PRICING DISCRETELY SAMPLED PATH-DEPENDENT EXOTIC OPTIONS USING REPLICATION METHODS

M. S. JOSHI

Abstract. A semi-static replication method is introduced for pricing discretely sampled path-dependent options. It depends upon buying and selling options at the reset times of the option but does not involve trading at intervening times. The method is model independent in that it only depends upon the existence of a pricing function for vanilla call options which depends purely on current time, time to expiry, spot and strike. For the special case of a discrete barrier, an alternative method is developed which involves trading only at the initial time and the knockout time or expiry of the option.

1. Introduction

In this note, we examine the problem of pricing path-dependent exotic options via the use of a new replication method. Our fundamental assumption is that the only state variable is spot and that it in combination with the current time determines the price of any specified call or put option. One advantage of this approach is that it naturally allows the inclusion of smile information in the pricing.

Thus we assume the existence of a deterministic function,

$$\text{Price}(S, K, t, T),$$

which gives the price of a call option with expiry T and strike K, at time t when spot is S. We shall say that a model which implies this property is a deterministic-smile model as it implies that smiles at future times are determined by the value of spot. Several popular models can be fit into this framework. In particular, the Black-Scholes model, the Dupire model, [9], Merton's jump-diffusion model, [17], and the variance gamma model, [14, 15, 16], all have this property. We remark that stochastic volatility models, eg [11], do not have this property as the volatility is a second state variable.

Whilst our method is quite general and applies to a large class of options, we develop it in detail in Sections 2, 3 for the discretely sampled arithmetic Asian option, and then indicate in Section 4 how to adapt it to many other options. Recall that a discretely sampled Asian option is an option on the average of the spot at certain specified times. Thus if the specified times are $t_1 < t_2 < \cdots < t_n$, and the strike of the option is K, then the Asian call and Asian put will pay at time t_n

$$\left(\frac{1}{n}\sum_{j=1}^{n} S_{t_j} - K\right)_+ \quad \text{and} \quad \left(K - \frac{1}{n}\sum_{j=1}^{n} S_{t_j}\right)_+, \tag{1.1}$$

respectively.

Our replication method is semi-static in that we only rehedge at the times t_j but dynamic in that at those times, we buy and sell portfolios of options depending on the value of spot and the realized path up to that point. We emphasize that we regard our trading strategy as a computational device rather than as a strategy which would actually be carried out by a trader.

The interesting fact about (1.1) is that whilst the payoff depends on the value of spot at the times $t_1, \dots, t_n$ it does so in a particularly simple way which is essentially one-dimensional. In particular, if we let

$$A_i = \frac{1}{i}\sum_{j=1}^{i} S_{t_j}, \tag{1.2}$$

then we have that

$$A_{i+1} = \frac{i}{i+1}A_i + \frac{1}{i+1}S_{t_{i+1}}. \tag{1.3}$$

Thus the value of an Asian option at time $t = t_i$, is a function of S_{t_i} and A_i – the precise values of $S_{t_1}, \dots, S_{t_{i-1}}$ are irrelevant as all relevant information has been encoded in A_i.

The upshot of this is that we can apply backwards methods to the pricing of Asian options, provided we introduce the auxiliary variables, A_i, the realized average of the path up to time t_i. In fact, this is a standard approach to implementing PDE methods for the pricing of Asian options. See [2, 6, 12, 20].

The fundamental difference between our approach and standard numerical PDE methods, is that we only solve for option values at the times t_i without looking at the intervening time steps. Indeed, if the reader is particularly wedded to PDE approaches, and the model being used has a PDE interpretation then one can view the replicating options as being a basis for the solutions of the PDE, [13]. It is also important to realize that this approach does not require the existence of a PDE describing the price evolution. The similarity here between PDE methods and replication methods suggests a vague general rule that if an option can be priced using PDEs in the Black-Scholes world then it can be priced using replication methods for any deterministic smile model.

