

Monte Carlo valuation of American options

L.C.G. ROGERS¹ *University of Bath*

Abstract. This paper introduces a ‘dual’ way to price American options, based on simulating the path of the option payoff, and of a judiciously-chosen Lagrangian martingale. Taking the pathwise maximum of the payoff less the martingale provides an upper bound for the price of the option, and this bound is sharp for the optimal choice of Lagrangian martingale. As a first exploration of this method, three examples are investigated numerically; the accuracy achieved with even very simple-minded choices of Lagrangian martingale is surprising. The method also leads naturally to candidate hedging policies for the option, and estimates of the risk involved in using them.

Key words: Monte Carlo, American option, duality, Lagrangian, martingale, Snell envelope.

Abbreviated Title. Monte Carlo valuation of American options.

AMS(MOS) subject classifications: 49K35, 60G40, 91B28.

1 Introduction

The pricing of American options by simulation techniques is an important and difficult task, as witnessed by the contributions of Tilley (1993), Barraquand & Martineau (1995), Carriere (1996) Broadie & Glasserman (1997), Broadie, Glasserman & Jain (1997), Raymar & Zwecher (1997), Carr (1998), and Longstaff & Schwartz (1999). Frequently, the payoff of an American-style derivative depends in a highly complex path-dependent fashion on many underlyings, which means that the traditional dynamic programming approach to computing the value and the optimal exercise policy is impossible, due to the dimension of the problem. This has prompted interest in possible simulation methods for pricing such derivatives, and the papers mentioned above offer a variety of approaches to the problem. In general terms, all use simulation in some way to derive a stopping rule, by comparing the current value of stopping with some estimate (based on simulated paths) of the value of waiting. It follows that the answers obtained will be lower bounds for the value of the option, since the value has been computed using an approximation to the optimal stopping rule.

In contrast, the approach adopted here makes no attempt to determine an approximately optimal exercise policy, and always comes up with an answer which is an *upper* bound for the true price. While it says little about how such an option should be *exercised*, it does give guidance on how the option should be *hedged*. Thus this approach should be of value to the party writing the option, and the other general approach would be of value to the party buying the option. The perceptive reader

¹Department of Mathematical Sciences, University of Bath, Bath BA2 7AY, England. Tel= +44 1225 826224; email= L.C.G.Rogers@ bath.ac.uk

may already have guessed that there are themes of convex duality at play here, and the germ of the method is in an interesting but little-appreciated paper of Davis & Karatzas (1994)². In Section 2, we show how the price of the American option may be expressed as the infimum of a family of expectations, the infimum being taken over the class of Lagrangian martingales. This expression immediately suggests how one might try to estimate the price of an American options, and in Section 4 we take this further with a numerical study of some examples: the standard American put; an American min-put (see Hartley (2000)); an American-Bermudan-Asian example of Longstaff & Schwartz (1999). The method requires a good choice of Lagrangian martingale to give good results, but it turns out that in the examples we study it is not too hard to find martingales which give reasonably close approximations to the true price.

2 The price of an American option.

We fix some finite time horizon $T > 0$, and suppose given on some filtered probability space³ $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0 \leq t \leq T}, P)$ two adapted processes $(r_t)_{0 \leq t \leq T}$ and $(\tilde{Z}_t)_{0 \leq t \leq T}$. The first is the spot rate of interest, and the second defines the amount paid to the holder of an American option at the moment of exercise. We shall also assume that the probability P is the (risk-neutral) pricing probability for the problem. Adopting the notational device that a random time denoted by τ (with or without superscripts or subscripts) should be understood to be a stopping time, standard arbitrage pricing theory gives the time-0 value of the American option to be

$$Y_0^* \equiv \sup_{0 \leq \tau \leq T} EZ_\tau, \quad (2.1)$$

where $Z_t \equiv \exp(-\int_0^t r_s ds) \tilde{Z}_t$ is the discounted exercise value of the option. To avoid trivialities, we need to assume that $Y_0^* < \infty$; in fact, for technical reasons we shall assume a little more, namely that for some $p > 1$, $\sup_{0 \leq t \leq T} |Z_t| \in L^p$, and also that the paths of Z are right continuous. Under this assumption, the *Snell envelope* process

$$Y_t^* \equiv \text{ess sup}_{t \leq \tau \leq T} E[Z_\tau | \mathcal{F}_t]. \quad (2.2)$$

is a supermartingale of class (D), and so has a Doob-Meyer decomposition

$$Y_t^* = Y_0^* + M_t^* - A_t^*, \quad (2.3)$$

where M^* is a martingale vanishing at zero, and A^* is an previsible integrable increasing process, also vanishing at zero. See, for example, Dellacherie & Meyer (1980), p432.

