x_i, T-t}(x)}{g(x)}. \quad (3.3)$$

The matrix $P^{(T-t)}$ is a stochastic matrix, that is, a matrix with nonnegative entries and unit row sums. We think of the entries as giving transition probabilities between points in the grid over the horizon $T - t$.

In the semidiscrete Markov chain setting, where A represents transition intensities, we note that the evolution of state probabilities is given by $p(s) = e^{A(s-t)}p(t)$ where $p(t)$ is the initial probability distribution at time t , for example it may be a delta function in the case where the initial state is known. The matrix $P^{(T-t)}$ thus gives us access to an approximation A to \mathcal{L} on \mathcal{X} as follows:

$$A \triangleq \frac{1}{T-t} \log P^{(T-t)}. \quad (3.4)$$

The matrix logarithm of $P^{(T-t)}$ certainly exists and is unique if the matrix is diagonalisable and has positive eigenvalues. We find these two properties hold

in our experiments; note however that $P^{(T-t)}$ is in general not symmetric. We shall see in Section 3.3 that instead of computing the matrix logarithm, one may alternatively compute the matrix root corresponding to the required time step.

3.3 Time discretisation

Let us now discretise (3.1) with respect to time. We denote the approximation at state x_i and time step t_k by $v_{i,k}$.

We use the θ -method, standard in the numerical solution to PDEs, to discretise (3.1). For PDE solutions, $\theta = 0$ corresponds to the explicit method, $\theta = 1$ corresponds to the implicit method and $\theta = \frac{1}{2}$ corresponds to the Crank-Nicolson method. The latter has δt^2 convergence for European problems, whereas the explicit and implicit methods exhibit δt convergence.

To implement the θ -method, we consider the vector $v^{(k)}$ of values at our grid points each at time t_k and discretise the first line of (3.1) as

$$\frac{v^{(k+1)} - v^{(k)}}{\delta t_k} + A \left((1 - \theta)v^{(k+1)} + \theta v^{(k)} \right) \leq 0 \quad (3.5)$$

where $\delta t_k \triangleq t_{k+1} - t_k$. Thus (3.1) becomes

$$\left\{ \begin{array}{l} (I + (1 - \theta)A\delta t_k) v^{(k+1)} - (I - \theta A\delta t_k) v^{(k)} \leq 0 \\ v^{(k)} - \psi \geq 0 \\ ((I + (1 - \theta)A\delta t_k) v^{(k+1)} - (I - \theta A\delta t_k) v^{(k)})' (v^{(k)} - \psi) = 0. \end{array} \right. \quad (3.6)$$

Now note that $I + A\delta t_k = \exp(A\delta t_k) + o(\delta t_k)$. We thus define the matrices

$$M_L = \exp \{-\theta A\delta t_k\} \quad (3.7)$$

$$M_R = \exp \{(1 - \theta)A\delta t_k\}. \quad (3.8)$$

The approximating complementarity problem to solve is then

$$\left\{ \begin{array}{l} M_R v^{(k+1)} - M_L v^{(k)} \leq 0 \\ v^{(k)} - \psi \geq 0 \\ (M_L v^{(k)} - M_R v^{(k+1)})' (v^{(k)} - \psi) = 0 \end{array} \right. \quad (3.9)$$

for $k = K - 1, \dots, 0$ where the inequalities are componentwise and $v^{(K)} = \psi$.

Numerically we must solve a LCP at each time step, for which the PSOR method of Cryer [8] has been used with much success in the past. Since the solution does not change greatly between time steps, a good starting guess for PSOR is

the solution at the previous time step. Various other methods may be used for solving (3.9); for example, see Dempster and Hutton [9] for American option pricing using linear programming in the one-dimensional case.

An error analysis of the discretisation in (3.5) may be undertaken along the lines of Glowinski et al. [14] on variational inequalities or that of Kushner and Dupuis [20] on stochastic control.

It turns out that the matrix logarithm does not have to be calculated explicitly in our method; instead we may calculate roots of the matrix $P^{(T-t)}$ corresponding to the time step and implicitness parameters. In particular we have

$$M_L = \left(P^{(T-t)} \right)^{-\theta \delta t_k / (T-t)} \quad (3.10)$$

$$M_R = \left(P^{(T-t)} \right)^{(1-\theta) \delta t_k / (T-t)}. \quad (3.11)$$

We prefer to use the matrix root because we have found it to be a quicker and more robust operation in Matlab than the matrix logarithm. If one would choose to compute the logarithm however, one would have access to a varying time step without performing any extra computations.

There are many methods available for evaluating matrix functions, as detailed in Golub and Van Loan [15]. The general method suggested involves Schur decomposition in combination with Parlett's algorithm, which computes general functions of an upper triangular matrix. Matrix functions can also be computed using eigen-decomposition, which is the method used by Matlab to compute general matrix powers. We note that the structure of the matrix $P^{(T-t)}$ may mean that more efficient methods are available for computing matrix roots and logarithms; it is not the purpose of the current research to investigate such methods however.

We now highlight the importance of using the matrix logarithm or root, as opposed to constructing $P^{(\delta t)}$ directly (the latter being more attractive computationally). The intuition for this importance is that $P^{(\delta t)}$ does not produce consistent transition probabilities over longer time horizons as in (3.12). We demonstrate the difference between the two constructions in Figure 3.1 for a one-dimensional example and a random grid. In particular, when δt is too small compared to the separation of grid points, the solutions become distorted. This problem is more pronounced in higher dimensions due to the larger average separation between grid points.

