A Monte Carlo Approach to Price American-Bermudan-Style Derivatives*

Javier Neves da Silva[1,2], Luis Fernández[3]

[1]HEC Lausanne, [2]Swissquote Bank, Qam@swissquote.ch
[3]Swissquote Bank, Luis.Fernandez@swissquote.ch

Abstract We study the performance of a simulation-based method to price American-Bermudan derivatives, using the approach of Longstaff and Schwartz (2001). This method, called Least Squares Monte Carlo, regresses the simulated future cash flows on a set of basis functions to compute the conditional expectation of the cash flows while holding the derivative. The computation is performed at each time step to decide whether to exercise the option or not by comparing it to the intrinsic value; this determines the estimation of the optimal stopping rule. We present results of the timing estimation performance and of the pricing accuracy for various derivative contracts—including path-dependent derivatives and multi-asset options—for different settings of the relevant parameters, and show that this method provides a reliable and efficient way to deal with American-Bermudan derivatives.

* This work is adapted from Javier Neves da Silva's Master thesis. The project was carried out at Swissquote Bank in 2008 under the supervision of Professor Micheal Rockinger, MScF Program Director, HEC, Lausanne.

1 Introduction

The aim of this study is to investigate the reliability and efficiency of a simulation-based method, called Least Squares Monte Carlo (LSMC), in order to price American-style options (see [5]). The main advantage of such method being its flexibility and robustness, it can be easily applied to price derivatives in a multi-factor setting. Moreover, simulation-based methods become much more efficient than numerical ones, e.g. finite difference, when several risk factors influence the value of a derivative. The value at time t of a European option maturing at T is given as:

$$V\left(S,t\right)=E\left[e^{-r(T-t)}h\left(S_T\right)\right],$$

where $h\left(S_T\right)$ is the payoff of the option at maturity, r is the constant risk-free rate, and E denotes the expectation with respect to the risk-neutral measure.

Closed-form formulas exist for some European-style options. This is not the case of American-style options, which are more complex to price, due to their early exercise feature. Bensoussan (1984) and Karatzas (1988) proved that, in order to eliminate any possibility of arbitrage, the value at time t of an American option maturing at T must be given by:

$$V\left(S,t\right)=\sup_{\tau}E\left[e^{-r(\tau-t)}h\left(S_\tau\right)\right],$$

where τ is the stopping time.

An American option must always be at least worth its intrinsic value, corresponding to the payoff from exercising immediately the option. Otherwise there would be an arbitrage opportunity.

We will focus on equity derivatives exclusively, though the method described can easily be adapted to other processes. The dynamics of the stock price is taken under the risk-neutral measure. Under that measure, all financial assets have the same expected return, which is the risk-free interest rate. The existence of a unique risk-neutral measure ensures a unique arbitrage-free price for each asset in the market, which is guaranteed under geometric Brownian motion.

In the case of a single asset, the stock price at a given time $t+s$ reads:

$$S\left(t+s\right)=S\left(t\right)\exp\left[\left(r-D-\frac{\sigma^2}{2}\right)s+\sigma\sqrt{s}Z\left(t\right)\right],$$

where r is the risk-free rate, D the dividend yield of the underlying, $S(t)$ the stock price at time t, σ the volatility of the underlying stock, and $Z\sim N(0,1)$ a normally distributed variable with zero mean and unit variance. In this way, the stock price is log-normally distributed.

In the case of multiple correlated assets, the dynamics of the stock price of each underlying asset is slightly modified as correlated geometric Brownian motions have to be generated.

The methodology to generate correlated stock prices relies on the fact that a linear combination of normal variables is still normal. The procedure to generate random samples of correlated variables $X \sim MN(0,\Sigma)$—$MN(0,\Sigma)$ being the multinormal distribution with covariance matrix Σ—is as follows:

- Generate a set of normally distributed random variables Z_i for each of the assets;
- Find the Choleski decomposition C of Σ;
- Compute the random sample of correlated variables X as: $X = C^T Z$.

