Monte Carlo valuation of American options

L.C.G. Rogers[1] *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, four 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 (1997a), (1997b), Broadie, Glasserman & Jain (1997), Raymar & Zwecher (1997), Carr (1998), Longstaff & Schwartz (2001) and Fu, Laprise, Madan, Su & Wu (2001). 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

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

reader 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 option, 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)); a Bermudan max-call (see Broadie & Glasserman (1997b)); an American-Bermudan-Asian example of Longstaff & Schwartz (2001). 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.

After the first draft of this paper was written, the author became aware of a working paper of Haugh & Kogan (2001), in which essentially the same dual approach to pricing of American options is advanced. Haugh & Kogan's numerical approach is to apply methods from neural nets to estimate the payoff function of continuing.

2 The price of an American option.

We fix some finite time horizon $T > 0$, and suppose given on some filtered probability space [3] $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0\leq t\leq T}, P)$ 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

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

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

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

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

$$Y_t^* = Y_0^* + M_t^* - A_t^*, \tag{2.3}$$

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

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

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.

Theorem 1

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

where H_0^1 is the space of martingales M for which $\sup_{0\le t\le T}|M_t| \in L^1$, 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 $E[\sup_{0\le t\le T}(Z_t - M_t)]$.* 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 $\sup_{0\le t\le T}|M_t^*| \le \sup_{0\le t\le T} z_t + |Y_0^*| + A_T$, proving that M^* is indeed in H_0^1.

Returning to the definition (2.1) of Y_0^*, we have for any $M \in H_0^1$ that

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

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 $Z_t \le Y_t^* = Y_0^* + M_t^* - A_t^*$,

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

as claimed.

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

Remark. Of course, a conditional form of Theorem 1 holds too.

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

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

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 $t \in [0, T]$

$$Z_t \le \eta + M_t,$$

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

$$Z_t \le E[\ \eta_T - \eta_0 \mid \mathcal{F}_t\ ] + (M_t + \eta_0) \tag{3.2}$$

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

$$E[\ \eta_T - \eta_0 \mid \mathcal{F}_t\ ]^+ \le E[\ (\eta_T - \eta_0)^+ \mid \mathcal{F}_t\ ] \tag{3.3}$$

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 $\sup_{u \le t}(Z_u - M_u)$. 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:

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

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.

(iii) The hedging martingales used in the various numerical examples in the next Section generally do not provide a good hedge in the sense that the mean absolute deviation is small; nevertheless, there are possible variants of the approach used here which it is intended will be explored in a subsequent paper.

4 Numerical examples.

In this section, we report the results of numerical studies of four examples, the standard American put, the American min-put (see Hartley (2000)), the Bermudan max-call (see Broadie & Glasserman (1997b)), and the Asian-Bermudan-American example studied by Longstaff & Schwartz (2001); firstly, though, we describe the general approach used in all four examples.

The first step is to simulate a relatively small number of sample paths (a few hundred, never more than 1000 in the current study) at a relatively coarse spacing of the time points (of the order of 40 time-steps). 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 three examples, but not for the last 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 $E[\, \sup_{0\le t\le T}(Z_t - \lambda \cdot M_t)\,]$, 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 (of the order of 10^4 here, but ideally more) and the corresponding martingales $\lambda^* \cdot M$, and then compute the average value of $E[\, \sup_{0\le t\le T}(Z_t - \lambda^* \cdot M_t)\,]$ 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 with little 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 $2n$ time-steps, we reduced the sample paths to observations at even-numbered times, recomputed the answer for these coarser paths with n 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

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

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 discounted value 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 first column gives the Black-Scholes values for the corresponding European option. The column of true American prices is quoted from the paper of Ait-Sahalia & Carr (1997), using their averaged binomial figures with 1000 time points. We next give the Monte Carlo values from the present method for comparison; there were

300 paths used in the optimisation step,and a further 5000 paths thereafter, with 50 time steps in each simulation, and Richardson extrapolation to give the figures in the Table. In every case, there is agreement to within 0.63%, with an average error of 0.34%. In contrast, the early exercise premium ranges from 2.1% up to 4.2%, and is in all but two cases over ten times the size of the error in our Monte Carlo value. The standard error 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. The final column presents the time taken by the entire calculation (on a 600MHz PC). The optimal value of λ (not reported) is very close to 1 in all cases, so 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 5000 simulations, the times taken were of the order of 10s. The calculations were performed throughout in Scilab[4] and can be expected to speed up considerably if coded throughout in a compiled language (typically one expects a speed-up of 5-10 times when going to compiled code, more if there is a lot of looping in the Scilab code).

