Valuing Average Rate ('Asian') Options

A P Carverhill

L J Clewlow

Financial Options Research Centre
University of Warwick

RISK, Vol 3, No 4

April 1990

Financial Options Research Centre
School of Industrial and Busines Studies
University of Warwick
Coventry
CV4 7AL
Phone: 0203 523606

FORC Preprint: 90/11

"It's complicated to make settlement of swaps coincide with futures. Swaps are settled on a monthly, quarterly or six-monthly average while futures are settled on the basis of a daily price."

Ultimately the project will depend on the size and liquidity of the oil swaps market, neither of which have yet been quantified. The locals who contribute liquidity in the near contracts would be unlikely to trade long-term crude futures at anything less than 30 cent spreads and would be subject to speculative position limits anyway, notes Marinchek. "Would you want to pay 30 basis points every time you did 1,000 lots?" he asks.

Marinchek admits to cautious optimism for a "wonderful idea" and sees potential arbitrage opportunities. "If I entered a three-year swap I could sell a two-year strip to Nymex and play with one year in the over-the-counter market," he notes.

Hampton, too, insists: "It will work. It is a way of adding a counterparty to your range. And it could be especially effective if, by having a visible price out there, it brings in new people." But he warns that the swaps market does not depend on the futures. "The swaps market already exists and already works. Long futures are just another way of laying off the risk."

Meanwhile, the International Petroleum Exchange (IPE) in London, whose crude futures go out only six months at the moment, is not commenting directly on the idea. Chief executive Peter Wildblood will only invoke the IPE's "good record of responding to change in the industry". So now it's up to Nymex to bring about the change.■

FLEXIBLE CONVOLUTION

Continuing the debate on the valuation of average-rate options, Andrew Carverhill **and** Les Clewlow **introduce the convolution method**

The convolution method is an efficient and flexible means of valuing average-rate (Asian) options. It allows us to answer the questions of how often the underlying should be read in taking the average and how the price is affected if underlying returns are not normal.

We will discuss two types of Asian option – floating-strike and fixed strike – but concentrate on the latter. As the name implies, the floating-strike option pays the difference – if positive – between the average value of the underlying on which it is written and the spot value of the underlying when it is exercised. The fixed-strike option pays the difference between the average and a previously agreed strike price. The floating-strike is easier to deal with, but less widely used. We will take all our Asian options to be European-style, and the average to be over the entire life of the option.

Our principal references will be Ingersoll[1], Kemna and Vorst[2] and the *RISK* articles by Krzyzak[3] and Ruttiens[4]. Of course, the elixir for the Asian option would be an explicit, exact formula for its value that is also flexible in the sense described above. None of these articles quite gives such a formula: Ingersoll claims that one can be constructed, but tantalises his reader by declining to do so, on the grounds that his option does not (yet!) exist in the market. (He is dealing with floating-strike options.) On the other hand, Kemna and Vorst (and Ruttiens) give an approximate formula, and claim that an exact one seems impossible in principle. (They are dealing with fixed-strike options.)

Note that Ingersoll and Kemna/Vorst could both be correct, because the floating-strike option is technically easier. We do not know of anyone who claims to have an exact formula for the Asian option.

[1] Ingersoll, J, *Theory of financial decision making*, Rowman & Littlefield

[2] Kemna, AGZ and ACF Vorst, 1986, *The value of an option based on an average security value*, Report 8611/B, Erasmus University, Rotterdam, and, 1987, *Options on average asset values*, Report 8717/B, Erasmus University, Rotterdam

[3] Krzyzak, K, 1989, *Asian elegance*, *RISK* Vol 3 No 1, December 1989–January 1990, page 30

[4] Ruttiens, A, 1990, *Classical replica*, *RISK* Vol 3 No 2, February 1990, page 33

A $/£ fixed-strike option

To fix our ideas, we will concentrate on the following situation: I am a corporate treasurer, and expect to earn $k over the coming year from my US activities. I expect these dollars to come in a continuous uniform stream, and my policy

Of course, the elixir for the Asian option would be an explicit, exact formula for its value that is also flexible

will be to change them to sterling at regular intervals, say on *n* occasions during the year. I am budgeting for a £/$ exchange rate of Y but am worried that the actual rate will be less than this.