Once this is done, the dynamics of the asset i is given by:

$$S_i\,(t+s) = S_i\,(t)\exp\left[\left(r - D_i - \frac{\sigma_i^2}{2}\right)s + \sqrt{s}X_i\,(t)\right].$$

Note that the Choleski decomposition is not possible in the extreme cases $\rho_{ij} = \pm 1$, where $\rho_{ij} = \frac{\Sigma_{ij}}{\sigma_i \sigma_j}$ is the correlation between two assets i and j, because the correlation matrix is not positive semi-definite.

Note that all simulations are run on a computer with 2.4 GHz and 3.6 GB RAM using the software Scilab.

2 Simulation-Based Method

The Monte Carlo approach to pricing European-style derivatives consists in approximating the European option value by the discounted average taken over all simulations of a given process:

$$V_0 = \exp(-r\,T)\sum_{i=1}^{M}\frac{V_{T,i}}{M},$$

where $V_{T,i}$ is the cash flow at maturity for the simulated path i, and M is the total number of simulations performed.

By the law of large numbers, the result obtained from Monte Carlo simulations converges to the theoretical value of the expectation as the number of simulations tends to infinity.

2.1 Least Squares Monte Carlo Approach

The valuation of American-style derivatives is one of the difficult challenges in finance. The difficulty comes from the necessity of determining the optimal exercise date of these derivatives. To do so, the holder of an American-style derivative compares, at each time t, the payoff from immediate exercise with the expected payoff from continuation, and she exercises the derivative only if the former is larger than the latter.

In 2001, Longstaff and Schwartz proposed a method relying on Monte Carlo, called Least Squares Monte Carlo (LSMC), for valuing American options. The idea consists in estimating the conditional expectation of the payoff from continuing to keep the derivative alive from the cross-sectional information contained in the simulation by using least squares. An important point is that only paths where the option is in-the-money should be considered to perform the regressions, since the exercise decision is only relevant in such cases as an out-of-the-money option would not be exercised anyway. This significantly increases the efficiency of the algorithm as it limits the region over which the conditional expectation must be estimated. Therefore, far fewer basis functions are needed to obtain an accurate approximation of the conditional expectation function.

More precisely, the approach is as follows:

1. Generate independent stock paths.

2. Start at maturity T and exercise the option only if it is in-the-money. Note that no regression is needed at the moment.

3. Start rolling back to the previous time steps.

4. For times t prior to maturity, regress the discounted subsequent realized cash flows from continuation, called Y_t, on a set of n basis functions B_i of the values of the relevant state variables, namely the actual stock prices S_t at time t. This will return an efficient unbiased estimate of the conditional expectation of the payoff from continuing to keep the derivative alive, called $E_t[Y_t|S_t]$, approximated as:

   $$E_t[Y_t|S_t] = \sum_{i=0}^{n} a_i \, B_i \left(S_t\right).$$

   A simple example would be to use:

   $$E_t[Y_t|S_t] = a + b \, S_t + c \, S_t^2.$$

5. Compare these conditional expectations to the intrinsic values $h(S_t)$ of the option at time t, corresponding to the payoff from immediate exercise, and exercise the derivative if the intrinsic value is larger than the conditional expectation from continuation. This determines the optimal stopping rule for each simulated path of the derivative at time t. Mathematically, the value of the derivative at each time step reads:

$$V_t(S_t) = \max\left(h(S_t)\,,\; E_t[Y_t|S_t]\right).$$

6. Once the derivative is exercised for a given path, set all subsequent cash flows to zero, as the derivative can only be exercised once.