One advantage of our approach is that the initial portfolio for the replication is independent of the initial value of spot. This means that the price, delta and gamma can be immediately read off from the portfolio for any value of spot. As the replicating portfolio does not change until the first averaging time, we can in fact get value, delta, gamma and theta for any value of spot and time before the first averaging time just by evaluating the relevant quantity for the replicating portfolio.

Another advantage from a more conceptual viewpoint is that we require no assumptions on the process of the underlying except the existence of the pricing function, and so we have proven our replication result simultaneously for all processes with deterministic pricing functions. This means that we do not need a process; if we wish to examine the effect of an arbitrarily specified pricing function then we can do so provided the function is arbitrage-free. Note that this means that the inclusion of smile information is natural in this model.

In the special case of a discrete barrier option, a simpler trading strategy can be employed which involves buying a portfolio of vanilla options at time zero and selling it when the option knockouts or expires. We develop this approach in Section 6. That a discrete barrier option can be hedged in this manner for the Black-Scholes model and Merton's jump-diffusion model has already been observed by Andersen, Andreasen and Eliezer, [1]. However, they do not appear to have developed the method as a pricing tool and their arguments are model dependent.

We place our results in historical context. Replication is one of the main ideas in modern derivatives pricing. We can think in terms of a hierachy of assumptions. We list some possible assumptions and then discuss which are needed for each trading strategy. We also make the general assumptions of no trading costs and heteroskedasticity throughout.

Assumptions on the underlying:

- A1 there exists a liquid market in the underlying (or forwards) at all times
- A2 the underlying follows a Markovian process
- A3 the underlying follows a continuous process
- A4 the underlying follows a diffusive process
- A5 the underlying follows a log-normal process

Assumptions on the vanilla options markets:

- B1 there exists a liquid market in calls and puts of all strikes and maturities today
- B2 there exists a liquid market in calls and puts of all strikes and maturities at all times
- B3 the prices of calls and puts satisfy "call-put symmetry" conditions at all times
- B4 the price of calls and puts are a known deterministic function of calendar time, spot, strike and maturity.

We give a classification of replication methods also.

- C1 strong static : the option pay-off can be perfectly replicated by a finite portfolio of calls, puts and the underlying set-up today with no further trading
- C2 mezzo static: the option pay-off can be perfectly replicated by a finite portfolio of calls, and puts set-up today together with a finite number of trades in the underlying
- C3 weak static: the option pay-off can perfectly replicated by setting up a finite portfolio of calls and puts today which may be sold before expiry
- C4 feeble static: the option pay-off can perfectly replicated by trading a finite number of calls and puts at a finite number of times
- C5 dynamic: the option pay-off can be perfectly replicated by continuous trading in the underlying

We shall also use the term "almost" to indicate that the pay-off can be replicated arbitrarily well with a finite portfolio rather than perfectly.

If the underyling satisfies A1-5 then C5 holds; this is the fundamental result of Black and Scholes, [3]. This still holds under A1-A4, Dupire [9].

Under assumption B1 then C1 holds for a straddle with no assumptions on the underlying. We also have that under the assumption B1 that digital European options can be almost strong statically replicated, by approximating using a call-spread. Unfortunately, strong static replication holds for very few options.

If we make the assumptions A1, A3 as well as B1 then more options can be mezzo statically replicated. For example, an up-and-in put option struck at K with up-barrier at K can be replicated. Purchase a call option struck at K. At the first time that the spot reaches K, go short a forward contract struck at K (which is of zero cost.) The forward turns the call into a put and the pay-off is replicated. If the spot never crosses K then the call option and the original both pay zero at expiry. See [4, 10].

Under B1, B2, B3, A3 it was shown in [5] that a class of barrier options can be weak statically replicated. For example, a down-and-out call can be replicated by holding a call option with the same expiry and going short a put option with strike below the barrier and the same expiry in such a way as to guarantee that the resultant portfolio has zero value on the boundary.

If we assume B1, B2 and B4 then we are in the situation of this paper. The method we present for hedging discrete barrier options under these assumptions is then an almost weak static replication. Note that we can also hedge continuous barrier options arbitrarily well by approximating with a discrete-barrier option with an arbitrarily large number of sampling dates.

