

Pricing and hedging American-style options: a simple simulation-based approach

Yang Wang

UCLA Mathematics Department, Box 951555, Los Angeles, CA 90095-1555, USA;
email: wangyang0905@math.ucla.edu

Russel Caflisch

UCLA Mathematics Department, Box 951555, Los Angeles, CA 90095-1555, USA;
email: caflisch@math.ucla.edu

This paper presents a simple yet powerful simulation-based approach for approximating the values of prices and Greeks (ie, derivatives with respect to the underlying spot prices, such as delta, gamma, etc) for American-style options. This approach is primarily based upon the least squares Monte Carlo (LSM) algorithm and is thus termed the modified LSM (MLSM) algorithm. The key to this approach is that with initial asset prices randomly generated from a carefully chosen distribution, we obtain a regression equation for the initial value function, which can be differentiated analytically to generate estimates for the Greeks. Our approach is intuitive, easy to apply, computationally efficient and, most importantly, provides a unified framework for estimating risk sensitivities of the option price to underlying spot prices. We demonstrate the effectiveness of this technique with a series of increasingly complex but realistic examples.

1 INTRODUCTION

In the past years, Monte Carlo simulation has emerged as the most popular approach in computational finance for determining the prices of American-style options. Some important contributions are those of Tilley (1993), Carriere (1996), Broadie and Glasserman (1997, 2004), Tsitsiklis and Van Roy (1999, 2001), Longstaff and Schwartz (2001), Rogers (2002) and Andersen and Broadie (2004).

However, while calculating prices is one objective of Monte Carlo simulation and tremendous progress has been made in this area, the accurate estimation of Greeks via simulation remains an equally important but much more difficult task. Therefore, Monte Carlo simulation plays a much more crucial role in the calculation of price sensitivities. Both first- and second-order derivatives are essential for hedging and risk analysis, and even higher order derivatives are sometimes used.

The authors would like to thank Professor Francis Longstaff for his helpful discussions and comments throughout.

They are known collectively as the “Greeks”. Important as it is, efficient calculation of price sensitivities, especially for American-style options, continues to be among the greatest practical challenges facing users of Monte Carlo methods in the derivatives industry. Naturally enough, much attention has been focused on this area in recent years.

As one of the efforts in this trend, we present a simple yet powerful new approach to approximate the values of prices and Greeks for American-style options. Our new approach is primarily based upon the well-known least squares Monte Carlo (LSM) approach proposed by Longstaff and Schwartz (2001), which makes use of least squares regression to estimate the conditional expected payout from continuation at each exercise date. The idea for our modified LSM (MLSM) approach can be seen as a natural application of the technique in Pelsser and Vorst (1994) to the LSM algorithm, where the binomial tree was extended in a similar way in order to obtain more accurate Greeks. The key insight is that, by generating random initial prices for stock price sample paths, we can exploit the cross-sectional information in the simulated paths at initial time to infer option value information over a range of initial asset prices. This is done by roughly equating the option value function with the additional conditional expectation function estimated at initial time. Simple manipulation of this function immediately yields the desired estimates for price and Greeks.

To illustrate this technique, we present a series of increasingly complex but realistic examples. Firstly, we value an American put option on a single asset. Secondly, we value Bermudan max-call options on multiple underlying assets. This option is a typical high-dimensional example and poses great computational challenges to traditional finite-difference and binomial techniques. In the third example, we consider an exotic American–Bermudan–Asian option. This option is more complex than the previous ones since it is both path-dependent and has multifactor features. In each of these cases, the MLSM approach is able to produce results that closely approximate the benchmark values we provide. Finally, we value American options on an asset that follows a jump-diffusion process. This option is not directly solvable using standard partial differential equation or binomial techniques but poses little difficulty for our MLSM algorithm.

A number of papers have addressed the issue of using Monte Carlo simulation to estimate sensitivities for European options. Important examples of this literature include Glynn (1989), Broadie and Glasserman (1996) and Fournié *et al* (1999). Glasserman (2004, Chapter 7) provides an overview of these methods, which broadly fall into two categories. The first category uses finite-difference approximations and is superficially easier to understand and implement; the second uses information about the simulated stochastic process to replace numerical differentiation with exact differentiation. The pathwise derivative method and likelihood ratio method belong to this second category, and are found to be computationally more efficient and capable of providing more robust results than the finite-difference approach.

Recently, an important contribution was made by Piterbarg (2004, 2005) to extend the pathwise method and likelihood method to handle Bermudan-style options. The author finds that extension of the likelihood method is quite straightforward and requires little extra effort; extension of the pathwise method turns out to be much harder and constitutes the main theoretical result presented in the paper. Another serious attempt at generalizing the methods in Glasserman (2004) to handle Bermudan-style options was made by Kaniel *et al* (2007). Their algorithm is based on a combination of the likelihood ratio method for calculating European option sensitivities and the duality formulation for pricing Bermudan options, thus termed the likelihood ratio and duality (LRD) algorithm. In the LRD algorithm the Bermudan option is treated as a European option that expires on the first exercise date of the Bermudan option. The likelihood ratio method is thus applied on this European option while the duality method is used to approximate prices for the new Bermudan option, which has one exercise date less.


Our work takes a fundamentally different approach by focusing directly on the conditional expected function. We feel that our MLSM algorithm has a few attributes that make it a promising candidate for estimating sensitivities in future practice. First, it is intuitive and easy to apply, since nothing more than simple regression is required. Second, it is computationally as efficient as the LSM algorithm, as it only involves one extra regression being conducted at initial time. Third, it is readily applicable to cases with complex price dynamics or arbitrary payout functions. To demonstrate the generality of our approach, we have studied a series of increasingly complicated examples in our paper. Fourth, it does not suffer the problem arising from increasing the number of exercise dates as the likelihood method and LRD method experience. It can be directly used to approximate Greeks for continuously exercisable options. To further put the MLSM approach to the test, we have run a detailed performance comparison with the pathwise method, likelihood ratio method and the LRD method throughout this paper.

The remainder of this paper is organized as follows. Section 2 presents a simple numerical example of the MLSM approach. Section 3 describes the valuation framework and MLSM algorithm within a general theoretical setting. Sections 4–7 provide specific examples of application for this approach. Section 8 discusses a number of numerical and implementation issues. Section 9 summarizes the results and discusses some possible future directions.

2 A NUMERICAL EXAMPLE

Let us briefly restate the methodology for this MLSM approach. First we need to generate a number of sample paths for the stock price process. However, instead of fixing the initial prices at one point S_0 as is required for implementation of the LSM algorithm, here we “perturb” these initial prices by randomly generating them from a carefully chosen distribution centered around S_0 . The entire sample paths are thus constructed, and we apply the LSM algorithm recursively to these paths to

FIGURE 1 A comparison of sample path generation for the LSM and MLSM methods. (a) The LSM algorithm. (b) The MLSM algorithm.



obtain the optimal stopping times for each path. Finally, at the initial point we obtain a regression equation for the initial value function by regressing all the pathwise discounted payouts on a set of basis functions of the initial prices. Amazingly, the fitted values of this final regression provide a good approximation to the American option value for a range of asset prices near S_0 , without having to perform a full Monte Carlo simulation each time the asset price changes slightly. In particular, this allows us to calculate the hedging parameters (in fact, any derivatives with respect to stock price) directly through differentiating the resultant analytic expression. Figure 1 clearly illustrates the difference between the LSM and MLSM algorithms in terms of the way initial stock prices are generated.

Perhaps the best way to convey the intuition of this MLSM approach is through a simple numerical example. Here we will just use the one in Longstaff and Schwartz (2001). Consider a three-year American put option on a share of non-dividend-paying stock that can be exercised at the end of year 1, year 2 and year 3. The current stock price is 1.00 and the strike price is 1.10. The risk-free rate is 6% per annum (continuously compounded). For simplicity, we illustrate the algorithm using only eight sample paths for the price of the stock and the initial prices are produced from a uniform distribution on $[0.90, 1.10]$. (This is for illustration only; in practice many more paths would be sampled and other distributions could be used to generate initial prices.) The entire sample paths are constructed in Table 1 (see page 99).

First, we apply the LSM method to these sample paths to obtain the optimal stopping rule that maximizes the value of the option at each point along each path. The LSM method is recursive in nature and we need to work backwards one step at a time. Conditional on not exercising the option before the final expiration date at time 3, the cashflows realized by the option holder at time 3 are given in Table 2 (see page 99).

TABLE 1 Stock price paths.

Path	$t = 0$	$t = 1$	$t = 2$	$t = 3$
1	1.05	1.09	1.08	1.34
2	1.07	1.16	1.26	1.54
3	1.02	1.22	1.07	1.03
4	0.99	0.93	0.97	0.92
5	1.01	1.11	1.56	1.52
6	0.91	0.76	0.77	0.90
7	1.00	0.92	0.84	1.01
8	0.95	0.88	1.22	1.34

TABLE 2 Cashflows if exercised only at time 3.

Path	$t = 1$	$t = 2$	$t = 3$
1	–	–	0.00
2	–	–	0.00
3	–	–	0.07
4	–	–	0.18
5	–	–	0.00
6	–	–	0.20
7	–	–	0.09
8	–	–	0.00

If the put is in-the-money at time 2, the holder must decide whether or not to exercise. There are only five paths for which the option is in-the-money at time 2. We use only these in-the-money paths since it allows us to better estimate the conditional expectation function.¹ Let X denote the stock prices at time 2 and Y the corresponding discounted cashflows from continuation. Our five observations on X are 1.08, 1.07, 0.97, 0.77 and 0.84, and the corresponding values for Y are $0.00 e^{-0.06 \times 1}$, $0.07 e^{-0.06 \times 1}$, $0.18 e^{-0.06 \times 1}$, $0.20 e^{-0.06 \times 1}$ and $0.09 e^{-0.06 \times 1}$. Regressing Y on a constant, X , and X^2 yields the estimated conditional expectation function: $E[Y | X] = -1.070 + 2.983X - 1.813X^2$. In fact, this specification is one of the simplest possible; more general specifications will be discussed later in the paper.