My solution is to buy a fixed-strike option, which will pay the difference (if positive):

$$kY - \left\{(k/n)X_1 + (k/n)X_2 + \ldots + (k/n)X_n\right\} = k(Y - A),$$

where X_1 ..., X_n are the spot rates at which I change my dollars, and A is the average of these. To value this option it is also necessary to know the UK and US interest rates (which we will call r and s), the current spot £/$ exchange rate (which we will call X_0), and the £/$ volatility (which we will call σ).

Valuation by the convolution method

Our starting point is the "expectation formula" for the value of the option. An expectation formula can be written down for any European option, but may not be easily tractable. The Monte Carlo and path enumeration (exploded tree) methods of valuing Asian options are also based on the expectation formula. To obtain the formula we first make a risk-neutral adjustment to the drift of the underlying process; the option value is the expected pay-off under this adjustment, discounted back to the present.

For our £/$ option the risk-neutral adjustment is simply to make the drift of the £/$ exchange rate equal to the difference between interest rates, r − s. (Note that we have not said, because it does not matter, what the drift actually is! Also note that the option pays off in sterling and the final discounting back to the present will be at the sterling interest rate r.)

The key to our method is formula (1), which gives the probability distribution of the average:

Put:

$$Y_i = \log X_i, \quad Z_i^{i+1} = Y_{i+1} - Y_i \text{ for } i = 0, 1, \ldots ,$$

so that:

$$X_i = \exp(Y_0 + Z^1{}_0 + \ldots + Z^i{}_{i-1}).$$

Note that the increments Z_1^{i+1} are all independent and normal, with mean $((r-s) - \frac{1}{2}\sigma^2)(T/n)$, where T is the time to maturity of the option.

Then the formula for the average A is:

$$A = {}^1/_n \left[X_1 + \ldots + X_n\right]$$

$$= {}^1/_n \Big[\exp(Y_0 + Z^1{}_0) + \exp(Y_0 + Z^1{}_0 + Z^2{}_1) + \ldots$$

$$\ldots + \exp(Y_0 + Z^1{}_0 + Z^2{}_1 + \ldots + Z^n{}_{n-1})\Big]$$

$$= {}^1/_n \exp Y_0 \Big[\exp Z^1{}_0 (1 + \exp Z^2{}_1 (1 + \exp Z^3{}_2$$

$$(1 + \ldots (1 + \exp Z^n{}_{n-1}) \ldots)\Big] \qquad (1)$$

(This key factorisation was discovered by Stewart Hodges.)

Our actual method is to obtain the density of the distribution of the average from the densities of its ingredients, as given by formula (1), these ingredients being the normal distribution $Z^1{}_0, \ldots, Z^n{}_{n-1}$. This is done in the inductive procedure:

$$\tilde{A}_{n-1} = Z^n{}_{n-i}$$

$$\tilde{A}_{i-1} = Z^i{}_{i-1} + \log(1 + \exp \tilde{A}_i) \text{ for } i = n-1, \ldots, 1, \qquad (2)$$

