

EXAMINATIONS

April 2004

Advanced Certificate in Derivatives: Further Mathematics, Principles and Practice

EXAMINERS' REPORT

QUESTION 1

Syllabus: (a) (ii) 7 & (v) 1

Reading: Hull (5th Edition) Ch. 19, Briggs et al Ch 6.

(i)

[More detail is covered in Hull Ch 19.]

Asian Option (or Average rate option) is an option where pay off at expiry is based on the average stock price during the life of the option (typically computed daily but other rules can apply) in relation to the strike price.

Barrier Option is an otherwise standard option, which deletes itself if the underlying stock touches a knock-out or activates if it touches a knock-in price.

The knock-out or knock-in price is usually out of the money but may be in the money.

Basket Option is an option where underlying "instrument" is a specified basket (or portfolio) of assets.

Bermuda Option is an otherwise standard option, which can be exercised only on a limited number of dates.

It lies in valuation between that of a European Option and an American Option.

Binary Option is an option with a discontinuous (digital) pay off.

The most common form is the "Cash or Nothing" in which the option pays out a fixed amount if the option is in the money at expiry.

Compound Option is an option on an option. The four main types are call on a call, put on a call, call on a put and put on a put.

The option has a first period and strike level, and a second period and strike level.

Forward Start Option is an option, which will start at an agreed date in the future. At this time the strike level will be set – here at the money. The option is then a standard one.

Lookback Option is an option where pay off at expiry uses the historical price evolution of the underlying stock over the life of the option to create the most favourable payout (i.e. the payout depends on the maximum or minimum stock price during the life of the option).

Shout Option is an otherwise Standard European option in which on an agreed set of dates or on the basis of a limited number of opportunities (including expiry) the option holder can lock in the then current stock price to create a positive or more positive payout at expiry.

It is in effect a limited form of lookback option and will lie in value between a standard European Option and a lookback option depending upon the number of "shouts" involved.

(ii)

Legal risk is defined to be the exposure to financial loss arising from adverse legal, legislative or regulatory action.

A major loss from legal risk occurred in the UK due to an adverse decision on local authority entering swaps, where they were deemed to be acting *ultra vires* (outside their powers).

Legal risks fall into the following four categories (Briggs et al split the fourth one):

(a) Enforceability

Whether derivatives can be deemed legal transactions, and not gambling ...

... several jurisdictions (e.g. France) have recently amended their law to clarify this.

(b) Capacity

Whether entity entering into the transaction, and the person actually arranging the deal, have the power and authority so to do ...

... which can be particularly problematic with government agencies

Capacity is a broad terms – it covers limitations on types of instruments, their use (e.g. in hedging or speculation), method of authorisation, accounting methods etc ...

(c) Documentation

The status of written documentation is important (e.g. ISDA agreement).

Some jurisdictions allow verbal confirmation

... and there can be uncertainty when verbal and written confirmations disagree.

(d) Credit risk mitigation

A collateral agreement would help reduce credit risk ...

... but the trade valuations will need to be agreed with the counterparty – not easy if they are complex.

The ability to net exposures across similar contracts with a counterparty when it defaults on one contract ...

... reduces credit risk loss potential by offsetting all payments due (both to and from the counterparty) against each other.

This should ideally apply across all derivative contracts between the counterparty and the life office ...

... and also across all branches and subsidiaries of the counterparty.

QUESTION 2

Syllabus: (b) (i) 6 & 7, (iii) 2

Reading: *Rebonato Ch 10 & 17, plus Hull (5th Edition) Ch 22.*

(i)

[The following is an example of one of a number of approaches which could be adopted. Similar marks could therefore be earned for alternative answers.]

An accrual swap can be decomposed into an ordinary fixed-floating swap plus a set of daily binary options (“cash-or-nothing” options) to cancel out the fixed payment for that day.

Hence the valuation would be that of an ordinary swap, less the value of the binary options which have effectively been sold by the holder of the accrual swap.