The following result is the theoretical basis of the paper.

²I am grateful to Mike Curran for drawing this paper to my attention, and persuading me that Monte Carlo pricing of American options could work.

³satisfying the usual conditions; see, for example, Rogers & Williams (2000).

Theorem 1

$$Y_0^* = \inf_{M \in H_0^1} E[\sup_{0 \leq t \leq T} (Z_t - M_t)], \quad (2.4)$$

where H_0^1 is the space of martingales M for which $\sup_{0 \leq t \leq T} |M_t| \in L^1$, and such that $M_0 = 0$. The infimum is attained by taking $M = M^*$.

Before proving Theorem 1, let us note how this leads to a method of pricing the American option: *we pick a suitable martingale M , and evaluate by simulation the expectation $E[\sup_{0 \leq t \leq T} (Z_t - M_t)]$.* In Section 4 we shall show how to go about finding a ‘suitable’ martingale. Obtaining the optimal martingale is of course a task of a similar complexity to finding the optimal exercise policy, but we can often find simple martingales which provide remarkably good (and quick) bounds.

PROOF OF THEOREM 1. Firstly, we note that Y^* is dominated by the L^p -bounded martingale $z_t \equiv E(\sup_s |Z_s| | \mathcal{F}_t)$, and so $\sup_{0 \leq t \leq T} |M_t^*| \leq \sup_{0 \leq t \leq T} z_t + |Y_0^*| + A_T$, proving that M^* is indeed in H_0^1 .

Returning to the definition (2.1) of Y_0^* , and letting σ denote a *general* random time with values in $[0, T]$, we have for any $M \in H_0^1$ that

$$\begin{aligned} Y_0^* &= \sup_{0 \leq \tau \leq T} EZ_\tau \\ &= \sup_{0 \leq \tau \leq T} E[Z_\tau - M_\tau] \\ &\leq \sup_{0 \leq \sigma \leq T} E[Z_\sigma - M_\sigma] \\ &= E[\sup_{0 \leq t \leq T} (Z_t - M_t)]; \end{aligned}$$

taking the infimum over all $M \in H_0^1$ proves that Y_0^* is bounded above by the right-hand side of (2.4). On the other hand, since $Z_t \leq Y_t^* = Y_0^* + M_t^* - A_t^*$,

$$\begin{aligned} \inf_{M \in H_0^1} E[\sup_{0 \leq t \leq T} (Z_t - M_t)] &\leq E[\sup_{0 \leq t \leq T} (Z_t - M_t^*)] \\ &\leq E[\sup_{0 \leq t \leq T} (Y_t^* - M_t^*)] \\ &= E[\sup_{0 \leq t \leq T} (Y_0^* - A_t^*)] \\ &= Y_0^* \end{aligned}$$

as claimed.

REMARK. Davis & Karatzas (1994) proved that $E[\sup_{0 \leq t \leq T} (Z_t + M_T^* - M_t^*)] = Y_0^*$, in the present notation.

3 Hedging and exercise.

Theorem 1 tells us that in order to find a good approximation to the price Y_0^* of the American option, it is necessary to find a ‘good’ martingale $M \in H_0^1$. We discuss later how this can be done in practice, but for the moment we suppose that we have a candidate martingale M , and interpret this in terms of hedging.

Holding M fixed, we have an upper bound for Y_0^* , namely, the mean of the random variable

$$\eta \equiv \sup_{0 \leq t \leq T} (Z_t - M_t). \quad (3.1)$$

Let us set $\eta_t \equiv E(\eta | \mathcal{F}_t)$ for the martingale closed on the right by η , so that $\eta \equiv \eta_T$. We now think of the martingale M as the discounted gains-from-trade process of some portfolio; thus if we started with wealth η_0 and used this portfolio, our discounted wealth at time t would just be $\eta_0 + M_t$. Now (3.1) implies the inequality for any $t \in [0, T]$

$$Z_t \leq \eta + M_t,$$

and taking conditional expectation given \mathcal{F}_t and rearranging gives the key inequality

$$Z_t \leq E[\eta_T - \eta_0 | \mathcal{F}_t] + (M_t + \eta_0) \quad (3.2)$$

The interpretation of this is immediate and illuminating; *the (discounted) amount Z_t which has to be paid out to the holder of the option if exercised at time t is almost hedged by the (discounted) value of our portfolio.* The shortfall is at worst

$$E[\eta_T - \eta_0 | \mathcal{F}_t]^+ \leq E[(\eta_T - \eta_0)^+ | \mathcal{F}_t] \quad (3.3)$$