With this conditional expectation function, we compare the value of immediate exercise with the value from continuation to find that it is optimal to exercise the

¹The reason for including only in-the-money paths in regression is primarily due to numerical considerations. This is clearly explained and numerically backed up in Longstaff and Schwartz (2001), “more than two or three times as many basis functions may be needed to obtain the same level of accuracy as obtained by the estimator based on in-the-money paths”. Further, for cases where the exercise value is not known immediately, this scheme of singling out in-the-money paths would fail and more specialized techniques need to be designed accordingly.

TABLE 3 Cashflows if exercised only at times 2 and 3.

Path	$t = 1$	$t = 2$	$t = 3$
1	—	0.00	0.00
2	—	0.00	0.00
3	—	0.00	0.07
4	—	0.13	0.00
5	—	0.00	0.00
6	—	0.33	0.00
7	—	0.26	0.00
8	—	0.00	0.00

TABLE 4 (a) Cashflows from the option; (b) option values at time 0.

Path	(a)			(b)	
	$t = 1$	$t = 2$	$t = 3$	Path	$t = 0$
1	0.00	0.00	0.00	1	0.02
2	0.00	0.00	0.00	2	0.00
3	0.00	0.00	0.07	3	0.07
4	0.17	0.00	0.00	4	0.13
5	0.00	0.00	0.00	5	0.09
6	0.34	0.00	0.00	6	0.33
7	0.18	0.00	0.00	7	0.11
8	0.22	0.00	0.00	8	0.23

option at time 2 for paths 4, 6 and 7. This leads us to the matrix in Table 3, which shows the cashflows received by the option holder conditional on not exercising prior to time 2.

Proceeding recursively, we next consider the paths that are in-the-money at time 1. These are paths 1, 4, 6, 7 and 8. Similarly, X represents the stock price at time 1 and Y the discounted value of subsequent option cashflows. The values of X for the paths are 1.09, 0.93, 0.76 and 0.92, and the corresponding values of Y are $0.00 e^{-0.06 \times 1}$, $0.13 e^{-0.06 \times 1}$, $0.33 e^{-0.06 \times 1}$, $0.26 e^{-0.06 \times 1}$ and $0.00 e^{-0.06 \times 1}$. Again linear regression gives us the estimated conditional expectation function: $E[Y | X] = 2.038 - 3.335X + 1.356X^2$.

This gives the value of continuing at time 1 for paths 1, 4, 6, 7 and 8 as 0.0139, 0.1092, 0.2866, 0.1175 and 0.1533, respectively. The value of immediate exercise is 0.01, 0.17, 0.34, 0.18 and 0.22. This means that we should exercise at time 1 for paths 4, 6, 7 and 8. Table 4(a) summarizes the cashflows assuming that early exercise is possible at all three times.

Having identified the cashflows generated by the American put at each date along each path, we can estimate the option value as a function of initial stock prices by

conducting a linear regression at time 0. For this we consider all the simulated paths (this time all the paths should be relevant for the regression²), and define X as initial stock prices for each path and Y the corresponding discounted payouts. Our eight observations on X are 1.05, 1.07, 1.02, 0.99, 1.01, 0.91, 1.00 and 0.95, and the values of Y are $0.00 e^{-0.06 \times 1}$, $0.00 e^{-0.06 \times 1}$, $0.07 e^{-0.06 \times 3}$, $0.17 e^{-0.06 \times 1}$, $0.00 e^{-0.06 \times 1}$, $0.34 e^{-0.06 \times 1}$, $0.18 e^{-0.06 \times 1}$ and $0.22 e^{-0.06 \times 1}$. Finally, regressing Y on a constant, X , and X^2 results in our desired option value function:

$$Y = 6.3828 - 10.5129X + 4.2437X^2 \quad (1)$$

Thus we can substitute the values of X into (1) to produce a table of estimates for option values at these initial prices, as shown in Table 4(b). Furthermore, this expression (1) provides us with a rough approximation to the option value for a continuous range of initial asset prices near S_0 ($=1.00$). With this analytical expression at hand, it is easy to calculate any derivatives with respect to asset price. $Y(S_0)$, $Y'(S_0)$ and $Y''(S_0)$ would immediately yield estimates for the price, Δ and Γ .

Since only eight sample paths are used here, the results provided above are by no means meant to accurately represent the true values. However, this simple example illustrates how least squares can use the cross-sectional information to estimate the conditional expected payout function as well as the initial value function. Like the original LSM algorithm, this MLSM algorithm is easily implemented since nothing more than simple regression is involved.

3 THE GENERAL MLSM ALGORITHM

In this section, we describe the general valuation framework and MLSM algorithm within a generic theoretical setting. We also discuss some related implementation issues and finally present a convergence result for the algorithm.

3.1 Valuation framework

The first step in implementing any numerical algorithm to price an American option is to assume that time can be discretized. Thus, we will assume that the derivative expires in L periods, and specify the exercise points as $t_0 = 0 < t_1 \leq t_2 \leq \dots \leq t_L = T$. In practice, of course, many American options are continuously exercisable;

²It should be noted that there is a fundamental difference between the regression at intermediate times and the regression at initial time. At intermediate times, we use only in-the-money paths for regression since we are interested in estimating the expectation conditional on the event that the option is in-the-money, in which case the comparison of exercise and continuation is relevant. At initial time, we use all the paths for regression because we are trying to estimate the option value function for all stock prices. In this case the comparison becomes irrelevant and hence there is no need for conditioning on the event that the option is in-the-money.

the MLSM algorithm can still be applied to these options by taking L to be sufficiently large.

We assume a complete probability space (Ω, \mathcal{F}, P) equipped with a discrete filtration $(\mathcal{F}(t_k))_{k=0}^L$. The underlying model is assumed to be Markovian, with state variables $(X(\omega, t_k))_{k=0}^L$ adapted to $(\mathcal{F}(t_k))_{k=0}^L$. We denote by $(Z(\omega, t_k))_{k=0}^L$ an adapted payout process for the derivative, satisfying $Z(\omega, t_k) = h(X(\omega, t_k), t_k)$, for a suitable function $h(\cdot, \cdot)$. As an example, consider the American put option from above, for which the only state variable of interest is the stock price, $X(\omega, t_k) = S(\omega, t_k)$. We have that $Z(\omega, t_k) = \max(K - S(\omega, t_k), 0)$.³

Here it is important to notice that in the LSM algorithm, $X(\omega, t_k) \in \mathcal{F}(t_k)$, $k = 0, 1, \dots, L$, and since $X(\omega, 0) = S_0$ is deterministic, $\mathcal{F}(0)$ is just a trivial σ -algebra. However, in our “modified” LSM algorithm, we make an important “modification”, that is, we randomly generate $X(\omega, 0)$ from some predetermined distribution $X_0(\omega)$ (ie, $X(\omega, 0) = X_0(\omega)$) and hence turn $\mathcal{F}(0)$ into a non-trivial σ -algebra.⁴

From the payout function $Z(\omega, t)$, we can define the function $C(\omega, \tilde{\tau}(t_k)) = e^{-r(\tilde{\tau}(t_k) - t_k)} Z(\omega, \tilde{\tau}(t_k))$ as the cashflow generated by the option, discounted back to t_k and conditional on no exercise prior to time t_k and on following a stopping strategy from t_k to expiration, written as $\tilde{\tau}(t_k)$ (essentially this corresponds to the $C(\omega, s; t_k, T)$ function from Longstaff and Schwartz (2001) defined in terms of stopping times). With this formulation we can specify the initial value function as:

$$V(X, 0) = \max_{\tilde{\tau}(0) \in \mathcal{T}(0)} E[C(\omega, \tilde{\tau}(0)) | X(\omega, 0)] \quad (2)$$

where the maximization is over stopping times $\tilde{\tau}(0) \in \mathcal{T}(0)$, with $\mathcal{T}(t_k)$ denoting the set of all stopping times with values in $\{t_k, \dots, t_K\}$. Here we suppress the randomness “ ω ” in the left-hand side. As we will see later, (2) is crucial to the formation of our MLSM algorithm.

3.2 The MLSM algorithm

Problems such as (2) are referred to as discrete time optimal stopping time problems and the preferred way to solve them is to use the dynamic programming principle. For the American option problem this can be written in terms of the optimal stopping

³Note that in some cases it might not be possible to determine the exercise value $Z(\omega, t_k)$ analytically. Therefore, we cannot focus on in-the-money paths when conducting the regressions. Consider an example where we have the option to enter into a (European) Asian option, in a model where an Asian option cannot be valued easily in closed form. We thank the referee for presenting such an example to elaborate this point.

⁴We are grateful to Professor Francis Longstaff for pointing out that this could be interpreted as starting the stock price process from some time before 0.

times $\tau(t_k)$ as follows:

$$\begin{cases} \tau(t_L) = T \\ \tau(t_k) = t_k \mathbf{1}_{\{Z(\omega, t_k) \geq E[C(\omega, \tau(t_{k+1})) | X(\omega, t_k)]\}} \\ \quad + \tau(t_{k+1}) \mathbf{1}_{\{Z(\omega, t_k) < E[C(\omega, \tau(t_{k+1})) | X(\omega, t_k)]\}}, \quad k \leq L - 1 \end{cases} \quad (3)$$

Thus the initial value function in (2) can be expressed in terms of the optimal stopping times in (3) as:

$$V(X, 0) = \max(F(\omega, 0), Z(\omega, 0)) \quad (4)$$

where $F(\omega, t_k) = E[C(\omega, \tau(t_{k+1})) | X(\omega, t_k)]$ following the notation in Longstaff and Schwartz (2001), which represents the expected payout from continuation at time t_k . Intuitively (4) makes sense because the option value at time 0 should be equal to the maximum of two things: the expected payout from continuation at time 0 and the payout from immediate exercise at time 0. By further restricting our attention to initial price regions where it is optimal to keep the option alive at time 0 (ie, $F(\omega, 0) \geq Z(\omega, 0)$),⁵ formula (4) reduces to:

$$V(X, 0) = F(\omega, 0) \quad (5)$$

The key contribution of Longstaff and Schwartz (2001) is that they provide a particularly useful method to approximate the conditional expectations ($F(\omega, t_k)$, $k = 0, 1, \dots, L - 1$) by using least squares regression. The theory on Hilbert spaces tells us that any function belonging to this space can be represented as a countable linear combination of basis vectors for the space. In particular, assuming that $F(\omega, t_k)$ belongs to Hilbert space, ie, is squarely integrable, we can write:

$$F(\omega, t_k) = \sum_{m=0}^{\infty} \phi_m(X(\omega, t_k)) a_m(t_k) \quad (6)$$

where $\{\phi_m(\cdot)\}_{m=0}^{\infty}$ form a set of basis functions.