$$A_0 = (Y_0 - \log n) + \tilde{A}_0, \qquad (3)$$

$$A = \exp A_0.$$

How convolution is used to calculate the density of the average

The convolution process

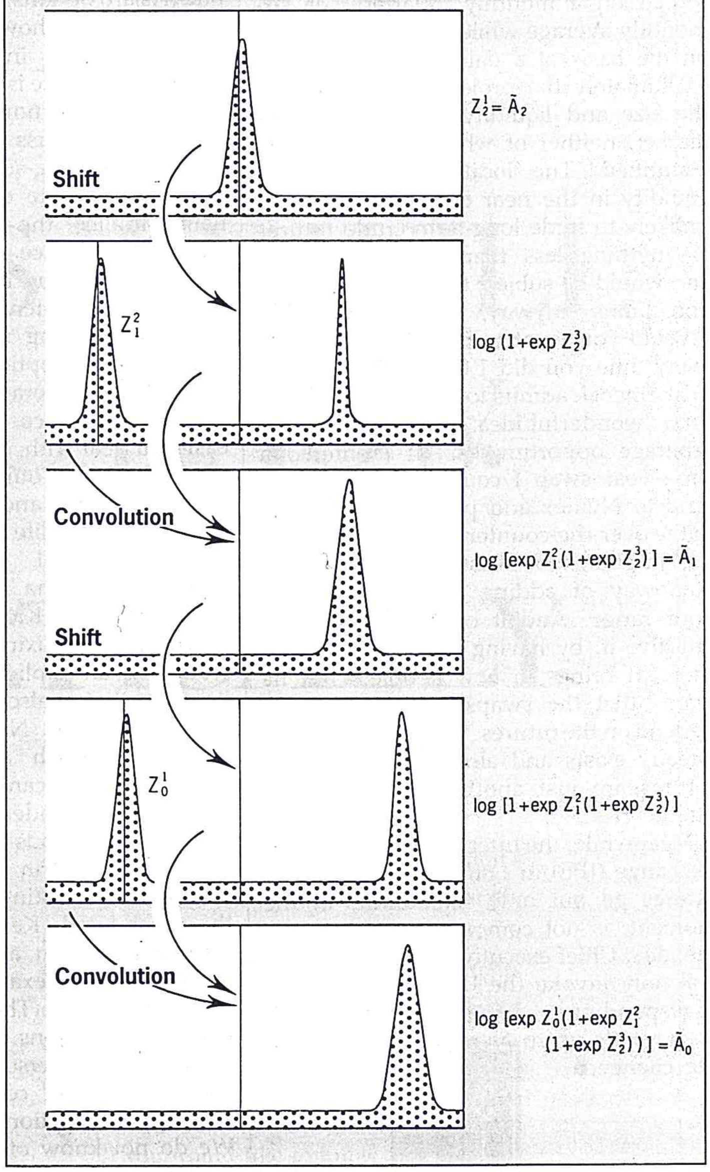

In the inductive step (2), the density $f_{\log(1+\exp \tilde{A}_i)}$ is obtained from the density $f_{\tilde{A}_i}$ by the easily verified transformation:

$$f_{\log(1+\exp X)}(x) = \frac{\exp x}{(\exp x) - 1} f_X(\log((\exp x) - 1)) \qquad (4)$$

The key to step (2) is to obtain the density for the sum of this and the normal density $Z^i{}_{i-1}$. This is obtained as a convolution using the fast Fourier transform.

(Mathematically, the density for the sum of random variable is the convolution of the individual densities. Also the Fourier transform of a convolution of densities is the product of the Fourier transforms of the individual densities. The fast Fourier transform is a technique to do Fourier transforms computationally in an efficient way.[5]).

Having obtained the density of the average, we can work out the expected pay-off simply by integrating the pay-off function against the density.

An intuitive way to understand the idea of valuation by convolution is as follows. The underlying process is built up from normal increments as time evolves, and the transformation (4) plus convolution, as represented by equation (2), specifies how each new normal increment is mixed in with the previous ones, to give the average of the underlying. In fact, transformation (4) can be understood as shifting the distribution to the right, so that $x \mapsto \log((\exp x) + 1)$, and the convolution of any two densities represents "smudging" them together. The diagram represents the calculation of the density of the average, if it is done as above, taking n=3.

Numerical results

We used the convolution method to value the Asian foreign exchange option discussed above, which is a put on the dollar. The value is given in the graphs, in terms of the number of iterations used by the method. The parameters for the option valuation were taken as:

Volatility σ = 10%
UK interest rate r = 13%
US interest rate s = 7%
Current £/$ rate = 0.625
Strike rate = 0.625
Time to maturity = 1 year

From the graphs we see that monthly averaging gives a value for the option within about 5% of the value with weekly averaging, and this last averaging is almost equivalent to continuous averaging. (Note that for any given number of iterations, the convolution should give the correct option value, with no error except round-off

[5] Press, WH, BP Flannery, SA Teukolsky and WT Vetterling, 1988, *Numerical recipes in C*, Cambridge University Press