In the accrual swap example given, there are two options on each day: one to cancel the fixed payment if rates go above 5%, the other to cancel it if rates go below 4%.

Binary options are valued via the Black formula, suitably adjusted for whether the option is a call or a put, using the $N(d_2)$ term – the formula book calls this $\Phi(d_2)$ [*which is actually the risk-neutral probability that exercise occurs*].

Theoretically, the binary option are valued daily and then summed over each accrual period, but in practice they are usually grouped weekly or fortnightly without any loss of accuracy.

The risk profile of the given accrual swap is that it behaves like an ordinary fixed-floating swap whilst rates lie well inside the 4% - 5% boundary, with small gamma.

However, as soon as rates approach either boundary, the gamma increases dramatically and the option effect becomes very pronounced. This makes accrual notes a bit “nasty” to hedge when rates get near the boundaries.

[Actually, on a boundary, the theoretical gamma of a binary option is infinite.]

(ii)

[Rebonato covers this point in Ch. 17.]

The “chooser” feature dramatically changes the valuation problem for accrual swaps.

Instead of having simple binary options dependent on market level, the market level is now virtually irrelevant, as the owner of the “chooser” feature can move the strike level as he wishes ...

... but the binary options now become very dependent on the slope of the yield curve, i.e. the imperfect correlation of forward rates.

Hence the valuation depends on more than one factor ...

... and thus needs a more complex valuation model, such as Monte-Carlo or the two-factor Hull-White model discussed in another question.

(iii)

Definition

The Monte Carlo method uses a random number generator to provide potential values for the stochastic term of an interest-rate model. One set of random numbers gives one simulated path.

Method: turn time horizon into discrete time by small time steps Δt ...

... sample known probability distribution to give a path over all the time steps ...

... gather large number of paths, say 10,000, of equal probability ...

... value the derivative (option or whatever) on each path, and take the average.

Technically it can be thought of as a forward induction method of calculating $\hat{E}[e^{-\bar{r}T} f_T]$, the (risk-neutral) expectation of the current value of the time T payoff f_T .

Used to value complex derivatives that have no analytic solution, such as exotic options which are path dependent or have non-linear payoffs.

Advantages

- Simple idea conceptually, albeit rather “brute force” in style.
- Extension of MC to multiple factors is simple, as the samples are taken e.g. from the multivariate Normal rather than the simple Normal distribution. For the same level of accuracy, the number of samples needed grows approximately linearly with the number of factors. In tree methods, it grows exponentially, so is much more costly.
- Binomial trees are not very efficient. They have large numbers of very extreme (large and small) nodes which contribute almost nothing to the value of the derivative.
- MC time steps can easily be tailor-made for the problem. Trees and lattices are much less flexible.
- Increasing the number of time steps in a tree increases the number of calculations by the square of that number, whereas for MC it is roughly linear.
- Successful for valuing many path-dependent options, because the path is always generated for each sample.

Disadvantages

- Needs complex and generally large-scale computer program to work effectively.
- Convergence is slow. The variance of the set of independent random variables decreases with as $\frac{1}{n}$, where n is the number of sample paths, so accuracy (as measured by the standard error) is proportional to $\frac{1}{\sqrt{n}}$.
- Deltas and gammas are hard to obtain directly, due to compounding of the approximation errors for these second and third order quantities.

- Often hard to get unbiased enough random or quasi-random sequences, so results nearly always have some hidden bias from the true answer.
- American options are tricky or impossible to value under MC because the method does not enable one path to know about other paths, so as to compare values for the American feature. Trees and lattices do this readily.
- Intermediate values along the paths (for e.g. short rates or bond prices) are almost impossible to calculate in MC, certainly for common log-normal distributions. Trees and lattices are much better for this.

(iv)

[There are several possible methods - a few are listed below. The aim of all of these is to improve the effectiveness of the sample distribution, by giving the same coverage of the sample space with fewer sample paths.]