However, the coefficients $\{a_m(t_k)\}$ in (6) are generally not known. Longstaff and Schwartz (2001) suggest in their algorithm a procedure for approximating $\{a_m(t_k)\}$ and thus $F(\omega, t_k)$ using the first M basis functions and N sample paths for stock price with $\widehat{F}_M^N(\omega, t_k)$ defined by:

$$\widehat{F}_M^N(\omega, t_k) = \sum_{m=0}^{M-1} \phi_m(X(\omega, t_k)) \widehat{a}_m^N(t_k) \quad (7)$$

⁵This is a reasonable adjustment to make since, for most options we deal with in practice, we are only interested in stock price regions where it is optimal not to exercise the option at initial time.

The algorithm is recursive, and at each point in time t_k the coefficients $\{\hat{a}_m^N(t_k)\}_{m=0}^{M-1}$ are calculated as the solution to the following least squares minimization problem:

$$\min_{\{\hat{a}_m^N\}_{m=0}^{M-1}} \sum_{n=1}^N (\hat{a}_0^N \phi_0(X(\omega_n, t_k)) + \dots + \hat{a}_{M-1}^N \phi_{M-1}(X(\omega_n, t_k)) - e^{-r(t_{k+1}-t_k)} C(\omega_n, \hat{\tau}_M^N(t_{k+1})))^2 \tag{8}$$

Eventually this procedure would give us $\{\hat{a}_m^N(0)\}_{m=0}^{M-1}$ and $\hat{F}_M^N(\omega, 0)$. Thus a natural approximation to the initial value function, in analogy to (5), can be designated as:

$$V(X, 0) \doteq \hat{F}_M^N(\omega, 0) = \sum_{m=0}^{M-1} \phi_m(X) \hat{a}_m^N(0) \tag{9}$$

The equation above tells us that the initial value function can be approximated by the conditional expected payout from continuation at time 0. This, we believe, is the essence of our MLSM algorithm. Compared to the original LSM algorithm, which evaluates the option value at only one point S_0 , Equation (9) turns out to be a significant step forward because it provides us with a direct estimation of the option values for a continuous range of stock prices near S_0 . In particular, we obtain $V(S_0, 0)$ simply by taking X to be S_0 in (9):

$$V(S_0, 0) \doteq \sum_{m=0}^{M-1} \phi_m(S_0) \hat{a}_m^N(0) \tag{10}$$

Equally important hedging parameters, such as Δ and Γ , are immediately produced by analytically differentiating the expression:

$$\Delta(S_0, 0) = \frac{\partial V}{\partial X}(S_0, 0) \doteq \sum_{m=0}^{M-1} \phi'_m(S_0) \hat{a}_m^N(0) \tag{11}$$

$$\Gamma(S_0, 0) = \frac{\partial^2 V}{\partial X^2}(S_0, 0) \doteq \sum_{m=0}^{M-1} \phi''_m(S_0) \hat{a}_m^N(0) \tag{12}$$

3.3 Initial distribution and convergence result

The concrete distribution $X_0(\omega)$ needs to be specified first for the initial prices to be generated. We have conducted a few experiments to investigate the relationship between initial distribution and the corresponding results, and we find that the results seem to be very robust to the choice of initial distribution. Specifically, for consistency and simplicity, we fix the initial distribution for all our examples to be:

$$X(\omega, 0) = S_0 e^{\alpha\sigma\sqrt{T}\omega}, \quad \omega \sim N(0, 1) \tag{13}$$

where σ and T are the corresponding stock volatility and time to expiration, and α could be used to adjust the variance of the distribution. The median of (13) is S_0 and it mimics the distribution for the underlying stock. One obvious advantage of this specification is that it allows the distribution to vary accordingly as problem parameters change and excludes extreme cases that might arise had the distribution been kept fixed.⁶ As repeated numerical experiments show, fixing α to be near 0.5 works reasonably well for all examples in this paper. Our experiments further show that too big or too small a value for α would lead to magnified variance and inaccuracy of the results.

However, it is worth noting that the initial distribution need not be random; a deterministic grid, spaced closely around the true initial value, would probably exhibit a comparable performance. One good candidate for such a deterministic grid is an evenly dispersed grid (with equal weights) over the interval:

$$(S(0) e^{-\alpha\sigma\sqrt{T}}, S(0) e^{\alpha\sigma\sqrt{T}})$$

where, once again, α is the parameter that could be used to vary the width of the grid. To further put it to test, we run a comparison between the results from an evenly dispersed grid and the ones from the random distribution (13). We found that the even grid would usually do as good a job in estimating the values as the random distribution, but tends to display a consistently bigger standard error. The results are presented in detail in Appendix A. Given the difference in producing standard errors displayed in these results, we recommend using a random distribution for generating starting points in future practice.

While the ultimate test of the MLSM algorithm is how well it performs on a set of realistic examples, it is also useful to examine what can be said about the theoretical convergence of the algorithm to the true option value function $V(X, 0)$. Fortunately, this topic has been extensively discussed by Stentoft (2004a, b), where the author has proved for the LSM algorithm the convergence in probability of the conditional expectation function $\widehat{F}_M^N(\omega, t_k)$ to the true function $F(\omega, t_k)$ under some general assumptions. This result answers virtually every question about the convergence of the MLSM algorithm, and we cite it below:

THEOREM 1 *Under Assumptions 1, 2 and 3,⁷ if M increases as N increases such that $M \rightarrow \infty$ and $M^3/N \rightarrow 0$, then $\widehat{F}_M^N(\omega, t_k)$ converges to $F(\omega, t_k)$ in probability,*

⁶As an example, consider an American option with a short maturity, on an asset with low volatility. If the initial distribution had been unchanged for all cases, the continuation values at future time steps would be estimated on too widely dispersed sample points, leading to less accurate results. However, (13) will generate a reasonable distribution for this case that is relatively concentrated near S_0 .

⁷Refer to Stentoft (2004a, b) for details of these assumptions. Basically these are regularity and integrability assumptions on the conditional expectation functions to ensure convergence in the result.

for $k = 0, 1, \dots, L - 1$, ie, for any $\epsilon > 0$:

$$\Pr(|\widehat{F}_M^N(\omega, t_k) - F(\omega, t_k)| > \epsilon) \longrightarrow 0 \quad (14)$$

This theorem states that in the limit $\widehat{F}_M^N(\omega, 0)$ converges to $F(\omega, 0)$, which is again equal to the true option value function $V(X, 0)$ in stock price regions where early exercise at time 0 is not optimal. Finally, this justifies the validity of formula (12), the foundation on which all of our calculations in this paper are made. Although it is required that both M and N increase to infinity for convergence of the algorithm, it turns out that a very small value for M would suffice for most cases in practice. For one-dimensional problems, $M = 5$ basis functions prove to be sufficient for a satisfactory level of accuracy.

4 AMERICAN PUT OPTION ON A SINGLE ASSET

Earlier in the second chapter we used a stylized example to illustrate how this approach could be applied to a standard American put option. In this section we present an in-depth example of the application of the MLSM algorithm to American put options.

Assume that we are interested in pricing an American-style put option on a single share of stock, whose risk-neutral process follows a geometric Brownian motion process:

$$dS = rS dt + \sigma S dZ \quad (15)$$

where r and σ are constants, Z is a standard Brownian motion and the stock does not pay dividends. Furthermore, assume that the option is continuously exercisable at a strike price K up to and including the final expiration date T of the option. No closed-form solution for the price and hedging parameters is known, but there are various existing numerical methods that give good approximations very rapidly, such as the well-known binomial/trinomial tree method, the finite-difference scheme for partial differential equations, etc.

In applying our MLSM approach, we use as the set of basis functions a constant and the first four power polynomials X, X^2, X^3, X^4 . For the initial distribution of $X(\omega, 0)$, we use the specification in (16), with the value of α set to be 0.5. Once the coefficients $\{\hat{a}_m^N(0)\}_{m=0}^4$ are calculated, we can approximate the option value function as:

$$V(X, 0) \doteq \hat{a}_0^N(0) + \hat{a}_1^N(0)X + \hat{a}_2^N(0)X^2 + \hat{a}_3^N(0)X^3 + \hat{a}_4^N(0)X^4 \quad (16)$$

Direct substitution and differentiation would yield estimates for price, Δ and Γ :

$$V(S_0, 0) \doteq \hat{a}_0^N(0) + \hat{a}_1^N(0)S_0 + \hat{a}_2^N(0)S_0^2 + \hat{a}_3^N(0)S_0^3 + \hat{a}_4^N(0)S_0^4 \quad (17)$$

$$\Delta(S_0, 0) \doteq \hat{a}_1^N(0) + 2\hat{a}_2^N(0)S_0 + 3\hat{a}_3^N(0)S_0^2 + 4\hat{a}_4^N(0)S_0^3 \quad (18)$$

$$\Gamma(S_0, 0) \doteq 2\hat{a}_2^N(0) + 6\hat{a}_3^N(0)S_0 + 12\hat{a}_4^N(0)S_0^2 \quad (19)$$

It is straightforward to add additional basis functions as explanatory variables in the regression if needed. Using more than five basis functions, however, causes little change to the numerical results; five basis functions are adequate to obtain effective convergence of the algorithm in this example.

To test the performance of the MLSM method, we compare our results of option prices and Greeks calculated from the MLSM method with the benchmark results, which in this case are produced by the standard binomial model with $N = 10,000$ time steps. To further compare the MLSM method with other existing simulation-based methods for computing Greeks, we also report estimates of deltas from the pathwise derivative method and the likelihood ratio method as well as estimates of gammas from the likelihood ratio method.⁸ They are labeled as “MLSM”, “Binomial”, “Pathwise” and “Likelihood” values respectively in Table 5 (see page 108). All simulation estimates are based on 150,000 sample paths for stock price process with 150 discretization points per year. Each estimate comes with a standard error, which is computed by independently running the procedure 15 times, and this is given in the parentheses immediately below.