Remark 3.1 It is clearly more efficient if the matrices M_L and M_R need be calculated only once; hence the choice of a constant time step $\delta t_k \equiv \delta t$ seems convenient. We also note that, given a small enough δt , M_L and M_R should be approximately sparse in that most elements can be set to zero without affecting the solution

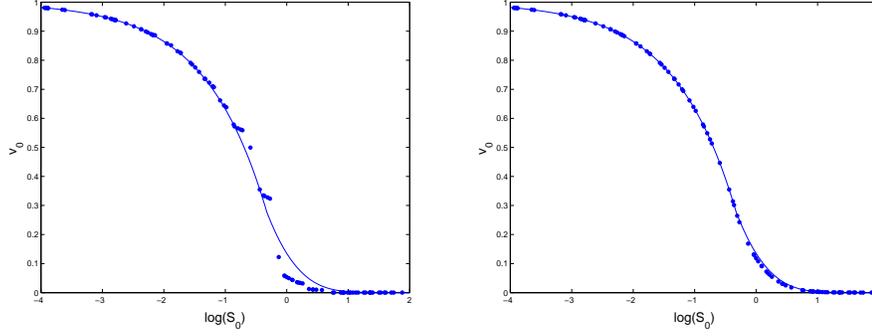


Figure 3.1: Random grid valuation of an American put option on a single asset with expiry 1, strike 1 and 100 asset points, using transition matrices $P^{(0.01)}$ and $(P^{(1)})^{0.01}$ respectively (*dots*). The plots are in the log domain. Also plotted are values computed to high accuracy (*solid lines*) using a standard finite difference method.

significantly. Using this observation can dramatically improve the efficiency of the solution procedure.

Remark 3.2 As already noted above, an important property of this construction in the European setting is that the time stepping reconstructs numerical integration with importance sampling function g . The reconstruction is realised as follows:

$$\begin{aligned}
 v_i(t) &= \left(\prod_{k=1}^K e^{-r\delta t} M_L^{-1} M_R \right) v^{(K)} \\
 &= e^{-r(T-t)} \exp(A) v^{(K)} \\
 &= e^{-r(T-t)} \sum_{j=1}^n \psi(x_j) p_{ij}^{(T-t)} \quad (3.12)
 \end{aligned}$$

where $v_i(t)$ is the value in state x_i at initial time t and $p_{ij}^{(T-t)}$ is defined in (3.2). The last line of the equation is precisely QMC integration of the payoff ψ with importance sampling function equal to the grid density g . Note that in case M_L and M_R are constructed from the matrix logarithm, (3.12) holds only asymptotically as $\delta t \rightarrow 0$.

Equation (3.12) also shows that the calculation of the European price on the grid may be carried out without time stepping, given that the transition probabilities $p_{ij}^{(T-t)}$ are available. Thus, using the European price as control variate is a faster operation than would normally be expected.

3.4 Randomisation

The QMC grids we have proposed are deterministic; however perturbing these points randomly allows us to observe the behaviour of solutions for a random selection of QMC grids, and thus to obtain estimates of the bias and standard error of solutions. The use of such methods is surveyed in Owen [23] for integration problems. The importance of randomised QMC is also emphasised in Glasserman [12].

When using Sobol' points and a normal density for example, one first generates the Sobol' points, then applies the inverse normal distribution function to the points. In order to realise randomised QMC points, one perturbs the Sobol' points modulo one by a random factor before applying the inverse normal distribution function.

Suppose $\mathcal{S} = (s_i)$ is our sequence of n Sobol' points, and U_j is a sequence of random variables uniformly distributed on the unit cube $[0, 1]^d$. We then realise the j th randomised Sobol' sequence as

$$\mathcal{S}_j = (s_{j,i})_{i \in \mathbb{N}} = (s_i + U_j \pmod{1})_{i \in \mathbb{N}}. \quad (3.13)$$

We refer to grids obtained in this way as randomised QMC (RQMC) grids.

3.5 Summary of procedure

We now present a concise statement of the proposed procedure as Algorithm 3.1. We let $\hat{v}_{i,j}$ denote the solution at initial time t and state x_i in the j th RQMC experiment. For the statement of the algorithm we assume a fixed number of grid points n and a constant time step $\delta t = (T - t)/K$ where K is the number of time steps.

4 Experimental setup and details

We now use the algorithm presented in Section 3.5 to estimate prices of multi-asset options. We first present a detailed exposition of the setting, experimental procedure and various considerations. Numerical results are presented in the next section.

4.1 Specification of dynamics

Suppose our American option is based on d assets following a correlated geometric Brownian motion where the risk-neutral dynamics in the log domain are given by

$$dX = (r\mathbb{1} - \delta - \frac{1}{2}\text{diag}(\Sigma)) dt + R'dW \quad (4.1)$$

Algorithm 3.1 The proposed irregular grid algorithm

for $j = 1, \dots, J$ **do**

 Generate a RQMC grid \mathcal{X}

 Compute the transition matrix for expiry $P^{(T-t)}$

 Compute the matrix root $(P^{(T-t)})^{1/2K}$ (Crank-Nicolson)

 Solve the LCPs (3.9)

 Let $\hat{v}_{i,j}$ be the solution at initial time t for state x_i

end for

for initial states of interest x_i **do**

 Estimate the solution as $\hat{v}_i = \frac{1}{J} \sum_j \hat{v}_{i,j}$

 Estimate the standard error as $\hat{\epsilon}_i = \left(\frac{1}{J-1} \sum_j (\hat{v}_{i,j} - \hat{v}_i)^2 \right)^{1/2}$.

end for

and r is the risk-free rate, $\mathbb{1}$ is the d -vector of ones, $\delta = (\delta_1, \dots, \delta_d)$ is the vector of dividend rates, $\Sigma = (\rho_{ij}\sigma_i\sigma_j)$ is the covariance matrix of the Brownian motions and $R'R$ is its Cholesky decomposition. The operator \mathcal{L} in this setting is just the multidimensional Black-Scholes operator given by

$$\mathcal{L} = \frac{1}{2} \sum_{i,j=1}^d \rho_{ij}\sigma_i\sigma_j \frac{\partial^2}{\partial x_i \partial x_j} + \sum_{i=1}^d (r - \delta_i - \frac{1}{2}\sigma_i^2) \frac{\partial}{\partial x_i} - r. \quad (4.2)$$

4.2 Elimination of drift

In order to facilitate reuse of the matrix roots, we first reformulate the problem so that the process has zero drift. We introduce the change of variables

$$X_0(s) = X(s) - (s - t)\mu, \quad (4.3)$$

where μ is the risk-neutral drift; for example in (4.1) we have $\mu = r\mathbb{1} - \delta - \frac{1}{2}\text{diag}(\Sigma)$. The new process X_0 has zero drift and the covariance Σ is unchanged:

$$dX_0(s) = RdW(s). \quad (4.4)$$

The payoff under the reformulation is

$$\psi_0(x_i, s) = \psi(x_i + (s - t)\mu). \quad (4.5)$$

One may also eliminate a deterministic, time-dependent risk-neutral drift by subtracting $\int_t^s \mu(u)du$ in (4.3).