Our method of replication for a general path-dependent exotic option makes the same assumptions, but it is almost feeble static in that it requires trading in options at multiple times.

Note that whilst these methods make no assumptions on the underlying, it is difficult to imagine a situation where B1, B2 and B4 hold but A2 does not.

If we make the additional assumption of continuity of sample paths, A3, then a simpler method can be used to almost weak replicate continuous barrier options. This method relies on dissolution of the portfolio at the instant the barrier is crossed, and therefore only requires the replicating portfolio to be of zero value on the barrier rather than behind it. See [19]. Under these assumptions, it is also possible to replicate American options, [13]. In fact, one does not really need continuity of sample paths, the crucial property is that the spot cannot jump across the barrier for knock-out options or into the exercise domain for American options. These techniques could therefore be applied in markets where only down jumps occur, for example equities, to the pricing of up-and-out barrier options and American call options.

This research was carried out while working in the Quantitative Research Centre of the Royal Bank of Scotland Group and I am grateful to my colleagues, Christopher Hunter, Peter Jäckel, and Riccardo Rebonato, for many helpful conversations.

2. The method for Asian options

In this section, we focus on an Asian call option for concreteness. We observed that the value of an Asian option at a time t_i depends only on the current value of spot and the value of the auxiliary variable A_i. Our approach is therefore to construct approximating portfolios of options with the composition of the portfolio depending upon spot and A_i. We do this by backwards iteration.

At the final time, $T = t_n$, the value of the option is $(A_n - K)_+$. We can rewrite this as

$$\left(\frac{n-1}{n}A_{n-1} + \frac{1}{n}S_{t_n} - K\right)_+ = \frac{1}{n}\left(S_{t_n} - (nK - (n-1)A_{n-1})\right)_+ . \tag{2.1}$$

This means that at time t_{n-1} a call option struck $(nK - (n-1)A_{n-1})$ of notional $1/n$ precisely replicates the final pay-off. Note that the replicating portfolio here depends upon A_{n-1} but not $S_{t_{n-1}}$.

We have assumed the existence of a pricing function so we immediately have that the value of this replicating call option is determined as a function of $S_{t_{n-1}}$ for each value of A_{n-1}. By no-arbitrage the value of the Asian call option must be equal to the value of this replicating option and we therefore know its value as a function of $S_{t_{n-1}}$ and A_{n-1}.

At time t_{n-2}, we now wish to construct portfolios of options expiring at time t_{n-1} whose payoffs precisely replicate the value of the Asian call at time t_{n-1}.

Recall that the value of A_{n-1} is equal to

$$\frac{n-2}{n-1}A_{n-2} + \frac{1}{n-1}S_{t_{n-1}}.$$

This means that the set of points in the $(S_{t_{n-1}}, A_{n-1})$ plane reachable from a point $(S_{t_{n-2}}, A_{n-2})$ is a line, and which line depends purely upon the value of A_{n-2} and not $S_{t_{n-2}}$. In particular, the line of points reachable is

$$\left(S_{t_{n-1}}, \frac{n-2}{n-1}A_{n-2} + \frac{1}{n-1}S_{t_{n-1}}\right).$$

Thus given a value of A_{n-2}, the value of $S_{t_{n-1}}$ determines a point in the $(S_{t_{n-1}}, A_{n-1})$ plane and a price for the Asian option at time t_{n-1}. Call this price $f_{A_{n-2}}(S_{t_{n-1}})$. This means that for each value of A_{n-2}, we can replicate the value of the Asian call at time t_{n-1} by a European option paying $f_{A_{n-2}}(S_{t_{n-1}})$ at time t_{n-1}.

This European option can be approximated arbitrarily well by using a portfolio of vanilla call and put options. The given pricing function can then be used to assess the value of this portfolio for any value of $S_{t_{n-2}}$, at time t_{n-2}. Thus by valuing the portfolio associated to each value of A_{n-2} for every value of $S_{t_{n-2}}$, we develop the price of the Asian call option as a function of $S_{t_{n-2}}$ and A_{n-2} at time t_{n-2}.