7. Once the optimal stopping rule is determined for each of the N paths, the initial value of the derivative V_0 is the average over all paths of the cash flows, denoted by CF_i, at the optimal stopping times τ_i, each discounted to time zero. The value of the derivative is then given by:

$$V_0 = \frac{1}{N}\sum_{i=1}^{N}\exp(-r\,\tau_i)\,CF_i.$$

Let the true value of the American derivative be $V(X)$, ω_i denote a simulated path, n the number of basis functions, K the number of exercise dates, $LSMC(\omega, n, K)$ the discounted cash flow resulting from following the LSMC rule of exercising, and N the number of paths. According to Longstaff and Schwartz (2001), the following inequality holds almost surely:

$$V(X) \geq \lim_{N\to\infty}\frac{1}{N}\sum_{i=1}^{N} LSMC(\omega_i,\, n,\, K).$$

This result implies that the LSMC algorithm provides a lower bound to the true value of an American derivative. The intuition from this result is that the true value of the American derivative is the one whose stopping rule maximizes the value of the derivative. Therefore, as the LSMC approximates the optimal stopping rule, the derivative value obtained with this method cannot be larger than the true value of the derivative, at least when the number of simulations is sufficiently large.

The main advantage of this approach is that it is based on simulation, which makes it readily applicable to path-dependent and multi-factor situations where traditional finite difference techniques are difficult to implement.

When several sources of risk are involved, the basis functions should depend on all sources of risk, including cross-products between them. Therefore, the number of basis functions can rise a lot as the number of sources of risk increases. In such cases, it is necessary to examine which basis functions are really important and which can be neglected.

2.2 Basis Functions

Besides simple powers of the state variables, it is always possible to regress the realized cash flows on other types of basis functions to approximate the conditional expectation function. Thus, several kinds of polynomials were considered: Hermite, Laguerre, simple powers of the explanatory variables, Chebyshev, and B-Splines.

According to Glasserman and Yu (2004), the number of basis functions—in the case of an asset following a geometric Brownian motion—is given by $O\left(\sqrt{\log(N)}\right)$. The number of basis functions should not be greater than 3 or 4 for a derivative affected by only one source of risk.

2.3 Alternative to the LSMC

Glasserman and Yu (2004) investigated a slightly different approach to the least squares method of Longstaff and Schwartz that consists in performing the regressions one step ahead. This implies that both the state variable (independent) and the response variable (dependent) are taken from one step ahead. This method implies that the values of the state variable are known at the next time step, and will thus yield an upper bound. However, the estimates produced would present a lower dispersion than those of the LSMC method.

3 Results of the LSMC Method

The accuracy and efficiency of the method is investigated using several Bermudan-style instruments, amongst which Vanilla put options, Asian options, and Max options on two and five underlying assets, whose results are presented hereafter.

3.1 American–Bermudan Put Options

Several tests were performed on Vanilla Put options. These tests have demonstrated that the LSMC method converges to the correct solution. The results are not sensitive to the kind of polynomial used. In addition, as noted by Glasserman and Longstaff and Schwartz, few basis functions are needed to approximate the conditional expectation of the payoff from continuation. The best combination used $50'000$ paths with 4 basis functions. However, considering the trade off between accuracy and timing performance, $10'000$ paths with 4 basis functions is the optimal combination.

The tests revealed that using polynomials other than B-Splines as basis functions is much more efficient. This is due to the fact that in the case of B-Splines the number of basis functions not only depends on the order of the polynomials, but also on the number of knots chosen. Moreover, B-Spline basis functions have to be determined iteratively, and each point in the matrix containing the basis functions has to be determined separately.

The most accurate results—lowest residuals in percent—are obtained for in-the-money options. This is intuitive as more information is available to perform the regressions. The computational time is lower for out-of-the-money options than for in-the-money options, because less paths are used to perform the regressions.

Results are not very sensitive to maturity or volatility, even if residuals tend to be somewhat lower in the case of short maturities as there is less uncertainty and less regressions have to be performed—implying less cumulative errors as the conditional expectation is only an approximation. Note that this is not true for residuals in percent, because options with short maturities are cheaper than those with longer maturities.

Except for B-Splines, the results are less accurate when using all the paths to determine the regression coefficients than those obtained when using only in-the-money paths, for a given number of basis functions. To get the same accuracy using all the paths as that obtained using only in-the-money paths, it is necessary to increase the number of basis functions by a factor 2 or 3. This is not optimal as it increases the computational time without improving significantly the results.