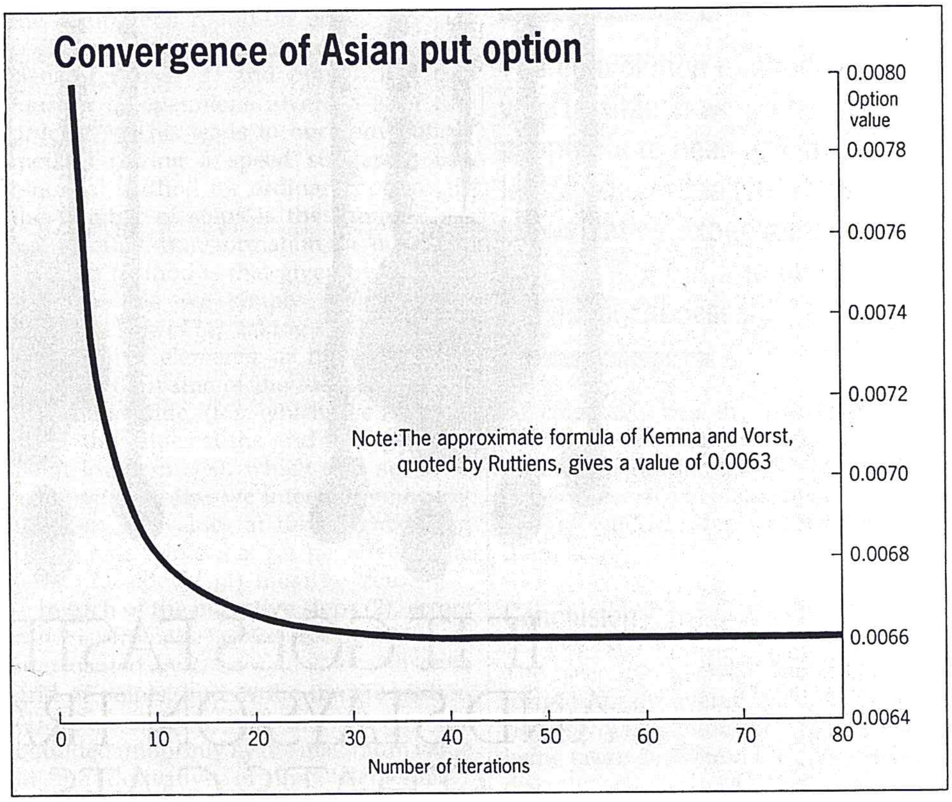

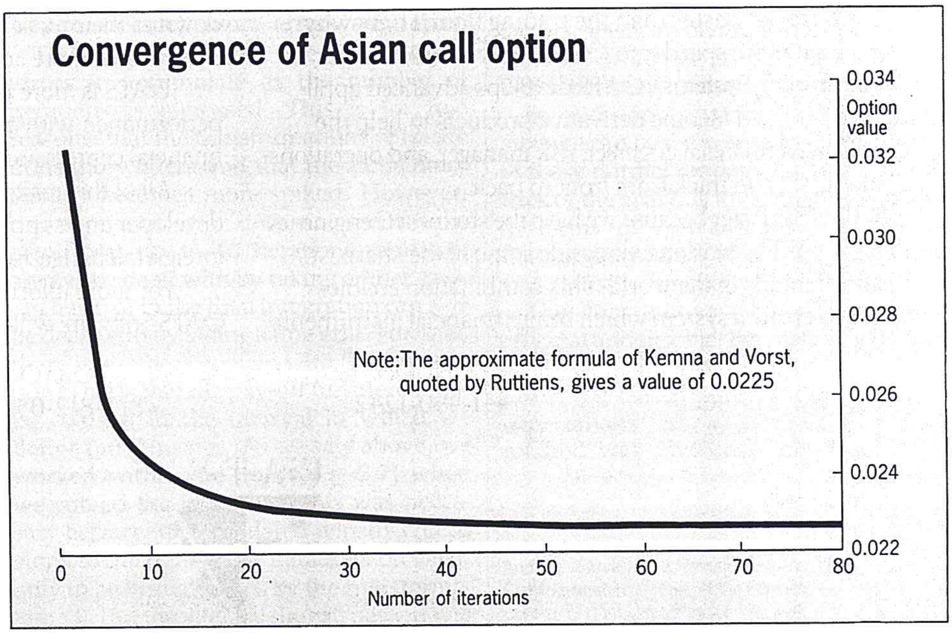

THE EXPLODED BINOMIAL TREE METHOD

This method involves approximating the underlying process by a binomial random walk – a branching process – which can step up or down, as time evolves, by discrete increments.