We can repeat this method to get the value of the call option across each plane (S_{t_j}, A_j) for $j = 1, \dots, n-1$.

At time t_1, we have, of course, that $A_1 = S_1$ and only the values along that line are relevant. Thus in setting up the initial replicating portfolio at time zero, we replicate the values along this line in the (S_1, A_1) at time t_1.

The value of this initial replicating portfolio is now the value of the Asian call option at time zero. Note that the initial value of spot was not used in the construction so we can get the value of the Asian call as a function of spot immediately just by revaluing the initial portfolio. This also means that the delta and gamma are equal to the delta and gamma of the initial portfolio. As the replicating portfolio does not change up to the first reset time, we can also read off the theta as the theta of the replicating portfolio. Note that this argument does not extend to Greeks with respect to the other model parameters as changing the other parameters will in general affect the composition of the initial portfolio.

Having used a replicating argument to price the derivative, what is the actual trading strategy? We set up the initial replicating portfolio and hold it until the first reset time. At the first reset time, we dissolve the portfolio by exercising all the options which are in the money, as all the options are at their expiry. The sum of money received is by construction precisely the cost of setting up a new portfolio which depends upon the value of A_1 and which replicates out to the second reset time. We then exercise again and use the money to buy a new portfolio out to the third reset time and so on. At all stages, the new portfolio set up will depend on the value of the auxiliary variable, and will be equal in value to exercised value of the previous portfolio.

The existence of a deterministic pricing function is crucial as any indeterminacy in the set-up cost of the later replicating portfolios will destroy our argument.

3. The implementation

The argument in the previous section implicitly assumed a perfect replication of the value of the Asian call option across all values of spot and the auxiliary variable. Clearly, if we are to implement the pricing method in a computer we need to use a discrete approximation. The approach tested was to use the same square two-dimensional grid at each time step. The grid was taken to have a uniform size of squares in log-space.

In the implementation of the method it was sometimes necessary to assess the value of a replicating portfolio at a point for which the auxiliary value was not on the grid. This was done by linearly interpolating the prices between the neighbouring auxiliary values.

Prices obtained were compared against quasi-Monte Carlo prices for Asian call options in Black-Scholes and jump-diffusion models and found to agree to high levels of accuracy. We give computed prices for a five-year Asian call option with yearly resets. Spot and strike were taken to be 100 and interest rates were taken to be zero.

We present numbers for the Black-Scholes model with volatility 10 percent. We present simulation results for various different step sizes. With step-size log(1.005), we get

Steps	Price	Delta	Gamma
10	5.100	0.2815	0.01662
15	5.520	0.3368	0.01896
20	5.535	0.3735	0.02041
25	5.715	0.4111	0.02232
30	5.721	0.4391	0.02320
35	5.839	0.4642	0.02462
40	5.838	0.4815	0.02506
45	5.890	0.4957	0.02590
50	5.887	0.5053	0.02606
55	5.909	0.5127	0.02650
60	5.906	0.5174	0.02653
65	5.914	0.5208	0.02673
70	5.912	0.5229	0.02673
75	5.914	0.5243	0.02681
80	5.913	0.5251	0.02680

With step-size log(1.01),

Steps	Price	Delta	Gamma
10	5.267	0.3570	0.01897
15	5.733	0.4310	0.02294
20	5.771	0.4743	0.02428
25	5.894	0.5026	0.02603
30	5.896	0.5156	0.02627
35	5.916	0.5223	0.02672
40	5.915	0.5249	0.02673
45	5.917	0.5259	0.02680
50	5.917	0.5262	0.02680
55	5.917	0.5263	0.02680
60	5.917	0.5264	0.02680
65	5.917	0.5264	0.02680
70	5.917	0.5264	0.02680
75	5.917	0.5264	0.02680
80	5.917	0.5264	0.02680

With step-size log(1.02), we obtain

Steps	Price	Delta	Gamma
10	3.101	0.172	0.006
15	3.814	0.253	0.007
20	4.155	0.297	0.008
25	4.651	0.350	0.010
30	4.927	0.381	0.011
35	5.260	0.422	0.014
40	5.429	0.445	0.016
45	5.617	0.472	0.019
50	5.706	0.486	0.021
55	5.798	0.501	0.023
60	5.839	0.508	0.024
65	5.880	0.516	0.025
70	5.897	0.519	0.025
75	5.913	0.522	0.026
80	5.920	0.524	0.026

We compare with a Quasi-monte-carlo method in which Greeks are computed by finite differencing.

