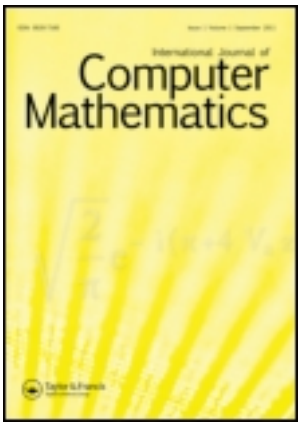

International Journal of Computer Mathematics

Publication details, including instructions for authors and subscription information:
http://tandfonline.com/loi/gcom20

Pricing high-dimensional Bermudan options using the stochastic grid method

Shashi Jain [a b] & Cornelis W. Oosterlee [a c]

[a] TU Delft, Delft Institute of Applied Mathematics, Delft, The Netherlands

[b] Nuclear Research Group, Petten, The Netherlands

[c] CWI-Centrum Wiskunde & Informatica, Amsterdam, The Netherlands

Version of record first published: 24 May 2012

To cite this article: Shashi Jain & Cornelis W. Oosterlee (2012): Pricing high-dimensional Bermudan options using the stochastic grid method, International Journal of Computer Mathematics, 89:9, 1186-1211

To link to this article: http://dx.doi.org/10.1080/00207160.2012.690035

International Journal of Computer Mathematics
Vol. 89, No. 9, June 2012, 1186–1211

Pricing high-dimensional Bermudan options using the stochastic grid method

Shashi Jain[a,b]* and Cornelis W. Oosterlee[a,c]

[a]*TU Delft, Delft Institute of Applied Mathematics, Delft, The Netherlands;* [b]*Nuclear Research Group, Petten, The Netherlands;* [c]*CWI-Centrum Wiskunde & Informatica, Amsterdam, The Netherlands*

(*Received 13 October 2011; revised version received 21 February 2012; accepted 23 April 2012*)

This paper considers the problem of pricing options with early-exercise features whose pay-off depends on several sources of uncertainty. We propose a *stochastic grid method* for estimating the optimal exercise policy and use this policy to obtain a low-biased estimator for high-dimensional Bermudan options. The method has elements of the *least-squares method* (LSM) of Longstaff and Schwartz [*Valuing American options by simulation: A simple least-squares approach*, Rev. Finan. Stud. 3 (2001), pp. 113–147], the *stochastic mesh method* of Broadie and Glasserman [*A stochastic mesh method for pricing high-dimensional American option*, J. Comput. Finance 7 (2004), pp. 35–72], and *stratified state aggregation along the pay-off* method of Barraquand and Martineau [*Numerical valuation of high-dimensional multivariate American securities*, J. Financ. Quant. Anal. 30 (1995), pp. 383–405], with certain distinct advantages over the existing methods. We focus on the numerical results for high-dimensional problems such as max option and arithmetic basket option on several assets, with basic error analysis for a general one-dimensional problem.

Keywords: American options; high dimensional; Monte Carlo; Gram Charlier; stochastic grid method; regression; stochastic mesh method; least squares method (LSM); Bermudan options

2000 AMS Subject Classifications: 65C05; 65C30; 62P05; 91B28; 60G40

1. Introduction

Pricing of Bermudan options[1] especially for multi-dimensional processes is a challenging problem owing to its path-dependent settings. The traditional valuation methods, such as lattice and tree-based techniques are often impractical in such cases due to the curse of dimensionality and hence are used only in the low-dimensional cases. In the recent years, many simulation-based algorithms have been proposed for pricing Bermudan options, most of which use a combination of Monte Carlo simulations and dynamic programming to estimate the option price.

Monte Carlo simulations for pricing options became popular after the pioneering works of Boyle [8], Bossaerts [7] and Tilley [30]. Regression-based approaches for pricing Bermudan options have been proposed by Carriere [16], Tsitsiklis and Van Roy [31] and Longstaff and Schwartz [25]. The Longstaff and Schwartz least-squares method (LSM) computes the option

*Corresponding author. Email: s.jain@cwi.nl, jain.shashi@gmail.com

ISSN 0020-7160 print/ISSN 1029-0265 online

http://dx.doi.org/10.1080/00207160.2012.690035
http://www.tandfonline.com

price by first determining the optimal exercise policy for a set of simulated paths and then finds the expected value of the discounted pay-off obtained by following this exercise policy. The option price obtained is the lower bound on the true option price, as the exercise policy obtained would either be inferior or equal to the optimal exercise policy. Eglof [18] and Zanger [32] analyse the convergence of the LSM. Belomestny *et al.* [4] compare local regression estimators which are popular for computing Greeks with global regression estimators, which is a generalization of the methods of Tsitsiklis and Van Roy [31] and Longstaff and Schwartz [25]. They also present an algorithm where instead of regressing continuation functions, the control and stopping times are backwardly constructed on a set of simulated trajectories.

Ibanez and Zapatero [21] compute at each exercise opportunity the fixed points of the optimal exercise frontier and obtain the parametric form of this frontier by regressing on quadratic or cubic function. They use the frontier obtained with plain vanilla Monte Carlo simulation to obtain a low-biased estimator of the true price.

Duality-based approaches for Bermudan option pricing are proposed by Haugh and Kogan [20] and Rogers [29] which can be used to construct an upper bound on the option value. Andersen and Broadie [1] improved the practical implementation of duality-based methods by proposing a simulation algorithm for obtaining the upper bounds from any given exercise policy. The duality-based algorithms work by first computing the lower bounds using some exercise policy (a sub-optimal policy) and then adding a non-negative quantity that penalizes potentially incorrect exercise decisions made by the sub-optimal policy.

The stochastic mesh method (SMM) of Broadie and Glasserman [14] approximates the option values using a dynamic programming-style backward recursion for approximating the price and optimal exercise policy. The continuation value at each mesh point is computed as the weighted sum of option values attained due to all possible transitions to mesh points in the next time step. In the original mesh method, the weights were computed from the transition density of the underlying process. In an improvement to the original stochastic mesh method, Broadie *et al.* [15] avoid the use of the transition density of the underlying process of asset prices and other state variables by choosing mesh weights through optimization of a convex objective function subject to known conditional expectations.

In an important attempt to circumvent the curse of dimensionality problem associated with pricing of multi-dimensional Bermudan options, Barraquand and Martineau [3] introduce the *state aggregation* technique, in which they partition the space of underlying assets (state space) into a tractable number of cells, and compute an approximate early-exercise strategy that is constant over those cells. They limit their search to strategies that depend upon a stratification map (a real-valued function mapping the state) rather than upon the entire state itself. Particularly in the case of Bermudan options, they use the pay-off as stratification map, and call this technique as *stratified state aggregation along the pay-off* (SSAP). Boyle *et al.* [10] draw attention to some drawbacks of using SSAP.

Berridge and Schumacher [6] introduced a hybrid method to price high-dimensional American options by first performing a discretization of the state space using quasi-Monte Carlo (QMC) points and then finding the approximation to the partial differential operator on this grid which is used to formulate linear complementarity problems at successive time points, working backwards from the option expiry.

The *stochastic grid method* (SGM) follows the dynamic programming style of SMM, by recursively computing the option price, moving backwards in time. The functional approximation, obtained using regression, of the option price at a given time step is used to compute the option price at the previous time step. The dimensionality of the problem is recursively reduced using the pay-off as a mapping function. Although numerical results are given for the high-dimensional problems, we show that error for SGM is bounded only for a one-dimensional problem.

The SGM has certain advantages over the existing methods. The LSM, although computationally fast and simple to implement, uses a large number of paths to obtain a good exercise policy. Also the number of basis functions required for regression grows almost exponentially with the dimensions of the problem. SGM on the other hand can be used to obtain a good exercise policy using far fewer paths. The number of basis functions used in the SGM is independent of the dimensions of the problem. SGM uses sub-simulation when moments required to approximate the transition density function are unavailable, which can make the method computationally expensive. SGM does not suffer from the limitations, pointed out by Boyle *et al.* [10], of the SSAP method of Barraquand and Martineau, making it an efficient algorithm for handling options with a large number of underlying assets.

The paper is organized as follows, Section 2 is devoted to the description of the SGM. In Section 3, we present a basic error analysis for a one-dimensional problem and discuss some of the results for the single asset case. In Section 4, we discuss and compare the results for high-dimensional problems with the other available models. In Section 5 we conclude, make observations about some existing open problems and directions in which the future research efforts can be made.

2. The method of stochastic grid

The SGM solves a general optimal stopping problem using a hybrid of dynamic programming and Monte Carlo methods. The method first computes the optimal exercise policy and a direct estimator of the true option price. The lower bound values are computed by discounting the pay-off obtained by following this exercise policy. We describe in detail how these bounds are obtained in the sections to follow.

2.1 *Problem formulation*

We assume complete probability space $(\Omega, \mathcal{F}, \mathcal{P})$ and finite time horizon $[0, T]$. Ω is the set of all possible realizations of the stochastic economy between 0 and T. $\mathcal{F}_T$ is the sigma field of distinguishable events at time T, and $\mathcal{P}$ is the risk-neutral probability measure on elements of $\mathcal{F}$. The information structure in this economy is represented by an augmented filtration $\mathcal{F}_t : t \in [0, T]$. We assume that $\mathcal{F}_t$ is generated by W_t, a d-dimensional standard Brownian motion, and the state of economy is represented by an $\mathcal{F}_t$-adapted Markovian process $S_t = (S_t^1, \ldots, S_t^d) \in \mathcal{R}^d$, where $t \in [t_0 = 0, \ldots, t_i, \ldots, t_k = T]$. Let $h_t = h(S_t)$ be a non-negative adapted process representing the pay-off of the option, i.e. the holder of the option receives h_t if the option is exercised at time t. Let the risk-less savings account process be $B_t = \exp(\int_0^t r_s \, \mathrm{d}s)$, where r_t denotes the instantaneous risk-free rate of return. We consider the special case where r_t is constant. The problem is then to compute

$$V_0 = \max_{\tau} \mathbb{E}\left[\frac{h(S_\tau)}{B_\tau}\right], \tag{1}$$

where τ is a stopping time taking values in the finite set $\{0, t_1, \ldots, t_k = T\}$. The value of the option at the terminal time T is equal to the products pay-off

$$V(T, x) = h(x). \tag{2}$$

The conditional continuation value $Q(t_i, S_{t_i} = x)$, i.e. the expected future pay-off at time t_i and state $S_{t_i} = x$ is given by

$$Q(t_i, S_{t_i} = x) = \frac{B_{t_i}}{B_{t_{i+1}}} \mathbb{E}[V(t_{i+1}, S_{t_{i+1}}) | S_{t_i} = x]. \tag{3}$$

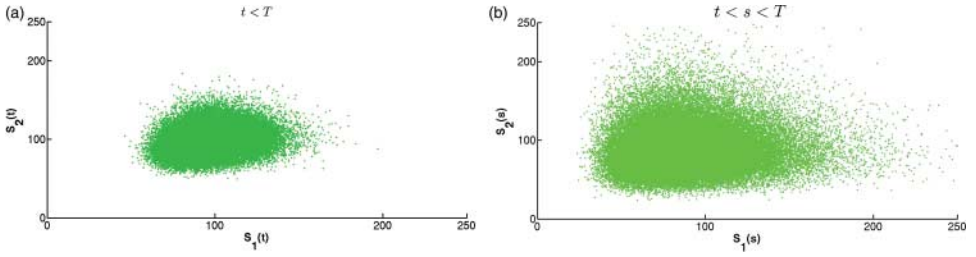

Figure 1. Grid points (30,000 × 30,000), figure (a) at t, figure (b) at s where $t < s < T$.

The Bermudan option value at time t_i and state $S_{t_i} = x$ is given by

$$V(t_i, S_{t_i}) = \max(h(S_{t_i}), Q(t_i, S_{t_i})). \tag{4}$$

We are interested in finding the value of the option at the initial state S_0, i.e. $V(0, S_0)$.