As shown, the differences between “MLSM”, “Pathwise” and benchmark results are quite small. The MLSM method generally performs equally as well as the pathwise method in estimating the deltas and their standard deviations. The likelihood method tends to exhibit much poorer results in that the “Likelihood” estimates usually have a much larger variance. The standard errors for the “MLSM” price, Δ and Γ , results are very small, usually accounting for less than 1% in proportion to the simulated values. All the benchmark results reported in Table 5 are within one standard error of the “MLSM” ones. In summary, these results suggest that the MLSM algorithm is able to approximate closely the binomial benchmark values.

4.1 LSM versus MLSM

Before ending the discussion of this section, we propose one more interesting diagnostic test between the LSM and MLSM methods by comparing their performances in computing the option prices for the same problem. Since both the LSM and MLSM methods can be used to calculate American option prices, we apply both algorithms to the same example with all parameters being identical.

To make the comparison more meaningful, a constant and the first four power polynomials are selected as common basis functions for both algorithms. The comparison of results and other implementation details are presented in Table 6 (see page 109). The simulations for the two algorithms are both based on 200,000 sample

⁸It is generally inapplicable to apply the pathwise method to estimating second derivatives for many important types of options due to the requirement of continuity in the discounted payout. Refer to Glasserman (2004) for a detailed explanation of this point. For this reason, we skip reporting estimates of gammas from the pathwise method for all examples in this paper.

TABLE 5 Standard American put option.

K	σ	T	MLSM price	Binomial price	MLSM Δ	Pathwise Δ	Likelihood Δ	Binomial Δ	MLSM Γ	Likelihood Γ	Binomial Γ
35	0.20	1/3	0.1991 (0.0026)	0.2004	-0.0903 (0.0012)	-0.0910 (0.0007)	-0.0909 (0.0025)	-0.0901	0.0367 (0.0012)	0.0368 (0.0069)	0.0357
35	0.20	7/12	0.4301 (0.0031)	0.4328	-0.1346 (0.0012)	-0.1349 (0.0012)	-0.1349 (0.0058)	-0.1338	0.0373 (0.0013)	0.0402 (0.0120)	0.0364
40	0.20	1/3	1.5786 (0.0071)	1.5798	-0.4434 (0.0029)	-0.4442 (0.0017)	-0.4414 (0.0077)	-0.4435	0.0930 (0.0024)	0.0951 (0.0187)	0.0923
40	0.20	7/12	1.9848 (0.0086)	1.9904	-0.4287 (0.0025)	-0.4280 (0.0027)	-0.4324 (0.0085)	-0.4287	0.0730 (0.0020)	0.0663 (0.0231)	0.0719
45	0.20	1/3	5.0942 (0.0073)	5.0883	-0.8848 (0.0039)	-0.8773 (0.0060)	-0.8792 (0.0099)	-0.8812	0.0811 (0.0015)	0.0652 (0.0554)	0.0827
45	0.20	7/12	5.2722 (0.0056)	5.2670	-0.7999 (0.0033)	-0.7925 (0.0040)	-0.7925 (0.0114)	-0.7948	0.0736 (0.0012)	0.0743 (0.0555)	0.0787
35	0.30	1/3	0.6972 (0.0059)	0.6975	-0.1745 (0.0014)	-0.1750 (0.0014)	-0.1741 (0.0060)	-0.1741	0.0377 (0.0013)	0.0357 (0.0087)	0.0376
35	0.30	7/12	1.2229 (0.0048)	1.2198	-0.2135 (0.0020)	-0.2144 (0.0009)	-0.2117 (0.0078)	-0.2126	0.0321 (0.0005)	0.0346 (0.0091)	0.0326
40	0.30	1/3	2.4808 (0.0090)	2.4825	-0.4414 (0.0029)	-0.4432 (0.0017)	-0.4455 (0.0113)	-0.4420	0.0591 (0.0016)	0.0585 (0.0195)	0.0597
40	0.30	7/12	3.1678 (0.0132)	3.1696	-0.4265 (0.0027)	-0.4277 (0.0022)	-0.4191 (0.0124)	-0.4256	0.0463 (0.0019)	0.0437 (0.0127)	0.0459
45	0.30	1/3	5.7012 (0.0152)	5.7056	-0.7266 (0.0042)	-0.7273 (0.0024)	-0.7256 (0.0097)	-0.7266	0.0576 (0.0014)	0.0720 (0.0251)	0.0572
45	0.30	7/12	6.2318 (0.0111)	6.2436	-0.6537 (0.0023)	-0.6535 (0.0029)	-0.6531 (0.0109)	-0.6520	0.0497 (0.0012)	0.0613 (0.0313)	0.0485

This table presents estimates of prices and sensitivities for standard American put options. The first three columns represent different values of the parameters K ($=35, 40, 45$), σ ($=0.2, 0.3$), T ($=1/3, 7/12$), and the other fixed parameters are $S_0 = 40$, $r = 4.88\%$. All simulation results are based on 150,000 sample paths for the stock-price process with 150 discretization points per year. Their respective standard errors are given in the parentheses immediately below them. The "Binomial" columns show the benchmark results for the corresponding values from the standard binomial method with $N = 10,000$ time steps. As shown, all the benchmark results are within one standard error away from those for "MLSM".

TABLE 6 American put option prices: LSM versus MLSM.

K	σ	T	LSM price	(s.e.)	MLSM price	(s.e.)	Binomial price
35	0.2	1/3	0.2012	(0.0021)	0.1991	(0.0026)	0.2004
35	0.2	7/12	0.4338	(0.0041)	0.4301	(0.0031)	0.4328
40	0.2	1/3	1.5806	(0.0073)	1.5786	(0.0071)	1.5798
40	0.2	7/12	1.9915	(0.0087)	1.9848	(0.0086)	1.9904
45	0.2	1/3	5.0909	(0.0048)	5.0942	(0.0073)	5.0883
45	0.2	7/12	5.2645	(0.0078)	5.2722	(0.0056)	5.2670
35	0.3	1/3	0.6995	(0.0049)	0.6972	(0.0059)	0.6975
35	0.3	7/12	1.2227	(0.0066)	1.2229	(0.0048)	1.2198
40	0.3	1/3	2.4846	(0.0109)	2.4808	(0.0090)	2.4825
40	0.3	7/12	3.1675	(0.0103)	3.1678	(0.0132)	3.1696
45	0.3	1/3	5.7034	(0.0146)	5.7012	(0.0152)	5.7056
45	0.3	7/12	6.2426	(0.0182)	6.2318	(0.0111)	6.2436
35	0.4	1/3	1.3483	(0.0073)	1.3455	(0.0103)	1.3460
35	0.4	7/12	2.1534	(0.0119)	2.1522	(0.0072)	2.1549
40	0.4	1/3	3.3906	(0.0120)	3.3840	(0.0182)	3.3874
40	0.4	7/12	4.3527	(0.0147)	4.3505	(0.0208)	4.3526
45	0.4	1/3	6.5125	(0.0122)	6.5119	(0.0112)	6.5099
45	0.4	7/12	7.3856	(0.0146)	7.3780	(0.0181)	7.3830

This table presents a comparison of price estimates for the American put option using both the LSM and MLSM algorithm. The first three columns represent different values for the parameters K , σ and T , and the other fixed parameters are $S_0 = 40$, $r = 4.88\%$. The simulations are all based on 200,000 sample paths for the stock-price process with 150 discretization points per year. Their respective standard errors are given in the parentheses immediately to the right. The "Binomial price" column shows the benchmark results for corresponding values from the standard binomial model with $N = 10,000$ time steps. All benchmark results are within one standard error of the simulated ones.

paths for the stock-price process with 150 discretization points per year. Again, each standard error in the table is computed by independently running the procedure 15 times.

As clearly shown in Table 6, the differences between the two algorithms in terms of computing option prices as well as their standard deviations are very slight. Both algorithms can be used to closely approximate the values of American options. It is thus safe to conclude that there is no need to run the LSM algorithm separately for prices and then the MLSM algorithm for Greeks in the same problem. We suggest that the MLSM algorithm be run only once to obtain price estimates as a by-product for any future application.

5 BERMUDAN OPTIONS ON MULTIPLE ASSETS

Like a typical simulation-based approach, the MLSM method is readily applicable in path-dependent and multifactor situations (particularly with five or more assets) where traditional lattice techniques usually suffer from serious numerical constraints.

In this section, we test its performance on the pricing of multi-asset equity options. Specifically, we price max-call equity options, a problem that has become a standard test case in the literature.

The payout of a max-call option at time t is equal to:

$$(\max(S_1(t), \dots, S_n(t)) - K)^+ \quad (20)$$

We denote $S(t) = (S_1(t), \dots, S_n(t))$ and assume that the risk-neutral dynamics for these n underlying assets follow correlated geometric Brownian motion processes:

$$dS_i = (r - \delta_i)S_i dt + \sigma_i S_i dZ_i \quad (21)$$

where $Z_i, i = 1, \dots, n$, are standard Brownian motion processes, and the instantaneous correlation of Z_i and Z_j is ρ_{ij} . For simplicity, in our numerical results we take $\delta_i = \delta, \sigma_i = \sigma$ and $\rho_{ij} = \rho$ for all $i, j = 1, \dots, n$ and $i \neq j$. The interest rate r is also assumed to be constant. Exercise opportunities are equally spaced at times $t_i = iT/d, i = 0, 1, \dots, d$. We test our MLSM method for $n = 2, 3, 5$, and the results are given in Tables 7 and 8 (see pages 111 and 113). The benchmark results are chosen to be the values produced from the classical multidimensional binomial routine by Boyle *et al* (1989).