So if we propose to use the martingale M as a hedging instrument, it will be highly desirable that the quantity $E|\eta_T - \eta_0|$, which bounds the mean of the shortfall, should be *small*. In the perfect solution, where $M = M^*$, the random variable η is constant, so we have a zero bound on the shortfall, but in general there is no reason why this quantity should be small. Notice in particular that it could be that a given martingale M gives a good bound on the *price* of the option (that is to say, $E(\eta) - Y_0^*$ is small), while having a large shortfall, and therefore being less desirable for hedging.

Remarks. (i) We can of course interpret the dual problem in a very concrete way; *we are trying to choose the hedging strategy to minimise the lookback value of $Z - M$.* In any Markovian example, we would typically have that Z were some function of time and a (possibly high-dimensional) Markov process X , and we would therefore expect the solution to be such that at any time the optimal hedging portfolio should be a function of t , X_t and $\sup_{u \leq t} (Z_u - M_u)$. In principle this could be solved by setting up the Hamilton-Jacobi-Bellman equations, but these are likely to be every bit as difficult to deal with as the original problem. Nevertheless, this suggests a much more refined approach to the choice of the hedging martingale than the very simple-minded approach of Section 4.

(ii) We may also use a candidate martingale M to suggest an exercise policy, namely, to stop when first Z exceeds the value of the hedging policy:

$$\tau_M \equiv \inf\{t \in [0, T] : M_t + \eta_0 \leq Z_t\} \wedge T.$$

In the case where the hedging policy was optimal, this stopping rule would also be optimal. However, it turns out in the examples studied in Section 4 that this rule was very poor, worse even than simply waiting until T and exercising then.

4 Numerical examples.

In this section, we report the results of numerical studies of three examples, the standard American put, the American min-put (see Hartley (2000)), and the Asian-Bermudan-American example studied by Longstaff & Schwartz (1999); firstly, though, we describe the general approach used in all the three examples.

The first step is to simulate a relatively small number of sample paths (300 for the first two examples, 1000 for the third) at a relatively coarse spacing of the time points (40 time-steps in the results here). Using these, we generate the corresponding sample paths of a small number of martingales, and the choice of these seems to be important. If, for example, the reward process Z is a semimartingale of the form $f(t, S_t)$ for some function of a (vector) log-Brownian price process S , then a natural choice to take is the martingale part of Z . This works well in our first two examples, but not for the third where Z is a finite-variation process, and so has no martingale part. Nonetheless, it is a fair guess that the martingale part of the corresponding European option should be close in some sense to the desired martingale. There are few general rules so far; the selection of the martingales appears to be more art than science.

Now take this vector M of martingales, and consider all linear combinations of them. By numerically minimising over λ the value of $E[\sup_{0 \leq t \leq T}(Z_t - \lambda \cdot M_t)]$, we make a presumably better martingale than any of those we began with. Using this minimising value λ^* , we now proceed to simulate a large number of sample paths (30,000 in most of the examples below) and the corresponding martingales $\lambda^* \cdot M$, and then compute the average value of $E[\sup_{0 \leq t \leq T}(Z_t - \lambda^* \cdot M_t)]$ over all the sample paths.

It is perhaps not surprising that the time-consuming part of this process is the numerical minimisation. Fortunately, high precision in the value of λ^* is not crucial, since we are finding a value where a convex function is minimised, and assuming the function is differentiable at the minimum, small departures from the exact minimising value will result in even smaller changes in value of the function. No attempt has been made to explore ways of speeding this part of the process up, though it would be worth finding out whether fewer sample paths would do an acceptable job, or whether there are rules for choosing the starting point for the minimisation

which improve the speed. The simulations themselves required relatively little time. Since the class of hedging martingales is in every case chosen without any attempt at refinement, it seems pointless to try to tighten up the minimisation step at this stage.

As one last improvement, having generated the sample paths with 40 time-steps, we reduced the sample paths to observations at even-numbered times, recomputed the answer for these coarser paths with 20 time-steps, and then used Richardson extrapolation.