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 $\sigma = 0.4$.

$S(0)$	European (true)	American (true)	American (MC)	Standard error	MAD	time (seconds)
80	20.6893	21.6059	21.6953	0.0037	0.2148	10.39
85	17.3530	18.0374	18.1008	0.0040	0.2367	10.45
90	14.4085	14.9187	14.9692	0.0038	0.2180	10.16
95	11.8516	12.2314	12.2685	0.0027	0.1413	10.30
100	9.6642	9.9458	9.9703	0.0027	0.1137	9.83
105	7.8183	8.0281	8.0439	0.0024	0.0977	9.73
110	6.2797	6.4352	6.4757	0.0054	0.2955	10.24
115	5.0113	5.1265	5.1363	0.0016	0.0548	10.42
120	3.9759	4.0611	4.0761	0.0036	0.1649	9.81

Theorem 1 tells us that if we use the martingale of the Doob-Meyer decomposition of the Snell envelope then the hedge should be exact. To check this out, we carried out the same estimation of the price as is reported in Table I, but using the martingale part of the (discounted) value of the American put in place of the discounted value of the European put; the results are in Table II. Of course, the value function has to be computed numerically, using once again 50 time steps and a Crank-Nicolson/SOR

[4]Scilab is a powerful computation package, similar in performance and ease of use to Matlab. It is available for most commonly used platforms by free download from the Scilab home page http://www-rocq.inria.fr/scilab/scilab.html

method. The results were even closer to the true values than the results in Table I; the maximum error was 0.27%, with an average error of 0.12%. Notice that sometimes the Monte Carlo value is several standard errors below the quoted true figure; this suggests that there may be small errors in the figures of Ait-Sahalia & Carr, and comparison with the values computed by the Crank-Nicolson/SOR method indicate that the fourth or fifth significant figure may be doubtful.

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

$S(0)$	American (true)	American (MC)	Standard error	MAD
80	21.6059	21.6270	0.0180	0.0879
85	18.0374	18.0534	0.0176	0.0535
90	14.9187	14.9366	0.0181	0.0964
95	12.2314	12.2234	0.0006	0.0369
100	9.9458	9.9398	0.0010	0.0626
105	8.0281	8.0202	0.0007	0.0444
110	6.4352	6.4268	0.0004	0.0194
115	5.1265	5.1403	0.0034	0.1947
120	4.0611	4.0546	0.0002	0.0125

Example 2: American min puts on 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 III and V 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 III we took the volatilities of both assets to be 0.6, whereas in Table V the volatilities were $\sigma_1 = 0.4$ and $\sigma_2 = 0.8$. In the example of Table III, we exploit the symmetry of the problem: since the assets are independent and have the same dynamics, we suppose that the weights on the two basic martingales are *the same*, so that in fact the calculation reported in Table III is a calculation based on the use of a single hedging martingale.

The tables give a European price (computed by numerical integration), 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 everywhere less than 1%. Times for the Monte Carlo pricing method are of the order of 180s; 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 III: Simulation prices of min-puts on two assets. Parameter values were $K = 100$, $T = 0.5$, $r = 0.06$, $\sigma_1 = \sigma_2 = 0.6$. The unique hedging martingale is whichever European put is in the lead. 1000 paths are used for the optimisation, followed by a further 10000 paths.