It is not hard to determine which basis functions to use for regression at intermediate time steps. For all three cases $n = 2, 3, 5$, we choose the set of basis functions to consist of a constant, the first five power polynomials in the highest price, the values and squares of values of the second through n th highest prices, the product of highest and second highest, second highest and third highest, etc, and finally, the product of all assets.

However, for higher-dimensional cases the choice of initial distribution and basis functions for regression at time 0 turns out to be a much more complicated issue than the previous one-dimensional case, and thus needs to be handled with care and treated differently. For $n = 2$ and 3, we report our results for three representative hedging parameters, $\Delta_1 (= \partial V / \partial S_1(0)), \Gamma_{11} (= \partial^2 V / \partial S_1(0)^2)$ and $\Gamma_{12} (= \partial^2 V / \partial S_1(0) \partial S_2(0))$. When calculating Δ_1 and Γ_{11} , we sample only $S_1(0)$ from the specification in (16) while keeping other $S_i(0)$ fixed at S_0 . For regression at time 0, we regress pathwise discounted payouts on a constant and the first four power polynomials of $S_1(0)$, and then differentiate the approximated function once and twice, respectively, to get estimates for Δ_1 and Γ_{11} . However, when it comes to calculating Γ_{12} , we sample both $S_1(0)$ and $S_2(0)$ from (16), and regress pathwise discounted payouts on a set of basis functions in $S_1(0)$ and $S_2(0)$ – a constant, the first four power polynomials of $S_1(0)$ and $S_2(0)$, their product, two terms of degree three ($S_1^2(0)S_2(0), S_1(0)S_2^2(0)$) and three terms of degree four ($S_1^3(0)S_2(0), S_1^2(0)S_2^2(0), S_1(0)S_2^3(0)$). Then the estimate for Γ_{12} is obtained by differentiating the approximated function with respect to both $S_1(0)$ and $S_2(0)$.

As shown in Table 7 (see pages 111–112), the pathwise method tends to exhibit a persistently smaller standard error in reporting Δ than the MLSM and likelihood

TABLE 7 Bermudan max-call option on two and three assets.

S_0	MLSM		Pathwise		Likelihood		Binomial		MLSM		Likelihood		Binomial		MLSM		Likelihood		Binomial		
	Δ_1	Γ_1	Δ_1	Γ_1	Δ_1	Γ_1	Δ_1	Γ_1	Δ_1	Γ_1	Δ_1	Γ_1	Δ_1	Γ_1	Δ_1	Γ_1	Δ_1	Γ_1	Δ_1	Γ_1	
$n = 2$																					
70	0.02265 (0.00304)	0.02362 (0.00040)	0.02376 (0.00152)	0.02348	0.00405 (0.00056)	0.00410 (0.00030)	0.00398	-0.00009 (0.00045)	-0.00016 (0.00033)	-0.00015	0.02265 (0.00304)	0.02362 (0.00040)	0.02376 (0.00152)	0.02348	0.00405 (0.00056)	0.00410 (0.00030)	0.00398	-0.00009 (0.00045)	-0.00016 (0.00033)	-0.00015	
80	0.08753 (0.00394)	0.08765 (0.00055)	0.08837 (0.00265)	0.08757	0.01027 (0.00028)	0.01032 (0.00038)	0.01019	-0.00106 (0.00053)	-0.00096 (0.00027)	-0.00105	0.08753 (0.00394)	0.08765 (0.00055)	0.08837 (0.00265)	0.08757	0.01027 (0.00028)	0.01032 (0.00038)	0.01019	-0.00106 (0.00053)	-0.00096 (0.00027)	-0.00105	
90	0.19878 (0.00545)	0.19992 (0.00150)	0.20078 (0.00333)	0.20026	0.01672 (0.00106)	0.01660 (0.00068)	0.01654	-0.00356 (0.00048)	-0.00388 (0.00033)	-0.00377	0.19878 (0.00545)	0.19992 (0.00150)	0.20078 (0.00333)	0.20026	0.01672 (0.00106)	0.01660 (0.00068)	0.01654	-0.00356 (0.00048)	-0.00388 (0.00033)	-0.00377	
100	0.32802 (0.00578)	0.32558 (0.00127)	0.32608 (0.00410)	0.32643	0.01981 (0.00079)	0.02004 (0.00066)	0.02018	-0.00835 (0.00068)	-0.00835 (0.00044)	-0.00844	0.32802 (0.00578)	0.32558 (0.00127)	0.32608 (0.00410)	0.32643	0.01981 (0.00079)	0.02004 (0.00066)	0.02018	-0.00835 (0.00068)	-0.00835 (0.00044)	-0.00844	
110	0.42077 (0.00710)	0.42269 (0.00100)	0.42128 (0.00430)	0.42304	0.02000 (0.00090)	0.02090 (0.00057)	0.02069	-0.01258 (0.00058)	-0.01379 (0.00068)	-0.01322	0.42077 (0.00710)	0.42269 (0.00100)	0.42128 (0.00430)	0.42304	0.02000 (0.00090)	0.02090 (0.00057)	0.02069	-0.01258 (0.00058)	-0.01379 (0.00068)	-0.01322	
120	0.47883 (0.00651)	0.47563 (0.00134)	0.47622 (0.00595)	0.47609	0.01905 (0.00081)	0.01910 (0.00095)	0.01957	-0.01497 (0.00067)	-0.01547 (0.00075)	-0.01604	0.47883 (0.00651)	0.47563 (0.00134)	0.47622 (0.00595)	0.47609	0.01905 (0.00081)	0.01910 (0.00095)	0.01957	-0.01497 (0.00067)	-0.01547 (0.00075)	-0.01604	
130	0.49789 (0.00814)	0.49856 (0.00131)	0.50131 (0.00666)	0.49868	0.01734 (0.00086)	0.01782 (0.00066)	0.01819	-0.01545 (0.00056)	-0.01633 (0.00074)	-0.01682	0.49789 (0.00814)	0.49856 (0.00131)	0.50131 (0.00666)	0.49868	0.01734 (0.00086)	0.01782 (0.00066)	0.01819	-0.01545 (0.00056)	-0.01633 (0.00074)	-0.01682	

TABLE 7 Continued.

S_0	MLSM Δ_1	Pathwise Δ_1	Likelihood Δ_1	Binomial Δ_1	MLSM Γ_{11}	Likelihood Γ_{11}	Binomial Γ_{11}	MLSM Γ_{12}	Likelihood Γ_{12}	Binomial Γ_{12}
$n = 3$										
70	0.02206 (0.00237)	0.02247 (0.00050)	0.02254 (0.00107)	0.02267	0.00403 (0.00044)	0.00398 (0.00034)	0.00388	-0.00019 (0.00021)	-0.00021 (0.00017)	-0.00014
80	0.07903 (0.00409)	0.08027 (0.00070)	0.08124 (0.00255)	0.08057	0.00979 (0.00082)	0.00973 (0.00058)	0.00966	-0.00094 (0.00060)	-0.00092 (0.00040)	-0.00088
90	0.17173 (0.00284)	0.17180 (0.00120)	0.17099 (0.00330)	0.17153	0.01512 (0.00107)	0.01500 (0.00080)	0.01510	-0.00295 (0.00060)	-0.00254 (0.00038)	-0.00275
100	0.26044 (0.00608)	0.25880 (0.00113)	0.25902 (0.00507)	0.25847	0.01791 (0.00087)	0.01753 (0.00030)	0.01781	-0.00518 (0.00066)	-0.00533 (0.00052)	-0.00523
110	0.31582 (0.00708)	0.31306 (0.00125)	0.31503 (0.00411)	0.31352	0.01789 (0.00100)	0.01815 (0.00088)	0.01801	-0.00662 (0.00066)	-0.00715 (0.00078)	-0.00708
120	0.34240 (0.00693)	0.33955 (0.00114)	0.33812 (0.00673)	0.33874	0.01728 (0.00114)	0.01694 (0.00084)	0.01713	-0.00743 (0.00057)	-0.00770 (0.00057)	-0.00780
130	0.35096 (0.00596)	0.34703 (0.00156)	0.34584 (0.00797)	0.34779	0.01577 (0.00110)	0.01603 (0.00096)	0.01607	-0.00740 (0.00050)	-0.00764 (0.00064)	-0.00777

This table presents estimates of sensitivities for Bermudan max-call options on two and three correlated assets. The first column represents different values for the initial price S_0 , where the initial price vector is $S(0) = (S_0, \dots, S_0)$. The other fixed parameters are $K = 100, T = 1, r = 5\%, q = 10\%, \sigma = 20\%, \rho = 0.3$. Exercise opportunities are equally spaced at times $t_i = iT/d, i = 0, 1, \dots, d$, with $d = 3$. All simulations are based on 150,000 sample paths for the stock-price process. Their respective standard errors are given in the parentheses immediately below them. The "Binomial" columns show the benchmark results for the corresponding values from the multidimensional binomial routine with $N = 600$ time steps. As shown, most of the simulated values are within one standard error of the benchmark results.

TABLE 8 Bermudan max-call option on five assets.

d	MLSM price	LRD price	MLSM Δ_1	LRD Δ_1	MLSM Γ_{11}	LRD Γ_{11}
5	25.695 to 25.805	25.735 to 25.795	0.196 to 0.209	0.197 to 0.208	0.0080 to 0.0092	0.0080 to 0.0090
10	26.083 to 26.211	26.157 to 26.233	0.198 to 0.209	0.199 to 0.212	0.0081 to 0.0094	0.0070 to 0.0088
15	26.206 to 26.372	26.304 to 26.392	0.196 to 0.210	0.191 to 0.211	0.0076 to 0.0088	0.0061 to 0.0093
20	26.307 to 26.452	26.370 to 26.480	0.194 to 0.210	0.190 to 0.222	0.0075 to 0.0093	0.0061 to 0.0117
30	26.344 to 26.517	26.445 to 26.571	0.196 to 0.211	0.175 to 0.224	0.0075 to 0.0095	0.0032 to 0.0150

This table presents 95% confidence intervals of prices and sensitivities for Bermudan max-call options on five uncorrelated assets using the MLSM and LRD methods. The number of exercise opportunities varies between $d = 5, 10, 15, 20, 30$. The initial price vector is $S(0) = (S_0, \dots, S_0)$, with $S_0 = 100$. The other fixed parameters are $K = 100, T = 3, r = 5\%, q = 10\%, \sigma = 20\%$. Our MLSM confidence intervals are given in the columns "MLSM price", "MLSM Δ_1 " and "MLSM Γ_{11} ", and the simulation is based on 200,000 sample paths for the stock-price process. They are constructed as 1.96 multiples of standard error around the MLSM estimates. As a comparison, LRD confidence intervals are constructed as $(L - 1.96\sigma_L, U + 1.96\sigma_U)$, where L, U, σ_L, σ_U are estimates for the lower and upper bounds and their associated standard errors, respectively.

method. Further, in this example the likelihood estimates do not suffer from huge variance as we might expect since there are three exercise opportunities. Other than these minor discrepancies, the differences between the “MLSM”, “Pathwise” and “Likelihood” estimates are quite small. From these results we conclude that the MLSM algorithm is able to approximate the benchmark values to a satisfactory level even for higher-dimensional cases.