Antithetic Variables

For each sample path from the distribution, create another valid path by taking the opposite sign on each random element. This creates an “antithetic” path.

The odd moments of the sample distribution are therefore zero, which works well with non-skew distributions like the Normal distribution.

Moment Matching

The n sample paths are generated and stored, then adjusted by a scaling factor to ensure that the variance (second moment) of the sample exactly matches that of the initial distribution.

Together with antithetic variables, this produces a tighter fit to the probability distribution without adding further samples, so should be more accurate.

The fourth moment can also be matched in the same way if kurtosis is important.

Stratified Sampling

Divide the range of outcomes into bands according to probability, and then take more sample paths from those with the higher probability.

This does not affect the overall accuracy, but removes the need to sample those outcomes where there is a low contribution to the valuation.

Contravariates (or Control variates)

If an exotic derivative is being valued, there is probably an analytical expression available for a similar, more simple, derivative (e.g. if calculating an average-rate caplet, the formula for the plain vanilla caplet is known).

Then use MC to calculate the simple and exotic derivatives, and adjust the exotic value by the same amount as the simple value needs to be adjusted to agree with the known formula.

Alternatively, the paths themselves can be adjusted (probably a few times) until the values agree.

Quasi-Random Sequences

The MC approach does not need random values to succeed - all it needs is to have representative enough paths that the approximation is not biased by the samples chosen.

There are some sequences which can be shown to be non-random, but which nevertheless will not bias the approximation. The advantage of this type of series is that its standard error is proportional to $\frac{1}{n}$, rather than $\frac{1}{\sqrt{n}}$ for the raw MC.

QUESTION 3

Syllabus: (b) (i) 6 & 7

Reading: *Rebonato (2nd Edition) Ch. 8, 12, 13.*

(i)

(a) Calculate 4-year cap and floor

The marking may need to be adjusted if candidates make a persistent slip. We assume the strike does not need to be adjusted for day-count, though candidates may do this if consistent.

The basic discounting process in BDT is used throughout. This assumes a probability of $\frac{1}{2}$ for up and down nodes, and $V(t) = \frac{1}{2} [V_{\text{up}}(t+1) + V_{\text{dn}}(t+1)] / [1 + r(t)]$.

Caplet prices: final payoff at node $j = \max \{r_j(t) - 2, 0\}$

$t = 1$:

	0.2923
0.1440	
	0.0000

$t = 2$:

		1.0957
	0.6030	
0.3305		0.1380
	0.0679	
		0.0000

$t = 3$:

			2.0577
		1.3642	
	0.8476		0.7560
0.5073		0.3701	
	0.1821		0.0000
		0.0000	
			0.0000

Floorlet prices: final payoff at node $j = \max \{2 - r_j(t), 0\}$

$t = 1$:

	0.0000
0.1920	
	0.3897

$t = 2$:

		0.0000
	0.0000	
0.1281		0.0000
	0.2600	
		0.5283

$t = 3$:

			0.0000
		0.0000	
	0.0324		0.0000
0.1370		0.0663	
	0.2456		0.1355
		0.4328	
			0.7427

Hence cap price = $0.1440 + 0.3305 + 0.5073 = 0.9818$ and floor price = $0.1920 + 0.1281 + 0.1370 = 0.4571$. These prices are in % of nominal.

(b) Check put-call parity of prices

If we take any caplet price minus the floorlet price, the final payoff is $r(t) - 2$ at each node.

Hence, since the tree creates the present value of the final payoff, the value at $t = 0$ will be the present value of the forward rate less the present value of the strike, which is put-call parity.

[Numerical examples could be given instead to show put-call parity for full marks.]

Delta

There are really several separate options, so we calculate for each caplet or floorlet, i.e. don't add the deltas.

Delta for any caplet or floorlet is $(P_{\text{up}} - P_{\text{down}})/(r_{\text{up}} - r_{\text{down}})$ at $t = 0$.