4.3 Grid specification

We consider normal RQMC grids as suggested in Section 3.4; thus the grid density is multivariate normal. We now discuss parameter selection for the grid density.

Importance sampling considerations tell us that the most efficient sampling is given by the density of the process itself; thus using a constant grid we cannot provide the most efficient importance sampling at all times. However, given the restriction to a constant grid, we can still provide an acceptable importance sampling.

As outlined in Evans and Swartz [10], the rate of convergence for importance sampling of normal densities using normal importance sampling functions is most damaged when the variance of the importance sampling function is less than that of the true density. Conversely, convergence rates are not greatly affected when the variance of the importance sampling function is greater than that of the true density.

The situation we should try to avoid is that the process has a significant probability of lying in the “tails” of the grid density. A further consideration is the minimisation of boundary effects on the solution. This suggests that the grid covariance should be larger than the covariance of the process.

These considerations lead us to set the grid mean to the initial state x_t and the grid covariance to be a multiple α of the grid density at expiry for some trial values $\alpha = 1.0, 1.5, 2.0$. Owing to the reformulation (4.3), this ensures that the grids are centered at the process mean for all times. We further ensure that the initial state is included in the grid.

Summarising, we suggest the parameters

$$\mu_g = x_t \tag{4.6}$$

$$\Sigma_g = \alpha \Sigma(T - t). \tag{4.7}$$

The first grid point in the j th RQMC experiment is $x_1 = \mu_g$ and the $(i + 1)$ th grid point is generated as

$$x_{i+1} = \mu_g + R'_g \left(\Psi^{-1}(s_{j,i,1}) \cdots \Psi^{-1}(s_{j,i,d}) \right)' \tag{4.8}$$

where Ψ^{-1} is the standard normal inverse function, $R'_g R_g$ is the Cholesky decomposition of Σ_g and $s_{j,i,k}$ is the k th component of $s_{j,i}$.

An example of a normal Sobol’ grid in two dimensions is shown in Figure 4.1. It should be noted however that the advantage of using an irregular grid is realised in dimensions of at least three.

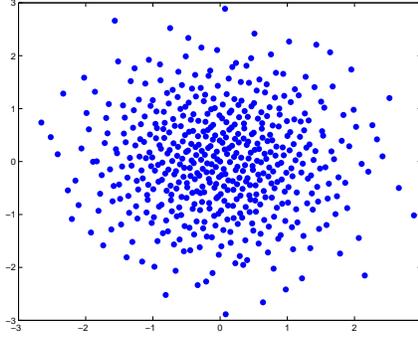


Figure 4.1: Example of normal QMC grid in 2 dimensions with 500 points.

4.4 Reuse of roots for similar processes

Given that generating matrix roots is an expensive operation compared to the final time stepping procedure, it is of interest to know under which conditions these matrix roots can be reused for related problems; for example, problems with different parameters.

Clearly a single matrix root can be reused for as many different payoff functions as required, but we also show how it can be reused for processes with different risk-neutral drifts and covariances. To answer this question for diffusion processes with zero drift we provide the following result.

Lemma 4.1 *Suppose that $P^{(T-t)}$ is the transition matrix corresponding, through (3.2) and (3.3), to the grid $\mathcal{X} = \{x_1, \dots, x_n\}$, respective importance sampling weights g_1, \dots, g_n , horizon $T - t$ and an d -dimensional Brownian motion with covariance I_d . Suppose further that $\tilde{P}^{(T-t)}$ is the transition matrix corresponding to the grid*

$$\mathcal{Y} = \{y_1, \dots, y_n\} = \{R'x_1, \dots, R'x_n\}, \quad (4.9)$$

importance sampling weights g_1, \dots, g_n , horizon $T - t$ and an d -dimensional Brownian motion with positive definite covariance $\Sigma = R'R$.

Then

$$\tilde{P}^{(T-t)} = P^{(T-t)}. \quad (4.10)$$

Proof. Let $f_{x,T-t}$ and $h_{x,T-t}$ be the densities at expiry from starting point x , expiry $T - t$ and corresponding to covariance I and Σ respectively. The densities

from x_i to x_j in grid \mathcal{X} and from $y_i = R'x_i$ to $y_j = R'x_j$ in grid \mathcal{Y} are respectively

$$\begin{aligned}
 f_{x_i, T-t}(x_j) &= |2\pi(T-t)I_d|^{-1/2} \exp \left\{ -\frac{1}{2(T-t)}(x_j - x_i)'(x_j - x_i) \right\} \\
 h_{y_i, T-t}(y_j) &= |2\pi(T-t)\Sigma|^{-1/2} \exp \left\{ -\frac{1}{2(T-t)}(x_j - x_i)'R\Sigma^{-1}R'(x_j - x_i) \right\} \\
 &= |2\pi(T-t)\Sigma|^{-1/2} \exp \left\{ -\frac{1}{2(T-t)}(x_j - x_i)'(x_j - x_i) \right\},
 \end{aligned}$$

which are equal up to the constant factor $|\Sigma|^{1/2}$. Given the latter observation and

that the weights are equal in both cases, we conclude that the normalised entries $p_{ij}^{(T-t)}$ and $\tilde{p}_{ij}^{(T-t)}$ obtained through (3.2) and (3.3) are also equal.

□

Remark 4.1 One may ask whether the weights g_i specified in Lemma 4.1 are indeed appropriate for the grid \mathcal{Y} . That is, whether (3.3) leads to a standard importance sampling procedure for \mathcal{Y} with these weights. We answer this question by comparing the grid densities.