3.2 American–Bermudan Asian Options

We considered an Asian option where the payoff is $\left(\overline{A_t} - K\right)^+$, where $\overline{A_t} = \frac{1}{0.25+t}\int_{-0.25}^{t} S_t\,dt$ is the average stock price computed over the period starting three month prior to the valuation date of the option up to time t. Moreover, this option can only be exercised after 3 months. The features of the Asian options are given in Table 1.

K	T	σ	r	D
100	2	0.2	0.06	0

Table 1: *Features of the Asian option.*

The number of exercise dates is taken to be 100 per year. Several Asian options will be priced with different initial average stock prices $\overline{A_0}$, computed over the period three months prior to the valuation date up to the valuation date, and different initial stock prices S_0. The initial average $\overline{A_0}$ can take three values, namely 90 for the out-of-the-money (OTM) case, 100 for the at-the-money (ATM) case, or 110 for the in-the-money (ITM) case, while the initial stock price S_0 can also take three values, namely 80, 100, and 120, which gives nine different options to be priced with the LSMC method.

In this case, there are two sources of risk. Therefore, the basis functions should depend on both the underlying stock price S and its average value $\overline{A}$.

Table 2 shows the reference values of each of these Asian options, as quoted in [5], and denoted by LSFD. These results were obtained using a finite difference method.

	$S_0 = 80$			$S_0 = 100$			$S_0 = 120$		
$\overline{A_0}$	90	100	110	90	100	110	90	100	110
LSFD	0.95	1.11	1.29	7.89	8.66	9.82	22.42	23.81	25.45

Table 2: *Asian option values from [5].*

Each LSMC simulation is repeated 50 times. The results displayed in Table 3 are the American values of the options A, and the standard deviations of the American values σ_A. Table 3 gives the results obtained when using 8 basis functions up to third order terms—$1, S, A, SA, S^2, A^2, S^2A, SA^2$—and $10'000$ paths.

	$S_0 = 80$			$S_0 = 100$			$S_0 = 120$		
$\overline{A_0}$	90	100	110	90	100	110	90	100	110
A	0.95	1.09	1.25	7.81	8.57	9.72	22.23	23.62	$\overline{A_0}$
σ_A	0.035	0.036	0.038	0.066	0.066	0.063	0.034	0.035	0.039

Table 3: *Asian option value from LSMC with 8 basis functions (10'000 paths).*

These 8 basis functions seem to be a good combination as they lead to results that are compatible with those of Table 2. Moreover, the results are always below those of Table 2. These simulations take about 45 seconds.

3.3 2D American–Bermudan Max Options

We next consider a 2D Max option with a payoff of $(\max(S_1, S_2) - K)^+$, where S_1 and S_2 are the values of both underlying stock prices. The features of the 2D Max options studied in this section are given in Table 4.

K	T	σ_1	σ_2	r	D_1	D_2
100	3	0.2	0.2	0.05	0.1	0.1

Table 4: *Features of the max option.*

Note that D_1 and D_2 are the dividend yields of each underlying. The initial values of each underlying stock price $S_{0,1}$ and $S_{0,2}$ are always considered to be equal and with same weights, but three different scenarios will be considered:

- $S_{0,1} = S_{0,2} = 90$ (OTM);
- $S_{0,1} = S_{0,2} = 100$ (ATM);
- $S_{0,1} = S_{0,2} = 110$ (ITM).

Concerning the correlation between both underlying assets, several cases are considered, namely $\rho = \pm 1$, $\rho = \pm 0.5$ and $\rho = 0$.

The cases with 3 and 50 possible exercise dates each year will be considered. Note that 20'000 paths are used for each of the 50 repeated simulations.

3.3.1 Optimal Basis Functions for the 2D Max Options

Let us first consider the case with 3 possible exercise dates each year. As quoted in [3], the reference values V of such American–Bermudan options, assuming no correlation between both underlying assets, are:

- $V_{ATM} = 13.9$;
- $V_{OTM} = 8.08$;
- $V_{ITM} = 21.34$.