[Numerically – not required for the solution:]

Given that $r_{\text{up}} - r_{\text{down}} = 2.299 - 1.604 = 0.695$, we get

Caplet Δ for $t = 1$: 0.9576 $t = 2$: 0.7699 $t = 3$: 0.4205

Floorlet Δ for $t = 1$: -0.3067 $t = 2$: -0.3740 $t = 3$: -0.5608]

(ii)

Both BDT and Black use a log-normal model so using the caplet volatilities for each time step should give very close answers.

BDT should give exact same answers as Black for caps and floors if properly calibrated and small Δt (i.e. $\Delta t \rightarrow 0$).

Our BDT tree is very coarse, so won't get very precise answers (though amazingly they are very close). Need a smaller time step to be accurate.

Need to be sure day-count methods are same.

Generally our deltas will be poor due to the large time step.

Black deltas are always w.r.t. the forward rate, not the short rate as we calculated.

BDT deltas can be converted into forward rate deltas by first calculating, then dividing by, the sensitivity of the forward rate to the short rate but this is a messy calculation and needs a small time step.

[It is not correct to say that BDT gives only an approximate answer in any circumstances.]

(iii)

Comparison with HW 1-factor model.

Both models can fit all shapes of yield curve perfectly ...

... but both are single factor models, so only have limited ability to model simultaneously the yield curve and volatility structure.

Both fail to give explicit imperfect instantaneous correlation of forward rates.

Both achieve a good fit by compromising on the future evolution of the term structure over time ...

... with HW generally having a poorer snapshot but better evolution.

HW uses explicit mean reversion ...

... whereas BDT this is achieved by a declining volatility.

BDT uses binomial tree to value ...

... HW a trinomial because of the extra degrees of freedom.

BDT fits the curve by adjusting the bottom rate ...

... whereas HW's trinomial is much more efficient because it avoids areas of low probability, thereby not wasting calculation on unimportant areas of the tree.

BDT is log normal, HW normal ...

... which makes BDT easier to calibrate to Black-derived caps and/or swaptions.

Also, theoretically HW could give negative rates, although only with extremely low probability.

In fact, very hard to calibrate HW, as at-the-money options are priced in both normal and log-normal distributions but the differences come away from the current level of rates. Any cap prices will include a range of at-the-money, out-of-the-money and in-the-money caplets.

Need to disentangle these effects, otherwise HW will not have correct balance between volatility and mean-reversion parameters. In some cases, declining volatility as observed in caplet pricing can lead to failure of the trinomial process further out in the tree.

QUESTION 4

Syllabus: (c) & (d) (ii) 2

Reading: Kemp 11, Dowd 12, GN25, DTI 94/6

[The report should at least cover these points. If a candidate makes a valid point not shown below, it can attract further marks.]

Internal management

- Senior management must decide what sort of risks they want to take on, e.g. how much do they want to be involved in derivatives?
- There has to be congruence between the intended risk appetite and the actual risks taken on. Sometimes this means encouraging more risk-taking, other times less.
- Ensure effective delegation of risk control at both local and global level.
- Beware hidden risks, for example those that cannot yet be reported on the firm's risk systems. These are often the biggest risks, as they usually involve illiquid or unquantifiable positions.
- Everything should be reported even if it can't be risk-measured (and if not, management may wish to ask why the position was taken in the first place).
- Check the accounting methodology, particularly if business looks very profitable.
- Check the documentation, that it is industry standard and robust.
- Check the models. Don't always believe the rocket scientists - must have reasonability check e.g. beware if they claim to be making money if other firms' models are "wrong".

Regulatory framework (banking)

Regulation for derivatives in banks in UK is covered by FSA (Financial Services Authority).

Their principles include:

- Settlement and documentation departments to be separated from trading, so a trader cannot settle his own trades (else e.g. he could "bottom draw" bad trades).
- An *independent* risk management department to be set up within the bank, reporting to senior management and *not* the head of trading ...
- ... to measure and report both market and credit risk ...
- ... and verify the financial models used in the trading areas, that they are robustly applied and, if not industry standard, then soundly based theoretically.
- Traders must be set trading limits expressed in risk-related terms, and any exceptions must be reported daily.
- Daily risk reporting to senior managers, who must be qualified to understand the risks involved and aware of what action to take to mitigate them.