$S_1(0)$	$S_2(0)$	European	FD price	MC price	SE	MAD	time
80	80	36.859	37.30	37.63	0.088	7.2894	177.96
80	100	31.639	32.08	32.30	0.078	6.3310	177.67
80	120	28.652	29.14	29.38	0.062	4.7030	177.70
100	100	24.728	25.06	25.17	0.079	6.3895	177.65
100	120	20.610	20.91	21.10	0.068	5.3851	198.74
120	120	15.704	15.92	16.02	0.062	4.8372	177.75

Although these results are already quite good, in view of experience with the n-max-call problem later, it seemed worth finding out whether including an exchange-type martingale in the hedging set made a significant improvement.

Table IV: Simulation prices of min-puts on two assets. Parameter values were $K = 100$, $T = 0.5$, $r = 0.06$, $\sigma_1 = \sigma_2 = 0.6$. The hedging martingales are the European put in the lead, and the exchange martingale specified in the text. 450 paths are used for the optimisation, followed by a further 10000 paths.

$S_1(0)$	$S_2(0)$	European	FD price	MC price	SE	MAD
80	80	36.859	37.30	37.50	0.075	5.95
80	100	31.639	32.08	32.19	0.038	2.61
80	120	28.652	29.14	29.25	0.037	2.46
100	100	24.728	25.06	25.08	0.079	6.24
100	120	20.610	20.91	20.99	0.050	3.80
120	120	15.704	15.92	16.09	0.053	4.35

In general, when there are n assets, we can for any $i, j \in \{1, \ldots, n\}$ explicitly express the value of a European option to exchange asset i for asset j at time T. The hedging martingale to be added to the set is the martingale which at any time when the option is in the money follows the (discounted) value of the option to exchange the second cheapest asset for the cheapest. The results of this more refined analysis are reported in Table IV; we find that the worst error is 1%, with an average error of less than 0.5%. The times taken (not reported) roughly double. The MAD values are quite a lot smaller.

Table V presents the results for an asymmetric example, where the volatilities of the two assets are different. The hedging martingales used here are just the (discounted) values of the two European puts, when the corresponding share is cheapest. The percentage errors (relative to Hartley's finite-difference prices) is never more than 1.6%, and averages 0.8%.

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

$S_1(0)$	$S_2(0)$	European	FD price	MC price	SE	MAD	time
80	80	37.5540	38.01	38.35	0.096	7.5528	275.45
80	100	31.8075	32.23	32.60	0.099	8.1361	298.29
80	120	28.0900	28.54	29.01	0.089	7.5012	309.36
100	80	32.8564	33.34	33.59	0.071	5.3996	271.04
100	100	25.4666	25.81	26.02	0.090	6.9387	259.08
100	120	20.4767	20.75	21.05	0.092	7.4975	271.19
120	80	30.6872	31.21	31.31	0.047	2.9212	328.28
120	100	22.4413	22.77	22.83	0.059	4.2130	289.55
120	120	16.7641	16.98	16.98	0.074	5.6365	298.78

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 VI. 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. The prices quoted are computed by his algorithm (50S) with 50 time steps calculating the drift and volatility separately at each node point. 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 close to the truth; the gap between his lower bounds and our upper bounds is of the order of 1–2%. The hedging martingales used were as for the example of Table IV.

Comparing the results for the one case common to Tables III, IV and VI shows small variations; in Table III, there was only one hedging martingale (compared with two in the other two cases), and Table IV used only 450 paths for the optimisation step, with a corresponding gain in speed and loss in the SE and MAD of the answer, compared with Table VI. The MC times taken are again satisfactory for some OTC product, but would ideally be shorter; Hartley's method is producing computation times which are much better. It is worth remarking that in some cases, the 50S result of Hartley is several standard errors below the Monte Carlo European price given in the table; this is consistent with the performance of the 50S algorithm in the case of two assets, where the values achieved were always below the finite-difference value, by at most 1%.