Suppose that the density g_X was used to generate the grid \mathcal{X} using random sampling. So that for every $S \subset \mathbb{R}^d$,

$$P(x \in S) = \int_S g_X(x) dx.$$

Applying the transformation $x \mapsto y = R'x$ leads us to conclude that the grid \mathcal{Y} consists of points generated randomly from some density g_Y satisfying, for each $S \subset \mathbb{R}^d$,

$$\begin{aligned} P(y \in R'S) &= \int_{R'S} g_Y(y) dy \\ &= \int_S g_Y(R'x) ||R|| dx \end{aligned} \quad (4.11)$$

by the multivariate substitution formula where $R'S = \{R'x : x \in S\}$ and $||R|| = \text{abs}(\det R)$. But since $y = R'x$, $P(y \in R'S) = P(x \in S)$, and since (4.11) holds for all $S \subset \mathbb{R}^d$, we conclude that

$$g_X(x) = ||R|| g_Y(R'x). \quad (4.12)$$

Finally, note that the averaging taking place in (3.3) implies that the weights $g_X(x)$, being proportional to $g_Y(R'x)$, are appropriate for importance sampling with respect to the grid \mathcal{Y} .

Remark 4.2 A time-dependent scaling of the covariance can also be incorporated by using the matrix logarithm, and constructing the time stepping matrices through (3.7) and (3.8) rather than (3.10) and (3.11).

4.5 Low-biased estimate

As is common practise in the American option pricing literature, a low biased estimate may be obtained by taking an exercise rule implied by the pricing method and determining the expected value of using this rule using out-of-sample paths.

A natural approximation to the optimal exercise rule using the pricing results of the irregular grid method is to take the implied rule of the nearest neighbour in the grid at the closest time. Specifically one may define the exercise rule for grid points to be

$$\tau(x_i, t_k) \triangleq \begin{cases} 1 & \text{if } v_i^{(k)} = \psi_i \\ 0 & \text{otherwise} \end{cases} \quad (4.13)$$

and for general points $x \in \mathbb{R}^d$

$$\tau(x, t) \triangleq \begin{cases} 1 & \text{if } v_i^{(k)} = \psi_i \\ 0 & \text{otherwise} \end{cases} \quad (4.14)$$

where $k = \operatorname{argmin}_j |t - t_j|$ and $i = \operatorname{argmin}_j \{ \|x - x_j\| : x_j \in \mathcal{X}_k \}$.

This rule is easily implemented and can also be adapted to the case where we have several different grids. In this case one could base the exercise rule on a vote between grids. One could also implement weighted schemes with respect to x and t rather than using nearest neighbour rules.

4.6 High-biased estimate

Whereas applying an exercise rule to out-of-sample paths leads to a low-biased estimate of the option value, simulating the cost of a hedging strategy leads to a high-biased estimate. The latter may be seen as follows: the optimal hedging strategy enables the seller of the option, given a cash amount equal to the value of the option at the initial time t , to perfectly reproduce the payoff. A suboptimal strategy however will on average require a larger initial cash amount, thus the cost of a suboptimal hedging strategy is on average higher than the true option value.

A formal demonstration can be given in terms of the dual formulation for American option pricing (see Rogers [24], Haugh and Kogan [16]) in which one minimises the cost of hedging by minimising an objective function over martingales. Since the value of our hedging strategy is a martingale, it corresponds in general to a suboptimal martingale, and thus a high-biased estimate.

In practise, obtaining an upper bound in the way we suggest requires knowledge of the optimal exercise rule. Since we only have an estimate of this, the cost of the hedging strategy may not be purely upward biased. We find however that one can approximate the optimal exercise strategy far more accurately than one can approximate the optimal hedging strategy. We shall see in Section 5 that experimental results support this statement.

In the literature on American options there is little said about the practicalities of hedging in a high dimensional setting. The difficulty with using an approach

such as LSM is that the method does not naturally form approximations to the value function from which derivatives can be estimated. One can form a hedging strategy by evaluating prices at states perturbed in each underlying; this demands the calculation of many additional option prices at each time step, each calculation being expensive in a high-dimensional setting. Furthermore one must be very careful with partial derivative estimates obtained from differencing stochastic point estimates; in particular the point estimates must be sufficiently accurate and the perturbations must be well-chosen with respect to the (unknown) curvature of the value function.

A solution provided by the irregular grid method involves estimates of the price not only at the current state, but at all states in the grid. This allows one to extract derivative estimates using value information from nearby points in the grid; for example using partial derivatives implied by a local linear regression. Indeed the irregular grid method provides derivative information as a by-product.

4.7 Benchmarks

There are few benchmark results for high-dimensional American options. Broadie and Glasserman [6] provide 90% confidence intervals for American call options on the maximum of five assets with nine exercise opportunities and the geometric average of five and seven assets with ten exercise opportunities using their stochastic mesh method. Longstaff and Schwartz [21] price the Broadie and Glasserman maximum options using the LSM method.

Stentoft [25] uses the binomial method of Boyle et al. [3] and the LSM method to price put options on the arithmetic average, geometric average, maximum and minimum of three and five assets. Broadie and Glasserman [5] and Fu et al. [11] provide benchmark results for options over five assets with three exercise opportunities. Finally, Rogers [24] and Haugh and Kogan [16] use the dual formulation to price a number of different American options.

A useful result involving options on the geometric average of several assets is that this problem can be easily reduced to an option pricing problem in one dimension. Suppose that the risk-neutral dynamics in the log domain are given by (4.1), and the payoff function

$$\psi(s) = \left(K - \left(\prod s_i \right)^{1/d} \right)^+ \quad (4.15)$$

where K is the strike price and d is the number of assets. Then using Itô's lemma one finds that the price is the same as that of a vanilla put on the asset with log

price Y where $Y(t) = \frac{1}{d} \sum_{i=1}^d X_i(t)$ and

$$dY(s) = \frac{1}{d} \sum_{i=1}^d dX_i(s) \quad (4.16)$$

$$= \tilde{\mu} ds + \tilde{\sigma} dW(s). \quad (4.17)$$

The parameters of the diffusion are given by

$$\tilde{\mu} = r - \frac{1}{2d} \sum_{i=1}^d \sigma_i^2 \quad (4.18)$$

$$\tilde{\sigma}^2 = \frac{1}{d^2} \sum_{i=1}^d \left(\sum_{j=1}^d R_{ij} \right)^2. \quad (4.19)$$

Using this we find that an accurate price for the geometric average European option in the Stentoft setting is 1.159, and the Bermudan and American prices are 1.342 and 1.355 respectively. Note that the difference in early exercise premium between the Bermudan, which allows ten exercise opportunities, and American prices is about 6%.

