
Valuing American Options by Simulation: A Simple Least-Squares Approach

Francis A. Longstaff
UCLA

Eduardo S. Schwartz
UCLA

This article presents a simple yet powerful new approach for approximating the value of American options by simulation. The key to this approach is the use of least squares to estimate the conditional expected payoff to the optionholder from continuation. This makes this approach readily applicable in path-dependent and multifactor situations where traditional finite difference techniques cannot be used. We illustrate this technique with several realistic examples including valuing an option when the underlying asset follows a jump-diffusion process and valuing an American swaption in a 20-factor string model of the term structure.

One of the most important problems in option pricing theory is the valuation and optimal exercise of derivatives with American-style exercise features. These types of derivatives are found in all major financial markets including the equity, commodity, foreign exchange, insurance, energy, sovereign, agency, municipal, mortgage, credit, real estate, convertible, swap, and emerging markets. Despite recent advances, however, the valuation and optimal exercise of American options remains one of the most challenging problems in derivatives finance, particularly when more than one factor affects the value of the option. This is primarily because finite difference and binomial techniques become impractical in situations where there are multiple factors.¹

We are grateful for the comments of Yaser Abu-Mostafa, Giovanni Barone-Adesi, Marco Avellaneda, Peter Bossaerts, Peter Carr, Peter DeCrem, Craig Fithian, Bjorn Flesaker, James Gammill, Gordon Gemmill, Robert Geske, Eric Ghysels, Ravit Efraty Mandell, Soetjo Tanudjaja, John Thornley, Bruce Tuckman, Pedro Santa-Clara, Pratap Sondhi, Ross Valkanov, and seminar participants at Bear Stearns, the University of British Columbia, the California Institute of Technology, Chase Manhattan Bank, Citibank, the Courant Institute at New York University, Credit Suisse First Boston, Daiwa Securities, Fuji Bank, Goldman Sachs, Greenwich Capital, Morgan Stanley Dean Witter, the Norinchukin Bank, Nikko Securities, the Math Week Risk Magazine Conferences in London and New York, Salomon Smith Barney in London and New York, Simplex Capital, the Sumitomo Bank, UCLA, the 1998 Danish Finance Association meetings, and the 1999 Western Finance Association meetings. We are particularly grateful for the comments of the editor Ravi Jagannathan and of an anonymous referee who made extensive and insightful suggestions for improving the article. All errors are our responsibility. Address correspondence to Francis A. Longstaff, The Anderson School at UCLA, Box 951481, Los Angeles, CA 90095-1481, or e-mail: francis.longstaff@anderson.ucla.edu.

¹ For example, this explains why virtually all Wall Street firms value and exercise American swaptions using a simple single-factor model despite clear evidence that the term structure is driven by multiple factors.

In this article, we present a simple, yet powerful new approach to approximating the value of American options by simulation. By its nature, simulation is a promising alternative to traditional finite difference and binomial techniques and has many advantages as a framework for valuing, risk managing, and optimally exercising American options. For example, simulation is readily applied when the value of the option depends on multiple factors. Simulation can also be used to value derivatives with both path-dependent and American-exercise features. Simulation allows state variables to follow general stochastic processes such as jump diffusions, as in Merton (1976) and Cox and Ross (1976), non-Markovian processes, as in Heath, Jarrow, and Morton (1992), and even general semimartingales, as in Harrison and Pliska (1981).² From a practical perspective, simulation is well suited to parallel computing, which allows significant gains in computational speed and efficiency. Finally, simulation techniques are simple, transparent, and flexible.

To understand the intuition behind this approach, recall that at any exercise time, the holder of an American option optimally compares the payoff from immediate exercise with the expected payoff from continuation, and then exercises if the immediate payoff is higher. Thus the optimal exercise strategy is fundamentally determined by the conditional expectation of the payoff from continuing to keep the option alive. The key insight underlying our approach is that this conditional expectation can be estimated from the cross-sectional information in the simulation by using least squares. Specifically, we regress the ex post realized payoffs from continuation on functions of the values of the state variables. The fitted value from this regression provides a direct estimate of the conditional expectation function. By estimating the conditional expectation function for each exercise date, we obtain a complete specification of the optimal exercise strategy along each path. With this specification, American options can then be valued accurately by simulation. We refer to this technique as the least squares Monte Carlo (LSM) approach.