2.2 *Method details of the SGM*

We use a (Markovian) discretization scheme which is easy to simulate, e.g. the Euler scheme, to generate N sample paths originating from the initial state S_0. When the diffusion process appears in a closed form, such as the case of the commonly used multi-dimensional Black and Scholes model, we can generate the sample paths directly. The stochastic grid points (t_i, S_{t_i}) can be interpreted as the intersections of the sample paths with a plane representing different intermediate time steps t_i. Figure 1 shows the grid points for an option with two underlying assets $S_{t_i} = (S_1, S_2)$ starting from the initial state $S_{t_0} = (100, 100)$ at two different time intervals t and s, where t is close to the initial time and s is closer to the final exercise time T. The number of grid points in the vicinity of the initial state $S_{t_0} = (100, 100)$, the point for which we are interested to find the option value, increases as we approach t_0, providing a natural refinement around the point of interest. This method of grid generation is closely related to the *binomial tree* approach, where only grid points associated with the initial state are generated.

This is the most basic method for generating grids to be used in SGM. It is possible to use a more advanced spatial discretization method like the *quantization tree method* of Bally *et al.* [2], where rather than settling the grids *a priori*, at each time step a grid Γ_k^* of size N_k is generated, which optimally fits to a large simulated sample of S_{t_k} among all grids with size N_k such that the closest neighbour rule projection of S_{t_k} onto the grid Γ_k^* is the best least-squares approximation of S_{t_k}.

The value of the option at the expiration time $t_k = T$ will be equal to its pay-off given by $h(S_T)$. We restrict our attention to financial derivatives with pay-off that are element of the space of square integrable or finite variance functions. Examples of pay-off functions on multiple assets include, for a basket call option, $h(S_t) = (a_1 S_t^1 + \cdots + a_n S_t^n - K)^+$, for an out-performance option $h(S_t) = (\max(a_1 S_t^1, \ldots, a_n S_t^n) - K)^+$, where the notation x^+ is short for $\max(x, 0)$.

2.3 *Computing the optimal exercise policy*

The main obstacle in pricing Bermudan options using Monte Carlo methods is the fact that we do not know the optimal exercise policy. SGM computes the continuation value at each grid point, starting from the grid points at the expiration time $t_k = T$ and moving backwards in time.

The option is exercised if the immediate pay-off is greater than the discounted continuation value.

The grid estimator is defined recursively starting with $\hat{V}(T, S_T) = h(S_T)$, and for $i = k - 1, \ldots, 1$, by

$$\hat{V}(t_i, S_{t_i} = x) = \max\left(h(S_{t_i} = x), \frac{B_{t_i}}{B_{t_{i+1}}}\mathbb{E}[\hat{Z}(t_{i+1}, g(S_{t_{i+1}}|S_{t_i} = x))|S_{t_i} = x]\right), \tag{5}$$

where

$$\hat{Z}(t_{i+1}, g(S_{t_{i+1}}|S_{t_i} = x)) = \mathbb{E}[\hat{V}(t_{i+1}, S_{t_{i+1}})|g(S_{t_{i+1}}|S_{t_i} = x)]. \tag{6}$$

Mapping function $g(\cdot)$ maps the high-dimensional $S_{t_{i+1}}$-space to a low-dimensional $g(S_{t_{i+1}})$-space. We use $g(S_{t_{i+1}}|S_{t_i} = x)$ to denote that mapping $g(\cdot)$ is applied to all grid points $S_{t_{i+1}}$ which are generated from source $S_{t_i} = x$. $\mathbb{E}[\hat{Z}(t_{i+1}, g(S_{t_{i+1}}|S_{t_i}))|S_{t_i}]$ represents the continuation value for the grid point S_{t_i}. Using iterated conditioning we can show,

$$\begin{aligned}\mathbb{E}[\hat{V}(t_{i+1}, S_{t_{i+1}})|S_{t_i} = x] &= \mathbb{E}[\mathbb{E}[\hat{V}(t_{i+1}, S_{t_{i+1}})|g(S_{t_{i+1}}|S_{t_i} = x)]|S_{t_i} = x] \\ &= \mathbb{E}[\hat{Z}(t_{i+1}, g(S_{t_{i+1}}|S_{t_i} = x))|S_{t_i} = x].\end{aligned} \tag{7}$$

In the sections to follow we discuss how to approximate $\hat{Z}(t_{i+1}, g(S_{t_{i+1}}|S_{t_i}))$ and the choice of the mapping function $g(\cdot)$. Once we have the functional approximation, $\hat{Z}(t_{i+1}, g(S_{t_{i+1}}|S_{t_i}))$, we can use it to compute the discounted continuation value at the grid points for t_i and thus make the optimal exercise decision, i.e. exercise if the discounted continuation value is less than the immediate pay-off.

2.4 *Parametrization of the option values*

The continuation value at time t_i and state $S_{t_i} = x$, i.e. $Q(t_i, S_{t_i} = x)$ can be computed from Equation (3). Instead of using the direct functional approximation of the option price at t_{i+1}, i.e. $\hat{V}(t_{i+1}, S_{t_{i+1}})$, we use the law of iterated conditioning, i.e. $\mathbb{E}[\mathbb{E}[X|\mathcal{G}]|\mathcal{H}] = \mathbb{E}[X|\mathcal{H}]$, where $\mathcal{H}$ is the sub- σ algebra of $\mathcal{G}$, to compute the continuation value. Then the continuation value can be written as (7).

In order to compute $Q(t_i, S_{t_i} = x)$, from Equation (7) we need to know the functional form of $\hat{Z}(t_{i+1}, g(S_{t_{i+1}}|S_{t_i}))$. At the expiration time, the option value is given by Equation (2).

In the examples to follow, the form of solution is simplified if we write the pay-off function in the following form:

$$h(S_t) = \max(g(S_t) + X, 0), \tag{8}$$

with $g : [0, T] \times \mathcal{R}^d \to \mathcal{R}$ explained before. In the case of a simple call on a single asset with strike K, $g(S_t) = S_t$ and $X = -K$, for a put on the maximum of d assets and strike K, $g(S_t^1, \ldots, S_t^d) = -\max(S_t^1, \ldots, S_t^d)$ and $X = K$. It should be noted however that this form of writing the pay-off function is not restrictive for SGM but is used as it simplifies the form of the solution.

We assume that the unknown functional form of $\hat{Z}(t_{i+1}, g(S_{t_{i+1}}|S_{t_i} = x))$ can be represented by a linear combination of a countable set of $\mathcal{F}_{t_{i+1}}$-measurable basis functions, where $\mathcal{F}_{t_{i+1}}$ is the information set at time t_{i+1}.

Similar to the regression-based algorithms [25,31] SGM approximates the unknown functional form of $\mathbb{E}[\hat{V}(t_{i+1}, S_{t_{i+1}})|g(S_{t_{i+1}}|S_{t_i} = x)]$ by projecting it on the first $M(< \infty)$ polynomial basis functions.

Remark 1 In the examples we approximate the function $\hat{Z}(t_{i+1}, g(S_{t_{i+1}}|S_{t_i}))$ by $\hat{Z}(t_{i+1}, g(S_{t_{i+1}}|S_{t_0}))$, as all the grid points at t_{i+1} generated from source S_{t_0} are used in the regression. The exercise policy obtained is still accurate as shown by the numerical results (lower bound values). To simplify the notations, we will be referring to $g(S_{t_{i+1}}|S_{t_0})$ by $g(S_{t_{i+1}})$ from here on. An improved approximation will be based on a more sophisticated regression scheme, where grid points at t_i are bundled based on proximity, and only those grid points at t_{i+1} are used for regression to approximate $\hat{Z}(t_{i+1}, g(S_{t_{i+1}}|S_{t_i}))$ that originate from the bundle containing S_{t_i}.

When we approximate $\hat{Z}(t_{i+1}, g(S_{t_{i+1}}|S_{t_i}))$ by $\hat{Z}(t_{i+1}, g(S_{t_{i+1}}|S_{t_0}))$, an accurate early-exercise policy is obtained when $g(\cdot)$ is of the form given by Equation (8). However, also other choices of $g(\cdot)$ can be made. For other choices, it becomes important that the grid points are bundled based on some nearest neighbour rules to get an accurate exercise policy. In the special case when $g(\cdot)$ is chosen to be constant, SGM with bundling would very closely resemble the state space partitioning method by Jin *et al.* [22].

We denote this approximation by $Z_M(t_{i+1}, g(S_{t_{i+1}}|S_{t_0}))$ or $Z_M(t_{i+1}, g(S_{t_{i+1}}))$. We approximate Equation (6) over a set of M polynomial basis functions, as

$$Z_M(t_{i+1}, g(S_{t_{i+1}})) = \mathbb{E}[\hat{V}(t_{i+1}, S_{t_{i+1}})|g(S_{t_{i+1}})] = \sum_{m=0}^{M-1} a_m \Psi_m(g(S_{t_{i+1}})), \tag{9}$$

such that at each time step

$$r = \min_{a_m} \sum_{1}^{N} |Z_M(t_{i+1}, g(S_{t_{i+1}})) - V(t_{i+1}, S_{t_{i+1}})|^2, \tag{10}$$

where $\{\Psi(\cdot)\}_{m=0}^{M-1}$ form a set of basis functions and r is the sum of squared residual errors.

This approximation can be justified if we assume that $V(t_i, S_{t_i})$ is an element of the L^2 space of square integrable functions relative to some measure and therefore can be written as the linear combination of basis functions. Rather than regressing over entire $g(S_{t_{i+1}})$-space, a better accuracy is obtained by *piecewise regression*, as explained in Section 3 and the specific examples to follow.

2.4.1 *Mapping high-dimensional state to single-dimensional $g(\cdot)$-space*

In an approach similar to Barraquand and Martineau's SSAP method [3], we reduce the dimensions of the problem by using $g(S_{t_{i+1}})$ rather than the cross-products of the underlying states (as in LSM) for regression.

Figure 2 shows in a schematic diagram how dimension reduction works in SGM. In order to compute the continuation value at S_{t_i} directly, a high-dimensional transition density function would be required, as shown in Figure 2(a). In SGM, however, we first project the option value at t_{i+1} over the $g(S_{t_{i+1}})$-space, see Figure 2(b). In other words, we compute the conditional expectation, $\mathbb{E}[V(t_{i+1}, S_{t_{i+1}})|g(S_{t_{i+1}})]$, using the least-squares regression. The continuation value is then computed using the tower property as explained in Equation (7), which involves a one-dimensional transition density function. When we use all the grid points at t_{i+1} for regression, we compute $\mathbb{E}[\hat{V}(t_{i+1}, S_{t_{i+1}})|g(S_{t_{i+1}}), S_{t_0}]$, instead of $\mathbb{E}[\hat{V}(t_{i+1}, S_{t_{i+1}})|g(S_{t_{i+1}}), S_{t_i}]$. A better approximation is obtained by bundling the grid points at t_i based on proximity and using only those grid points at t_{i+1} that originate from the bundle containing S_{t_i} for regression. However, in the present paper we find that in the case that all the grid points at t_i are in a single bundle, we still obtain a very satisfactory exercise policy (as is reflected in the lower bounds), when the mapping function $g(\cdot)$- is of the form of the pay-off function. We report on these latter results for higher-dimensional problems in the numerical section.

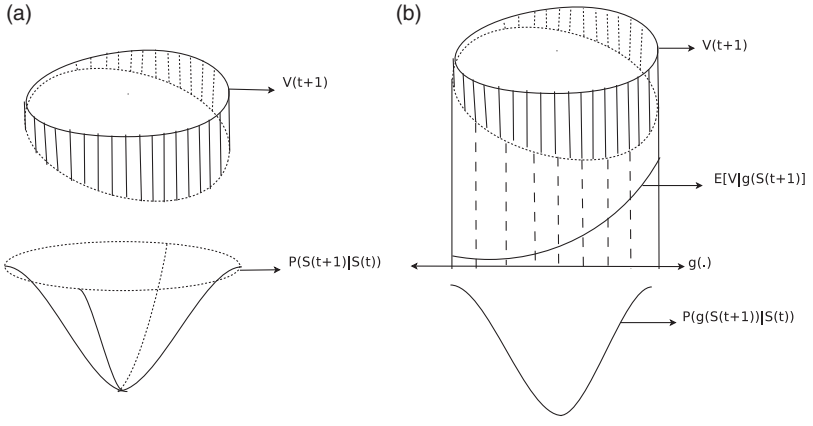

Figure 2. Schematic diagram showing how dimension reduction works in SGM. The option value at step $t+1$ is given, figure (a) shows the conventional way of computing the continuation value at $S(t)$, based on $\mathbb{P}(S_{t+1}|S_t)$; figure (b) shows how the continuation value is computed in SGM by means of projection $\mathbb{E}[V|g(S_{t+1})]$ onto $g(\cdot)$ and $\mathbb{P}(g(S_{t+1})|S_t)$.