Table VI: 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	50S price	MC price	SE	MAD	European price	SE	MC time	50S time
2	24.87	25.16	0.057	4.57	24.80	0.027	794.05	5.64
3	31.21	31.76	0.095	7.88	31.22	0.026	742.31	8.39
4	35.72	36.28	0.081	6.69	35.80	0.024	713.93	11.31
5	39.01	39.47	0.095	7.93	39.18	0.022	725.30	14.05
10	47.99	48.33	0.100	8.44	48.02	0.017	873.49	25.64
15	52.23	52.14	0.108	9.09	52.11	0.014	1175.86	41.58

Example 3: Bermudan max calls on n assets.

A benchmark example in this subject appears to be the Bermudan n-max-call, studied in Broadie & Glasserman (1997b), and used as a test example by Haugh & Kogan (2001), Fu, Laprise, Madan, Su & Wu (2001), and Andersen & Broadie (2001). Here we apply the methods developed in this paper to the pricing of the example where at any time $t = iT/d$, $i = 0, \ldots, d$, the holder of the option may exercise and receive the payoff

$$(\max(S_t^1, \ldots, S_t^n) - K)^+.$$

As usual, the assets S^i are log-Brownian motions, which are assumed to be independent and identically-distributed for the purposes of this example (though this is not of course necessary for the operation of the method.) There is a continuous dividend payout at rate δ.

In applying the methods developed in this paper, the obvious first guess at the martingales in the hedging set was to use the (discounted) value of the European call on whichever asset was in the lead at the time. This turned out to give bounds which were quite high, of the order of 3% - 5% too high, so something else was needed. This was the exchange martingale described above; at any time t, the increments of the martingale were the increments of the (discounted) value of the European option to exchange asset i for asset j at time T, where asset j was the most expensive and asset i was the second most expensive at time t. The increments were turned off when the option was out of the money (that is, all the assets were worth less than the strike K). Finally, one other martingale was added to the hedging set; this was the European call in the lead, but only switched on once the asset went into the money. This could only make a difference to the pricing when $S(0)$ was less than or equal to the strike.

The values reported in Table VII are based on simulations of 1000 paths for the optimisation, then another 8000 paths thereafter to refine the estimate. The agreement between the values of Broadie & Glasserman (1997b) and the upper bounds produced by the dual simulation method are quite good, within 2% on average, and at most 2.6%.

Table VII: Simulation prices of max-calls on 5 assets. The strike is $K = 100$ throughout, and all shares have the same start value, $S(0) = 90, 100, 110$, and $\sigma_i = 0.2$, with $\delta = 0.1$. The expiry is $T = 3$, interest rate is $r = 0.05$. Exercise can occur at any of the times $t = iT/d$, $i = 0, \ldots, d$, where $d = 3, 6, 9$ are the values used in the table

d	S_0	BG price	MC price	SE	MAD	MC time
3	90	16.006	16.24	0.060	4.48	337.73
	100	25.284	25.70	0.072	5.39	227.47
	110	35.695	36.19	0.060	4.98	208.07
6	90	16.474	16.91	0.057	4.28	299.63
	100	25.92	26.40	0.060	4.38	329.52
	110	36.497	37.18	0.065	4.74	345.98
9	90	16.659	16.98	0.061	4.54	710.71
	100	26.158	26.75	0.061	4.40	419.56
	110	36.782	37.61	0.066	4.84	431.86

Example 4: American-Bermudan-Asian option.

This is one of the examples studied by Longstaff & Schwartz (2001), 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

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

The positive value δ is incorporated to prevent wild fluctations 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, in their Monte Carlo approach, Longstaff & Schwartz 'use 100 discretization points per year to approximate the continuous exercise feature of the option,' and are therefore technically pricing a Bermudan option. Thus 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 are based on a discretization using 10,000 time steps per year, and 200 space steps in each of two dimensions.

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

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

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 obervations. Firstly, there would never be exercise at a time when $A_t \le K$; and secondly, there would never be exercise at a time when

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

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

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

For the second, we take the closely-related martingale

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

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

$$M_3'(t) = E[e^{-rT}(A_T - K)^+|\mathcal{F}_t] = \frac{E[e^{-rT}\{\int_t^T S_u\,du - (K(T+\delta) - \int_{-\delta}^t S_u\,du)\}|\mathcal{F}_t]}{T+\delta}. \tag{4.4}$$

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 $t > t^*$.