5 Experimental results

Our experiments are conducted in a Matlab environment and are based on the five-dimensional examples of Stentoft [25]. Specifically we consider five stock processes driven by correlated Brownian motions for put options with four different payoff functions. The method we use for valuation is that of Section 3. Our programs are mostly script-based but some computationally intensive routines, for example the PSOR code, have been implemented in C.

We are given initial stock prices $S_i(0) = 40$ for each i , the correlations between log stock prices are $\rho_{ij} = 0.25$, $i \neq j$, and volatilities are $\sigma_i = 0.2$ for all i , the risk-free interest rate is fixed at $r = 0.06$, the expiry is $T = 1$ and we use $K = 10$ time-steps.

We generate 50 RQMC normal grids as detailed in Section 3.4 using the parameter values $\alpha = 1, 1.5, 2$ respectively (these were found to give the best rates of convergence). The number of grids need not be so high in practise, depending on the accuracy required. The vector of initial stock prices x_0 was always included in the grid.

The payoff functions considered correspond to put options on the arithmetic mean, geometric mean, maximum and minimum respectively,

$$\begin{aligned} \psi_1(s) &= \left(K - \frac{1}{d} \sum s_i\right)^+ & \psi_2(s) &= \left(K - (\prod s_i)^{1/d}\right)^+ \\ \psi_3(s) &= \left(K - \max(s_i)\right)^+ & \psi_4(s) &= \left(K - \min(s_i)\right)^+ \end{aligned} \quad (5.1)$$

where x^+ denotes the positive part of x .

Figures 5.1–5.3 show the convergence behaviour of the irregular grid method where the implicitness parameter is $\theta = 0, \frac{1}{2}, 1$ respectively, and for grid sizes up to 1000. For the constrained solutions, we see that convergence is usually fastest for $\alpha = 1.5$, the algorithm reaching a fairly stable value for $n = 1000$ for all but the maximum option.

The solutions for the arithmetic and geometric average options appear to converge to Stentoft's solutions for Bermudan options with ten exercise opportunities in the explicit case. For the Crank-Nicolson and implicit cases, the solutions appear to converge to a higher value.

The previous observation may be explained as follows. In the explicit case, our method calculates the price of a Bermudan option with ten exercise opportunities, just as in the case of Stentoft (provided we use ten equal time steps). This is because the explicit formulation takes the maximum of the intrinsic and continuation values at each exercise opportunity, and because we use exactly ten time steps. One can still see this Bermudan price as an approximation to the true American price, which we calculated previously to have an early exercise premium approximately 6% higher than the Bermudan price in the case of the geometric average put option over five assets.

In the Crank-Nicolson and implicit cases however, we cannot interpret the solution as approximating a Bermudan option due to the implicitness of the formulation. We can only say that as $\delta t \rightarrow 0$, the solution should converge to the American price. In the Crank-Nicolson case we suspect that the convergence is faster than in the implicit case (drawing a parallel with the unconstrained problem), and so we can think of our Crank-Nicolson solution as being close to our best possible approximation to the true American option value, given that we use ten equal time steps. We thus stress that the convergence of the Bermudan price does not require $\delta t \rightarrow 0$, but the convergence of the American price does.

The fastest convergence rate in Figures 5.1–5.3 is achieved with α (the ratio of grid density to process density) being 1.5. We thus present in Tables 5.1 and 5.2 some results and comparisons for Bermudan and American option prices respectively using the irregular grid method with a normal grid and $\alpha = 1.5$. Given the previous discussion, we take our explicit solutions to be approximations to the

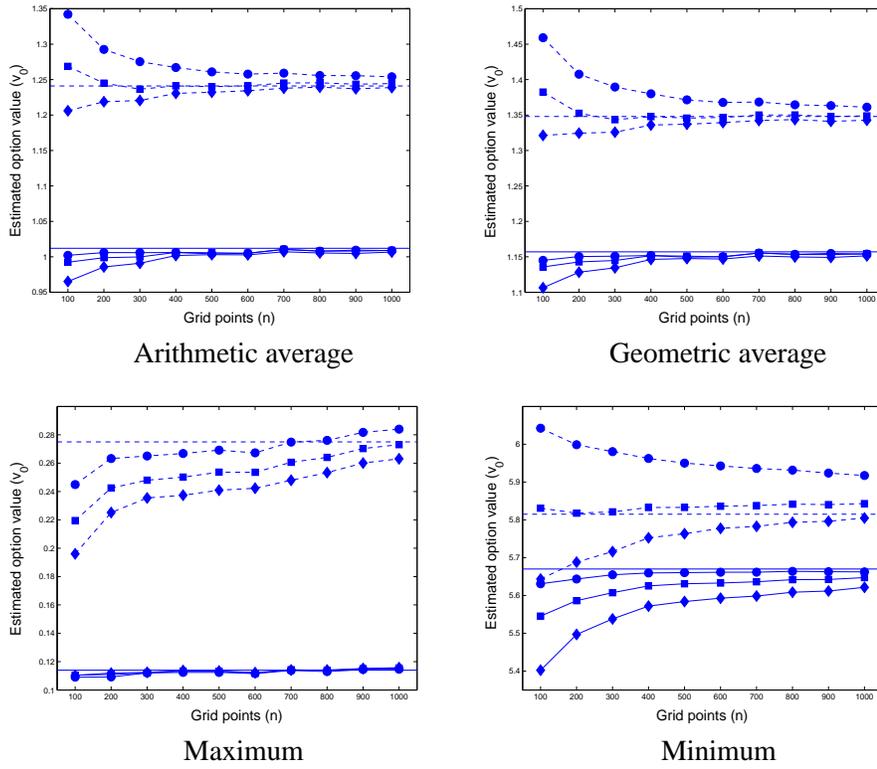


Figure 5.1: Average QMC grid valuation over 50 normal grids with $\alpha = 1.0$ (circles), $\alpha = 1.5$ (squares), $\alpha = 2.0$ (diamonds) of European (solid lines) and American (dotted lines) put options over five assets using the explicit method ($\theta = 0.0$) and ten time steps. Stentoft's Bermudan LSM solutions are drawn as horizontal lines.