Boyle *et al.* [10] and Broadie and Detemple [12] show that the pay-off value is not a sufficient statistic for determining the optimal exercise decision for options on the maximum of several assets for SSAP. This argument, however, is specific to the SSAP and would not apply to SGM. In the SSAP method, the state space is first mapped to the partitions (cells) along the pay-off space $h(S_t)$ and then the same exercise decision is applied for all underlying states that fall into a particular cell or partition. This results in seemingly far off state points (like (100,90), (100,100) and (100,50)) to have the same exercise decision. In SGM *first* the exercise decision is made for each underlying state S_{t_i} (or grid point) at time step t_i and *then* the state space is reduced to $g(S_{t_i})$.

In order to give a better intuition about our method and allay the concerns raised by Boyle *et al.* [10], we use the same example given by them. Figures 3–6 show the evolution of two asset prices $S_t = (S_t^1, S_t^2)$ with two exercise time steps. The option pay-off, $h(S_t = (S_t^1, S_t^2)) = g(S_t^1, S_t^2) = \max(S_t^1, S_t^2)$ and for convenience the risk-free interest rate is taken to be zero. The steps followed at each time step starting from the final expiration time t_2 are

- *Step 1*: Compute the continuation value at each state point.
- *Step 2*: Make the exercise decision, based on the greater of immediate exercise $h(S_t = x)$ or continuation value $Q(t, S_t = x)$.
- *Step 3*: Regress the option value obtained over $g(S_t^1, S_t^2) = \max(S_t^1, S_t^2)$ to be used in the previous exercise time step (as we move backwards in time) to compute the continuation value.
- *Step 4*: In the previous exercise time step, compute the transition probability from each state point to the $g(\cdot)$-space in the next time step, i.e. $\mathbb{P}(g(S_{t_{i+1}})|S_{t_i} = x)$.
- *Step 5*: Compute the continuation value $\hat{Q}(t_i, S_{t_i})$ and the option value $\hat{V}(t_i, S_{t_i})$ using Equation (5).

Focusing on the example, Figure 3 shows that at time t_2 the option values $V(t_2, S_{t_2} = (14, 2))$ and $V(t_2, S_{t_2} = (2, 14))$ are 14 and $V(t_2, S_{t_2} = (4, 2))$ is 4. On regressing these values over $\max(S_t^1, S_t^2)$ we obtain $\hat{Z}(t_2, g(S_{t_2}) = 14) = 14$ and $\hat{Z}(t_2, g(S_{t_2}) = 4) = 4$, as shown in Figure 4. Moving to exercise time step t_1 we first compute the transition probability for each state point (grid point) at t_1 to the $g(\cdot)$-space in t_2. In the present example, the state $S_{t_1} = (8, 8)$ transitions to $g(S_{t_2}) = 14$ with probability 1. Similarly, the conditional transition probability for $S_{t_1} = (8, 4)$ equals $\mathbb{P}(g(S_{t_2}) = 4|S_{t_1} = (8, 4)) = 1$. Together with these conditional transition probabilities and the approximation of the option values at t_2, we compute the continuation value for the state points

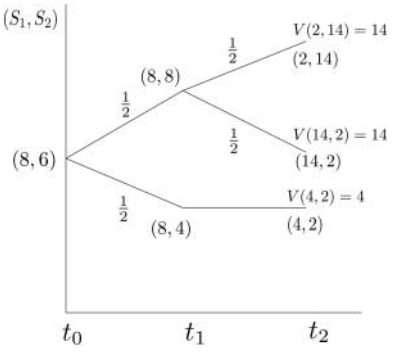

Figure 3. Step I: compute the option values at t_2 as a function of (S_1, S_2).

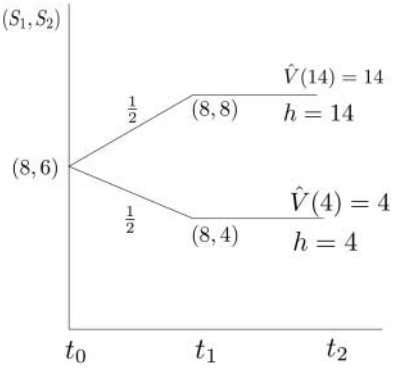

Figure 4. Step II: map the option prices to $\max(S_1, S_2)$.

at t_1. The continuation value at $S_{t_1} = (8,8)$ equals 14, computed by

$$\hat{Q}(t_1, S_{t_1}) = \sum_i \hat{Z}(t_2, g(S_{t_2}) = i) \cdot \mathbb{P}(g(S_{t_2}) = i | S_{t_1} = (8,8)).$$

The continuation value at $S_{t_1} = (8,4)$ is 4, determined as

$$\hat{Q}(t_1, S_{t_1}) = \sum_i \hat{Z}(t_2, g(S_{t_2}) = i) \cdot \mathbb{P}(g(S_{t_2}) = i | S_{t_1} = (8,4)).$$

Figure 5 shows that the option value at S_{t_1} is the *maximum* of immediate exercise and continuation, i.e. $\max(8,14)$ for $S_{t_1} = (8,8)$ and $\max(8,4)$ for $S_{t_1} = (8,4)$. Thus, it is optimal to exercise in state $S_{t_1} = (8,4)$ and to continue in the state $S_{t_1} = (8,8)$. On regressing these values over $\max(S_1, S_2)$, we obtain $\hat{V}(t_1, g(S_{t_1}) = 8)$ is 11, as shown in Figure 6. Finally, for time step t_0 state $(8,6)$ evolves to $g(S_{t_1}) = 8$ with probability 1. Therefore, the conditional continuation value is 11,

$$\sum_i \hat{Z}(t_1, g(S_{t_1}) = i) \cdot \mathbb{P}(g(S_{t_1}) = i | S_{t_0} = (8,6)),$$

and the option value $\hat{V}(t_0, (8,6)) = \max(8,11)$, which gives the correct value.

Although this example is over simplified, it gives a basic understanding of our approach. In Figure 7 we plot the shape of typical exercise regions ε^X for an Bermudan call option on

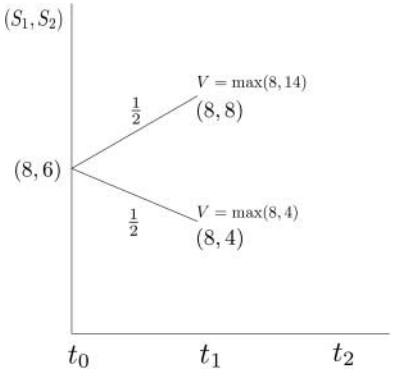

Figure 5. Step III: compute the option values at t_1 as function of (S_1, S_2).

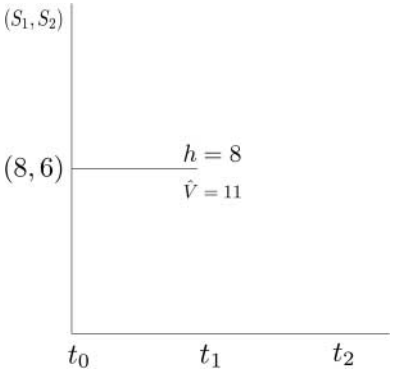

Figure 6. Step IV: map the option price to $\max(S_1, S_2)$.

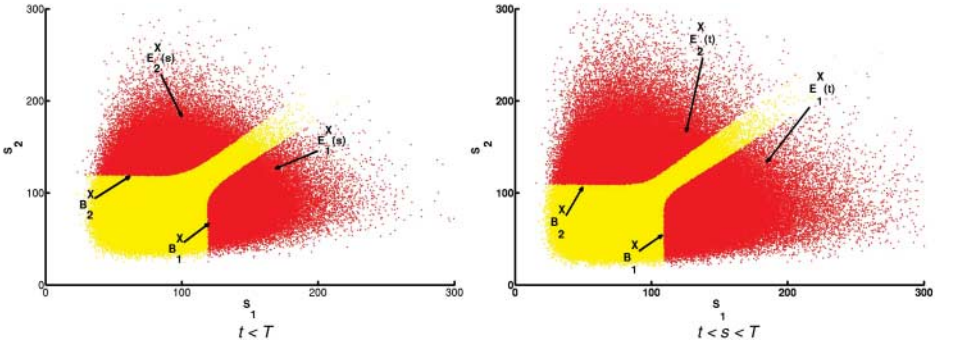

Figure 7. Exercise regions for a max-call option.

the max of two underlying assets obtained using SGM. The figures are in agreement with those deduced by Broadie and Detemple [12]. Interestingly we can see, as was found by Broadie *et al.* that prior to maturity exercise is not optimal when the prices of the underlying assets are equal.

2.5 *Computing the continuation value*

The continuation value for grid point S_{t_i} is the discounted conditional expectation of the option values in the next time step t_{i+1} given S_{t_i}. This can be written as

$$Q(t_i, S_{t_i} = x) = \frac{B_{t_i}}{B_{t_{i+1}}} \mathbb{E}[V(t_{i+1}, S_{t_{i+1}}) | S_{t_i} = x].$$

As mentioned in Section 2.4, we first approximate the conditional expectation of the option values at t_{i+1} given $g(S_{t_{i+1}})$ as a polynomial function of $g(S_{t_{i+1}})$, Equation (9). The continuation value can then be approximated using iterated conditioning as

$$\hat{Q}(t_i, S_{t_i} = x) = \frac{B_{t_i}}{B_{t_{i+1}}} \mathbb{E}[\hat{Z}(t_{i+1}, g(S_{t_{i+1}})) | S_{t_i} = x]. \tag{11}$$

Here, $\hat{Z}$ is a polynomial function of the adapted process $g(S_{t_{i+1}})$ and hence we need to determine the conditional probability density function $\mathbb{P}(g(S_{t_{i+1}}) | S_{t_i} = x)$ in order to compute its expectation. Using Equation (9), Equation (11) can be written as

$$\hat{Q}(t_i, S_{t_i} = x) = \frac{B_{t_i}}{B_{t_{i+1}}} \int_{S_{t_i} \in \mathcal{R}^d} \left(\sum_{m=0}^{M-1} a_m \Psi_m(g(S_{t_{i+1}})) \right) \mathrm{d}\mathbb{P}(g(S_{t_{i+1}}) | S_{t_i} = x). \tag{12}$$

There are three possibilities for computing the distribution of $g(S_{t_{i+1}})$ given state S_{t_i}:

(1) The exact transition probability density function $\mathbb{P}(g(S_{t_{i+1}}) | S_{t_i} = x)$ is known, for example for a call or put on a single asset in the Black–Scholes framework, a call or put on the geometric mean of d assets.
(2) The transition probability density function $\mathbb{P}(g(S_{t_{i+1}}) | S_{t_i} = x)$ is unknown; however, the moments of the distribution are known, for example for a call or put on the *Max* or *Min* of d assets in the Black– Scholes framework.
(3) The transition probability density function $\mathbb{P}(g(S_{t_{i+1}}) | S_{t_i} = x)$ and its moments are unknown.

Case 1 is the trivial case where the density function is already known. This case can also be handled efficiently by Fourier techniques, particularly when the conditional density function is not known but when the characteristic function (the Fourier transform of the conditional density) is [19]. Case 3 can be reduced to Case 2, by computing the moments with the help of Monte Carlo sub-simulations. For each grid point at time step t_i, we generate sub-paths until time t_i and compute the first four non-central moments (denoted by 'prime'), $\mu_1' = \mu, \mu_2' = \mu^2 + \sigma^2, \mu_3', \mu_4'$, of $g(S_{t_{i+1}})$. The computational effort required for such a sub-simulation is of order $O(N_G \times N_S)$, where N_G are the number of grid points and N_S are the number of sub-paths simulated. In the examples we considered, when sub-simulation was required, the computational time was a few minutes. The computational time can further be reduced by using GPUs and generating sub-paths for a group of nearest neighbour grid points, rather than for each one of them.