While we could have tried various antithetic variable and control variate techniques to reduce the variance of the estimate of the price, we avoided this, since the estimated value of the mean absolute deviation from the mean (MAD) has an important interpretation, and using such variance reduction techniques would have distorted the values of the MAD.

Example 1: an American put on a single asset.

Our first study is of the example of an American put on a single log-Brownian asset, whose price process is given by

$$S_t = S_0 \exp(\sigma W_t + (r - \sigma^2/2)t), \quad (4.1)$$

with r denoting as usual the riskless rate of interest, assumed constant, and σ denoting the constant volatility. No closed-form solution for the price is known, but there are various numerical methods which give good approximations to the price very rapidly. See the papers of Broadie & Detemple (1997), and Ait-Sahalia & Carr (1997) for surveys and comparisons of some of the methods proposed.

In applying the present method, the choice of martingales was almost the crudest possible; there was just one martingale in the hedging set, namely, the martingale part of the corresponding European put, started when the option goes in the money, at the first time that S_t falls below the strike K . The results of the simulation are presented in Table I; parameter values are $K = 100$, $\sigma = 0.4$, $r = 0.06$, and $T = 0.5$, with S_0 varying as shown in the table.

The column of true prices is quoted from the paper of Ait-Sahalia & Carr (1997), using their averaged binomial figures with 1000 time points. We give the Monte Carlo values from the present method for comparison; in every case, there is agreement to within 1%. The standard deviation of the estimates of the price is reported, and the mean absolute deviation from the mean, $E|\eta_T - \eta_0|$, one half of which bounds the expected hedging loss, is reported in the next column. Though they are not reported, the program output the simulation values and Black-Scholes values for the corresponding European puts; the agreement between European values was generally no better than the agreement between the American values. The final two columns present the optimal value of λ , and the time taken by the entire calculation

(on a 600MHz PC). Since the optimal value of λ is very close to 1 in all cases, one could obtain a very quick estimate of the price by just using $\lambda = 1$, thereby cutting out the slow numerical minimisation. Even including the numerical minimisation and performing 30000 simulations, the times taken were of the order of 20-25s.

Other runs were tried with a variety of martingales in the hedging set, but none of them showed any marked improvement over this simplest situation.

Table I: Simulation prices of standard American puts. Parameter values were $K = 100$, $r = 0.06$, $T = 0.5$, and $\sigma = 0.4$.

$S(0)$	American price (true)	American price (MC)	Standard deviation	MAD	λ^*	time (seconds)
80	21.6059	21.8057	0.0084	0.9698	1.0542	24.5400
85	18.0374	18.2105	0.0097	1.1796	1.0485	23.4500
90	14.9187	15.0702	0.0105	1.3226	1.0372	23.3200
95	12.2314	12.3598	0.0111	1.4236	1.0361	23.0900
100	9.9458	10.0550	0.0264	3.2487	1.0152	21.9000
105	8.0281	8.1256	0.0294	4.0260	1.0124	22.4000
110	6.4352	6.4915	0.0307	4.6116	1.0262	23.1600
115	5.1265	5.1799	0.0304	4.7150	1.0371	20.7400
120	4.0611	4.1051	0.0289	4.4798	1.0409	21.9200

Example 2: American puts on the cheapest of n assets.

This study takes n log-Brownian assets (which are assumed independent so that we can compare with the results of Hartley (2000)), given by

$$S_i(t) = S_i(0) \exp(\sigma_i W_i(t) + (r - \sigma_i^2/2)t), \quad i = 1, \dots, n.$$

The reward process Z is simply

$$Z_t = \max_{i=1, \dots, n} e^{-rt} (K - S_i(t))^+.$$

The set of hedging martingales for this example is again almost as rudimentary as one could imagine: we use the martingale parts of each of the corresponding European puts, started once the process Z first goes positive, but only while that share is the cheapest to deliver. Tables II and III both report a range of numerical values for different parameter choices. Throughout, we used $K = 100$, $T = 0.5$, $r = 0.06$, and in the Table II we took the volatilities of both assets to be 0.6, whereas in Table III the volatilities were $\sigma_1 = 0.4$ and $\sigma_2 = 0.8$.

The two tables give a price computed by finite-difference methods, quoted from Hartley (2000), alongside the simulation values. The differences are somewhat larger

than in the first example, but are still in the range 1–2%. Times are of the order of 80s; this is too high for a real-time trading environment, but perfectly acceptable for pricing an OTC product. Reassuringly, the price estimates coming out of the present method are all higher than the prices from Hartley’s finite-difference calculation.