Steps	Price	Delta	Gamma
2097152	5.910	0.5263	0.0166
1048576	5.910	0.5264	0.0184
524288	5.910	0.5263	0.0141
262144	5.910	0.5262	0.0091
131072	5.910	0.5262	0.0060
65536	5.909	0.5262	0.0000
32768	5.908	0.5265	0.0000
16384	5.910	0.5277	0.0000
8192	5.910	0.5272	0.0000
4096	5.899	0.5270	0.0000

The Gamma here has clearly not converged even after two million paths.

4. Other options

So far we have concentrated on the study of the discrete Asian call option, however the method can be used to price other path-dependent exotic options. We also remark that the addition of Bermudan-style features causes no new difficulties.

What properties are necessary for implementing this method? The pay-off should be dependent on the value of spot on a finite set of times $t_1, \dots, t_n$. There should exist a sequence of functions A_j such A_j is a function of $(S_{t_1}, \dots, S_{t_j})$, and a second sequence f_j of updating function such that

$$A_{j+1} = f_{j+1}(S_{t_{j+1}}, A_j). \tag{4.1}$$

The crucial property is that the final pay-off should depend only on A_n and S_{t_n}.

The geometric-average Asian option is easily fit into the framework. It pays on the geometric average of the spot at the sampling dates instead of the arithmetic average. We take

$$A_j = \left(\prod_{i=1}^{j} S_{t_i} \right)^{\frac{1}{j}}. \tag{4.2}$$

The updating function is then

$$f_{j+1}(S_{t_{j+1}}, A_j) = \left(A_j^j S_{t_{j+1}} \right)^{\frac{1}{j+1}}. \tag{4.3}$$

A look-back option pays the positive part of the maximum value of $S_{t_j}, j = 1, \dots, n$, minus the strike. In this case, A_j is the maximum of S_{t_i} for $i \le j$. The updating function f_j is just the maximum of A_{j-1} and S_{t_j}.

A forward-starting option pays the sum

$$(S_{t_2} - \alpha S_{t_1})_+$$

for some specified α. In this case,

$$A_1 = A_2 = S_{t_1}$$

and the pay-off is

$$(S_{t_2} - \alpha A_{t_2})_+.$$

This option is sometimes called a cliquet. A related option is the ratio cliquet which pays

$$\left(\frac{S_{t_2}}{S_{t_1}} - 1 \right)_+.$$

In this case, we take $A_1 = A_2 = S_{t_1}$, and the pay-off is trivially purely-dependent on A_2 and S_{t_2}.

We can also price a growth certificate which pays a multiple of the stock price at time t_3 if the price at time t_2 is greater than the price at time t_1, and zero otherwise. Here we take $A_1 = S_{t_1}$. We set A_2 equal to one or zero according to whether S_{t_2} is greater than A_1. We set $A_3 = A_2$. The pay-off is then $S_{t_3} A_3$. Note that there are only two possible auxiliary values for this option which massively speeds up implementations.

Discrete barrier options can similarly be fit into this framework. Suppose the option pays $f(S_{t_n})$ at time t_n, provided the spot lies in a given interval I_j at each of the times t_j. Let $g_j(x)$ equal one if $x \in I_j$ and zero otherwise. We let A_1 equal $g_1(S_{t_1})$ and A_j equal $g_j(S_{t_j})A_{j-1}$, for each j greater than 1. The final pay-off is then $f(S_{t_n})A_n$. As with the growth certificate, we require only two auxiliary values for this implementation.

In conclusion, the method presented here allows the rapid pricing of many path-dependent options using alternative models which allow the incorporation of smile information.