- Daily capital calculation, according to FSA rules, usually involving some sort of Value-at-Risk methodology and stress testing for large market moves.
- Adequate systems resources to be applied to valuation, P&L and risk reporting.
- Regular discussions with FSA and occasional visits.

Regulatory framework (insurance)

[The candidate may choose to list some of the previous section's points under this heading, so can receive those marks here instead.]

Regulation for derivatives in insurance subsidiaries is covered by DTI, based on the Insurance Companies Act 1982.

Their principles include:

- Prudential management controls must be in place, similar to the FSA's.
- Appointed Actuaries must ensure that suitable reserves or provisions exist ...
- ... and report on these at appropriate intervals - "continuous review" ...
- ... including a daily schedule of derivative transactions undertaken, exposures and sensitivities to large market movements.

Global vs. local considerations

Whenever businesses are managed in other locations, there is scope for unacceptable risk and unexpected losses, especially if a crisis occurs in one timezone whilst other timezones are closed.

Need to ensure adequate local supervision by trusted employees, who are empowered to intervene swiftly if the local situation demands, but who will also provide more general feedback to Head Office.

Comments on the specific businesses

(a) Swaps and MM in New York and Tokyo:

- The business described is a fairly standard Treasury operation.
- Documentation, settlement and accounting are crucial. If the business is "fairly large scale", these will need to be fully automated and robust.
- Management will need to assess which parts of the business will require mark-to-market valuation and which will require accruals valuation. Usually, Swaps are market-to-market and MM instruments are accrued.
- The setting and monitoring of risk limits for proprietary traders is crucial.
- Risks should be approximately linear, as there are no options traded, so provided limits are adhered to there should be no surprises.

(b) Capital market trading in Hong Kong:

- The key problem here is assessing the overall risk of the firm from all the different asset classes ...
- ... and return on capital employed.
- Value-at-risk would be useful here, to benefit from imperfect correlation between the asset classes, hence reducing capital requirements ...
- ... but integrating equities and debt into the same framework needs careful thought, and it may not be worth it as the business is small scale.

(c) Equity options in the insurance business:

- Extent of hedging of liabilities important - should in theory be exact. There should be little need to “churn” any of the positions.
- Biggest risk is credit risk with respect to investment banks - don't want to be too concentrated.
- Careful adherence to schedule of positions with regards exercise of options.
- Policy on managing lapses, i.e. closing out of options.

QUESTION 5

Syllabus: (b) (i) 7

Reading: Rebonato Ch 17, plus Hull (5th Edition) Ch. 24

(i)

[The question does not ask for all the algebra, but some is helpful. Different approaches could also be valid here.]

Consider the variable $y(t, T) = \ln P(t, T)$. Then, using Ito's lemma,

$$dy = d(\ln P) = \left[r - \frac{1}{2} \frac{\sigma^2}{P^2} \right] dt + \frac{\sigma}{P} dz = \left[r - \frac{1}{2} v^2 \right] dt + v dz \quad (*)$$

where $v = \sigma / P$ depends (like σ) on t , T and P .

But the time t forward rate spanning period T_1 to T_2 is given by:

$$f(t, T_1, T_2) = - \frac{y(t, T_2) - y(t, T_1)}{T_2 - T_1} \quad (**)$$

and, substituting equation (*) into the differential of equation (**), gives an expression for $df(t, T_1, T_2)$ of the form [...] dt + [...] dz involving v , the normalised volatility ...

... but not r or P directly, i.e. only on these variables in so far as v depends on them.

Letting $T_2 \rightarrow T_1 = T$, the LHS of the above equation becomes the instantaneous forward rate $F(t, T)$:

$$F(t, T) = - \frac{\partial y(t, T)}{\partial T}$$

which leads to a risk-neutral process of the form [...] dt + [...] dz also for $dF(t, T)$...