As we stated in the first section, we now compare our MLSM algorithm with the LRD algorithm for the case of $n = 5$ underlying assets. 95% confidence intervals for price, Δ_1 , Γ_{11} are reported in Table 8. In applying the MLSM method, we sample only $S_1(0)$ from (16) while keeping the other $S_i(0)$ fixed at S_0 and, at the end, regress pathwise payouts on a constant and the first four power polynomials of $S_1(0)$. Due to the exponential computational complexity involved with traditional lattice methods, the LRD confidence intervals are provided here as an alternative comparison. Since the LRD algorithm treats the Bermudan option as a European option that expires on the first exercise date, its results are expected to deteriorate and have greater discrepancies as the number of exercise dates increases. This is rather clear from Table 8, where we report the values for a series of increasing exercise opportunities. It is also readily inferred from Table 8 that the MLSM algorithm dominates the LRD algorithm in that its performance remains relatively stable for various numbers of exercise dates.

6 AMERICAN-BERMUDAN-ASIAN OPTION

In this section we apply the MLSM algorithm to a more exotic path-dependent option. In particular, we consider a call option on the average price of a stock over some horizon, where the call option can be exercised at any time after some initial lockout period. Thus this option is an Asian option since it is an option on an average, and has both a Bermudan and American exercise feature. This is one of the examples studied by Longstaff and Schwartz (2001), of an American-Bermudan-Asian option, specified as follows.

Define the current valuation date as time 0. We assume that the option has a final expiration date of $T = 2$, and that the option can be exercised at any time after $t^* = 0.25$ by payment of the strike price K . The underlying average A_t , $0.25 \leq t \leq T$, is the continuous arithmetic average of the underlying stock price during the period $t_0 = -0.25$ to time t .⁹

$$A_t = \frac{\int_{-0.25}^t S_u \, du}{t + 0.25}, \quad 0.25 \leq t \leq T \quad (22)$$

Thus the cashflow from exercising the option at time t is $\max(A_t - K, 0)$. The risk-neutral dynamics for the stock price are the same as in (18).

⁹Note that $t^* = 0.25$ and $t_0 = -0.25$ represent different times and should not be confused with each other.

In Table 8, we report our results for price, $\Delta (= (\partial V / \partial S(0))(S_0, A_0, 0))$ and $\Gamma (= (\partial^2 V / \partial S(0)^2)(S_0, A_0, 0))$. For regression at intermediate times, we use a total of eight basis functions: a constant, the first two power polynomials in the stock price and the average stock price, and the cross products of these power polynomials up to third-order terms. At time 0, we again sample $S(0)$ from (16) while keeping $A(0)$ fixed at A_0 and, in the final regression, regress pathwise payouts on a constant and the first four power polynomials in $S(0)$. Thus the desired estimates are obtained by substituting and differentiating the approximated function accordingly.

To provide a benchmark result in this case, we resort to the standard finite-difference techniques for solving the partial differential equation that models the option. In general, this type of problem is very difficult to solve by finite-difference techniques since the cashflow from exercise depends on the past stock price over the averaging window. However, in this particular case, we can transform the problem from a path-dependent one to a Markovian problem by introducing the average to date A_t as a second state variable in the problem. Consequently, the option price $V(S, A, t)$ is the solution to the following two-dimensional partial differential equation:

$$(\sigma^2 S^2 / 2) V_{SS} + r V_S + \frac{1}{0.25 + t} (S - A) V_A - r V + V_t = 0 \quad (23)$$

along with the early exercise constraint:

$$V(S, A, t) \geq \max(A - K, 0), \quad 0.25 \leq t \leq T \quad (24)$$

subject to the expiration condition $V(S, A, T) = \max(A - K, 0)$. Note that the path dependence of the option payout does not pose any difficulties to the simulation-based MLSM algorithm.

As shown in Table 9 (see page 116), the pathwise method again tends to exhibit a persistently smaller standard error in reporting Δ , while the likelihood method suffers from a much bigger standard error for both Δ and Γ due to the large number of exercise dates used. In contrast, the MLSM method is relatively stable for the results throughout the table. The simulation estimates are typically within two standard errors of the finite-difference benchmark values. This is a good indication that the MLSM algorithm can also be effective in closely approximating the benchmark values for an exotic path-dependent problem.

7 AMERICAN PUT OPTION IN A JUMP-DIFFUSION MODEL

In this section, we illustrate how the MLSM approach can be applied to American options when the underlying asset follows a jump-diffusion process. In particular, we revisit the American put option considered in Section 6.

To simplify the illustration, we focus on the basic jump-to-ruin model presented in Merton (1976). In this model, the stock price follows a geometric Brownian motion as in (18) until a Poisson event occurs, at which point the stock price becomes zero.

TABLE 9 American–Bermudan–Asian option.

A_0	S_0	MLSM price	Finite-difference price	MLSM Δ	Pathwise Δ	Likelihood Δ	Finite-difference Δ	MLSM Γ	Likelihood Γ	Finite-difference Γ
90	90	3.3361 (0.0256)	3.2956	0.3451 (0.0058)	0.3463 (0.0013)	0.3377 (0.0119)	0.3446	0.0227 (0.0007)	0.0225 (0.0104)	0.0232
90	100	7.9112 (0.0296)	7.8940	0.5687 (0.0051)	0.5786* (0.0010)	0.5678 (0.0215)	0.5720	0.0206 (0.0007)	0.0171 (0.0146)	0.0212
90	110	14.5052 (0.0511)	14.5511	0.7405 (0.0049)	0.7511* (0.0011)	0.7465 (0.0290)	0.7447	0.0125 (0.0008)	0.0196 (0.0212)	0.0126
100	90	3.7229 (0.0357)	3.7066	0.3793 (0.0037)	0.3806* (0.0017)	0.3738 (0.0135)	0.3739	0.0241 (0.0007)	0.0291 (0.0126)	0.0244
100	100	8.6854 (0.0373)	8.6713	0.6099 (0.0045)	0.6259* (0.0021)	0.6039 (0.0221)	0.6146	0.0210 (0.0008)	0.0148 (0.0149)	0.0218
100	110	15.6724 (0.0417)	15.7367	0.7702* (0.0043)	0.7843* (0.0013)	0.7974 (0.0300)	0.7789	0.0110 (0.0005)	0.0117 (0.0244)	0.0103
110	90	4.2207 (0.0261)	4.1872	0.4255 (0.0043)	0.4314* (0.0020)	0.4111 (0.0177)	0.4199	0.0272 (0.0007)	0.0308 (0.0134)	0.0284
110	100	9.7935 (0.0413)	9.8367	0.6882 (0.0044)	0.6929 (0.0026)	0.6844 (0.0166)	0.6950	0.0193 (0.0008)	0.0243 (0.0144)	0.0196
110	110	17.3073* (0.0411)	17.4147	0.7917 (0.0043)	0.7952 (0.0012)	0.7931 (0.0209)	0.7949	0.0067* (0.0008)	0.0010 (0.0215)	0.0042

This table presents estimates of prices and sensitivities for American–Bermudan–Asian options. The first two columns represent different values for the initial parameters A_0 and S_0 , and the other fixed parameters are $K = 100$, $T = 2$, $r = 6\%$, $\sigma = 0.20$. All simulations are based on 100,000 sample paths for the stock-price process with 100 discretization points per year. Their respective standard errors are given in the parentheses immediately right to them. As a benchmark, the “finite-difference” columns show the numerical results for corresponding values from the finite-difference scheme with 10,000 time steps per year and 400 steps in both A and S . The simulated values are typically within two standard errors of the benchmark results; those that are two standard errors away we mark with an “*”.

The resultant risk-neutral dynamics are given by:

$$dS = (r + \lambda)S dt + \sigma S dZ - S dq \quad (25)$$

where q is an independent Poisson process with intensity λ . When a Poisson event occurs, the value of q jumps from zero to one, implying that $dq = 1$, and the stock price jumps downward from S to zero. Merton (1976) shows that the price of an American option in such a model is given by a complex mixed differential-difference equation, which is difficult to solve. As usual, this does not pose any difficulty to the simulation-based MLSM approach and, furthermore, the MLSM approach can be readily applied to much more complex jump-diffusion processes than in this example or the other examples in Merton.

To put the results into perspective, we compare the prices of the American put option for the cases where there is no possibility of a jump $\lambda = 0$ and when a jump can occur with intensity $\lambda = 0.05$. To make the comparison more meaningful, we adjust the parameters in the two cases so that the means and variances are equal. Because of the martingale restriction implied by the risk-neutral framework, the mean of the risk-neutral distribution for stock price is $S_0 \exp(rT)$ and it is the same across cases. The variance of the stock price is:

$$S_0^2 \exp(2r)(\exp((\lambda + \sigma^2)T) - 1) \quad (26)$$

Therefore, in order to equalize their first two moments, we set the parameter values to be $\lambda = 0$, $\sigma = 0.3$ and $\lambda = 0.05$, $\sigma = 0.2$. Other parameters being identical, a comparison of prices for the two cases is presented in Table 9 on page 116. To provide an additional comparison, we also report estimates for Δ from the pathwise method; however, it is in general quite difficult to apply the likelihood approach here as this example fails to have a closed-form density function for the underlying asset price. All simulations are based on 150,000 sample paths with 100 discretization points per year. Longstaff and Schwartz (2001) points out that the value of early exercise premium is typically less in the “jump” case where $\lambda = 0.05$, $\sigma = 0.2$, so there is less incentive to keep the option alive. They explain that this is because “the windfall gain to the option holder from a jump does not offset the effects of a lower diffusion coefficient λ ”.