5. Additional Simplifications

If our pricing function comes from a model that admits certain homogeneity properties, and the option is also homogeneous then we can greatly reduce the number of computations necessary. In particular, suppose our pay-off is that of an Asian call option. The class of pay-offs of call options is homogeneous in that we have

$$(\alpha x - K)_+ = \alpha \left(x - K/\alpha\right)_+ . \tag{5.1}$$

This means to price for all values of K and α, we only need to price for one value of α and all values of K. We can adapt this to speed up the pricing of Asian options for certain pricing functions.

Our assumption on the pricing function is that it is the discounted expectation of a Markovian process in which for some a, the distribution of $\log(S_t + a) - \log(S_s + a)$ for $t > s$ is independent of the value of S_s. With $a = 0$, this assumption is satisfied by the Black-Scholes model with determinsistic time-dependent volatility, Merton's jump-diffusion model and the Variance Gamma model. For a non-zero, we can obtained the displaced diffusion model.

Given this assumption, at time t_j we have that the value of the option is

$$\mathbb{E}\left(\left(\frac{1}{n}\left(\sum_{j=i+1}^{n} S_{t_j} + iA_i\right) - K\right)_+\right) P(t_j, T)$$

where A_i is as in Section 2 and $P(t_j, T)$ is the value of a zero coupon bond expiring at time T at time t_j. Our assumption implies that

$$(S_{t_j} + a) = (S_{t_i} + a)X_j, \tag{5.2}$$

where $\{X_j\}$ is a collection of random variables (which will not be independent.) We can therefore rewrite the value of the Asian option as

$$\begin{aligned}(5.3)\quad & \mathbb{E}\left(\left(\frac{1}{n}\left(\sum_{j=i+1}^{n} (S_{t_i} + a)X_i - a + iA_i\right) - K\right)_+\right) P(t_j, T) \\ &= (S_{t_i} + a)\mathbb{E}\left(\left(\frac{1}{n}\sum_{j=i+1}^{n} X_i + \frac{iA_i - nK - (n-i)a}{n(S_{t_i} + a)}\right)_+\right) P(t_j, T).\end{aligned}$$

The interesting thing about the right-hand-side of (5.3) is that putting

$$\theta(A_i, S_i) = \frac{iA_i - nK - (n-i)a}{n(S_{t_i} + a)}, \tag{5.4}$$

we see that the dependence of the expectation upon A_i and S_{t_i} is now one-dimensional.

Here we have adapted a method of Rogers and Shi, [18], developed for pricing arithmetic continuous average options in a Black-Scholes world using PDE methods. This has previously been applied to PDE pricing in a Black-Scholes setting for discrete arithmetic average options by Benhamou and Duguet, [7]. Note that the Black-Scholes case would have $a = 0$ and the variables X_j all being log-normal.

6. Replicating Discrete Barriers

Suppose we wish to price a discrete-barrier knock-out option. The method we have already presented works, however it requires buying and selling options at each knock-out date. In this section, we present a method that requires buying and selling only at the initial time and at the time of knock-out or expiry. Recall that a discrete barrier option pays off at some time T, unless the price of the underlying is outside a specified range on a predetermined finite set of dates, $t_1 < t_2 < \cdots <$

$t_n < T$, in which case the option pays zero (or possibly a fixed rebate at the time of knockout.)

The technique we present here is related to that in [1] and will produce the same replicating portfolio. The essential difference is that our approach is algorithmic and relies on the existence of a deterministic pricing function for options rather on a specific process. A consequence of this is that our approach can be used for pricing whereas the arguments presented in [1] relied on the price of the knock-out option already being known by different means.