The results are presented in Table 4. The agreement between the Monte Carlo prices and the finite-difference prices is impressive, getting closer as A_0 gets smaller. In just one place the Monte Carlo price is less than the finite-difference price ($A_0 = 110$, $S_0 = 80$), but the difference is about one standard error.

Table VIII: 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	SE	MAD	λ_1	λ_2	λ_3	time
90	80	0.949	0.952	0.018	1.5681	1.0000	0.0000	0.0000	151.64
90	90	3.267	3.297	0.031	3.8778	-1.2691	2.9607	0.0000	136.09
90	100	7.889	7.892	0.040	5.8087	2.9753	2.977	0.0000	144.10
90	110	14.538	14.575	0.052	7.3824	4.2962	2.5357	-0.4190	163.14
90	120	22.423	22.513	0.055	7.7752	-3.7813	2.7438	7.4299	164.38
100	80	1.108	1.094	0.019	1.7583	1.0000	0.0000	0.0000	113.52
100	90	3.710	3.697	0.035	4.3349	-0.8784	2.049	0.0000	158.80
100	100	8.658	8.752	0.040	5.6947	2.2548	3.464	0.0000	161.07
100	110	15.717	15.913	0.054	7.7434	8.0858	2.3646	-4.5773	163.37
100	120	23.811	23.924	0.056	7.9883	0.5154	2.5132	3.3145	164.15
110	80	1.288	1.265	0.021	1.9862	-0.4028	1.2406	0.0000	158.22
110	90	4.136	4.409	0.029	3.9949	-2.3167	3.6033	2.6306	160.86
110	100	9.821	10.359	0.038	5.2434	3.0337	3.5928	-1.6677	163.05
110	110	17.399	17.684	0.047	6.6736	5.4616	3.0411	-3.345	163.08
110	120	25.453	25.661	0.055	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 four 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.

As befits a new development, there remain many interesting and important questions. It appears very easy in the few examples studied here to pick a small family of hedging martingales which will get the upper bound within 2% of the 'true' value; more ingenuity is required to pull the error down to 1%, and it seems that getting a smaller error in any multi-dimensional example gets increasingly difficult with the crude methods proposed here. Another issue is the large MAD figures resulting from these hedging policies - so large that calling them 'hedging' policies is a bit of a misnomer. Better methods are required both for hedging and for pricing. The approach used here can indeed be developed further and it is intended to present the results of such developments in a later paper; the current paper should be seen more as a beginning than as an end.

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

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

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

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

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

$$\sup_{C \in \mathcal{A}} \{\Phi(C) - (M, C)\} = E[\sup_{0 \le t \le T} (Z_t - M_t)]. \tag{A.1}$$

The easy minimax inequality gives us

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

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

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

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

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

the last equality being evident.

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

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

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

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

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.

ANDERSEN, L. & BROADIE, M. (2001). A primal-dual simulation algorithm for pricing multi-dimensional American options. Columbia University working paper.

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. (1997a). 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.

BROADIE, M., & GLASSERMAN, P. (1997b). A stochastic mesh method for pricing high-dimensional American options. Columbia University working paper.

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.

FU, M. C., LAPRISE, S. B., MADAN, D. B., SU, Y., and WU, R. (2001). Pricing American options; a comparison of Monte Carlo simulation approaches. *Journal of Computational Finance* **4**, 39–88.

HARTLEY, P. M. (2000). Pricing a multi-asset American option. Working paper, University of Bath.

HAUGH, M., & KOGAN, L. (2001). Pricing American options: a duality approach. Working paper.

LEVY, E. (1990). Asian arithmetic. *RISK* **3**(5), 7–8.

LONGSTAFF, F. A., & SCHWARTZ, E. A. (2001). Valuing American options by simulation: a simple least-squares approach. *Review of Financial Studies* **14**, 113-147.

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.