As with the other examples, we provide a benchmark for the MLSM values to demonstrate the effectiveness of our MLSM algorithm. Among the many established techniques available in the literature, we adopt the generalized binomial model for jump-diffusion processes proposed by Amin (1993) to produce the desired benchmark values. These values are also reported in Table 10 (see page 118) and closely approximated by the corresponding MLSM values as expected.

8 NUMERICAL AND IMPLEMENTATION ISSUES

In this section, we discuss in detail a number of numerical and implementation issues that are associated with the MLSM algorithm. These are divided into three major

TABLE 10 American put options: jump-diffusion model.

S_0	T	No-jump price	MLSM price	Binomial price	MLSM Δ	Pathwise Δ	Binomial Δ	MLSM Γ	Binomial Γ
36	1	5.7292 (0.0158)	5.0689 (0.0233)	5.0837	-0.5641 (0.0087)	-0.5478 (0.0044)	-0.5526	0.0861 (0.0034)	0.0791
36	2	6.6262 (0.0272)	6.2901 (0.0347)	6.3132	-0.4104 (0.0080)	-0.3968 (0.0051)	-0.3988	0.0619 (0.0033)	0.0554
38	1	4.6806 (0.0210)	4.1087 (0.0279)	4.1271	-0.4134 (0.0083)	-0.4057 (0.0030)	-0.4098	0.0689 (0.0042)	0.0638
38	2	5.6861 (0.0149)	5.5938 (0.0309)	5.6168	-0.3127 (0.0108)	-0.2977 (0.0040)	-0.3014	0.0474 (0.0025)	0.0426
40	1	3.8100 (0.0159)	3.4122 (0.0291)	3.4256	-0.2971 (0.0098)	-0.2951 (0.0025)	-0.2964	0.0520 (0.0036)	0.0500
40	2	4.8819 (0.0142)	5.0785 (0.0520)	5.0917	-0.2278 (0.0080)	-0.2247 (0.0024)	-0.2266	0.0357 (0.0032)	0.0326
42	1	3.0741 (0.0163)	2.9200 (0.0311)	2.9240	-0.2097 (0.0085)	-0.2081 (0.0022)	-0.2090	0.0383 (0.0028)	0.0377
42	2	4.1990 (0.0209)	4.6912 (0.0409)	4.6982	-0.1726 (0.0076)	-0.1670 (0.0016)	-0.1694	0.0259 (0.0026)	0.0248
44	1	2.4751 (0.0160)	2.5563 (0.0162)	2.5740	-0.1441 (0.0057)	-0.1429 (0.0017)	-0.1440	0.0285 (0.0024)	0.0276
44	2	3.6009 (0.0132)	4.3988 (0.0319)	4.4049	-0.1244 (0.0090)	-0.1250 (0.0012)	-0.1260	0.0195 (0.0015)	0.0188

This table presents a comparison of prices and hedging parameters for the American put option under a jump-diffusion model. We also report price estimates for the non-jump model as a comparison. We equalize the means and variances in the two cases by setting $\lambda = 0$, $\sigma = 0.3$ and $\lambda = 0.05$, $\sigma = 0.2$. The first two columns represent different values for the parameters S_0 and T , and the other fixed parameters are $K = 40$, $r = 6\%$. The simulation is based on 150,000 sample paths for the stock-price process with 100 discretization points per year. Their respective standard errors are given in the parentheses below them. We report "Binomial" values as benchmarks in the table, which are calculated from the binomial model for the jump-diffusion processes proposed by Amin (1993). The binomial model is built with 600 time steps per year. As shown, all simulated values are within two standard errors of the corresponding benchmark results.

categories and discussed individually below. Much like its successful predecessor, the LSM method, the MLSM method shares virtually all the good qualities that the LSM method possesses.

8.1 Choice of basis functions

Extensive numerical tests indicate that the results from the MLSM algorithm are remarkably robust to the choice of basis functions. For instance, we use the first four power polynomials as basis functions in the American put example in Section 6. We obtain results that are virtually identical to those reported in Table 5 when we use the first four Laguerre polynomials as basis functions, when we use the first four Hermite polynomials as basis functions or when we use four trigonometric functions as basis functions.¹⁰ This is also true for all the other examples presented in this paper.

Here we point out that the MLSM algorithm is primarily intended to calculate option price derivatives that are more sensitive to the approximated option value function than the option prices themselves. Therefore, we recommend using one or two more basis functions in implementing the MLSM method than would usually be necessary in implementing the LSM method to allow for a better approximation of the option value function. For example, Longstaff and Schwartz (2001) report that, in applying the LSM method to the American put example, using only the first three polynomials would be sufficient, while we choose to use the first four polynomials in applying the MLSM method.

We also recommend normalizing appropriately to avoid potential numerical errors resulting from scaling problems. This is because for certain types of polynomials, such as weighted Laguerre polynomials, directly applying them to the problems could result in a scaling mismatch, further leading to computational underflows. To avoid this problem, we renormalize the regressions by dividing all the cashflows and prices by the strike price, and estimating the conditional expectation function in the renormalized space; all of the results reported in this paper are based on this renormalization.

8.2 Choice of initial distribution

As we have argued in Section 5, the conditional expectation we are trying to estimate, $F(\omega, 0) = E[C(\omega, \tau(t_1)) | X(\omega, 0)]$, is actually a deterministic function of X , and should be independent of the distribution for $X(\omega, 0)$ we use in the regression. For concerns on simplicity and consistency, we reuse the initial distribution specified in (16) for all of our examples in this paper. Extensive experiments show that a range of values for α from (0.1, 1) all do very well in producing a reasonably satisfactory

¹⁰A detailed illustration and comparison of these different basis functions can be found in Moreno and Navas (2003).

result based on a sufficiently large sample of stock price paths. However, significant biases in the results do arise when too small or too large a value for α is chosen. We believe this is because too small a value for α would produce a layout of initial prices too concentrated near S_0 and hence do a poor job in extracting the information of a neighborhood of S_0 for evaluation of the price sensitivities; meanwhile, too large a value for α would cause the initial prices to spread out excessively, thus undermining the resolution of the final regression.

8.3 Computational speed

First, we point out that the MLSM method differs little from the LSM method in terms of computational cost and speed, or we could say that the MLSM method is as fast as the LSM method. For example, in applying the LSM method, suppose we use 10,000 paths for a stock price with 100 discretization points for each path. That is, we have to simulate $10,000 \times 100$ normal variates to construct the paths, and do 99 regressions for the algorithm in order to calculate one estimate of the option value. To obtain the same level of accuracy in implementing the MLSM method, we need to use 10,000 paths for a stock price with 101 discretization points for each path (one additional dimension for the random initial prices). Then we have to simulate $10,000 \times 101$ normal variates to construct the paths, and do 100 regressions (one additional regression is conducted at time 0) for the algorithm in order to calculate one estimate of option value and other derivatives. Our experiments show that the differences in computational effort between the two algorithms are so small as to be negligible.

One important advantage of Monte Carlo simulation techniques is that they lend themselves well to parallel computing architecture. From the perspective of the MLSM method, the only constraint on parallel computation is that the regression needs to use the cross-sectional information in the simulation. To overcome this bottleneck, there are many ways in which regressions could be estimated using individual CPUs, and then aggregated across CPUs to form a composite estimate of the conditional expectation function. Furthermore, it may be promising to use quasi-Monte-Carlo techniques in conjunction with the MLSM algorithm to significantly improve its computational speed and efficiency. Important recent efforts in this direction include Cafilisch and Chaudhary (2004) and Chaudhary (2005), where they have developed various approaches to incorporate quasi-Monte-Carlo techniques into the LSM algorithm.

8.4 Calculating other risk sensitivities

At present the MLSM method is focused on estimating risk sensitivities of the option price to initial spot prices. Computing Greeks with respect to other parameters, like vega, falls out of the current scope of our MLSM method. However, we believe that the MLSM idea could be slightly extended to handle relevant issues like this.

Actually, one possible strategy for extending the MLSM method to estimate vega we now have in mind is to simulate a different value of volatility for each stock path from a carefully chosen distribution, and eventually compute a regression function of volatility values by regressing pathwise payouts against pathwise volatilities. We feel that this would be one interesting direction to explore for future research.

9 CONCLUSION AND FUTURE DIRECTIONS

This paper presents a simple new approach for approximating the values and hedging parameters of American-style options by simulation. This approach is intuitive, accurate, easy to implement and computationally efficient. We illustrate this technique using a number of increasingly complex but realistic examples, including a Bermudan max-call option on multiple assets, an exotic American–Bermudan–Asian option and an American put option when the underlying asset follows a jump-diffusion process.

At present, three future directions seem promising with the MLSM algorithm. First, as Chaudhary (2005) has already successfully incorporated the use of quasi-random sequences and Brownian bridges into the LSM algorithm, we believe it will also be fruitful to try similar ideas to greatly speed up the MLSM algorithm. Second, the MLSM algorithm is ready to be applied to more complex options than the ones we have presented in this paper. For example, it could be used to calculate the price sensitivities of a Bermudan swaption to different forward rates on the initial term structure.¹¹ Third, it looks promising to modify the MLSM approach further to be able to estimate risk sensitivities for American-style options under a stochastic volatility setting. Partial results can be found in Stentoft (2005).

¹¹Longstaff and Schwartz (2001) conduct an in-depth analysis of this example using the LSM algorithm.