Once we have these moments for $g(S_{t_{i+1}})$ corresponding to the grid points at t_i, we approximate the conditional density function $f(x)$ using the *Gram Charlier Series* (see [24]). Given the moments of a distribution, the Gram Charlier series approximates the density function $f(x)$ as

$$\hat{f}(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(x-\mu)^2}{2\sigma^2} \right] \left[1 + \frac{\kappa_3}{3!\sigma^3} H_3\left(\frac{x-\mu}{\sigma} \right) + \frac{\kappa_4}{4!\sigma^4} H_4\left(\frac{x-\mu}{\sigma} \right) \right], \tag{13}$$

where $H_3(x) = x^3 - 3x$ and $H_4(x) = x^4 - 6x^2 + 3$ are Hermite polynomials. $\kappa_1 = \mu$, $\kappa_2 = \sigma^2$, $\kappa_3 = \mu_3, \kappa_4 = \mu_4 - 3\mu_2^2$ are the first four cumulants. More details about computing the probability

density function are given in the specific examples in the sections to follow. In Appendix 3, we discuss the convergence of Gram Charlier Series and also show some numerical results for its error analysis.

2.5.1 *Need for peripheral paths*

We notice that in the high-dimensional problems, the exercise policy obtained is better if we generate additional paths from points on the periphery of source point S_0. This idea is not new and was originally proposed by Rasmussen [28] as an improvement for the LSM, which he calls *initial state dispersion* where instead of using the original initial state S_0 for generating the state variables one starts with some fictitious initial time point $-T_D < 0$ and the original state for generating the state variables. More recently Kan *et al.* [23] propose a scheme to disperse the points around the initial source point without starting from a fictitious initial time point. In our examples, however, we use two additional point sources around the initial point and generate an equal number of paths from these three source points.

2.6 *Lower bound values*

The solution from the SGM can be validated by computing the lower bound on the option price, using the exercise policy obtained from it. To compute the lower bound on the option price, we simulate a number of sample paths (fresh set of paths should be used) originating from S_0 using the same discretization scheme. The continuation value at the new grid points is then obtained using

$$Q(t_i, S_{t_i} = x) = \frac{B_{t_i}}{B_{t_{i+1}}} \mathbb{E}[\hat{Z}(t_{i+1}, g(S_{t_{i+1}}))|S_{t_i} = x],$$

where the functional approximation of the conditional option values $\hat{Z}(t_{i+1}, g(S_{t_{i+1}}))$ is obtained from the SGM algorithm. For each sample path, we find the first exercise period t_i, if it exists, for which $h(S_{t_i}) \geq \hat{Q}(t_i, S_{t_i})$. The option is then exercised and its discounted pay-off is given by $h(S_{t_i})/B_{t_i}$. The lower bound on the option price is then obtained as

$$\underline{V_0} = \mathbb{E}_0 \left[\frac{h_{\tilde{\tau}}}{B_{\tilde{\tau}}} \right], \tag{14}$$

where $\tilde{\tau} = \min\{t \in [0, T] : \widehat{Q}_t \leq h_t\}$. The option value obtained by following any exercise strategy is dominated by the optimal strategy. In other words, as the option value is obtained by following a stopping rule $\tilde{\tau}$ it gives the lower bound on the true price (see [1]).

2.6.1 *Algorithm*

We briefly summarize the SGM algorithm.

- *Step I*: Generate N sample paths $\{S_{t_0}, \ldots, S_{t_k}\}$, where $[t_0 = 0, \ldots, t_k = T]$ and $S_{t_i} \in \mathcal{R}^d$, starting from $S_{t_0} = S_0$. The paths are discretized in time using some discretization scheme (e.g. *Euler's discretization scheme*). Each of the N asset prices S_{t_i} represents the grid points in t_i.
- *Step II*: Compute the option value for grid points in $t_k = T$ as

$$V(T, S_T) = h(S_T) = \max(g(S_T) + X, 0).$$

- *Step III*: Compute the approximate functional form,

$$\hat{Z}(T, g(S_T|S_0)) = \mathbb{E}[\hat{V}(T, S_T)|g(S_T)],$$

by *regressing* the option value at the grid points over polynomial basis functions of $g(S_T)$;
- *Step IV*: Perform the following steps for each exercise time t_i moving backwards in time, starting from t_{k-1} until we reach t_0 to obtain the direct SGM estimator value $V(t_0, S_{t_0} = S_0)$:
  (1) Compute the continuation value for grid points at t_i using the functional approximation of $\hat{Z}(t_{i+1}, g(S_{t_{i+1}}))$,

$$\hat{Q}(t_i, S_{t_i}) = \frac{B_{t_i}}{B_{t_{i+1}}}\mathbb{E}[\hat{Z}(t_{i+1}, g(S_{t_{i+1}}))|S_{t_i})].$$

  (2) Compute the option value for grid points at t_i as

$$\hat{V}(t_i, S_{t_i}) = \max(g(S_{t_i}) + X, \hat{Q}(t_i, S_{t_i})).$$

  (3) Compute the functional approximation for the conditional expectation, i.e.

$$\hat{Z}(t_i, g(S_{t_i})) = \mathbb{E}[\hat{V}(t_i, S_{t_i})|g(S_{t_i})]$$

  by *regressing* the option value obtained at each grid point in t_i over a set of polynomial basis function of $g(S_{t_i})$.
  (4) Go to the previous time step ($i \Rightarrow i - 1$).
- *Step V*: Using the exercise strategy obtained while computing the direct SGM estimator, for each path (from a set of new paths) determine the earliest time to exercise $\tilde{\tau} = \min\{t \in [0, T] : \widehat{Q}_t \leq h_t\}$. Obtain the lower bound option value as $\mathbb{E}_0[h_{\tilde{\tau}}/B_{\tilde{\tau}}]$.

3. Error analysis for the single asset case

We perform a basic error analysis for a single asset case. SGM has two main sources of error in the penultimate exercise opportunity, i.e. when $t_{i+1} = T$. They are,

- $\epsilon_z(t_{i+1}, g(S_{t_{i+1}}))$: error in the approximation of

$$Z(t_{i+1}, g(S_{t_{i+1}})) = \mathbb{E}[V(t_{i+1}, S_{t_{i+1}})|g(S_{t_{i+1}})],$$

- $\epsilon_f(g(S_{t_{i+1}})|S_{t_i})$: error in the approximation of the transition density function,

$$f(g(S_{t_{i+1}})|S_{t_i} = x).$$

The approximation of the continuation value at t_i, is given by,

$$\hat{Q}(t_i, S_{t_i}) = \int (Z(t_{i+1}, x) + \epsilon_z(t_{i+1}, x))(f(x|S_{t_i}) + \epsilon_f(x|S_{t_i}))\,\mathrm{d}x, \tag{15}$$

error in the estimation of the continuation value, $\epsilon_Q(t_i, S_{t_i})$, comes from error in the approximation of $Z(t_{i+1}, S_{t_{i+1}})$ and transition density function $f(g(S_{t_{i+1}})|S_{t_i})$. Error, $\epsilon_Q(t_i, S_{t_i})$, can be split

into, error due to approximation of the transition density function, $\epsilon_{Q_f}(t_i, S_{t_i})$, and error due to approximation of $Z(t_{i+1}, g(S_{t_{i+1}}))$, i.e. $\epsilon_{Q_z}(t_i, S_{t_i})$.

$$\begin{aligned}\epsilon_Q(t_i, S_{t_i}) &\approx \int \epsilon_f(t_{i+1}, x) Z(t_{i+1}, x)\, \mathrm{d}x + \int \epsilon_z(t_{i+1}, x) f(x|S_{t_i})\, \mathrm{d}x \\ &\leq \int |\epsilon_f(t_{i+1}, x)| Z(t_{i+1}, x)\, \mathrm{d}x + \int |\epsilon_z(t_{i+1})| f(x|S_{t_i})\, \mathrm{d}x \qquad (16) \\ &= \epsilon_{Q_f}(t_i, S_{t_i}) + \epsilon_{Q_z}(t_i, S_{t_i}). \qquad (17)\end{aligned}$$

We show these two errors are bounded.

3.1 *Error due to Gram Charlier approximation*

Milne [27] showed that if $f(x)$ satisfies a condition of the form

$$|\mathrm{e}^{x_1^2/4} f(x_1) - \mathrm{e}^{x_2^2/4} f(x_2)| < L|x_1 - x_2|, \qquad (18)$$

and if

$$|x\, \mathrm{e}^{x^2/4} f(x)| < L, \qquad (19)$$

with L constant, then the error of a Gram Charlier series as in (13) with n terms is bounded by

$$|f(x) - f_n(x)| = |\epsilon_f(x)| < BLn^{-1/2}\, \mathrm{e}^{-x^2/4}, \qquad (20)$$

where B is a constant independent of n. Assuming that the conditions above are satisfied, the error in the continuation value due to the Gram Charlier approximation can be bounded by

$$\epsilon_{Q_f}(t_i, S_{t_i}) < BLn^{-1/2} \int \mathrm{e}^{-x^2/4} Z(t_{i+1}, x)\, \mathrm{d}x. \qquad (21)$$

3.2 *Error due to parametrization of option price*

We approximate $Z(t_{i+1}, g(S_{t_{i+1}}))$ by piecewise interpolation. If we use a single high-degree polynomial regression, it can lead to significant errors if one of the derivatives of $Z(t_{i+1}, g(S_{t_{i+1}}))$ is discontinuous. A robust alternative is to replace the single high-degree polynomial for regression in $[x_0, x_n]$, here $x_0 = \min(g(S_{t_{i+1}}))$ and $x_n = \max(g(S_{t_{i+1}}))$ by several low-degree polynomials by appropriately dividing the regression domain $[x_0, x_n]$. An extreme case of this would be to use a linear polynomial to interpolate between adjacent data points. In such a case, the maximum error due to regression is bounded by

$$\max_{x \in [x_0, x_n]} |Z(t_{i+1}, x) - \hat{Z}(t_{i+1}, x)| = |\epsilon_z|_{\max} \leq \max_{x \in [x_0, x_n]} \frac{1}{2} \left| \frac{\partial^2 Z(t_{i+1}, x)}{\partial x^2} \right| \Delta^2, \qquad (22)$$

where Δ denotes the largest space between interpolation points.

In practice, however, dividing the domain upto six regions with four polynomial basis functions for each region already gives a small regression error. The break points for dividing the domain $[x_0, x_n]$ are chosen as the early-exercise point and the critical points for $\partial^2 Z(t_{i+1}, x)/\partial x^2$, Figure 8 compares the maximum and mean regression error with different numbers of pieces (keeping the number of grid points constant) and with different numbers of grid points (keeping the number of pieces constant) for a call option on single asset. It can be seen that for the same number of grid points, significantly smaller errors in regression can be obtained using more partitions.