This approach is easy to implement since nothing more than simple least squares is required. To illustrate this, we present a series of increasingly complex but realistic examples. In the first, we value an American put option in a single-factor setting. In the second, we value an exotic American-Bermuda-Asian option. This option is path dependent and has multifactor features. In the third, we value a cancelable index amortizing swap where the term structure is driven by several factors. This standard fixed-income derivative product has almost pathological path-dependent properties. In each case, the simulation algorithm gives values that are indistinguishable from those given by more computationally intensive finite difference techniques. In the fourth example, we value American options on an asset which follows a jump-diffusion process. This option cannot be valued using standard finite difference techniques. To illustrate the full generality of this approach, the fifth

² Semimartingales are essentially the broadest class of processes for which stochastic integrals can be defined and standard option pricing theory applied.

example values a deferred American swaption in a 20-factor string model where each point on the interest-rate curve is a separate factor. We also show how the algorithm can be used in a risk-management context by computing the sensitivity of swaption values to each point along the curve.

A number of other recent articles also address the pricing of American options by simulation. In an important early contribution to this literature, Bossaerts (1989) solves for the exercise strategy that maximizes the simulated value of the option. Other important examples of this literature include Tilley (1993), Barraquand and Martineau (1995), Averbukh (1997), Broadie and Glasserman (1997a,b,c), Broadie, Glasserman, and Jain (1997), Raymar and Zwecher (1997), Broadie et al. (1998), Carr (1998), Ibanez and Zapatero (1998), and Garcia (1999). These articles use various stratification or parameterization techniques to approximate the transitional density function or the early exercise boundary. This article takes a fundamentally different approach by focusing directly on the conditional expectation function.

Several recent articles that use an approach similar to ours include Carriere (1996) and Tsitsiklis and Van Roy (1999). Our work, however, differs in a number of ways. For example, neither of these articles take the approach to the level of practical implementation we do in this article. Furthermore, we include in the regression only paths for which the option is in the money. This significantly increases the efficiency of the algorithm and decreases the computational time. In addition, we demonstrate the application of the methodology to complex derivatives with many underlying factors and evaluate the accuracy of the algorithm by comparing our solutions to finite difference approximations.³

The remainder of this article is organized as follows. Section 1 presents a simple numerical example of the simulation approach. Section 2 describes the underlying theoretical framework. Sections 3–7 provide specific examples of the application of this approach. Section 8 discusses a number of numerical and implementation issues. Section 9 summarizes the results and contains concluding remarks.

1. A Numerical Example

At the final exercise date, the optimal exercise strategy for an American option is to exercise the option if it is in the money. Prior to the final date, however, the optimal strategy is to compare the immediate exercise value with the expected cash flows from continuing, and then exercise if immediate exercise is more valuable. Thus, the key to optimally exercising an American option is identifying the conditional expected value of continuation. In this approach, we use the cross-sectional information in the simulated paths to

³ Another related article is Keane and Wolpin (1994), which uses regression in a simulation context to solve discrete choice dynamic programming problems.

identify the conditional expectation function. This is done by regressing the subsequent realized cash flows from continuation on a set of basis functions of the values of the relevant state variables. The fitted value of this regression is an efficient unbiased estimate of the conditional expectation function and allows us to accurately estimate the optimal stopping rule for the option.

Perhaps the best way to convey the intuition of the LSM approach is through a simple numerical example. Consider an American put option on a share of non-dividend-paying stock. The put option is exercisable at a strike price of 1.10 at times 1, 2, and 3, where time three is the final expiration date of the option. The riskless rate is 6%. For simplicity, we illustrate the algorithm using only eight sample paths for the price of the stock. These sample paths are generated under the risk-neutral measure and are shown in the following matrix.

Stock price paths				
Path	$t = 0$	$t = 1$	$t = 2$	$t = 3$
1	1.00	1.09	1.08	1.34
2	1.00	1.16	1.26	1.54
3	1.00	1.22	1.07	1.03
4	1.00	.93	.97	.92
5	1.00	1.11	1.56	1.52
6	1.00	.76	.77	.90
7	1.00	.92	.84	1.01
8	1.00	.88	1.22	1.34

Our objective is to solve for the stopping rule that maximizes the value of the option at each point along each path. Since the algorithm is recursive, however, we first need to compute a number of intermediate matrices. Conditional on not exercising the option before the final expiration date at time 3, the cash flows realized by the optionholder from following the optimal strategy at time 3 are given below.