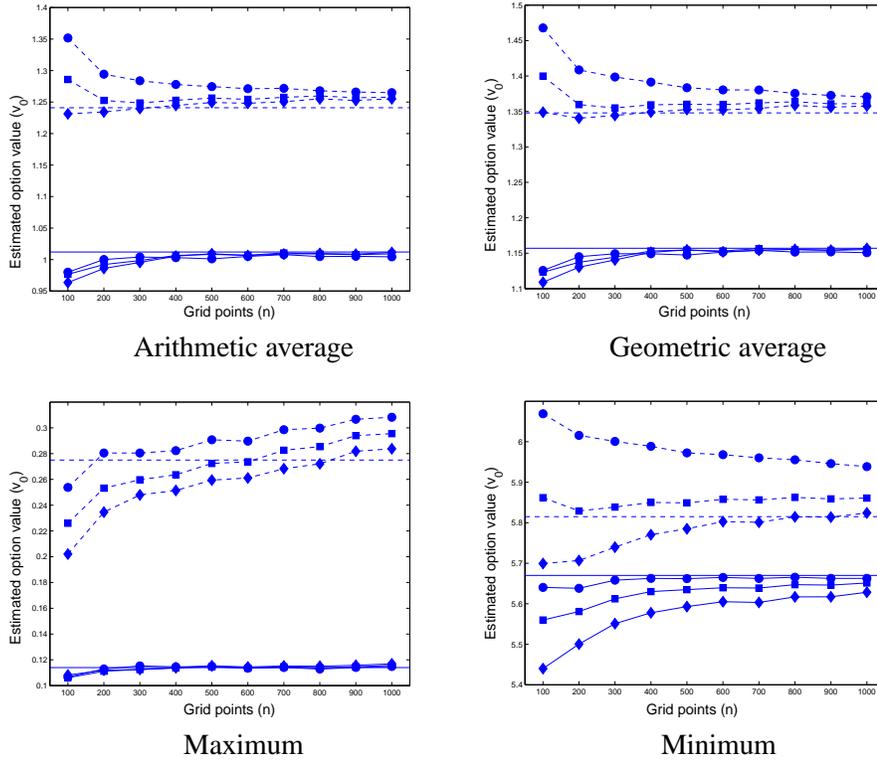


Figure 5.2: Average QMC grid valuation over 50 normal grids with $\alpha = 1.0$ (circles), $\alpha = 1.5$ (squares), $\alpha = 2.0$ (diamonds) of European (solid lines) and American (dotted lines) put options over five assets using the Crank-Nicolson method ($\theta = 0.5$) and ten time steps. Stentoft's Bermudan LSM solutions are drawn as horizontal lines.

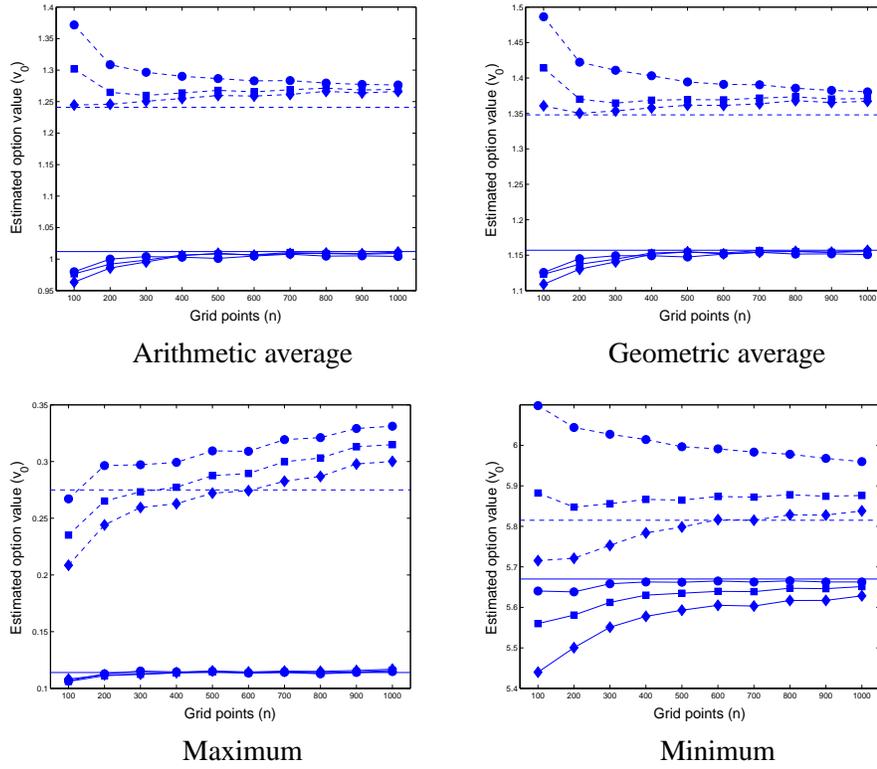


Figure 5.3: Average QMC grid valuation over 50 normal grids with $\alpha = 1.0$ (circles), $\alpha = 1.5$ (squares), $\alpha = 2.0$ (diamonds) of European (solid lines) and American (dotted lines) put options over five assets using the implicit method ($\theta = 1.0$) and ten time steps. Stentoft's Bermudan LSM solutions are drawn as horizontal lines.

Option type	Exact	Binomial	LSM	LSM OS (low-biased)	Normal grid	Normal grid OS (low-biased)
Arith. Average	-	1.235	1.241 (0.0006)	1.231 (0.0006)	1.246 (0.004)	1.238 (0.005)
Geom. Average	1.342	1.340	1.348 (0.0006)	1.335 (0.0007)	1.350 (0.004)	1.345 (0.005)
Maximum	-	0.230	0.275 (0.0004)	0.268 (0.0004)	0.276 (0.008)	0.233 (0.002)
Minimum	-	5.841	5.815 (0.0012)	5.816 (0.0014)	5.847 (0.009)	5.821 (0.013)

Table 5.1: Comparison of Bermudan price estimates ($\theta = 0$) with ten exercise opportunities. The grid estimates are the average price taken over 50 normal RQMC grids with size 1000, with $\alpha = 1.5$ and using ten time steps. The binomial method of Boyle et al. [3] was used with Richardson extrapolation. The OS (out-of-sample) columns give the value of the exercise strategy implied by the 50 grid solutions, calculated by taking the mean value over 100,000 simulated paths. The binomial and LSM prices are given in [25] and the OS prices for LSM are computed by running the LSM method 20 times each with 100,000 out-of-sample paths. The exact price given in the first column is the numerical solution to the equivalent one-dimensional problem. Standard errors are shown in brackets.