Table II: Simulation prices of min-puts on two assets. Parameter values were $K = 100$, $T = 0.5$, $r = 0.06$, $\sigma_1 = \sigma_2 = 0.6$.

$S_1(0)$	$S_2(0)$	FD price	MC price	SD	MAD	time
80	80	37.30	37.94	0.0779	10.4550	80.83
80	100	32.08	32.52	0.0648	8.3611	81.53
80	120	29.14	29.51	0.0519	6.0384	83.28
100	100	25.06	25.48	0.0703	9.2609	79.02
100	120	20.91	21.24	0.0649	8.1642	81.67
120	120	15.92	16.23	0.0635	7.8660	78.62

Table III: Simulation prices of min-puts on two assets. Parameter values were $K = 100$, $T = 0.5$, $r = 0.06$, $\sigma_1 = 0.4$, $\sigma_2 = 0.8$.

$S_1(0)$	$S_2(0)$	FD price	MC price	SD	MAD	time
80	80	38.01	38.72	0.0812	10.8557	82.25
80	100	32.23	32.98	0.0752	10.0912	81.82
80	120	28.54	29.13	0.0664	8.6926	83.08
100	80	33.34	33.89	0.0610	7.4415	80.23
100	100	25.81	26.30	0.0750	9.9030	80.66
100	120	20.75	21.18	0.0725	9.7199	77.28
120	80	31.21	31.64	0.0431	4.7429	82.21
120	100	22.77	23.16	0.0636	7.8909	74.30
120	120	16.98	17.37	0.0678	9.0109	77.57

The approach used by Hartley is an ingenious attempt to base a stopping rule on the minimum value of all the shares; of course, the stopping rule should ideally depend on the values of each of the shares, but Hartley shows in the case of two assets that his approximate method delivers numerical results within 1% of the finite-difference values. He also uses the method on some examples with more shares, and we quote his results in Table IV. For these calculations, the shares all have volatility $\sigma_i = 0.6$ and initial value $S_i(0) = 100$, with the other parameters as before. Of course, in Hartley’s paper it is only possible to say that the prices obtained are lower bounds on the price, and there is no way of knowing whether the true value is 1% larger, or 40% larger. Hartley presents upper bounds based on the mean of the maximum of Z along the path (corresponding to what we would get using $M = 0$), but these bounds are not at all close. However, when we look at the upper bounds derived by

the present simulation method, we see that in fact Hartley's values are very close to the truth; the gap between his lower bounds and our upper bounds is of the order of 1–2%. The times taken are again satisfactory for some OTC product, but would ideally be shorter. The size of the standard deviation achieved is a lot smaller than the error we are finding in the prices, so we could probably get prices close enough with far fewer than 30000 simulations. In any case, errors of the order of 1–2% are likely to be present in any estimate of the volatilities, and the assumption of constant interest rates is likely to introduce errors of similar size, so in summary, we have a method accurate to observational error.

Table IV: Simulation prices of min-puts on n assets. All shares have $S_i(0) = 100$ and $\sigma_i = 0.6$, with $T = 0.5$, $K = 100$ and $r = 0.06$.

n	FD price	MC price	SD	MAD	European (MC)	SD	time
2	24.87	25.3858	0.0856	9.2157	18.5185	0.1075	53.79
3	31.21	31.8906	0.1044	11.6200	23.6787	0.1036	90.80
4	35.72	36.3456	0.1192	13.2661	27.1762	0.0978	155.97
5	39.01	39.7567	0.1273	14.2813	29.8910	0.0920	209.67
10	47.99	48.7954	0.1356	15.2721	37.3026	0.0737	609.84
15	52.23	52.9456	0.1351	15.3063	40.8255	0.0649	1263.28

Example 3: American-Bermudan-Asian option.

This is one of the examples studied by Longstaff & Schwartz (1999), of an American-Bermudan-Asian option, specified as follows. There is a single risky log-Brownian asset, with dynamics (4.1), in terms of which is defined the cumulative average

$$A_t = \frac{\int_{-\delta}^t S_u du}{t + \delta}, \quad (t \geq 0).$$

The positive value δ is incorporated to prevent wild fluctuations near $t = 0$. There is an initial lockout period t^* during which the option may not be exercised, but at any time between t^* and T the holder may exercise the option and receive the option payoff $(A_t - K)^+$. In fact, Longstaff & Schwartz allow only Bermudan exercise, with 100 equally-spaced exercise dates per year, so a better description of what they have computed would be Bermudan-Asian; we shall compute values for the American-Asian option, where there is unrestricted exercise between t^* and T . The values (reported below in Table 4) are in any case very close; the Longstaff-Schwartz figures quoted there for the finite-difference value of the option would rise very slightly if the full American exercise were allowed, thereby closing the already small gap.