APPENDIX A INITIAL REGRESSION: RANDOM DISTRIBUTION VERSUS DETERMINISTIC GRID

K	σ	T	Random distribution		Deterministic grid		Binomial price		Random distribution		Deterministic grid		Binomial Δ		Random distribution		Deterministic grid		Binomial Γ	
			price	Γ	price	Δ	price	Δ	price	Δ	price	Γ	price	Δ	price	Γ	price	Γ	price	Γ
35	0.20	1/3	0.1991 (0.0026)	0.1996 (0.0028)	0.2004	-0.0903 (0.0012)	-0.0910 (0.0030)	-0.0901	0.0367 (0.0012)	0.0378 (0.0071)	0.0357	0.0378 (0.0071)	0.0367 (0.0012)	0.0378 (0.0071)	0.0357	0.0367 (0.0012)	0.0378 (0.0071)	0.0357	0.0367 (0.0012)	0.0378 (0.0071)
35	0.20	7/12	0.4301 (0.0031)	0.4333 (0.0058)	0.4328	-0.1346 (0.0012)	-0.1348 (0.0043)	-0.1338	0.0373 (0.0013)	0.0385 (0.0063)	0.0364	0.0373 (0.0013)	0.0385 (0.0063)	0.0364	0.0373 (0.0013)	0.0385 (0.0063)	0.0364	0.0373 (0.0013)	0.0385 (0.0063)	
40	0.20	1/3	1.5786 (0.0071)	1.5755 (0.0108)	1.5798	-0.4434 (0.0029)	-0.4441 (0.0128)	-0.4435	0.0930 (0.0024)	0.0966 (0.0155)	0.0923	0.0930 (0.0024)	0.0966 (0.0155)	0.0923	0.0930 (0.0024)	0.0966 (0.0155)	0.0923	0.0930 (0.0024)	0.0966 (0.0155)	
40	0.20	7/12	1.9848 (0.0086)	1.9881 (0.0113)	1.9904	-0.4287 (0.0025)	-0.4292 (0.0091)	-0.4287	0.0730 (0.0020)	0.0771 (0.0164)	0.0719	0.0730 (0.0020)	0.0771 (0.0164)	0.0719	0.0730 (0.0020)	0.0771 (0.0164)	0.0719	0.0730 (0.0020)	0.0771 (0.0164)	
45	0.20	1/3	5.0942 (0.0073)	5.0881 (0.0092)	5.0883	-0.8848 (0.0039)	-0.8819 (0.0117)	-0.8812	0.0811 (0.0015)	0.0835 (0.0158)	0.0827	0.0811 (0.0015)	0.0835 (0.0158)	0.0827	0.0811 (0.0015)	0.0835 (0.0158)	0.0827	0.0811 (0.0015)	0.0835 (0.0158)	
45	0.20	7/12	5.2722 (0.0056)	5.2670 (0.0135)	5.2670	-0.7999 (0.0033)	-0.7931 (0.0110)	-0.7948	0.0736 (0.0012)	0.0784 (0.0124)	0.0787	0.0736 (0.0012)	0.0784 (0.0124)	0.0787	0.0736 (0.0012)	0.0784 (0.0124)	0.0787	0.0736 (0.0012)	0.0784 (0.0124)	
35	0.30	1/3	0.6972 (0.0059)	0.6991 (0.0082)	0.6975	-0.1745 (0.0014)	-0.1744 (0.0056)	-0.1741	0.0377 (0.0013)	0.0353 (0.0072)	0.0376	0.0377 (0.0013)	0.0353 (0.0072)	0.0376	0.0377 (0.0013)	0.0353 (0.0072)	0.0376	0.0377 (0.0013)	0.0353 (0.0072)	
35	0.30	7/12	1.2229 (0.0048)	1.2168 (0.0121)	1.2198	-0.2135 (0.0020)	-0.2139 (0.0063)	-0.2126	0.0321 (0.0015)	0.0333 (0.0061)	0.0326	0.0321 (0.0015)	0.0333 (0.0061)	0.0326	0.0321 (0.0015)	0.0333 (0.0061)	0.0326	0.0321 (0.0015)	0.0333 (0.0061)	

APPENDIX A Continued

K	σ	T	Random distribution price	Deterministic grid price	Binomial price	Random distribution Δ	Deterministic grid Δ	Binomial Δ	Random distribution Γ	Deterministic grid Γ	Binomial Γ
40	0.30	1/3	2.4808 (0.0090)	2.4835 (0.0145)	2.4825	-0.4414 (0.0029)	-0.4418 (0.0098)	-0.4420	0.0591 (0.0016)	0.0612 (0.0105)	0.0597
40	0.30	7/12	3.1678 (0.0132)	3.1711 (0.0154)	3.1696	-0.4265 (0.0027)	-0.4233 (0.0077)	-0.4256	0.0463 (0.0019)	0.0437 (0.0077)	0.0459
45	0.30	1/3	5.7012 (0.0152)	5.7032 (0.0191)	5.7056	-0.7266 (0.0042)	-0.7288 (0.0084)	-0.7266	0.0576 (0.0014)	0.0558 (0.0129)	0.0572
45	0.30	7/12	6.2318 (0.0111)	6.2390 (0.0142)	6.2436	-0.6537 (0.0023)	-0.6513 (0.0086)	-0.6520	0.0497 (0.0012)	0.0493 (0.0104)	0.0485

This table presents a comparison of price estimates for the American put option under two different specifications for the spot price layout in the initial regression: random distribution and deterministic grid. The random distribution is specified in (13) with the coefficient α set to be 0.5, while an evenly dispersed grid (with equal weights) over the interval $(S(0) e^{-\alpha\sigma\sqrt{T}}, S(0) e^{\alpha\sigma\sqrt{T}})$ is used as the deterministic grid. The first three columns represent different values for the parameters K , σ and T , and the other fixed parameters are $S_0 = 40$, $r = 4.88\%$. The simulations are all based on 150,000 sample paths for the stock-price process with 150 discretization points per year. Their respective standard errors are given in the parentheses immediately below. The "Binomial" column shows the benchmark results for corresponding values from the standard binomial model with $N = 10,000$ time steps.

REFERENCES

- Amin, K. I. (1993). Jump diffusion option valuation in discrete time. *Journal of Finance* **48**(5), 1833–1863.
- Andersen, L. (2000). A simple approach to the pricing of Bermudan swaptions in the multi-factor Libor market model. *The Journal of Computational Finance* **3**(2), 1–32.
- Andersen, L., and Broadie, M. (2004). Primal-dual simulation algorithm for pricing multidimensional American options. *Management Science* **50**(9), 1222–1234.
- Broadie, M., and Glasserman, P. (1996). Estimating security price derivatives using simulation. *Management Science* **42**, 262–285.
- Broadie, M., and Glasserman, P. (1997). Pricing American-style securities using simulation. *Journal of Economic Dynamics and Control* **21**(8–9), 1323–1352.
- Broadie, M., and Glasserman, P. (2004). A stochastic mesh method for pricing high-dimensional American options. *The Journal of Computational Finance* **7**(4), 35–72.
- Boyle, P., Evnine, J., and Gibbs, S. (1989). Numerical evaluation of multivariate contingent claims. *Review of Financial Studies* **2**, 241–250.
- Caflisch, R., and Chaudhary, S. (2004). Monte Carlo methods for American options. *Proceedings of the 2004 Winter Simulation Conference*, Vol. 2, pp. 1656–1660.
- Carriere, J. (1996). Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics* **19**, 19–30.
- Chaudhary, S. (2005). American options and the LSM algorithm: quasi-random sequences and Brownian bridges. *The Journal of Computational Finance* **8**(4), 101–115.
- Fournié, E., Lasry, J. M., Lebuchoux, J., Lions, P. L., and Touzi, N. (1999). Applications of Malliavin calculus to Monte Carlo methods in Finance. *Finance and Stochastics* **3**, 391–412.
- Giles, M. B. (2007). Monte Carlo evaluation of sensitivities in computational finance. Working Paper, The Oxford-Man Institute, University of Oxford.
- Giles, M. B., and Glasserman, P. (2006). Smoking adjoints: fast evaluation of Greeks in Monte Carlo calculations. *Risk* **19**(1), 88–92.
- Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer-Verlag, New York.
- Glynn, P. W. (1989). Optimization of stochastic systems via simulation. *Proceedings of 1989 Simulation Conference*. Society for Computer Simulation, San Diego, CA, pp. 90–105.
- Kaniel, R., Tompaidis, S., and Zemliyanov, A. (2007). Efficient computation of hedging parameters for discretely exercisable options. *Operations Research* **55**(6), 811–826.
- Longstaff, F. A., and Schwartz, E. S. (2001). Valuing American options by simulation: a simple least-square approach. *Review of Financial Studies* **14**, 113–147.
- Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* **3**, 125–144.
- Moreno, M., and Navas, J. (2003). On the robustness of least-squares Monte Carlo (LSM) for pricing American derivatives. *Review of Derivatives Research* **6**(2), 107–128.

- Pelsser, A. A. J., and Vorst, A. C. F. (1994). The binomial model and the Greeks. *Journal of Derivatives* **1994**(Spring), 45–49.
- Piterbarg, V. V. (2004). Computing deltas of callable Libor exotics in forward Libor models. *The Journal of Computational Finance* **7**(3), 107–144.
- Piterbarg, V. V. (2005). Pricing and hedging callable Libor exotics in forward Libor models. *The Journal of Computational Finance* **8**(2), 65–117.
- Rogers, L. C. G. (2002). Monte Carlo valuation of American options. *Mathematical Finance* **12**, 271–286.
- Stentoft, L. (2004a). Assessing the least squares Monte-Carlo approach to American option valuation. *Review of Derivatives Research* **7**(2), 129–168.
- Stentoft, L. (2004b). Convergence of the least squares Monte Carlo approach to American option valuation. *Management Science* **50**(9), 1193–1203.
- Stentoft, L. (2005). Pricing American options when the underlying asset follows GARCH processes. *Journal of Empirical Finance* **12**(4), 576–611.
- Tilley, J. A. (1993). Valuing American options in a path simulation model. *Transactions of the Society of Actuaries* **45**, 83–104.
- Tsitsiklis, J., and Van Roy, B. (1999). Optimal stopping of Markov processes: Hilbert space theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives. *IEEE Transactions on Automatic Control* **44**(10), 1840–1851.
- Tsitsiklis, J., and Van Roy, B. (2001). Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks* **12**(4), 694–703.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.