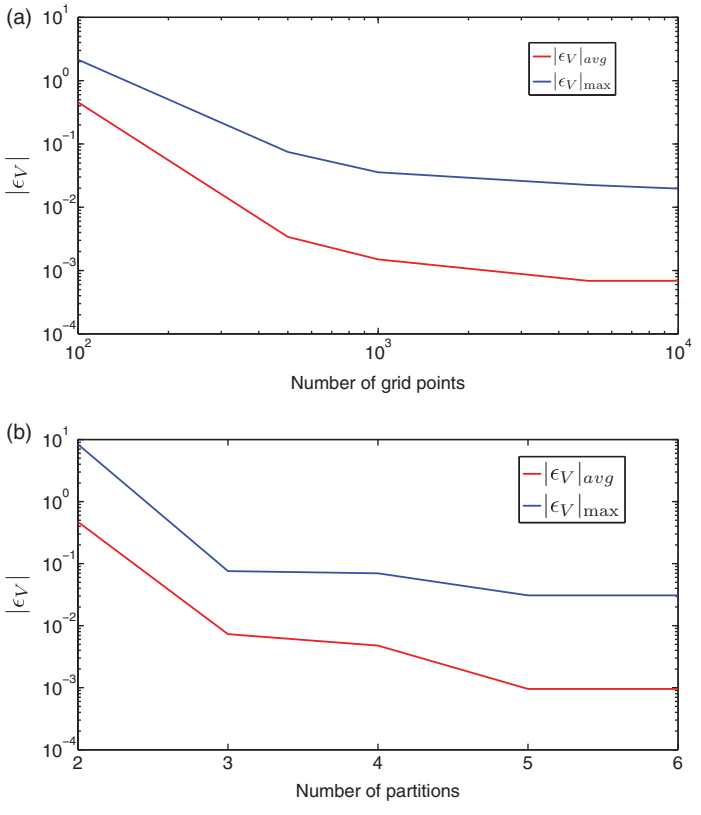

Figure 8. Maximum and average-squared residual errors due to parametrization of the option price when (a) the number of segments in the piecewise regression is constant = 6 and (b) the number of grid points used in the regression is constant = 10,000.

Assuming that the conditions above are satisfied, the error in continuation value due to parametrization of the option price is then bounded by

$$\epsilon_{Q_z}(t_i, S_{t_i}) \leq |\epsilon_z|_{\max} \int f(x|S_{t_i})\,\mathrm{d}x. \tag{23}$$

Under the assumption that the conditions for convergence of Gram Charlier series expansion are satisfied and we use large number of local regression functions, the error in the continuation value is bounded by

$$\epsilon_Q(t_i, S_{t_i}) \leq BLn^{-1/2} \int \mathrm{e}^{-x^2/4} Z(t_{i+1}, x)\,\mathrm{d}x + |\epsilon_z|_{\max} \int f(x|S_{t_i})\,\mathrm{d}x. \tag{24}$$

Here, we assume that $\int \mathrm{e}^{-x^2/4} Z(t_{i+1}, x)\,\mathrm{d}x$ is bounded.

3.3 *Error due to recursion*

From (24), the error in continuation value at t_i is bounded. At t_i , the error in the option price $V(t_i, S_{t_i})$ can be determined using

$$\begin{aligned} \hat{V}(t_i, S_{t_i}) &= \max(Q(t_i, S_{t_i}) + \epsilon_Q(t_i, S_{t_i}), h(S_{t_i})) \\ &\leq \max(Q(t_i, S_{t_i}), h(S_{t_i})) + |\epsilon_Q(t_i, S_{t_i})|. \end{aligned} \tag{25}$$

The continuation value at t_{i-1} will have error described by

$$\hat{Q}(t_{i-1}, S_{t_{i-1}}) \leq \int (Z(t_i, x) + |\epsilon_z(t_i, x)| + |\epsilon_Q(t_i, x)|)(f(x|S_{t_{i-1}}) + \epsilon_f(t_i, x))\,\mathrm{d}x. \tag{26}$$

The additional term in Equation (26) when compared to (15), is the error due to recursion, ϵ_R :

$$\epsilon_R \leq \int |\epsilon_Q(t_i, x)|(f(x|S_{t_{i-1}}) + \epsilon_f(t_i, x))\,\mathrm{d}x,$$

which is bounded by

$$\epsilon_R \leq \max_{S_{t_i}}(|\epsilon_Q(t_i, S_{t_i})|).$$

It can be shown that the error due to recursion at time step t_0 is bounded by

$$\epsilon_{R_0} \leq \sum_i \max_{S_{t_i}}(|\epsilon_Q(t_i, S_{t_i})|).$$

3.4 *Numerical results for Bermudan put on a single asset*

We illustrate the error analysis using numerical results for a put on a single asset, where the risk-neutral asset price follows the stochastic differential equation:

$$\mathrm{d}S = rS\,\mathrm{d}t + \sigma S\,\mathrm{d}W, \tag{27}$$

r being the continuously compounded risk-free interest rate, σ the annualized volatility. Here, we assume r and σ to be constant. W is the standard Brownian motion. We assume that the option is exercisable a finite number of times (k) per year, at a strike price of K, up-to and including the final expiration time T. We generate N sample paths $\{S_{t_0}, \ldots, S_{t_i}\}$, using the closed-form solution for the SDE (27). The asset values S_{t_i} represent the grid points in t_i.

3.4.1 *Parametrization of the option value for a single asset*

The option price at any time t_i prior to the expiration time T is given by

$$V(t_i, S_{t_i}) = \max(g(S_{t_i}) + X, Q(t_i, S_{t_i})).$$

To compute the functional approximation of the option value at time t_i, we regress the option values obtained at the grid points on polynomial basis functions of $g(S_{t_i}) = S_{t_i}$. We perform a piecewise least-squares regression with one of the break points at $\mathcal{X}_t^* = S_{t_i}^*$, where $S_{t_i}^*$ is the early-exercise point. For better approximation, the continuation region can be further divided into pieces with break points selected at the critical points for $Z''(t_i, S_{t_i})$. For the two segment case, we regress the option value as

$$\hat{Z}(t_i, S_{t_i}) = \mathbf{1}_{\{g(S_{t_i}) < \mathcal{X}_t^*\}} \sum_{m=0}^{M-1} a_m(|g(S_{t_i})|)^m + \mathbf{1}_{\{g(S_{t_i}) \geq \mathcal{X}_t^*\}} \sum_{m=0}^{M-1} b_m(|g(S_{t_i})|)^m, \tag{28}$$

with the coefficients a_m and b_m chosen so that residuals r_1 and r_2 are minimized,

$$r_1 = \min_{a_m} \left(\mathbf{1}_{\{g(S_{t_i}) < \mathcal{X}_t^*\}} \sum |V(t_i, S_{t_i}) - \hat{Z}(t_i, |g(S_{t_i})|)|^2 \right),$$

$$r_2 = \min_{b_m} \left(\mathbf{1}_{\{g(S_{t_i}) \geq \mathcal{X}_t^*\}} \sum |V(t_i, S_{t_i}) - \hat{Z}(t_i, |g(S_{t_i})|)|^2 \right).$$

We choose the first four polynomials (including the constant) as basis functions. Increasing the number of basis functions does not significantly improve the approximation; however, increasing the number of pieces does improve the solution.

3.4.2 *Continuation value for the single asset case*

In order to compute the continuation value for the grid points at t_i using Equation (11), we need the transition probability density function $\mathbb{P}(g(S_{t_{i+1}})|S_{t_i})$. For a single asset following a stochastic process given by Equation (27), the conditional transition density function is given by

$$\mathbb{P}(|g(S_{t_{i+1}})| = x|S_{t_i}) = S_{t_i}\,\mathrm{e}^{((r-\sigma^2/2)\Delta t+\sigma\sqrt{\Delta t}Y)}\mathbb{P}(Y = x^*), \tag{29}$$

where $\Delta t = t_{i+1} - t_i$, $Y \sim \mathcal{N}(0,1)$ and

$$x^* := \frac{1}{\sigma\sqrt{\Delta t}}\left[\log\left(\frac{x}{S_{t_i}}\right) - \left(r - \frac{\sigma^2}{2}\right)\Delta t\right].$$

Equation (12) can then be written as

$$\hat{Q}(t_i, S_{t_i}) = \frac{B_{t_i}}{B_{t_{i+1}}}\left(\int_{-\infty}^{K^*}\sum_{m=0}^{M-1} a_m(f(Y))^m\,\mathrm{d}\mathbb{P}(Y) + \int_{K^*}^{\infty}\sum_{m=0}^{M-1} b_m(f(Y))^m\,\mathrm{d}\mathbb{P}(Y)\right), \tag{30}$$

where

$$K^* = \frac{1}{\sigma\sqrt{\Delta t}}\left[\log\left(\frac{|\mathcal{X}^*_{t_{i+1}}|}{S_{t_i}}\right) - \left(r - \frac{\sigma^2}{2}\right)\Delta t\right],$$
$$f(Y) = S_{t_i}\,\mathrm{e}^{((r-\sigma^2/2)\Delta t+\sigma\sqrt{\Delta t}Y)},$$
$$\mathrm{d}\mathbb{P}(Y) = \frac{1}{\sqrt{2\pi}}\,\mathrm{e}^{-Y^2/2}\,\mathrm{d}Y.$$

Solving Equation (30) we obtain the continuation value at each grid point as

$$\hat{Q}(t_i, S_{t_i}) = \frac{B_{t_i}}{B_{t_{i+1}}}\left[\sum_{m=0}^{M-1}\varphi^m_{t_i}((a_m - b_m)\Phi(K^* - m\sigma\sqrt{\Delta t}) + b_m)\right], \tag{31}$$

where

$$\varphi^m_{t_i} = (S^m_{t_i}\,\mathrm{e}^{m((r-\sigma^2/2)+(m/2)\sigma^2)\Delta t}),$$

and

$$\Phi(x) = \frac{1}{2}\left[1 + erf\left(\frac{x}{\sqrt{2}}\right)\right].$$

In order to compute the value of $\mathcal{X}^*_{t_i} = |g((S_{t_i})|$, we need to solve the non-linear equation

$$g(S_{t_i}) = K - Q(t_i, S_{t_i}), \tag{32}$$

where the value of $Q(t_i, S_{t_i})$ is obtained from Equation (31). The value of $\mathcal{X}^*_{t_i}$ can be approximated as

$$\mathcal{X}^*_{t_i} = \max(|g(S_{t_i})|\mathbf{1}_{g(S_{t_i})\geq Q(t_i,S_{t_i})-X}),$$

i.e. we find the maximum value of the asset price for the grid points lying in the *early-exercise* region or alternatively the minimum value of the asset price for grid points in the *continuation* region.

3.4.3 *Results for single asset put option*

To illustrate the results, Table 1 reports the value of the early-exercise option implied by both the COS method and SGM. We use the COS method with $N = 2^{10}$ terms in the Fourier expansion, as our reference. The lower bound values, which are obtained by following the exercise policy from SGM on a fresh set of paths, are sometimes greater than the true option price. The lower bound values are taken as the mean of 30 simulation results. True lower bound values can be obtained by computing the mean over a large number of simulation results.

The SGM estimates are based on 10,000 (5000 plus and 5000 antithetic) paths using 50 exercise points per year, while the LSM estimates are based on 100,000 (50,000 plus and 50,000 antithetic) paths. Figure 9 compares the SGM direct estimator with the true option price for different numbers of grid points. Figure 10 compares the lower bound values obtained from SGM with lower bound from the LSM algorithm for different numbers of paths. The exercise policy obtained using SGM is better and stable compared with the one obtained using LSM, as can be deduced from the standard errors for the lower bounds for the two algorithms.The direct estimator value converges fast to the true price, as the number of partitions and grid points increases. The standard errors of the direct estimator are small compared with that of SGM lower bound values and much lower than that of LSM values.

The time taken for each simulation is few seconds on a system with Intel(R) Duo-Core 2.13 GHz processors and 2 GB RAM.

Table 1. Comparison of the SGM direct estimator and lower bound values with the LSM and COS method results for an Bermudan put option on a single asset, where the option is exercisable 50 times per year.

S_0	σ	T	COS method Bermudan	SGM lower bound (s.e.)	SGM direct estimator (s.e.)	Closed-form LSM (s.e.)	European
36	0.4	2	8.508	8.512 (0.56)	8.509 (0.010)	8.488 (0.51)	7.700
38	0.4	2	7.670	7.665 (0.53)	7.670 (0.011)	7.669 (0.50)	6.979
40	0.4	2	6.920	6.913 (0.59)	6.919 (0.011)	6.921 (0.55)	6.326
42	0.4	2	6.248	6.252 (0.59)	6.246 (0.013)	6.243 (0.51)	5.736
44	0.4	2	5.647	5.632 (0.66)	5.642 (0.014)	5.622 (0.51)	5.202

The strike price of the put is 40, the short-term interest rate is 0.06. The simulation for SGM is based on 10,000 (5000 plus 5000 antithetic) paths for the asset price process, and for LSM is based on 100,000 (50,000 plus and 50,000 antithetic) paths.
The standard error for the simulation (s.e.) is in cents while the option values are in dollars.