Cash-flow matrix at time 3			
Path	$t = 1$	$t = 2$	$t = 3$
1	—	—	.00
2	—	—	.00
3	—	—	.07
4	—	—	.18
5	—	—	.00
6	—	—	.20
7	—	—	.09
8	—	—	.00

These cash flows are identical to the cash flows that would be received if the option were European rather than American.

If the put is in the money at time 2, the optionholder must then decide whether to exercise the option immediately or continue the option's life until the final expiration date at time 3. From the stock-price matrix, there are only five paths for which the option is in the money at time 2. Let X denote the stock prices at time 2 for these five paths and Y denote the corresponding discounted cash flows received at time 3 if the put is not exercised at time 2. We use only in-the-money paths since it allows us to better estimate the conditional expectation function in the region where exercise is relevant and significantly improves the efficiency of the algorithm. The vectors X and Y are given by the nondashed entries below.

Regression at time 2		
Path	Y	X
1	$.00 \times .94176$	1.08
2	—	—
3	$.07 \times .94176$	1.07
4	$.18 \times .94176$	.97
5	—	—
6	$.20 \times .94176$	.77
7	$.09 \times .94176$	.84
8	—	—

To estimate the expected cash flow from continuing the option's life conditional on the stock price at time 2, we regress Y on a constant, X , and X^2 . This specification is one of the simplest possible; more general specifications are considered later in the article. The resulting conditional expectation function is $E[Y \mid X] = -1.070 + 2.983X - 1.813X^2$.

With this conditional expectation function, we now compare the value of immediate exercise at time 2, given in the first column below, with the value from continuation, given in the second column below.

Optimal early exercise decision at time 2		
Path	Exercise	Continuation
1	.02	.0369
2	—	—
3	.03	.0461
4	.13	.1176
5	—	—
6	.33	.1520
7	.26	.1565
8	—	—

The value of immediate exercise equals the intrinsic value $1.10 - X$ for the in-the-money paths, while the value from continuation is given by substituting X into the conditional expectation function. This comparison implies that

it is optimal to exercise the option at time 2 for the fourth, sixth, and seventh paths. This leads to the following matrix, which shows the cash flows received by the optionholder conditional on not exercising prior to time 2.

Cash-flow matrix at time 2			
Path	$t = 1$	$t = 2$	$t = 3$
1	—	.00	.00
2	—	.00	.00
3	—	.00	.07
4	—	.13	.00
5	—	.00	.00
6	—	.33	.00
7	—	.26	.00
8	—	.00	.00

Observe that when the option is exercised at time 2, the cash flow in the final column becomes zero. This is because once the option is exercised there are no further cash flows since the option can only be exercised once.

Proceeding recursively, we next examine whether the option should be exercised at time 1. From the stock price matrix, there are again five paths where the option is in the money at time 1. For these paths, we again define Y as the discounted value of subsequent option cash flows. Note that in defining Y , we use actual realized cash flows along each path; we do not use the conditional expected value of Y estimated at time 2 in defining Y at time 1. As is discussed later, discounting back the conditional expected value rather than actual cash flows can lead to an upward bias in the value of the option.

Since the option can only be exercised once, future cash flows occur at either time 2 or time 3, but not both. Cash flows received at time 2 are discounted back one period to time 1, and any cash flows received at time 3 are discounted back two periods to time 1. Similarly X represents the stock prices at time 1 for the paths where the option is in the money. The vectors X and Y are given by the nondashed elements in the following matrix.

Regression at time 1		
Path	Y	X
1	$.00 \times .94176$	1.09
2	—	—
3	—	—
4	$.13 \times .94176$	.93
5	—	—
6	$.33 \times .94176$	.76
7	$.26 \times .94176$	.92
8	$.00 \times .94176$	.88

The conditional expectation function at time 1 is estimated by again regressing Y on a constant, X and X^2 . The estimated conditional expectation function is $E[Y \mid X] = 2.038 - 3.335X + 1.356X^2$. Substituting the values of X into this regression gives the estimated conditional expectation function. These estimated continuation values and immediate exercise values at time 1 are given in the first and second columns below. Comparing the two columns shows that exercise at time 1 is optimal for the fourth, sixth, seventh, and eighth paths.