The simulation used linear combinations of three Lagrangian martingales, which needed to be chosen with some care, in the light of the derivative being hedged. In

the notation of Section 2, the discounted exercise value of the option is

$$Z_t = e^{-rt}(A_t - K)^+ I_{\{t^* \leq t\}},$$

and unlike the previous two examples, this has no martingale part; the paths are absolutely continuous, except possibly at t^* . This means that we cannot simply follow the recipe of taking the martingale part of Z as one of the candidate hedging martingales, but we can still make certain observations. Firstly, there would never be exercise at a time when $A_t \leq K$; and secondly, there would never be exercise at a time when

$$G_t \equiv e^{-rt} \left[\frac{S_t - A_t}{t + \delta} - r(A_t - K) \right]$$

were positive: the interpretation of G is that it is the derivative of Z with respect to t , and it is clear that if the exercise value were increasing, then optimal exercise requires the holder to wait to exercise, since the value will assuredly rise in the next small instant of time.

The payoff of the European-style analogue would be the positive part of $e^{-rT}(A_T - K)$, and it is easy to work out that

$$M_0(t) \equiv E[e^{-rT}(A_T - K) | \mathcal{F}_t] = e^{-rT} \left\{ \frac{\int_{-\delta}^t S_u du + S_t(e^{r(T-t)} - 1)/r}{T + \delta} - K \right\}.$$

It follows from this that

$$dM_0(t) = H_1(t) d\tilde{S}_t,$$

where $\tilde{S}_t = e^{-rt} S_t$ is the discounted share price process, and

$$H_1(t) = \frac{1 - e^{-r(T-t)}}{r(T + \delta)}.$$

Guided by this, we choose the first martingale in the hedging set to be

$$dM_1(t) = I_{\{G_t < 0, t \geq t^*, M_0(t) > 0\}} dM_0(t). \quad (4.2)$$

For the second, we take the closely-related martingale

$$dM_2(t) = I_{\{t \geq t^*, M_0(t) > 0\}} dM_0(t). \quad (4.3)$$

As for the third martingale, we consider the European-style problem, whose value at time t will be

$$M'_3(t) = E[e^{-rT}(A_T - K)^+ | \mathcal{F}_t] = \frac{E[e^{-rT} \{ \int_t^T S_u du - (K(T + \delta) - \int_{-\delta}^t S_u du) \} | \mathcal{F}_t]}{T + \delta}. \quad (4.4)$$

Now there is no closed-form expression for this, but it is known (see Levy (1990)) that by approximating the conditional distribution of $\int_t^T S_u du$ by a log-normal distribution with matching first two moments, we get quite similar numerical values. With this simplifying assumption, the conditional expectation in (4.4) can be expressed as a Black-Scholes-like formula; even though this new expression will not be a martingale, we take for M_3 its martingale part, when $G_t < 0$ and $t > t^*$.

Table V: American-Bermudan-Asian option prices. The parameters were $\sigma = 0.2$, $K = 100$, $t^* = 0.25 = \delta$, and $T = 2$. The optimisation was based on 1000 simulated paths with 40 time-steps, and the subsequent simulation used a further 30000 simulated paths. The martingales used are specified in the text.

A_0	S_0	FD price	MC price	SD	MAD	λ_1	λ_2	λ_3	time
90	80	0.949	0.9525	0.0179	1.5681	1.0000	0.0000	0.0000	151.64
90	90	3.267	3.297	0.0312	3.8778	-1.2691	2.9607	0.0000	136.09
90	100	7.889	7.8916	0.0404	5.8087	2.9753	2.977	0.0000	144.10
90	110	14.538	14.5746	0.0521	7.3824	4.2962	2.5357	-0.4190	163.14
90	120	22.423	22.5133	0.0554	7.7752	-3.7813	2.7438	7.4299	164.38
100	80	1.108	1.0942	0.0191	1.7583	1.0000	0.0000	0.0000	113.52
100	90	3.710	3.6967	0.0348	4.3349	-0.8784	2.049	0.0000	158.80
100	100	8.658	8.7524	0.0399	5.6947	2.2548	3.464	0.0000	161.07
100	110	15.717	15.9126	0.0539	7.7434	8.0858	2.3646	-4.5773	163.37
100	120	23.811	23.9239	0.0565	7.9883	0.5154	2.5132	3.3145	164.15
110	80	1.288	1.2655	0.0207	1.9862	-0.4028	1.2406	0.0000	158.22
110	90	4.136	4.4089	0.0286	3.9949	-2.3167	3.6033	2.6306	160.86
110	100	9.821	10.3594	0.0376	5.2434	3.0337	3.5928	-1.6677	163.05
110	110	17.399	17.6839	0.0467	6.6736	5.4616	3.0411	-3.345	163.08
110	120	25.453	25.6614	0.0551	7.9072	13.6087	2.8399	-11.266	128.31