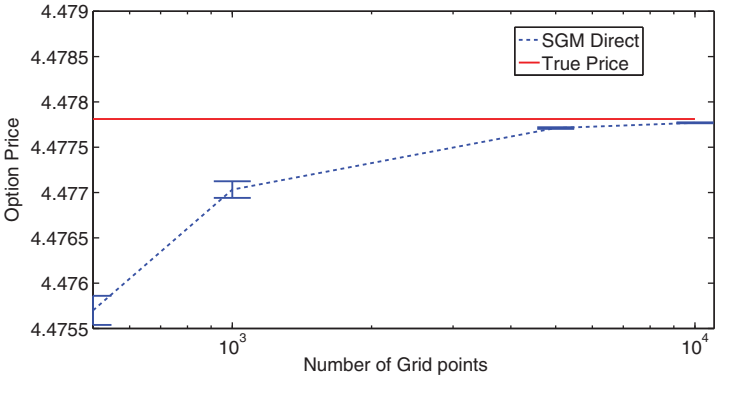

Figure 9. SGM direct estimator with confidence interval for different number of grid points. The regression is performed on six different pieces.

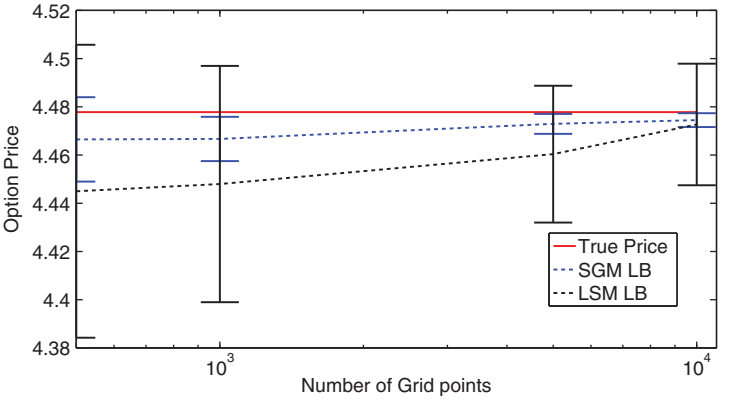

Figure 10. Comparison between lower bounds and confidence interval obtained using the exercise policy from SGM and LSM for different number of grid points (paths for latter).

4. Numerical results for high dimensions

In this section we illustrate our methodology by pricing Bermudan options on the max of two, three and five assets, and a basket option on an arithmetic mean of four and five assets. The underlying assets are assumed to follow the standard single and multi-asset Black–Scholes model (geometric Brownian motion, GBM).

4.1 *Bermudan call on maximum of d assets*

A Bermudan max-option is a discretely exercisable option on multiple underlying assets whose pay-off depends on the maximum among all asset prices. We assume that the asset prices follow correlated GBM processes, i.e.

$$\frac{\mathrm{d}S_t^i}{S_t^i} = (r - q_i)\,\mathrm{d}t + \sigma_i\,\mathrm{d}W_t^i, \tag{33}$$

where each asset pays a dividend at a continuous rate of q_i. $W_t^i, i = 1, \ldots, d$, are standard Brownian motions and the instantaneous correlation between W_t^i and W_t^j is ρ_{ij}. We assume that the option expires at time T and there are k equally spaced exercise dates in the interval $[0, T]$. If we use K to denote the strike price of the option, then the pay-off for d underlying assets is $\max(g(S_t^i) + X, 0)$, where $X = -K$ and, $g(S_t^i) = \max(S_t^1, \ldots, S_t^d)$. We start by generating N sample grid points $(S_{t_i}^1, \ldots, S_{t_i}^d)$ at each time step t_i, using the discretization scheme

$$S_{t_i}^j = S_{t_{i-1}}^j \exp\left(\left(r - q_i - \frac{1}{2}|\sigma_i|^2\right)\Delta t + \sum_{1\le k\le d} \sigma_{jk} W_{\Delta t}^k\right), \quad 1 \le j \le d, \tag{34}$$

where $\Delta t = t_i - t_{i-1}$. As explained in Section 2.5, for high-dimensional options additional peripheral paths are required to obtain better lower bound values. In the present example, we generate additional sample paths from two points around initial source point S_0, the points selected as $S_0\,\mathrm{e}^{0.3\sigma\sqrt{\Delta t}}$ and $S_0\,\mathrm{e}^{-0.1\sigma\sqrt{\Delta t}}$, which already significantly improves the lower bound values. The peripheral paths are used only to obtain the exercise-policy from the direct SGM estimator and are not used to obtain the lower bound values. Additional peripheral paths are required because in their absence the regressed function values around peripheral grid points become a source of

error. In the subsequent section, we discuss the scheme of parametrization and computing the continuation values specific to the Bermudan max-call option.

4.1.1 *Parametrization of the option value for max options*

In order to compute the functional form of the option value at t_{i+1}, we regress the option values obtained at the grid points over the polynomial basis functions of $g(S_{t_{i+1}})$. We use piecewise regression, with the break points at $\mathcal{X}_t^* = g(S_{t_{i+1}}^*)$, where $g(S_{t_{i+1}}^*) + X = Q(t_{i+1}, S_{t_{i+1}}^*)$. The regression scheme can be written as

$$\begin{aligned}\hat{Z}(t_{i+1}, g(S_{t_{i+1}})) = \mathbf{1}_{\{g(S_{t_{i+1}}) < \mathcal{X}_t^*\}} \sum_{m=0}^{M-1} a_m(\Psi_m(g(S_{t_{i+1}}))) \\ + \mathbf{1}_{\{g(S_{t_{i+1}}) \geq \mathcal{X}_t^*\}} \sum_{m=0}^{M-1} b_m(\Psi_m(g(S_{t_{i+1}}))),\end{aligned} \tag{35}$$

where Ψ are the basis functions. The coefficients a_m and b_m are chosen such that residuals r_1 and r_2 are minimized,

$$\begin{aligned} r_1 &= \min_{a_m} \left(\mathbf{1}_{\{g(S_{t_{i+1}}) < \mathcal{X}_t^*\}} \sum |V(t_{i+1}, S_{t_{i+1}}) - \hat{Z}(t_{i+1}, g(S_{t_{i+1}}))|^2 \right), \\ r_2 &= \min_{b_m} \left(\mathbf{1}_{\{g(S_{t_{i+1}}) \geq \mathcal{X}_t^*\}} \sum |V(t_{i+1}, S_{t_{i+1}}) - \hat{Z}(t_{i+1}, g(S_{t_{i+1}}))|^2 \right). \end{aligned}$$

We use a set of four (including the constant) *Hermite polynomial* basis functions of $g(S_{t_{i+1}})$ for regression in our example.

4.1.2 *Computing the continuation value for max options*

In order to compute the continuation value for grid points at t_i using Equation (11), we need to know the transition probability density function $\mathbb{P}(g(S_{t_{i+1}})|S_{t_i})$. For a call on the max of d underlying assets, $(g(S_{t_{i+1}}) = \max(S_{t_{i+1}}^1, \ldots, S_{t_{i+1}}^d))$, it is difficult to compute the exact transition density function. Like Boyle and Tse [9], we use Clark's algorithm to compute the first four moments of this distribution. The approximation of the transition probability density function can be obtained from these moments using the *Gram Charlier expansion*. Clark's algorithm [17] gives the exact expression for the first four moments of the maximum of a pair of jointly normal variates as well as the correlation coefficient between the maximum of the pair and the third normal variate. The details of Clark's algorithm are given in Appendix 1. S_{t_i} being a log-normal process given by Equation (34), we can write

$$\begin{aligned}\mathbb{P}(g(S_{t_{i+1}}) = X|S_{t_i}) &= \mathbb{P}\left(\max_{1 \leq j \leq d} (S_{t_{i+1}}^j) = X|S_{t_i} \right) \\ &= \mathbb{P}\left(\max_{1 \leq j \leq d} (Y_{t_{i+1}}^j) = \log(X)|S_{t_i} \right),\end{aligned} \tag{36}$$

where $Y_{t_{i+1}}^j$, $1 \leq j \leq d$ has a multivariate normal distribution. Using Clark's algorithm we can obtain the first four moments of the random variable $Y = \max(Y_{t_{i+1}}^1, \ldots, Y_{t_{i+1}}^d)$. If ${\kappa_i}^2$ $(1 \leq i \leq 4)$ are the first four cumulants of Y then using the *Gram Charlier Expansion*, we can write the

Table 2. Bermudan max-call option on 2, 3 and 5 underlying asset: the results are compared with Andersen and Broadie [1] and Broadie and Cao [11].

S_0	SGM LB (s.e.)	SGM direct (s.e.)	Binomial value	95% CI AB	95 % CI BC
$n = 2$ assets					
90	8.069 (0.026)	8.088 (0.003)	8.075	[8.053, 8.082]	–
100	13.892 (0.024)	13.900 (0.004)	13.902	[13.892, 13.934]	–
110	21.282 (0.028)	21.290 (0.003)	21.345	[21.316, 21.359]	–
$n = 3$ assets					
90	11.228 (0.023)	11.253 (0.003)	11.29	[11.265, 11.308]	–
100	18.665 (0.031)	18.625 (0.005)	18.69	[18.661, 18.728]	–
110	27.463 (0.036)	27.413 (0.006)	27.58	[27.512, 27.663]	–
$n = 5$ assets					
90	16.527 (0.028)	16.644 (0.005)	–	[16.602, 16.655]	[16.620, 16.653]
100	25.992 (0.033)	26.141 (0.006)	–	[26.109, 26.292]	[26.115, 26.164]
110	36.590 (0.047)	36.725 (0.005)	–	[36.704, 36.832]	[36.710, 36.798]

The parameters are: $K = 100$, $r = 5\%$, $q = 10\%$, $\rho = 0$, $T = 3$, $\sigma = 20\%$. There are 10 exercise opportunities equally spaced in time. Values in parentheses are standard errors. The total number of grid points at each time step was 30,000 with an equal number of paths generated from the three source grid points (two peripheral and one initial point).

approximate probability density function of Y as Equation (13). The continuation value given by Equation (12) can then be written as

$$\hat{Q}(t_i, S_{t_i}) = \frac{B_{t_i}}{B_{t_{i+1}}} \left(\int_{-\infty}^{K^*} \sum_{m=0}^{M-1} a_m \Psi_m(\mathrm{e}^x)\, \mathrm{d}\mathbb{P}(Y = x|S_{t_i}) + \int_{K^*}^{\infty} \sum_{m=0}^{M-1} b_m \Psi_m(\mathrm{e}^x)\, \mathrm{d}\mathbb{P}(Y = x|S_{t_i}) \right), \tag{37}$$

where $K^* = \log(\mathcal{X}^*_{t_i})$ The solution of Equation (37) is given in Appendix 2.

4.1.3 *Results for Bermudan call on max of several assets*

To illustrate the results, Table 2 compares the result of a Bermudan max option on 2, 3 and 5 underlying assets. The results reported in Table 2 are fairly remarkable given the simplicity of the method. The values obtained from the SGM are close to the values reported in the literature. The number of paths required to obtain an accurate exercise policy (as reflected by the lower-bound values) is far less than required to obtain the exercise policy for the duality-based methods. Also the time for each simulation is less than a minute on a system with Intel(R) Duo-Core 2.13 GHz processors and 2 GB RAM. The number of basis functions required for regression, irrespective of the dimensions of the problem, is upto 4 (including the constant).

4.2 *Bermudan put on arithmetic mean of* d *assets*

A Bermudan basket option is a discretely exercisable option on multiple underlying assets whose pay-off depends on the weighted average of the underlying asset prices. We assume that the asset prices follow correlated GBM processes given by Equation (33). The pay-off for d underlying assets is $\max(g(S_t^i) + X, 0)$, where $X = K$ and,

$$g(S_t^i) = -(w_1 S_t^1 + \cdots + w_d S_t^d), \tag{38}$$

such that

$$\sum_{i=1}^{d} w_i = 1.$$

The discretization and parametrization scheme for a Bermudan put on a basket is the same as that for Bermudan call on max of several assets. However, we use the case of *basket option* to show how the conditional continuation value can be computed in the general case.