The results obtained with the LSMC method are given in Table 5—σ represents the standard deviation of the results of the 50 repeated simulations. Two cases are considered, namely the case with 8 or 9 basis functions. The combinations of basis functions are:

- 8 basis functions: $1, S_1, S_2, S_1S_2, S_1^2, S_2^2, S_1^2S_2, S_1S_2^2$.
- 9 basis functions: $1, S_1, S_2, S_1S_2, S_1^2, S_2^2, S_1^2S_2, S_1S_2^2, S_1^2S_2^2$.

Note that using the maximum between the two options as basis functions does not yield more accurate results.

The combination with 8 basis functions already gives pretty accurate results. Adding more basis functions does not improve the results significantly. Concerning the computational time, it takes less than a second to compute the American–Bermudan option value with either kind of polynomials, which shows that the LSMC can be very efficient to price Bermudan options with a few exercise dates.

	OTM		ATM		ITM	
	8 b. f.	9 b. f.	8 b. f.	9 b. f.	8 b. f.	9 b. f.
V	8.11	8.11	13.93	13.93	21.38	21.37
σ	0.061	0.06	0.087	0.087	0.087	0.088

Table 5: *Max option values V with 8 and 9 basis functions ($\rho = 0$ and 3 exercise dates per year).*

3.3.2 2D LSMC vs 1D LSMC

The idea is to compare the results obtained when using both underlying assets and their cross-products as basis functions—2D LSMC with 8 basis functions—to price the 2D American Max option, or when simple powers of an index, containing both underlying assets, are used as basis functions—1D LSMC with 4 basis functions.

The index I is determined as follows:

- Simulate several paths for each underlying separately, as in the usual Monte Carlo framework;
- The index for path j and time k is simply: $I\,(j,k) = \max\,[S_1(j,k), S_2(j,k)]$.

In this 1D LSMC setting, the combination of basis functions used is:

- 4 basis functions: $1, I, I^2, I^3$.

The features of the Bermudan 2D Max options are the same as those of the Max options studied in the previous section. The only difference is that the option will have 50 possible exercise dates per year instead of 3. In addition, several scenarios with different correlations ρ are considered. Recall that the extreme cases $\rho = \pm 1$ have to be approximated by $\rho = \pm 0.9999$, so that the correlation matrix has a Cholesky decomposition. The results obtained are presented in Tables 6, 7, and 8 for respectively the at-the-money, in-the-money, and out-of-the-money cases when using $20'000$ simulations and repeating each simulation 50 times. Note that these results will then be compared to those obtained with a 2-dimensional explicit finite difference method.

Concerning the accuracy of the results, at least for a Max option on two assets, the difference when using or not the index is most of the times lower than 1%. The time needed to compute the value of the American–Bermudan Max option is about 30 seconds (2D LSMC). The time gain when using the index is not negligible—about 10%—and therefore using an index is more efficient in this case.

		2D LSMC	1D LSMC
$\rho = 0.9999$	V_{ATM}	8.1781	8.1658
	σ	0.0404	0.0439
$\rho = 0.5$	V_{ATM}	12.3	12.1905
	σ	0.0488	0.0335
$\rho = 0$	V_{ATM}	13.8994	13.8774
	σ	0.0494	0.04186
$\rho = -0.5$	V_{ATM}	15.0579	15.1301
	σ	0.07227	0.0496
$\rho = -0.9999$	V_{ATM}	15.9887	15.9092
	σ	0.0773	0.0898

Table 6: *ATM Max option values V_{ATM} from 2D LSMC and 1D LSMC (20'000 paths, 50 exercise dates per year) for different correlations.*

Comparison with the 2D explicit finite difference method results

The results obtained from the 2-dimensional explicit finite difference method are given in Table 9 for $dS = 5$ in both directions (S_1 and S_2).

The results obtained with this method are quite similar to those obtained from the 2D LSMC. The differences between both methods are around 1%, and should converge to almost the same solution as the number of simulations is increased and the size of the grid is decreased. Note that the differences between both methods increase with the moneyness of the option.