Optimal early exercise decision at time 1		
Path	Exercise	Continuation
1	.01	.0139
2	—	—
3	—	—
4	.17	.1092
5	—	—
6	.34	.2866
7	.18	.1175
8	.22	.1533

Having identified the exercise strategy at times 1, 2, and 3, the stopping rule can now be represented by the following matrix, where the ones denote exercise dates at which the option is exercised.

Path	Stopping rule		
	$t = 1$	$t = 2$	$t = 3$
1	0	0	0
2	0	0	0
3	0	0	1
4	1	0	0
5	0	0	0
6	1	0	0
7	1	0	0
8	1	0	0

With this specification of the stopping rule, it is now straightforward to determine the cash flows realized by following this stopping rule. This is done by simply exercising the option at the exercise dates where there is a one in the above matrix. This leads to the following option cash

flow matrix.

Option cash flow matrix			
Path	$t = 1$	$t = 2$	$t = 3$
1	.00	.00	.00
2	.00	.00	.00
3	.00	.00	.07
4	.17	.00	.00
5	.00	.00	.00
6	.34	.00	.00
7	.18	.00	.00
8	.22	.00	.00

Having identified the cash flows generated by the American put at each date along each path, the option can now be valued by discounting each cash flow in the option cash flow matrix back to time zero, and averaging over all paths. Applying this procedure results in a value of .1144 for the American put. This is roughly twice the value of .0564 for the European put obtained by discounting back the cash flows at time 3 from the first cash flow matrix.

Although very stylized, this example illustrates how least squares can use the cross-sectional information in the simulated paths to estimate the conditional expectation function. In turn, the conditional expectation function is used to identify the exercise decision that maximizes the value of the option at each date along each path. As shown by this example, the LSM approach is easily implemented since nothing more than simple regression is involved.

2. The Valuation Algorithm

In this section we describe the general LSM algorithm. The valuation framework underlying the LSM algorithm is based on the general derivative pricing paradigm of Black and Scholes (1973), Merton (1973), Harrison and Kreps (1979), Harrison and Pliska (1981), Cox, Ingersoll, and Ross (1985), Heath, Jarrow, and Morton (1992), and others. We also present several convergence results for the algorithm.

2.1 The valuation framework

We assume an underlying complete probability space $(\Omega, \mathcal{F}, P)$ and finite time horizon $[0, T]$, where the state space Ω is the set of all possible realizations of the stochastic economy between time 0 and T and has typical element ω representing a sample path, $\mathcal{F}$ is the sigma field of distinguishable events at time T , and P is a probability measure defined on the elements of $\mathcal{F}$. We define $F = \{\mathcal{F}_t; t \in [0, T]\}$ to be the augmented filtration generated by the relevant price processes for the securities in the economy, and assume that $\mathcal{F}_T = \mathcal{F}$. Consistent with the no-arbitrage paradigm, we assume the existence of an equivalent martingale measure Q for this economy.

We are interested in valuing American-style derivative securities with random cash flows which may occur during $[0, T]$. We restrict our attention to

derivatives with payoffs that are elements of the space of square-integrable or finite-variance functions $L^2(\Omega, \mathcal{F}, Q)$. Standard results such as Bensoussan (1984) and Karatzas (1988) imply that the value of an American option can be represented by the Snell envelope; the value of an American option equals the maximized value of the discounted cash flows from the option, where the maximum is taken over all stopping times with respect to the filtration F . We introduce the notation $C(\omega, s; t, T)$ to denote the path of cash flows generated by the option, conditional on the option not being exercised at or prior to time t and on the optionholder following the optimal stopping strategy for all $s, t < s \leq T$. This function is analogous to the intermediate cash-flow matrices used in the previous section.

The objective of the LSM algorithm is to provide a pathwise approximation to the optimal stopping rule that maximizes the value of the American option. To convey the intuition behind the LSM algorithm, we focus the discussion on the case where the American option can only be exercised at the K discrete times $0 < t_1 \leq t_2 \leq t_3 \leq \dots \leq t_K = T$, and consider the optimal stopping policy at each exercise date. In practice, of course, many American options are continuously exercisable; the LSM algorithm can be used to approximate the value of these options by taking K to be sufficiently large.