The essential idea is that we construct a portfolio of plain vanilla options which has the same value as the option being modelled at pay-off time, and has zero value at the points in spot-time where the original option knocks-out. So the option knocks-out if the spot passes below the value B at any of the times t_1 through t_n. Our portfolio should therefore have zero value on the set

$$[0, B] \times \{t_1\} \cup [0, B] \times \{t_2\} \cup \cdots \cup [0, B] \times \{t_n\}.$$

The idea here is that if the option knocks-out then the replicating portfolio would be immediately liquidated at zero cost. Unlike our construction for a general path-dependent exotic, the strategy here involves selling options before their expiry which is the crucial point where we use the existence of a deterministic pricing function. In what follows, we concentrate on the case of a down-and-out option for concreteness. However, the same techniques apply with little change to pricing a double-barrier option or an up-and-out option.

To construct this portfolio, we induct backwards. First, we choose a portfolio of vanilla options with expiry T which approximates the final pay-off as accurately as we desire. Of course, if we were modelling a knock-out call or put, this would just be the call or put without the knock-out condition. Call this initial portfolio, P_0. For a portfolio P, we denote its value at the point (S, t) by $P(S, t)$.

As we know the price of any unexpired vanilla option for any value of spot and time. We can value P_0 along the last barrier $[0, B] \times \{t_n\}$. We can kill the value of P_0 at the point (B, t_n) by shorting a digital option with value equal to $P_0(B, t_n)$, below B and zero above. In order to retain the property that the portfolio consist solely of vanilla options, we approximate the digital by a tight put-spread. Call our new portfolio P_1^0. This portfolio then has correct final pay-off profile and has zero value at (B, t_n) but may have non-zero value along $[0, B) \times \{t_n\}$. We partition $[0, B)$ into $[0, x_1], [x_1, x_2], \ldots [x_{k-1}, B)$. We then approximate the value of P_1^0 along $[0, B) \times \{t_n\}$ by assuming it is affine on each of these subintervals. We now remove this value by moving successively inward. The portfolio P_1^0 has zero value at (B, t_n) so we if short put options struck at B with expiry at t_n and notional $P_1^1(x_{k_1}, t_n)$, we obtain a portfolio P_1^1 which has zero value at both (B, t_n) and (x_{n-1}, t_n). If our partition is suitably small, the linear approximation will be close in value to the original and the value will be small on the interval $[x_{n-1}, B] \times \{t_n\}$. We now iterate along the barrier at each stage shorting put options struck at x_{n-j} with expiry t_n and notional $P(x_{n-j-1}, t_n)$ to obtain a sequence of portfolios, P_1^j. Note that the put options only affect value for spot below their strike so will not affect the value of the portfolio at the points already fixed. The portfolio P_1^{k-1} will then have close to zero value along barrier as desired. Let P_1 be the portfolio P_1^{k-1}.

We can repeat this procedure along the barrier $[0, B] \times \{t_{n-1}\}$ using the portfolio P_1 instead of P_0 to obtain a portfolio P_2. Repeating, we obtain a sequence of portfolios, P_l. If we regard an option as having zero value after expiry then the portfolio P_n then has the property of having close to zero value along all the barriers $[0, B] \times \{t_j\}$ and approximates the final payoff profile. As we immediately liquidate the replicating portfolio when the original options knocks out, options will only expire in the money at a time of liquidation. This means that we do not need to keep track of any pay-offs from options maturing. The value of P_n at time zero and today's spot will therefore be the value of the knock-out option.

A similar procedure would be effective for up barriers, simply replacing puts by calls and inducting upwards instead of downwards. To do double barriers, we simply do each independently at each barrier time as there is no interaction between the two pieces at a given knock-out time, although the portfolio is of course affected by both barriers at previous times. We note that our procedure does not require the barrier level to be constant.

One disadvantage of this approach over that presented in Section 2 is that for each option added the entire portfolio of already added options has to be valued for a new value of spot and time. This means that the time required to evaluate using replication will be proportional to the square of the total number of options. The total number of options will be the number of steps per barrier times the number of barriers. On the other hand, the fact that there is only one rehedging makes the approach conceptually nicer.

We present the results of some simulations for a jump-diffusion model.