The method is essentially to average the pay-off of the option over every possible path that this random walk could take. This is easy to programme and gives a rough-and-ready estimation of the option value.

However, it is extremely inefficient to implement, and the answer can only be approximate; for a realistic answer one would have to represent the underlying process by at least 20 steps of the random walk, but to take many more steps would be computationally intractable. (Twenty steps corresponds to 2^{20} = 1 million branches in the tree. Our valuation for the foreign exchange option in this article with 20 steps is 0.0479 for the call and 0.00657 for the put.)

Also, it is not clear that 20 steps is appropriate for 20 averaging points; the exploded tree method is really too insensitive to distinguish between Asian options with different numbers of averaging points.■

error, for that number of averaging points.)

The parameters for the convolution method itself were taken as follows. We represented the density function by their values at 4,096 (=2^{12}) equally spaced points in the interval [−7,7]. Throughout, we worked with the logarithmic transformation of the underlying process.

The fast Fourier transform requires that the function to be transformed is represented in terms of its values at an equally spaced grid of points in an interval, the number of points in this grid being a power of 2. The technique of the fast Fourier transform is essentially to reverse the digits in the binary expression for the grid point numbers; this reversed expression represents the decomposition of the function in terms of cycles. To convolve we take the fast Fourier transform of the two functions, multiply the answer, and then take the inverse transform of this product. This procedure actually calculates the (discrete) convolution without error except for the computer's round-off error.

The fast Fourier transform is very efficient; if it uses 2^N grid points then the number of operations involved is of the order $2^N N$. This leads to our convolution method having a speed similar to the binomial method for ordinary options, if the number of steps is the same.

The other transformation in our convolution method is that given by formula (4). For this we simply calculate the left-hand side of (4), taking x to be each of the positive elements of the grid. This involves the value of the function on the right-hand side (for which we already know the values at the grid points) at the point log ((exp x)−), which will not be a grid-point; for this we interpolate linearly between the values at the neighbouring grid points. Note that for negative x, the left-hand side of (4) must be zero.

In each of the inductive steps (2), errors can arise from the fact that the densities are represented by their values at a discrete grid of points, and in the transformation (4). One expects that these errors are bounded uniformly by the maximum value of the derivative of the "elementary" normal density function for Z_i^{i+1} multiplied by the distance between the grid points. Also, when we finally integrate the pay-off function against the density (using the trapezium rule), more errors may arise.

Thus, for a fixed grid, one expects these errors to accumulate as the number of iterations is increased. This is for two reasons: that the transformations of (2) are done more often, and that the elementary density becomes more spiked. However, this source of error does not seem to be significant up to 40 iterations, and can always be dealt with by taking a finer grid.

The error in the final integration can also be dealt with by taking a fine grid. For calls, there is another problem with this integration. This is that the pay-off function might be colossal in the interval in which we define our densities. As we said above, we worked within the interval [−7,7] when we valued the FX option; this was necessary because the calculated density creeps along to the right as the iterations proceed, only to be brought back by the transformation (3). But the value of the call pay-off will be colossal when the underlying is e^7. (Note that we work in logarithmic transformation throughout.) The remedy for this is simply to cut off the final value in the integral when the density becomes smaller than, say, 1/1,000.

The convolution method seems very flexible. It would be adaptable to dealing with leptokurtic (fat-tailed) returns; or to evaluating expectations for which a Monte Carlo method was previously necessary

Conclusion

The basic ingredient of this article is the formula for the average value of a geometric Brownian motion (GBM), the average being taken over the value of the GBM at a discrete set of times. The value of the option involves the expectation of the pay-off function with respect to the distribution of this average. The key is that this formula can be decomposed into factors, each involving a normal random variable, which gives the evolution of the GBM between consecutive time points.

By judiciously taking logarithms, the average can be expressed in terms of sums of these normal random variables; and the trick of the article is to calculate these sums in terms of convolutions (via the fast Fourier transform) of the associated normal density functions.

The convolution method seems very efficient and flexible. For instance, it would easily be adaptable to dealing with leptokurtic (fat-tailed) returns; or to evaluating expectations for which a Monte Carlo method was previously necessary.■

Andrew Carverhill is senior research fellow, and Les Clewlow research fellow, at the Financial Options Research Centre of the University of Warwick, UK