The 2-dimensional finite difference method is not applicable in the extreme cases $\rho = \pm 1$. Concerning $\rho = 1$, as correlation is perfect then the stock prices of both underlying assets behave exactly the same and either can be used to compute the option price with a 1-dimensional finite difference method, because all parameters of both underlying assets are the same. Table 9 shows the value of the option when $\rho = 1$, computed with the 1-dimensional Crank–Nicholson method, called 1D, with $dS = 1$ and $dt = 10^{-2}$. As expected, this is not the result obtained from the 2-dimensional method. Note that the LSMC technique converges to the correct solution in such cases of extreme correlation.

Figure 1 represents the value of the Max option at time zero for the case $\rho = 0.5$, while Figure 2 corresponds to the Max option value at time zero for $S_2 = 120$ depending on the correlation. These two figures correspond to the results obtained with the finite difference method. Figure 2 shows that the value of the Max option increases when the correlation between the underlying assets decreases.

		2D LSMC	1D LSMC
$\rho = 0.9999$	V_{ITM}	13.4866	13.4767
	σ	0.0484	0.0485
$\rho = 0.5$	V_{ITM}	19.0007	18.7113
	σ	0.053	0.0429
$\rho = 0$	V_{ITM}	σ	21.1524
	σ	0.0688	0.0554
$\rho = -0.5$	V_{ITM}	23.4018	23.2189
	σ	0.092	0.0705
$\rho = -0.9999$	V_{ITM}	25.1762	25.0724
	σ	0.0865	0.1176

Table 7: *ITM Max option values V_{ITM} from 2D LSMC and 1D LSMC (20'000 paths, 50 exercise dates per year) for different correlations.*

3.4 5D American–Bermudan Max Options

The LSMC is now used to price a Bermudan Max option on 5 underlying assets. Again, one will test if using an index, instead of the 5 underlyings separately, leads to the correct solution as in the 2-dimensional case.

The payoff of a 5D Max option is $(\max(S_1, S_2, S_3, S_4, S_5) - K)^+$, where S_i are the underlying stock prices.

Again, the simulations are repeated 50 times to have a standard deviation of the results.

Concerning the 5-dimensional LSMC, several combinations of basis functions are tested:

- 6 basis functions: $1, S_1, S_2, S_3, S_4, S_5$.
- 11 basis functions: $1, S_1, S_2, S_3, S_4, S_5, S_1^2, S_2^2, S_3^2, S_4^2, S_5^2$.
- 21 basis functions:

$$1, S_1, S_2, S_3, S_4, S_5,\; S_1^2, S_2^2, S_3^2, S_4^2, S_5^2,$$
$$S_1S_2, S_1S_3, S_1S_4, S_1S_5, S_2S_3, S_2S_4, S_2S_5, S_3S_4, S_3S_5, S_4S_5.$$

- 46 basis functions:

The 21 basis functions above plus

$$S_1^3, S_2^3, S_3^3, S_4^3, S_5^3,\; S_1^2S_2, S_1^2S_3, S_1^2S_4, S_1^2S_5, S_2^2S_1, S_2^2S_3, S_2^2S_4, S_2^2S_5,$$
$$S_3^2S_1, S_3^2S_2, S_3^2S_4, S_3^2S_5, S_4^2S_1, S_4^2S_2, S_4^2S_3, S_4^2S_5, S_5^2S_1, S_5^2S_2, S_5^2S_3, S_5^2S_4.$$

		2D LSMC	1D LSMC
$\rho = 0.9999$	V_{OTM}	4.5301	4.5118
	σ	0.03115	0.0252
$\rho = 0.5$	V_{OTM}	7.1849	7.1687
	σ	0.0282	0.0352
$\rho = 0$	V_{OTM}	8.0205	8.1632
	σ	0.0439	0.0368
$\rho = -0.5$	V_{OTM}	8.5873	8.7164
	σ	0.0576	0.0466
$\rho = -0.9999$	V_{OTM}	8.9614	8.9302
	σ	0.047	0.0504

Table 8: *OTM Max option values V_{OTM} from 2D LSMC and 1D LSMC (20'000 paths, 50 exercise dates per year) for different correlations.*