... where the [...] terms again depend only on v and its partial derivatives with respect to P and T .

Therefore, once the volatility function v has been determined, this is enough to define both the drift (dt) and the random (dz) term ...

... and so there is an explicit link between the volatility of zero-coupon bond prices and the drifts of the forward rates.

(ii)

Benefits of two-factor HJM:

- HJM is a generalised framework for modelling the yield curve evolution, enabling a wide range of valuations from one model, i.e. efficient and consistent.
- Two factors enable the model to describe greater variations in the evolution of the yield curve, since forward rates can be imperfectly correlated.

- Two factors also enable valuation of options which depend not just on the market level but also on the yield curve slope, i.e. include correlation between forward rates on the curve \Rightarrow valuing a wider range of options.
- HJM uses MC, which only needs twice the iterations of a single factor model, compared with square of that number for e.g. tree or lattice methods.

Choice of factors:

The method of principal components can be used to identify the main drivers of the yield curve ...

... which generally are the direction and the slope.

Implementation and calibration:

It is far from trivial to implement the HJM model based on market prices.

Would like to use prices of caps or swaptions to drive model parameters, but ...

... the problem is that these are valued in the market using the Black model, which is a log-normal distribution ...

... and specifying a log-normal volatility for forward rates in HJM severely restricts the forms of the HJM model that can be used.

Choice of correlation function is difficult - hard to be precise about it since so few reliable option values to base it on.

Swaptions and caplets imply certain correlation effects but, considering the calibration difficulties, are not likely to suggest a good correlation function.

(iii)

Equilibrium models are a particular class of models where a simple form of an entire economy is described by the model.

To obtain a parsimonious model (few parameters) the economy is often simplified to consist only of a single production source and a single “good” which can be consumed. Deferred consumption (investing) enables higher future consumption, and is governed by a utility function.

All securities and contingent claims are priced endogenously in this model. This gives a world of “absolute” pricing.

However, in practice, the simplifications of the economy are so great that only certain yield curve shapes are possible, and hence any attempt to map more complicated shapes at regular intervals results in unstable parameters.

No-arbitrage models are a class of models which allow recovery of market prices of one set of securities given prices of another set.

This gives a world of “relative” pricing.

In a non-arbitrage-free model, securities could be priced using the model and then traded at a different price in the real world, leading to persistent profit. In simplest terms, no-arbitrage is the absence of a “free lunch”.

[Note: Hull describes equilibrium models as providing security prices as an output of the model, whereas no-arbitrage models take security prices as an input.]

No-arbitrage is very important in yield curve models, since most complex structures are limiting cases of simpler structures (such as swaps, caps, floors) and hence ideally the model should recover the prices of the latter exactly.

Also, hedging is done using the simpler structures, so the absence of no-arbitrage would mean the accounting process would be distorted by imaginary gains and losses.

Equilibrium models are no-arbitrage as far as the stylized economy they describe is concerned, but the inadequacy of their formulation usually makes them poor at pricing complicated structures (and sometimes simple structures). Hence the possibility of arbitrage is introduced amongst instruments traded in the more complex (“real”) economy.

QUESTION 6

Syllabus: (d) (i) 2 & (v) & (vi)

Reading: Dowd Ch 12, Hull (5th Edition) Ch 16

(i)

Value at risk (VaR) for a portfolio is a number which indicates the maximum loss which the portfolio can sustain over a given timeframe (say 1 day) for a given confidence level (say 95% or 99% certainty).

It is normally calculated by taking into account:

- future price movements (volatility)
- the interdependencies of the portfolio constituents (correlations).

If normal returns are assumed, the VaR of the portfolio can be synthesized easily from the VaR of the constituent individual holdings. If the portfolio has very many holdings, this can lead to a large correlation matrix, so some “benchmarking” is necessary to first express positions in terms of representative indices (FTSE100, FTSE250, SmallCap, AllShare etc) before calculating VaR.