At the final expiration date of the option, the investor exercises the option if it is in the money, or allows it to expire if it is out of the money. At exercise time t_k prior to the final expiration date, however, the optionholder must choose whether to exercise immediately or to continue the life of the option and revisit the exercise decision at the next exercise time. The value of the option is maximized pathwise, and hence unconditionally, if the investor exercises as soon as the immediate exercise value is greater than or equal to the value of continuation.⁴

At time t_k , the cash flow from immediate exercise is known to the investor, and the value of immediate exercise simply equals this cash flow. The cash flows from continuation, of course, are not known at time t_k . No-arbitrage valuation theory, however, implies that the value of continuation, or equivalently, the value of the option assuming that it cannot be exercised until after t_k , is given by taking the expectation of the remaining discounted cash flows $C(\omega, s; t_k, T)$ with respect to the risk-neutral pricing measure Q . Specifically, at time t_k , the value of continuation $F(\omega; t_k)$ can be expressed as

$$F(\omega; t_k) = E_Q \left[\sum_{j=k+1}^K \exp \left(- \int_{t_k}^{t_j} r(\omega, s) ds \right) C(\omega, t_j; t_k, T) \mid \mathcal{F}_{t_k} \right], \quad (1)$$

⁴ For a discussion of optimal exercise policies for American options, see Duffie (1996) or Lamberton and Lapeyre (1996). Bossaerts (1989) directly uses this maximization property in developing simulation estimates of American option prices by parameterizing the stopping rule and then solving for the parameters that maximize the value of the option.

where $r(\omega, t)$ is the (possibly stochastic) riskless discount rate, and the expectation is taken conditional on the information set $\mathcal{F}_{t_k}$ at time t_k . With this representation, the problem of optimal exercise reduces to comparing the immediate exercise value with this conditional expectation, and then exercising as soon as the immediate exercise value is positive and greater than or equal to the conditional expectation.

2.2 The LSM algorithm

The LSM approach uses least squares to approximate the conditional expectation function at $t_{K-1}, t_{K-2}, \dots, t_1$. We work backwards since the path of cash flows $C(\omega, s; t, T)$ generated by the option is defined recursively; $C(\omega, s; t_k, T)$ can differ from $C(\omega, s; t_{k+1}, T)$ since it may be optimal to stop at time t_{k+1} , thereby changing all subsequent cash flows along a realized path ω . Specifically, at time t_{K-1} , we assume that the unknown functional form of $F(\omega; t_{K-1})$ in Equation (1) can be represented as a linear combination of a countable set of $\mathcal{F}_{t_{K-1}}$ -measurable basis functions.

This assumption can be formally justified, for example, when the conditional expectation is an element of the L^2 space of square-integrable functions relative to some measure. Since L^2 is a Hilbert space, it has a countable orthonormal basis and the conditional expectation can be represented as a linear function of the elements of the basis.⁵ As an example, assume that X is the value of the asset underlying the option and that X follows a Markov process.⁶ One possible choice of basis functions is the set of (weighted) Laguerre polynomials:

$$L_0(X) = \exp(-X/2), \quad (2)$$

$$L_1(X) = \exp(-X/2) (1 - X), \quad (3)$$

$$L_2(X) = \exp(-X/2) (1 - 2X + X^2/2), \quad (4)$$

$$L_n(X) = \exp(-X/2) \frac{e^X}{n!} \frac{d^n}{dX^n} (X^n e^{-X}). \quad (5)$$

With this specification, $F(\omega; t_{K-1})$ can be represented as

$$F(\omega; t_{K-1}) = \sum_{j=0}^{\infty} a_j L_j(X), \quad (6)$$

where the a_j coefficients are constants. Other types of basis functions include the Hermite, Legendre, Chebyshev, Gegenbauer, and Jacobi polynomials.⁷

⁵ For a discussion of Hilbert space theory and Hilbert space representations of square-integrable functions, see Royden (1968).

⁶ For Markovian problems, only current values of the state variables are necessary. For non-Markovian problems, both current and past realizations of the state variables can be included in the basis functions and the regressions.

⁷ These functions are described in Chapter 22 of Abramowitz and Stegun (1970).

以上内容仅为本文档的试下载部分，为可阅读页数的一半内容。如要下载或阅读全文，请访问：<https://d.book118.com/757166200015006121>