We first present the pricing of a down-and-out call option with the following contract,

Barrier Level	90
Strike	100
Notional	1
Expiry	1

and with knock-out dates $0.1, 0.2, 0.3, \ldots, 0.9$. We use a log-normal model for stock evolution with log-normal jumps as follows,

Interest Rate	0.05
Dividend Rate	0
Initial Spot	100
Diffusive Volatility	0.1
Jump Intensity	0.2
Jump Mean	0.8
Jump Sigma	0.2

The prices produced by the replication method given above were,

Steps per barrier	Price
8	8.685
16	8.686
32	8.687
64	8.687
128	8.687
256	8.687

These compare with prices produced by Monte Carlo as follows,

Number of Paths	Price
1000000	8.687
500000	8.689
250000	8.691
125000	8.687
62500	8.693
31250	8.691
15625	8.649

It is also interesting to compare the implied price with that obtained from pricing with a Black-Scholes model. Pricing the same option using

Monte Carlo but in a Black-Scholes world with the at-the-money implied volatility, the price was 8.42 and using the at-the-barrier implied volatility was 9.25.

We now present prices with the more extreme parameters,

Diffusive Volatility	0.2
Jump Intensity	0.5
Jump Mean	0.6
Jump Sigma	0.3

For the replication method, we get

Steps per barrier	Price
8	17.377
16	17.389
32	17.393
64	17.391
128	17.391
256	17.391

and for Monte Carlo,

Paths	Price
1000000	17.394
500000	17.400
250000	17.406
125000	17.392
62500	17.397
31250	17.436
15625	17.374

Pricing the same option using Monte Carlo but in a Black-Scholes world with the at-the-money implied volatility, the price was 13.89 and using the at-the-barrier implied volatility was 14.47.

References

[1] L. Andersen, J. Andreasen, D. Eliezer, Static Replication of Barrier Options: Some General Results, preprint 2000

[2] Y.Z. Bergman, Pricing Path Contingent Claims, Research in Finance, 5, 229-241

[3] F. Black, M. Scholes, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy*, 81, 637-654

[4] H. Brown, D. Hobson, L.C.G. Rogers, Robust Hedging of Barrier Options, preprint 2000

[5] P. Carr, K. Ellis, V.Gupta, Static Hedging of Exotic Options, *Journal of Finance,* LIII, 1165-1191, 1998
[6] J.N. Dewynne, P. Wilmott, Partial to the Exotic, 6 (3) 38-46
[7] E. Benhamou, A. Duguet, Small Dimension PDE for Discrete Asian Options, preprint 2000
[8] E. Dinenis, D. Flamouris, J. Hatgioannides, Implied Valuation of Asian Options, preprint 2000
[9] B. Dupire, Pricing with a smile, *Risk*, 1994
[10] S. Hodges, A. Neuberger, Rational Bounds for Exotic Options, preprint
[11] J. Hull, A. White, The pricing of options on assets with stochastic volatilities, *The Journal of Finance,* Vol XLII no 2, 1987, 281-300
[12] J.E. Ingersoll, Theory of Financial Decison Making, Rowman and Littlefield 1987
[13] P. Jäckel, R. Rebonato, An efficient and general method to value American-style equity and exotic options in the presence of user-defined smiles and time-dependent volatility, preprint 2000
[14] D. Madan, P. Carr, E. Chang, The Variance Gamma Process and Option Pricing, *European Finance Review* Vol. 2, No. 1 1998
[15] D. Madan, F. Milne, Option Pricing with V.G. Martingale Components, *Mathematical Finance,* 1, 4, 39-55
[16] D. Madan, E. Seneta, The Variance Gamma Model for Share Market Returns, *Journal of Business,* 63, 4, 511-524
[17] R. Merton, Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* 3, 1973, 141-83
[18] L. Rogers, Z. Shi, *The Value of an Asian Option, J. Appl. Probability* **32**, 1995, 1077-1088
[19] R. Rebonato, Volatility and Correlation in the Pricing of Equity, FX and Interest-rate Options, Wiley Chichester New York Weinheim Brisbane Singapore Toronto 1999
[20] P. Wilmott, Derivatives, The Theory and Practice of Financial Engineering, Wiley 1998

Royal Bank of Scotland Group Risk, 135 Bishopsgate, London EC2M 3UR