VaR is essentially a linear measure of risk, i.e. assumes the loss is always proportional to the exposure, which is true except for option-type instruments. VaR can be adjusted for these convex securities, but not if the convexity is too pronounced (e.g. very “out-of-the-money” or exotic options).

[Note: as with many such questions, a shorter answer could attract full marks. More details are given here to illustrate the range of possible points that could be covered.]

(ii)

In note format:

(a) matrix

- standard method - statistically sound
- too few factors not precise - too many and matrix too big so cumbersome to calculate
- correlations not stable or even accurate
- not forward looking - past not guide to future
- problem in choosing correct amount of history to compute vols and correlations as far past may not be relevant (could weight nearer observations but then have to choose weights!)
- normal model the only practical one - not realistic for outliers (kurtosis effect) ...
- ... hence VaR does not really give the extremes - bit of a problem since it is supposed to be risk capital
- easy to calculate based on benchmarks (eg FTSE index) - not so easy for individual securities

(b) historical simulation

- fast to compute
- lots of data for a multi-currency portfolio
- very easy to obtain distribution of losses from the data without relying on normal model (but may be a weakness since distribution is implicit)
- no problem with correlation assessment - all implicit (again, may be a weakness)
- not forward looking
- very affected by particular major events e.g. '87 crash
- allows for specific risk as well as sectoral
- problem in knowing how far back to go (as for (a))
- difficult to cope with new securities since not available throughout history ...
- ... so have to use sectors and hence thrown back on correlation problem/specific vs index

(c) Monte-Carlo & GARCH time series

- Monte Carlo and GARCH are actually totally separate methodologies, but the question has linked them for convenience
- Monte Carlo - uses random number generator to provide potential values for the stochastic term of an interest-rate model - one set of random numbers gives one simulated path
- can tackle path-dependent and non-linear instruments easily
- extendible to many assets as can cope with several dimensions without excessive calculation, since number of calculations is proportional to m , not m^2 as in matrix method (m = number of factors)
- however, is slow to converge to solution (rate of $\frac{1}{\sqrt{n}}$ for n iterations), unless random numbers "tweaked" to give better coverage of domain by variance-reduction techniques
- needs an explicit underlying model, which may have inadequacies
- technically quite complicated to set up - needs skilled programming, especially to get a truly random series
- GARCH - predicts volatility and correlations based on evolution of previous observations, using an autoregressive (i.e. implicitly fitting to itself) relationship.
- good to try to look forward - VaR really should use implied vols not historical vols
- problem is we are not really sure what model we are using
- forecasting is unstable and maybe completely erroneous if trying to predict in a range outside its "experience"
- not so useful for predicting correlations - number of parameters required is enormous and problems with keeping output values in range -1 to +1.

[There are many other points which could be raised. The above is only a selection.]

(iii)

Forward currency positions

Create two payoffs, one in each currency for the appropriate amounts at the agreed conversion rate. These then act as zero coupon bonds without purchase price.

Dowd assumes that there is only one currency yield curve, but most banks run the DVaR process for several currencies at once.

Structured notes

Mostly these can be decomposed into more basic securities, usually with a bond underlying.

For example, an inverse floater can be transformed into a FRN plus twice the amount of a fixed-floating interest-rate swap.

Certain types of structured note have options which are very difficult to include within the VaR framework, especially if they have barriers or path-dependent behaviour. Also, additional factors such as dual-currency bonds can be hard to value and measure risk sensitivities.

At the linear level, one could use option delta to give exposure at current level and enter this into the VaR model.

However, non-linear characteristics of options (gamma) make VaR's linear approximation only work for small range about the current market level – hence no use for capital on extreme move.

Exotic options also have very variable gamma and vega (volatility) profiles so risk profile at current level is misleading. Some exotics are virtually riskless at current levels, but “blow up” somewhere far away (e.g. barrier options or cancellation options).

END OF REPORT