5 Conclusions.

This paper presents a simple method for evaluating the prices of American-style options by a direct simulation approach, based on a dual characterisation of the optimal exercise problem. The method involves the choice of a suitable Lagrangian hedging martingale, which can be thought of as a hedging strategy designed to minimise the lookback value of the excess of the option exercise value over the chosen hedging strategy. A choice of the hedging strategy gives bounds on expected shortfall (evaluated through simulation).

Even using very primitive choices for the hedging martingales, the agreement with other numerical methods in the three examples considered is remarkably good, usually in the range 1%–2%, or better in the first two examples. Errors of this order are already present in the problem, in the estimates one would need of volatilities, or in the assumption of constant interest rates.

Further work remains to be done in deriving better hedging martingales, and in reducing the number of simulations required, but already the method looks promising.

Appendix: Duality.

There is a short but sweet convex duality story to be told about the main result, Theorem 1, which appears as an example of the minimax principle when suitably interpreted. Many of the ideas are already present in Davis & Karatzas (1994). Recall that we assume that

$$\text{for some } p > 1, \sup_{0 \leq t \leq T} |Z_t| \in L^p,$$

and that the paths of Z are right continuous with left limits.

The convex duality story requires two convex sets: for the first we use $H_0^1 \equiv \{M \in H^1 : M_0 = 0\}$, and for the other we take the collection

$$\mathcal{A} = \{\text{right-continuous increasing processes } C, C_0 = 0, C_T = 1\}.$$

Notice that the processes in \mathcal{A} are assumed jointly measurable, but *not* adapted. There is a pairing on $H_0^1 \times \mathcal{A}$ defined by

$$(M, C) \mapsto (M, C) \equiv E \int_0^T M_t dC_t.$$

In view of our assumptions, this pairing is finite-valued. Now we define a function on \mathcal{A} by

$$C \mapsto \Phi(C) \equiv E \left[\int_0^T Z_s dC_s \right] :$$

evidently, Φ is convex (in fact, linear) and in view of our assumptions, finite-valued. Now for any $M \in H_0^1$,

$$\sup_{C \in \mathcal{A}} \{\Phi(C) - (M, C)\} = E \left[\sup_{0 \leq t \leq T} (Z_t - M_t) \right]. \quad (\text{A.1})$$

The easy minimax inequality gives us

$$\inf_{M \in H_0^1} \sup_{C \in \mathcal{A}} \{\Phi(C) - (M, C)\} \geq \sup_{C \in \mathcal{A}} \inf_{M \in H_0^1} \{\Phi(C) - (M, C)\} \quad (\text{A.2})$$

To understand the infimum on the right-hand side of (A.2), suppose that C is held fixed, and that \tilde{C} denotes the dual optional projection of C (see, for example, Rogers & Williams (2000), Chapter VI for definition and properties of the dual optional projection.) We have

$$\begin{aligned} (M, C) &\equiv E \int_0^T M_t dC_t \\ &= E \int_0^T M_t d\tilde{C}_t \\ &= E(M_T \tilde{C}_T) \end{aligned}$$

and if we now seek to take the supremum of this expression over all $M \in H_0^1$, we obtain an infinite value unless \tilde{C}_T is almost surely constant; and that constant must be 1, since $C_T = 1$. If we *do* have that $\tilde{C}_T = 1$, then $(M, C) = 0$; thus the right-hand side of (A.2) is

$$\sup_{\tilde{C} \in \tilde{\mathcal{A}}} \Phi(\tilde{C}),$$

where

$$\tilde{\mathcal{A}} = \{\text{right-continuous adapted increasing processes } \tilde{C}, \tilde{C}_0 = 0, \tilde{C}_T = 1\}.$$