Bermudan problem, and the Crank-Nicolson solutions to be approximations for the American problem.

Tables 5.1 and 5.2 also show out-of-sample results for LSM and the irregular grid methods. These are estimates of the expected value, under the risk-neutral measure, of using the implied exercise strategy. We implement the LSM method ourselves, as specified in Stentoft [25], to obtain out-of-sample values for the LSM algorithm (these results are not given in [25]). Our LSM implementation also reproduced (up to a statistically insignificant difference) the in-sample LSM results given in [25]. For details of how out-of-sample paths are used in the LSM method to obtain low-biased estimators, we refer the reader to Longstaff and Schwartz [21].

We remark that the values obtained from the irregular grid method are higher than those produced by the LSM algorithm, although this is not statistically significant except in the case of the minimum option. The OS results are also higher for all but the maximum option.

Option type	Exact American	Normal grid American	Normal grid American OS (low-biased)	Hedged American OS (high-biased)
Arith.	-	1.257	1.243	1.363
Average		(0.004)	(0.004)	(0.004)
Geom.	1.355	1.360	1.348	1.462
Average		(0.004)	(0.005)	(0.004)
Maximum	-	0.295	0.267	0.504
		(0.009)	(0.002)	(0.006)
Minimum	-	5.862	5.789	6.355
		(0.009)	(0.012)	(0.010)

Table 5.2: Comparison of American price estimates ($\theta = 0.5$). The grid estimates in the third and fourth columns are the average price taken over 50 normal RQMC grids with size 1000, with $\alpha = 1.5$ and using ten time steps. The OS (out-of-sample) column gives the estimated value of the implied exercise strategy, calculated by taking the mean value over 100,000 simulated paths and using 50 time steps. The hedged column gives the average cost of the hedging strategy obtained as a by-product of a single price computation; it is implemented using the results of a single grid solution, 50 time steps and computes the hedge as detailed in Section 4.6. In particular note that we have used a different time step in the OS exercise and hedging simulations than in computing the grid solutions. The exact price given in the first column is the numerical solution to the equivalent one-dimensional problem. Standard errors are shown in brackets.

For the more problematic cases of the maximum and minimum options, we see that convergence is slower. In the case of the maximum it is not clear with 1000 grid points what an appropriate estimate should be. It is also not clear whether the convergence in our case for the explicit method agrees with the value obtained by Stentoft. These are cases where the grid could be adapted to the payoff function as well as to the process itself; such extensions are left for future investigation.

In Table 5.1 it is encouraging to see that the irregular grid prediction for the geometric average option is very accurate as compared to the benchmark. The exercise strategy performs well for the arithmetic and geometric average options, but not for the more problematic maximum and minimum payoffs.

As detailed in Section 4.6, our method yields a hedging strategy as a by-product; thus simulation of a hedging strategy can be done quickly and efficiently. Using the implied hedging strategy of a single grid, and taking 20 nearest neighbours for the delta estimation, we obtain the results shown in the last column of Table 5.2. It is clear that the hedging errors are much larger than the exercise errors; this may be expected given that the exercise rule is a function having only two possible values, whereas the hedging rule takes values in \mathbb{R}^d . The hedging strategy used is naive in that the results of only a single grid solution are used. It could probably be improved by using information from different grid solutions.

The most time-consuming operation in the irregular grid method is the computation of the matrix root. Some timings for computing matrix roots in Matlab 6.1 on a PIII 866MHz machine are presented in Table 5.3. It should be noted that the time does not depend strongly on the order of the root, so that square root and tenth root operations for example take about the same amount of time. The time taken for the construction of the matrix $P^{(T-t)}$ is seen to be small compared to the root operation.

Although the matrix root operation is time-consuming for large values of n , it should be noted that once a root has been computed for a single normally distributed grid, it can subsequently be used for valuing options on a large class of diffusion processes with arbitrary payoff functions without the need for recomputation.

6 Conclusions

We have proposed a new method for finding the value of American and Bermudan options in a high-dimensional setting. Central to this method is the use of an irregular grid over the state space and an approximation of the partial differential operator on this grid.

In our analysis we allow any grid which is generated using MC or QMC trials

Size P (n)	Memory full (MB)	Memory sparse (MB)	Time for P (sec)	Time for $P^{1/10}$ (sec)	Prop $P^{1/10}$ nonzero	Time- stepping (sec)
500	2.0	0.6	1	22	0.190	0.5
1000	8.0	1.8	5	200	0.147	1.3
1500	16.0	3.3	12	750	0.123	2.0
2000	32.0	5.1	22	2000	0.106	2.9
2500	50.0	7.1	37	4000	0.094	3.8
3000	72.0	9.1	55	7200	0.084	4.9

Table 5.3: Timings and storage requirements for the irregular grid method using Matlab 6.1 with a PIII 866MHz processor with 512 MB RAM, matrix entries stored in double precision (8 bytes per entry). The sparse matrices are formed by eliminating all entries smaller than 5×10^{-4} and renormalising. The time stepping column gives the total time to complete 10 time steps, using the sparse matrix and the explicit method. Note that sparse matrices were not used for any experiments in this paper, the information provided rather serves to illustrate the complexity of the method as n increases. The second-to-last column gives the proportion of nonzero entries in the sparse matrices, an important consideration for computational complexity. Note that MB denotes 10^6 bytes in this context.

with respect to a known density function. Once the Markov chain approximation has been obtained, we use the transition probability matrix to form a semidiscrete approximation to the partial differential operator corresponding to this Markov chain. This is done through taking a logarithm of the transition probability matrix; however solving the fully discrete problem only requires computing a certain root of the matrix related to the time step and implicitness parameters, at the cost of an extra approximation error.

An important aspect of the proposed method is the absence of any requirement to specify basis functions for approximating the value function or exercise strategy. Indeed the only specification needed is a grid density, although asymptotically even this choice is not critical. Furthermore, convergence in the Bermudan case should require asymptotics in only one parameter, namely the number of grid points. In the American case one also requires $\delta t \rightarrow 0$. These aspects set the root method aside from the LSM method where the specification of basis functions plays a critical role in the success of the method, and convergence involves asymptotics in two parameters in the Bermudan case, namely the number of basis functions and the number of paths, these two parameters producing opposite biases.