4.2.1 *Computing the continuation value for Bermudan basket options*

In order to compute the continuation value for grid points at t_i using Equation (11), the transition probability density function $\mathbb{P}(g(S_{t_{i+1}})|S_{t_i})$ is required. For a put on the weighted mean of d underlying assets, the exact transition density function is unknown. The moments for the distribution of $g(S_{t_{i+1}})$ can be obtained using *sub-simulations*, which can be used to approximate the density function using the *Gram Charlier series* (Equation (13)). For each grid point at t_i, sub-paths are generated until the next time step t_{i+1}, and the first four non-central moments of the distribution of $g(S_{t_{i+1}})$ so obtained are computed. In order to re-use the results obtained for the Bermudan max option, we find the distribution $\mathbb{P}(\log(|g(S_{t_{i+1}})|) = x|S_{t_i})$, rather than determining $\mathbb{P}(|g(S_{t_{i+1}})| = x|S_{t_i})$. The continuation value is then given by Equation (37).

4.2.2 *Results for Bermudan basket option*

To illustrate the results, Table 3 compares the result of a Bermudan put option on four underlying assets. In order to compute the continuation value, we generate $N_S = 1000$, sub-paths for each of the underlying assets. The computational effort increases linearly with the number of exercise opportunities k, the number of paths N_S in the sub-simulations, and the dimension of the problem d. Although computationally more expensive than the case where the moments of the distribution can be computed analytically, this example shows a generic case when its not easy to compute the transition probability density or its moments directly. The time taken for each simulation was a few (<5) minutes. Table 4 compares the results of a Bermudan put option on five underlying assets with those reported by Bender *et al.* [5]. The LSM values and confidence intervals reported by Bender *et al.* our close to our values.

Table 3. Bermudan put option on arithmetic mean of four underlying assets: the results are compared with the CONV method of Lord *et al.* [26] and the LSM values.

S_0	SGM LB (s.e)	SGM direct (s.e)	FFT value	LSM (s.e.)
40	1.739 (0.37)	1.740 (0.16)	1.739(0.08)	1.739

The parameters are: $K = 40$, $r = 6\%$, $q = 2\%$, $\rho = 0.25$, $T = 1$, $\sigma = 20\%$. There are 10 exercise opportunities equally spaced in time. Values in parentheses are standard errors. The total number of grid points at each time step were 30,000 with equal number of paths generated from the three source grid points (two peripheral and one initial point). For the LSM algorithm there were 300,000 paths for each asset, and 18 set of basis functions.

Table 4. Bermudan put option on arithmetic mean of five underlying assets: the results are compared with the intervals reported by Bender *et al.* [5] and the LSM values.

S_0	SGM LB (s.e.)	SGM direct (s.e.)	LSM (s.e.)	BKS 95% CI
90	10.000 (0.00)	10.000 (0.00)	10.000 (0.00)	[10.000, 10.004]
100	2.134 (0.012)	2.141 (0.008)	2.163 (0.001)	[2.154, 2.164]
110	0.540 (0.010)	0.550 (0.006)	0.540 (0.001)	[0.535, 0.540]

The parameters are: $K = 100$, $r = 5\%$, $\rho = 0$, $T = 3$, $\sigma = 20\%$. There are four exercise opportunities (including t_0) equally spaced in time. Values in parentheses are standard errors. The total number of grid points at each time step were 3000 with equal number of paths generated from the 3 source grid points (two peripheral and one initial point). For the LSM algorithm, there were 120,000 paths for each asset and 24 sets of basis functions.

5. Conclusion

In this paper, we have presented the SGM for pricing and exercising Bermudan options. Our approach is based on dynamic programming and linear least-squares regression. One of the main achievements of the algorithm is its ability to reduce a multi-dimensional problem to a single-dimensional case, and yet avoid some of the associated short comings as were discussed by Boyle *et al.* [10].

The SGM method can be seen as a hybrid of the *stochastic mesh method* of Broadie and Glasserman [13], the *LSM* of Longstaff and Schwartz [25] and the *stratified sampling along the pay-off* method of Barraquand and Martineau (1996). It is similar to SMM as we follow the same dynamic programming approach, by approximating the option price at exercise times t_{i+1} and moving backwards in time using the information at t_{i+1}, to approximate the continuation value and hence the option price at exercise step t_i. We use the regression approach of LSM for approximating the conditional expectation $\mathbb{E}[\hat{V}(t_{i+1}, S_{t_{i+1}})|g(S_{t_{i+1}})]$. Similar to the approach of SSAP, we use the pay-off function to reduce the dimensions of the problem.

The regression in SGM differs from the LSM algorithm, as SGM does not approximate the functional form of the *continuation value*, rather it uses regression to approximate the functional form of $\mathbb{E}[V(t_{i+1}, S_{t_{i+1}})|g(S_{t_{i+1}})]$ at the exercise dates t_{i+1}.

In SSAP before pricing the option, the entire (i.e. at all time steps) state space is reduced to a one-dimensional state (using pay-off as the mapping function). The option pricing is then carried out for this reduced state space. This scheme can result in wrong exercise policy as was shown by Boyle *et al.* [10]. In SGM, option price and exercise policy for grid points in t_{i+1} are first computed in the high-dimensional space. This is followed by reducing the state space at t_{i+1}. The continuation value at t_i is then computed (by iterated conditioning) using one-dimensional probability density function, rather than multivariate distributions for transition $S_{t_{i+1}}|S_{t_i}$.

SGM can be computationally expensive when sub-simulations are required, especially when there are many early-exercise dates. The computational time when sub-simulations are required can be reduced by running SGM on parallel GPUs. SGM algorithm is well suited for parallel processing as computing the moments using sub-simulation can run independent of other processes. Another possible improvement could be using sub-simulation for a group of grid points (nearest neighbours) rather than for each grid points independently.

Although we did not look into the computation of price sensitivities for hedging purposes here, we believe that the method is well suited for computing the Greeks. Once the functional approximation of the option values at the next time step is available, it can be used to compute the continuation values and option values in the close neighbourhood of the initial point, thus allowing for fast approximation of some of the Greeks.

An aspect not yet explored is the choice of more refined methods to generate the initial grid points. The stochastic method for grid generation is a convenient method of grid generation, such that the density of the generated grid points at different time steps is closely related to the transition probability density function, when the transitions happen from the initial point. Since once the grid points are generated, the method employed for generating the grid points has no effect on the solution, it allows using more sophisticated methods of grid generation, like the *quantization tree method* of Bally *et al.* [2] in future work.

In order to show convergence of the SGM algorithm in high dimensions with an increasing number of paths, the method should include bundling for more accurately computation of the conditional expectation given by Equation (6).

Bundling can be interpreted as the *punishment for the dimension reduction* introduced in this paper, when high accuracy is required. However, in a follow-up paper we will show the strength of a bundling algorithm once it is implemented.

Acknowledgement

The second author thanks CWI-Centrum Wiskunde & Informatica, Amsterdam.

Notes

1. A Bermudan option is an option where the buyer has the right to exercise at a set (discretely spaced) of times. This is intermediate between a European option which allows exercise at a single time, namely expiry and an American option, which allows exercise at any time. With an increasing number of exercise opportunities Bermudan option values approach the value of an American option.
2. $$\kappa_1 = \mu = \mu_1'$$
$$\kappa_2 = \sigma^2 = \mu_2' - \mu_1'^2$$
$$\kappa_3 = \mu_3' - 3\mu_2'\mu_1' + 2\mu_1'^3$$
$$\kappa_4 = \mu_4' - 4\mu_3'\mu_1' - 3\mu_2'^2 + 12\mu_2'\mu_1'^2 - 6\mu_1'^4$$
where μ_i' is the ith non-central moment

References

[1] L. Andersen and M. Broadie, *A primal–dual simulation algorithm for pricing multi-dimensional American options*, Manag. Sci. 50 (2004), pp. 1222–1234.
[2] V. Bally, G. Pages, and J. Printems, *A quantization tree method for pricing and hedging multi-dimensional American options*, Math. Finance 15 (2005), pp. 119–168.
[3] J. Barraquand and D. Martineau, *Numerical valuation of high-dimensional multivariate American securities*, J. Financ. Quant. Anal. 30 (1995), pp. 383–405.
[4] D. Belomestny, A. Kolodko, and J. Schoenmakers, *Regression methods for stochastic control problems and their convergence analysis*, SIAM J. Control Optim. 48 (2010), pp. 35–62.
[5] C. Bender, A. Kolodko, and J. Schoenmakers, *Policy iteration for American options: Overview*, Monte Carlo Methods Appl. 12(5–6) (2006), pp. 347–362.
[6] S.J. Berridge and J.M. Schumacher, *An irregular grid approach for pricing high-dimensional American options*, J. Comput. Appl. Math. 222(1) (2008), pp. 94–111.
[7] P. Bossartes, *Simulation estimators of optimal early exercise*. Working Paper, Carnegie Mellon University, 1989.
[8] P. Boyle, *Options: A Monte Carlo approach*, J. Financ. Econ. 4 (1977), pp. 323–338.
[9] P. Boyle and Y. Tse, *An algorithm for computing values of options on the maximum or minimum of several assets*, J. Financ. Quant. Anal. 25 (1990), pp. 215–227.
[10] P. Boyle, M. Broadie, and P. Glasserman, *Monte Carlo methods for security pricing*, J. Econ. Dynam. Control 21 (1997), pp. 1267–1321.
[11] M. Broadie and M. Cao, *Improved lower and upper bound algorithms for pricing American options by simulation*, Quant. Finance 8 (2008), pp. 845–861.
[12] M. Broadie and J. Detemple, *American option valuation: New bounds, approximations, and a comparison of existing methods*, Rev. Financ. Stud. 9 (1996), pp. 1211–1250.
[13] M. Broadie and P. Glasserman, *Pricing American-style securities using simulation*, J. Econ. Dynam. Control 21(8–9) (1997), pp. 1323–1352.
[14] M. Broadie and P. Glasserman, *A stochastic mesh method for pricing high-dimensional American option*, J. Comput. Finance 7 (2004), pp. 35–72.
[15] M. Broadie, P. Glasserman, and Z. Ha, *Pricing American options by simulation using a stochastic mesh with optimized weights*, in *Probabilistic Constrained Optimization: Methodology and Applications*, S. Uryasev, ed., Kluwer Academic Publishers, Norwell, MA, 2000, pp. 32–50.
[16] J.F. Carriere, *Valuation of the early-exercise price for derivative securities using simulation and splines*, Insurance: Math. Econ. 19 (1996), pp. 19–30.
[17] C. Clark, *The greatest of a finite set of random variables*, Oper. Res. 9 (1961), pp. 145–162.
[18] D. Eglof, *Monte Carlo algorithms for optimal stopping and statistical learning*, Ann. Appl. Probab. 15(2) (2005), pp. 1396–1432.
[19] F. Fang and C.W. Oosterlee, *A novel pricing method for European options based on fourier-cosine series expansions*, SIAM J. Sci. Comput. 31 (2008), pp. 826–848.
[20] M. Haugh and L. Kogan, *Approximating pricing and exercising of high-dimensional American options: A duality approach*, Oper. Res. 52(2) (2004), pp. 258–270.
[21] A. Ibanez and F. Zapatero, *Monte Carlo valuation of American options through computation of the optimal exercise frontier*, J. Financ. Quant. Anal. 39(2) (2004), pp. 253–275.