Now for any $\tilde{C} \in \tilde{\mathcal{A}}$ and any $t \in [0, 1)$ we may define the stopping time

$$\tau_t \equiv \inf\{s : \tilde{C}_s > t\},$$

and we may rewrite

$$\Phi(\tilde{C}) = E\left[\int_0^T Z_s d\tilde{C}_s\right] = \int_0^1 EZ(\tau_t) dt = EZ(\tau^*),$$

say, where τ^* is the *randomised* stopping time τ_U , where U is chosen uniformly from $[0, 1]$ independently of everything else. Thus if \mathcal{T}^* denotes the class of randomised stopping times, the right-hand side of (A.2) becomes

$$\sup_{C \in \mathcal{A}} \inf_{M \in H_0^1} \{\Phi(C) - (M, C)\} = \sup_{\tilde{C} \in \tilde{\mathcal{A}}} \Phi(\tilde{C}) = \sup_{\tau^* \in \mathcal{T}^*} EZ(\tau^*) = \sup_{\tau} EZ_{\tau}, \quad (\text{A.3})$$

the last equality being evident.

Turning to the left-hand side of (A.2), we have that

$$\sup_{C \in \mathcal{A}} \{\Phi(C) - (M, C)\} = \sup_{C \in \mathcal{A}} E\left[\int_0^T (Z_s - M_s) dC_s\right] = E\left[\sup_{0 \leq t \leq T} (Z_t - M_t)\right],$$

so taking the infimum over $M \in H_0^1$ transforms the inequality (A.2) into

$$\begin{aligned} \inf_{M \in H_0^1} E\left[\sup_{0 \leq t \leq T} (Z_t - M_t)\right] &= \inf_{M \in H_0^1} \sup_{C \in \mathcal{A}} \{\Phi(C) - (M, C)\} \\ &\geq \sup_{C \in \mathcal{A}} \inf_{M \in H_0^1} \{\Phi(C) - (M, C)\} = \sup_{\tau} EZ_{\tau}. \end{aligned}$$

The reverse inequality is part of Theorem 1, and the statement of that Theorem can be reinterpreted as a minimax equality.

References

AIT-SAHALIA, F., & CARR, P. (1997). American options: a comparison of numerical methods. In *Numerical Methods in Finance*, eds. L. C. G. Rogers and D. Talay, Cambridge University Press, Cambridge.

- BARRAQUAND, J. & MARTINEAU, D. (1995). Numerical valuation of high dimensional multivariate American securities, *Journal of Financial and Quantitative Analysis* **30**, 383–405.
- BROADIE, M., & DETEMPLE, J. (1997). Recent advances in numerical methods for pricing derivative securities. In *Numerical Methods in Finance*, eds. L. C. G. Rogers and D. Talay, Cambridge University Press, Cambridge.
- BROADIE, M., & GLASSERMAN, P. (1997). Pricing American-style securities using simulation, *Journal of Economic Dynamics and Control* **21**, 1323–1352.
- BROADIE, M., GLASSERMAN, P., & JAIN, G. (1997). Enhanced Monte Carlo estimates for American option prices, *Journal of Derivatives* **5**, 25–44.
- CARR, P. (1998). Randomization and the American put, *Review of Financial Studies* **11**, 597–626.
- CARRIERE, J. (1996). Valuation of early-exercise price of options using simulations and nonparametric regression, *Insurance: Mathematics and Economics* **19**, 19–30.
- DAVIS, M. H. A., and KARATZAS, I. (1994). A deterministic approach to optimal stopping, with applications. In *Probability, Statistics and Optimisation: a Tribute to Peter Whittle*, ed. F. P. Kelly, 455–466, John Wiley & Sons, New York & Chichester.
- DELLACHERIE, C., & MEYER, P.-A. (1980). *Probabilités et Potentiel: Théorie des Martingales*, Hermann, Paris.
- HARTLEY, P. M. (2000). Pricing a multi-asset American option. Working paper, University of Bath.
- LEVY, E. (1990). Asian arithmetic. *RISK* **3**(5), 7–8.
- LONGSTAFF, F. A., & SCHWARTZ, E. A. (1999). Valuing American options by simulation: a simple least-squares approach. Working paper, University of California, Los Angeles.
- RAYMAR, S., & ZWECHER, M. (1997). A Monte Carlo valuation of American call options on the maximum of several stocks, *Journal of Derivatives* **5**, 7–23.
- ROGERS, L. C. G., & WILLIAMS, D. (2000). *Diffusions, Markov Processes and Martingales, Volume 2*, Cambridge University Press, Cambridge.