ρ	OTM	ATM	ITM
1D	4.49	8.17	13.47
1	5.32	9.41	14.97
0.5	7.24	12.39	19.04
0	8.22	14.14	21.65
-0.5	8.75	15.32	23.64
-1	8.91	16.01	25.21

Table 9: *Max option values for different correlations and moneyness (with $dS = 5$).*

The results from the 5-dimensional LSMC method are compared to those obtained with a 1-dimensional LSMC whose basis functions depend on an index. The index I is determined as follows:

- Simulate several paths for each underlying separately, as in the usual Monte Carlo framework;
- The index for path j and time k is simply: $I\,(j,k) = \max\,[S_i(j,k)]$, i=1,...,5.

Again, in the 1-dimensional LSMC method, one combination of basis functions is used:

- 4 basis functions: $1, I, I^2, I^3$.

The features of the 5D Max options studied in this section are given in Table 10.

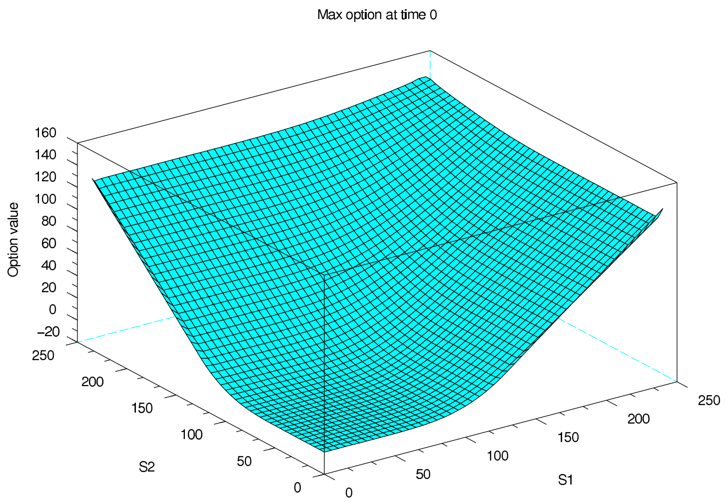

Figure 1: *Value of the max option at time zero (for $dS = 5$).*

K	T	σ_i	r	D_i
100	3	0.2	0.05	0.1

Table 10: *Features of the max option.*

Note that D_i are the dividend yields of each underlying i. The initial values of each underlying stock prices $S_{0,i}$ are always considered to be equal and with same weights, but three different scenarios are considered:

- $S_{0,j} = 90$ (OTM);
- $S_{0,j} = 100$ (ATM);
- $S_{0,j} = 110$ (ITM).

The correlation between both underlying assets is taken to be $\rho = 0$. The option has 3 possible exercise dates per year.

The reference values, as quoted in [6], are given in Table 11. Table 12 gives the values of the 5D Bermudan Max option obtained from the LSMC method for the different combinations of basis functions.

Table 12 shows that 21 basis functions give very accurate results. However, 11 basis functions already give quite accurate results. The value of the option (with the 5-dimensional LSMC with 5'000 paths) is returned in about 30 seconds for most combinations of basis functions.

Table 12 also shows that the results obtained when using an index are not accurate at all. Increasing the number of basis functions of the number of paths does not improve the accuracy of the results. Therefore, using an index on many underlyings leads to an incorrect solution.