The irregular grid solution gives price estimates at all points in the grid. This is useful if one requires partial derivative information, for example when hedging. Partial derivatives can be easily estimated from the solution by performing a linear regression using values from neighbouring points.

Our experiments suggest that the irregular grid method has very good convergence properties, especially when the grid density is related to the density of the process itself. In particular, the grid density should have a larger variance than the process; for a geometric Brownian motion process in five dimensions it was found that a ratio of 1.5 gave a good rate of convergence, although (slower) convergence was also observed for ratios of 1.0 and 2.0. Convergence of estimates for the maximum option was not clear with grids of up to 1000 points.

The numerical results obtained largely agree with those of Stentoft [25]. We find that the early exercise premium is increased by about 6% for the examples he considers when allowing a continuum of exercise opportunities rather than only ten. We also find that the exercise strategies implied by the LSM method produce significantly lower values (statistically) than the LSM price implies, except in the case of the minimum option; this is an indication that out-of-sample paths should be used in simulation methods — in this way the price obtained corresponds directly to the average value of the implied exercise strategy. This suggests that one should be careful in higher dimensions when applying the recommendation of Longstaff and Schwartz [21] to save time by only using in-sample paths.

A possible variance reduction technique is to adjust the transition probabilities according to the empirical density of the grid points rather than the density used for

generation of the grid. Adjustment may be done after constructing the transition matrix for example using quadratic programming to improve local consistency in the sense of Kushner and Dupuis [20], but may also take inspiration from the literature on nonparametric analysis. These and other possible refinements are reserved for future investigation.

References

- [1] Jérôme Barraquand and Didier Martineau. Numerical valuation of high dimensional multivariate American securities. *Journal of Financial and Quantitative Analysis*, 30:383–405, 1995.
- [2] Phelim Boyle, Mark Broadie, and Paul Glasserman. Monte Carlo methods for security pricing. *Journal of Economic Dynamics and Control*, 21:1267–1321, 1997.
- [3] Phelim P. Boyle, Jeremy Evnine, and Stephen Gibbs. Numerical evaluation of multivariate contingent claims. *Review of Financial Studies*, 2:241–250, 1989.
- [4] Phelim P. Boyle, Adam W. Kolkiewicz, and Ken Seng Tan. An improved simulation method for pricing high-dimensional American derivatives. *Mathematics and Computers in Simulation*, 62:315–322, 2003.
- [5] Mark Broadie and Paul Glasserman. Pricing American-style securities using simulation. *Journal of Economic Dynamics and Control*, 21:1323–1352, 1997.
- [6] Mark Broadie and Paul Glasserman. A stochastic mesh method for pricing high-dimensional American options. *Working Paper, Paine Webber Series, Columbia Business School*, 1997.
- [7] E. Clément, D. Lamberton, and P. Protter. An analysis of a least squares regression method for American option pricing. *Finance and Stochastics*, 6:449–471, 2002.
- [8] Colin W. Cryer. The solution of a quadratic programming problem using systematic overrelaxation. *SIAM Journal of Control*, 9:385–392, 1971.
- [9] M.A.H. Dempster and J.P. Hutton. Pricing American stock options by linear programming. *Mathematical Finance*, 9:229–254, 1999.

- [10] M. Evans and T. Swartz. *Approximating Integrals via Monte Carlo and Deterministic Methods*. Oxford University Press, Oxford, 2000.
- [11] Michael C. Fu, Scott B. Laprise, Dilip B. Madan, Yi Su, and Rongwen Wu. Pricing American options: a comparison of Monte Carlo simulation approaches. *Journal of Computational Finance*, 4:39–88, 2001.
- [12] Paul Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer, New York, 2003.
- [13] Paul Glasserman and Bin Yu. Number of paths versus number of basis functions in American option pricing. *Working Paper, Columbia University*, 2003.
- [14] R. Glowinski, J.L. Lions, and R. Trémolières. *Numerical Analysis of Variational Inequalities*. North-Holland, Amsterdam, 1981.
- [15] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, 1996.
- [16] Martin B. Haugh and Leonid Kogan. Pricing American options: a duality approach. *Working Paper, The Wharton School, University of Pennsylvania*, 2001.
- [17] Patrick Jaillet, Damien Lamberton, and Bernard Lapeyre. Variational inequalities and the pricing of American options. *Acta Applicandae Mathematicae*, 21:263–289, 1990.
- [18] Farshid Jamshidian. Minimax optimality of Bermudan and American claims and their Monte-Carlo upper bound approximation. *Working Paper, NIB Capital Bank*, 2003.
- [19] Anastasia Kolodko and John Schoenmakers. An efficient dual Monte Carlo upper bound for Bermudan style derivatives. *Preprint 877, Weierstrass Institute*, 2003.
- [20] Harold J. Kushner and Paul G. Dupuis. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag, New York, 1992.
- [21] Francis A. Longstaff and Eduardo S. Schwartz. Valuing American options by simulation: a simple least squares approach. *Review of Financial Studies*, 14:113–147, 2001.
- [22] Manuel Moreno and Javier F. Navas. On the robustness of least-squares Monte Carlo (LSM) for pricing American derivatives. *Review of Derivatives Research*, 6:107–128, 2003.

- [23] Art B. Owen. Randomly permuted (t, m, s) -nets and (t, s) -sequences. In *Monte Carlo and Quasi-Monte Carlo Methods in Scientific Computing* (H. Niederreiter and P. J.-S. Shiue, eds.), pages 299–317, 1995.
- [24] L.C.G. Rogers. Monte Carlo valuation of American options. *Mathematical Finance*, 12:271–286, 2002.
- [25] Lars Stentoft. Assessing the least squares Monte-Carlo approach to American option valuation. *CAF Working Paper No. 90, Aarhus School of Business*, 2001.
- [26] Lars Stentoft. Convergence of the least squares Monte-Carlo approach to American option valuation. *CAF Working Paper No. 113, Aarhus School of Business*, 2002.
- [27] John N. Tsitsiklis and Benjamin Van Roy. Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks*, 12:694–703, 2000.