[22] X. Jin, H.H. Tana, and J. Sun, *A state-space partitioning method for pricing high-dimensional American-style options*, Math. Finance 17 (2007), pp. 399–426.
[23] K.H.F. Kan, G. Frank, V. Mozgin, and R.M. Reesor, *Optimized least-squares Monte Carlo for measuring counterparty credit exposure of American-style options*, Math.-in-Industry Case Stud. J. 2 (2010), pp. 1–21.
[24] M.G. Kendall and A. Stuart, *The Advanced Theory of Statistics*, Vol. 1, Charles Griffin, London, 1969.
[25] F.A. Longstaff and E.S. Schwartz, *Valuing American options by simulation: A simple least-squares approach*, Rev. Financ. Stud. 3 (2001), pp. 113–147.
[26] R. Lord, F. Fang, F. Bervoets, and C.W. Oosterlee, *A fast and accurate FFT-based method for pricing early-exercise options under Lévy processes*, SIAM J. Sci. Comput. 30(4) (2008), pp. 1678–1705.
[27] W.E. Milne, *On the degree of convergence of the Gram-Charlier series*, Trans. Amer. Math. Soc. 31 (1929), pp. 422–443.
[28] N.S. Rasmussen, *Control variates for Monte Carlo valuation of American options*, J. Comput. Finance 9(1) (2005), pp. 84–102.
[29] L.C.G. Rogers, *Monte Carlo valuation of American options*, Math. Finance 12 (2002), pp. 271–286.
[30] J.A. Tilley, *Valuing American options in a path simulation model*, Trans. Soc. Actuaries 45 (1993), pp. 83–104.
[31] J. Tsitsiklis and B. Van Roy, *Optimal stopping of Markov processes: Hilbert space theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives*, IEEE Trans. Automat. Control 44 (1999), pp. 1840–1851.
[32] D. Zanger, *Convergence of a least-squares Monte Carlo algorithm for bounded approximating sets*, Appl. Math. Finance 16(2) (2009), pp. 123–150.

Appendix 1. The Clark algorithm

The Clark algorithm [17], calculates the first four moments of the maximum of a pair of jointly normal variates as well as the correlation coefficient between the maximum of the pair and a third normal variate. Let X_1 and X_2 have a bivariate normal distribution, with means μ_1 and μ_2, and standard deviations σ_1 and σ_2, respectively. The correlation coefficient between the two is ρ. Y denotes the maximum of (X_1,X_2). Let ν_i denote the ith non-central moment for the distribution of Y, then

$$\nu_1 = \mu_1\Phi(\alpha) + \mu_2\Phi(-\alpha) + a\phi(\alpha), \tag{A1}$$

$$\nu_2 = (\mu_1^2 + \sigma_1^2)\Phi(\alpha) + (\mu_2^2 + \sigma_2^2)\Phi(-\alpha) + (\mu_1 + \mu_2)a\phi(\alpha), \tag{A2}$$

$$\begin{aligned}\nu_3 = {} & (\mu_1^3 + 3\mu_1\sigma_1^2)\Phi(\alpha) + (\mu_2^3 + 3\mu_2\sigma_2^2)\Phi(-\alpha) \\ & + [(\mu_1^2 + \mu_1\mu_2 + \mu_2^2)a + (2\sigma_1^4 + \sigma_1^2\sigma_2^2 + 2\sigma_2^4 \\ & - 2\sigma_1^3\sigma_2\rho - 2\sigma_1\sigma_2^3\rho - \sigma_1^2\sigma_2^2\rho^2)a^{-1}]\phi(\alpha),\end{aligned} \tag{A3}$$

$$\begin{aligned}\nu_4 = {} & (\mu_1^4 + 6\mu_1^2\sigma_1^2 + 3\sigma_1^4)\Phi(\alpha) + (\mu_2^4 + 6\mu_2^2\sigma_2^2 + 3\sigma_2^4)\Phi(-\alpha) \\ & + \Bigg\{(\mu_1^3 + \mu_1^2\mu_2 + \mu_1\mu_2^2 + \mu_2^3)a - 3\alpha(\sigma_1^4 - \sigma_2^4) \\ & + 4\mu_1\sigma_1^3\left[3\left(\frac{\sigma_1 - \sigma_2\rho}{a}\right) - \left(\frac{\sigma_2 - \sigma_1\rho}{a}\right)^3\right] \\ & + 4\mu_2\sigma_2^3\left[3\left(\frac{\sigma_2 - \sigma_1\rho}{a}\right) - \left(\frac{\sigma_1 - \sigma_2\rho}{a}\right)^3\right]\Bigg\}\phi(\alpha).\end{aligned} \tag{A4}$$

If X_3 is a random variable with normal distribution, and the correlation coefficients between X_3 and X_1, X_2 are ρ_1, ρ_2, respectively, then the correlation coefficient ρ_{X_3Y} between X_3 and $Y = \max(X_1, X_2)$ is given by

$$\rho_{X_3Y} = [\sigma_1\rho_1\Phi(\alpha) + \sigma_2\rho_2\Phi(-\alpha)]/(\nu_2 - \nu_1^2)^{1/2}, \tag{A5}$$

where

$$a^2 = \sigma_1^2 + \sigma_2^2 - 2\sigma_1\sigma_2\rho,$$

$$\alpha = \frac{\mu_1 - \mu_2}{a},$$

$$\phi(x) = (2\pi)^{-1/2}\exp\left(-\frac{x^2}{2}\right),$$

$$\Phi(x) = \int_{-\infty}^{x}\phi(t)\,\mathrm{d}t.$$

Clark's method can be used to obtain the exact moments of Y and its correlation with X_3; however, as the distribution of Y is not exactly normal, the method can be only used to obtain the approximation of the first four moments of the maximum of a set of d normal variates. If $X_1, \ldots, X_d$ are the d jointly normal variates, and Y is the maximum of these d variates then by using the recursive scheme

$$Y_i = \max(X_1, X_2, \ldots, X_{i+1}) = \max(Y_{i-1}, X_{i+1})$$

and applying Clark's approximation at each step we can compute the approximation of the first four moments for the distribution of Y. It is easy to deduce how Clark's method can be used to obtain the moments for the *minimum* of d assets as well (see Boyle *et al.* [9]).

Appendix 2. Solution for continuation value

The solution to equation

$$\int_{-\infty}^{K^*} \sum_{m=0}^{M-1} a_m \Psi_m(\mathrm{e}^x)\, \mathrm{d}\mathbb{P}(Y = x|S_{t_{i-1}}),$$

where $\Psi(\cdot)$ is polynomial basis function, can be written as the linear combination of,

$$\int_{-\infty}^{K^*} (\mathrm{e}^{mx})\, \mathrm{d}\mathbb{P}(Y = x|S_{t_{i-1}}),$$

$m = (0, \ldots, (M-1))$, and $\mathbb{P}(Y = x|S_{t_{i-1}})$ is

$$\phi\left(\frac{x-\mu}{\sigma}\right)\left[1 + \frac{\kappa_3}{3!\sigma^3} H_3\left(\frac{x-\mu}{\sigma}\right) + \frac{\kappa_4}{4!\sigma^4} H_4\left(\frac{x-\mu}{\sigma}\right)\right].$$

Here $\phi(x)$ is,

$$\phi(x) = \frac{1}{\sqrt{2\pi}}\, \mathrm{e}^{-x^2/2}.$$

We need to solve,

$$\int_{-\infty}^{K^*} (\mathrm{e}^{mx})\phi\left(\frac{x-\mu}{\sigma}\right)\left[1 + \frac{\kappa_3}{3!\sigma^3} H_3\left(\frac{x-\mu}{\sigma}\right) + \frac{\kappa_4}{4!\sigma^4} H_4\left(\frac{x-\mu}{\sigma}\right)\right] \mathrm{d}x. \tag{A6}$$

Equation (A6) can be written as

$$A \int_{-\infty}^{K^*} \phi\left(\frac{x-\theta}{\sigma}\right)\left[1 + \frac{\kappa_3}{3!\sigma^3} H_3\left(\frac{x-\mu}{\sigma}\right) + \frac{\kappa_4}{4!\sigma^4} H_4\left(\frac{x-\mu}{\sigma}\right)\right] \mathrm{d}x, \tag{A7}$$

where

$$A = \mathrm{e}^{(\mu m + m^2\sigma^2/2)},$$

and

$$\theta = (\mu + m\sigma^2).$$

This can be written in a form easy to integrate,

$$A \int_{-\infty}^{K^*} \phi\left(\frac{x-\theta}{\sigma}\right)\left[1 + \frac{\kappa_3}{3!\sigma^3} \sum_{j=0}^{3} \binom{3}{j} \left(\frac{\theta-\mu}{\sigma}\right)^j H_{3-j}\left(\frac{x-\theta}{\sigma}\right)\right.$$
$$\left. + \frac{\kappa_4}{4!\sigma^4} \sum_{k=0}^{4} \binom{4}{k} \left(\frac{\theta-\mu}{\sigma}\right)^k H_{4-k}\left(\frac{x-\theta}{\sigma}\right)\right] \mathrm{d}x, \tag{A8}$$

using the property,

$$\int_{-\infty}^{x} \phi(y) H_n(y)\, \mathrm{d}y = -\phi(x) H_{n-1}(x).$$

Appendix 3. Gram Charlier series

The convergence of the Gram Charlier series has been discussed by Milne [27]. If a distribution satisfies the conditions given by Equations (18) and (19) then Milne shows the order of convergence for Gram Charlier series approximation of the distribution is $O(n^{-1/2})$, where n is number of terms in the series expansion. Here, we give some numerical results to show the effect of

- error in the moment estimates in the case of sub-simulation,
- non-random error in the Gram Charlier series approximation on the SGM estimator.

The effect of error in the moment estimates from sub-simulation is illustrated in Figure A1. It plots the standard error for the direct estimator when an increasing number of sub-paths is used. We plot the standard error *vs.* $1/\sqrt{N_S}$, where N_S is number of sub-paths used. When the exact values of the higher moments are used (as can be computed using the Clark algorithm), the standard error should be Y-intercept of the fitted function (as it is the case when $N_S \to \infty$). We find that this is indeed the standard error for the direct SGM estimator when we use exact moments from the Clark algorithm. Also we find that the mean of the direct SGM estimator and SGM lower bound values obtained with sub-simulation are close to those obtained using the exact moments.

As the error from the Gram Charlier series is independent of the regression error, we look at the case of a European option price for a max option. In this case, the error in the approximation of $\mathbb{E}[V(T, S_T)|g(S_T)]$ is zero. Then the error in the option price is only due to the Gram Charlier series. We compare the results with those from Boyle [9] as reference values.

Table A1 gives the max European call option values for a three-dimensional case, when the first two, three and four moments are used in the Gram Charlier series. We found that the error due to non-inclusion of higher moments while approximating the Gram Charlier series in the case of max option is significant only when the volatilities of the underlying assets are not the same.

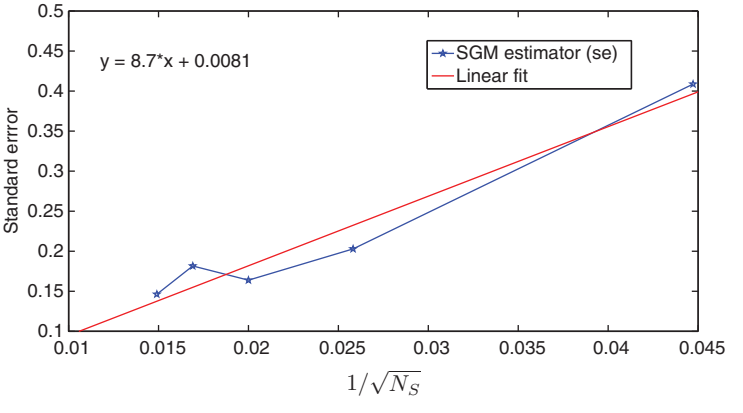

Figure A1. Standard error *vs.* $1/\sqrt{N_S}$ for a Bermudan max call option with 10 equally spaced exercise opportunity and 30,000 grid points. The parameters are same as Table 2 for 2 assets and $S_0 = [100, 100]$.

Table A1. European call option on max of three underlying assets: the results are compared with Boyle (1990).

Strike	Boyle	2 moments GC	3 moments GC	4 moments GC
30	16.703	16.705	16.700	16.705
40	9.235	9.249	9.251	9.237
50	4.438	4.375	4.458	4.439

The parameters are: $S_0 = [40, 40, 40]$, $r = 10\%$, $\rho = 0.9$, $T = 1$, $\sigma = [25, 30, 35]\%$.