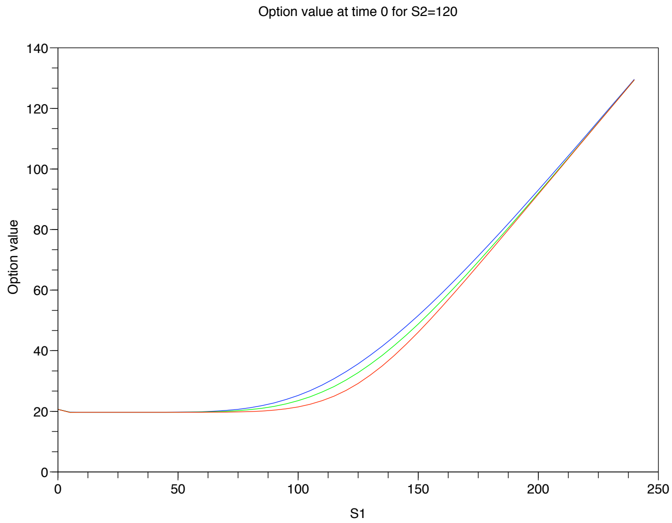

Figure 2: *Value of the max option at time zero for $S_2 = 120$ (for $dS = 1$). The red curve represents the case $\rho = 0.5$, the green curve the case $\rho = 0$, and the blue curve the case $\rho = -0.5$.*

95% interval		
V_{OTM}	V_{ATM}	V_{ITM}
[16.602, 16.655]	[26.109, 26.292]	[36.704, 36.832]

Table 11: *Reference Max option values.*

4 Conclusions

This study has clearly verified the reliability of the LSMC algorithm using different kinds of derivatives.

When several risk factors affect the value of the option, simulation-based techniques are far easier to implement and are much more efficient than numerical methods. However, the efficiency of the LSMC method decreases as the number of exercise dates increases. Therefore, this method is more efficient to price Bermudan rather than pure American-style derivatives. When the LSMC method involves many dimensions—typically more than 3—it is crucial to choose the optimal set of basis functions, by performing some tests, in order to have an acceptable amount of basis functions. Choosing too many basis functions would increase significantly the computational time, while choosing too few would lead to an inaccurate result.

Concerning the LSMC method, when a few underlying assets influence the value of the derivative—this is true at least for the case with two underlying assets—it is possible to use basis functions on a pseudo-index containing the underlying assets, which is somewhat more

	OTM		ATM		ITM	
b. f.	V_{OTM}	σ_{OTM}	V_{ATM}	σ_{ATM}	V_{ITM}	σ_{ITM}
6 (5D)	16.15	0.2	25.76	0.25	35.96	0.25
11 (5D)	16.6	0.2	26.03	0.23	36.55	0.31
21 (5D)	16.66	0.21	26.16	0.24	36.77	0.26
46 (5D)	16.89	0.2	26.4	0.22	37.05	0.27
4 (1D)	6.56	0.07	12.55	0.1	22.9	0.11

Table 12: *Bermudan Max option values with the Least Squares Monte Carlo method using 5'000 paths (repeated 50 times). V is the option value and σ is the standard deviation of V. 5D means 5-dimensional LSMC, while 1D means 1-dimensional LSMC.*

efficient. However, when more underlying assets are involved, this is not possible anymore as the results do not converge to the correct solution. Another advantage of the simulation-based techniques is that, unlike numerical methods, they can easily be adapted to consider different dynamics of the stock price. Both numerical and simulation-based methods are very useful tools for pricing complex derivatives, but each of these methods is applicable in different settings. Note however that, due to their inefficiency, the application range of the numerical methods is restricted.

The next step would be to apply this method to derivatives with even higher dimensions. In addition, this study did not consider transaction costs. Moreover, in order to obtain more realistic dynamics of the stock price, some of the assumptions of the Black–Scholes model could be relaxed. For instance, jumps or stochastic volatility could be introduced in the model.

References

[1] Bensoussan, A. On the Theory of Option Pricing, *Acta Applicandae Mathematicae 2 139*, 158 (1984).

[2] Glasserman, P., and Yu, B. Number of paths versus number of basis functions in American option pricing, *The Annals of Applied Probability 14*, 4 (2004), 2090–2119.

[3] Glasserman, P. *Monte Carlo Methods in Financial Engineering*, Springer, 2003.

[4] Karatzas, I. On the Pricing of the American Option, *Appl. Math. and Optim. 17*, (1988), 37–60.

[5] Longstaff, F.A., and Schwartz, E.S. Valuing American Options by Simulation: A Simple Least-Squares Approach, *The Review of Financial Studies 14*, 1 (2001).

[6] Powell Firth, N., *High dimensional American Options*, Ph.D. Thesis, University of Oxford 2005.