

EXAMINERS' REPORT

April 2010 examinations

Subject ST6 — Finance and Investment Specialist Technical B

Introduction

The attached subject report has been written by the Principal Examiner with the aim of helping candidates. The questions and comments are based around Core Reading as the interpretation of the syllabus to which the examiners are working. They have however given credit for any alternative approach or interpretation which they consider to be reasonable.

R D Muckart
Chairman of the Board of Examiners

July 2010

Examiners' comments

Question 1

This was a straightforward bookwork question that was generally well answered.

In part (i), whilst demonstrating knowledge of the right concepts, several candidates did not write their definitions “in the context of a binomial model” as requested, and therefore did not receive full marks.

Question 2

This question involved using a different form of probability distribution to value a Call option, namely the gamma distribution. (This is not to be confused with the Gamma sensitivity measure, i.e. one of the “Greeks”.)

In part (i), several candidates worked on the assumption that the distribution was lognormal, when a more generic response was required.

For part (ii), due to a typographical error, the solution presented in the examination paper omitted a second bracketed term. As a result of this, candidates who had shown the right approach were given appropriate marks whether or not they had reached the correct answer or the printed answer in the paper.

Part (iii), required candidates to understand the contrasting features of the gamma and lognormal distributions. The key point was that distributions can be found which fit the fat tails in observed stochastic market movements, theoretically providing more accurate option pricing, but then they are often less tractable to deal with in practice.

Question 3

In part (i), candidates were asked about the “market price of risk”, which most were aware of as the measure that marks the trade off between risk and return for an investor in an asset in the real world. However, several candidates defined it in terms of the μ and σ of the stochastic differential equation for r given at the start of the question, failing to observe that r is not an asset, so such a representation is meaningless.

Part (ii), comparing the binomial and finite difference methods, was a familiar concept to many and was generally well answered.

Question 4

Part (i) referred to some of the recent additions to the Core Reading on the subject of longevity and inflation risk, and was well answered by most.

For part (ii), the list of counterparty credit risk mitigants given in the model solution is very comprehensive and candidates would not be expected to include them all – they are mentioned in full to indicate the range of possible answers. Counterparty credit risk has

become a major focus for bank regulators since the “credit crunch”, and some additional material will be added to the Core Reading on this topic in due course.

Part (iii) tried to apply the bookwork in the first two parts to a real-life situation. A well prepared candidate was able to contribute some response to these application-type questions (and hence garner some valuable marks) even when they did not have extensive knowledge in that area.

Question 5

This question was based on profit and loss (P&L) attribution. Several candidates tended to treat the tables as providing values for the Greeks, not P&L outcomes, which would have been a more familiar treatment but was not what the question asked.

Attribution of P&L is a very important tool for option traders, as it shows clearly how well a particular hedge is working in practice. A key concept is the way in which Gamma losses (or profits) offset Theta (time decay) – a short option position gives positive Theta profits but market movements then lead to Gamma losses on each re-hedging. Also, marking the book to a chosen overall volatility level will affect the hedging outcome if volatility during the hedging period turns out to be different.

Question 6

This question was about the difference between two methods of fitting swap curves at a point in time to create a pricing tool.

In part (i), quite a few candidates appeared to want to treat the fitting exercise as a calibration of an interest rate model, which is something completely different. Part (ii) was better answered as most candidates recognised the concept of “bootstrapping”.

Part (iv) formed to some extent a separate question on calculating the value of a swap, and here candidates were in more familiar territory.

Question 7

Although this question involved the brand new concept of a “quadrinomial” valuation tree, it was very well answered as candidates immediately recognised how to extend the binomial and trinomial valuation methodologies to four branches.

The skew graph in part (iii) was also generally well drawn, although a few candidates did not make their attempts distinctive enough even for a sketch. Graphical responses, such as this one and Q8 part (iii), have improved dramatically since ST6 was first introduced and it is pleasing that candidates are getting good marks for these now.

Question 8

The application of an option strategy to an insurance company's hedging requirement was simple in concept and elicited some good attempts. Although questions on practical situations where options are used will vary from paper to paper, most are based around a

straightforward application, so a little thought should enable the candidate to gain a large proportion of the available marks.

Where candidates could have improved was in their attention to detail about the reasons for the hedge – the first two parts of the question amount to seven marks, so it is clearly not enough just to make one basic point for part (i) and another for part (ii). The secret here (as elsewhere) is to make important and relevant points succinctly, not to scramble down anything that might possibly be linked to the topic. Despite the time pressure, which is acknowledged, some brief pause for thought should yield an advantage.

QUESTION 1

Syllabus section: (h)(i)-(iii)

Core reading: 8, 9

(i)

(a) Stochastic process

A stochastic process can be defined as a mathematical model designed to follow the progress over time of a random phenomenon such as the price of a stock.

In the simple one-period binomial model, the value of the process is S_0 at time 0. The stock price at the end of period 1 is a random variable S_1 capable of taking only two values, uS_0 and dS_0 , and u and d are up- and down-ratios.* In period 2, each of these nodes has two further possibilities, and so on – this creates the tree effect.

[*This is a multiplicative process; an arithmetic process could go to $S_0 + u$ or $S_0 + d$.]

(b) Filtration

A filtration is the history of the stochastic process movements up until a particular time. The price process S_i , $0 \leq i \leq T$, generates a filtration F_i , $0 \leq i \leq T$, where F_i is the collection of all the events that depend only on S_0, S_1, \dots, S_i .

In a binomial tree, the filtration at node i is the history of movements up to node i .

(c) Previsible process

We will say that a process S_i , $0 \leq i \leq T$ is *previsible* if S_i depends only on the filtration F_{i-1} , i.e. up to the previous time step.

A previsible process is a binomial tree in its own right, but compared to an ordinary stochastic process is known one node in advance.

(d) Probability measure

A probability measure \mathbf{P} for a binomial stochastic process describes the likelihood of the stochastic process increasing or decreasing.

In a binomial tree, it is specified by assigning a probability p_i to the up move at node i and $1 - p_i$ to the down move.

(e) Contingent claim

A contingent claim on the tree is a function of the nodes up to some claim time horizon T , say, i.e. it is a function of the filtration F_T .

In a binomial tree, the claim payoff is specified at relevant nodes up to the final node of the tree at time T .

[It is then valued by discounting back from time T to 0.]

(f) Risk Neutral Valuation

Risk neutral valuations are valuations using the risk-neutral measure, which gives an expected return equal to the risk-free rate on any security (or zero market price of risk).

In a binomial tree, it is the valuation of a claim using the risk-neutral probabilities, obtained by equating the total return of the stochastic process to the risk-free rate.

(ii)

Stochastic processes underlying derivatives normally have non-zero drift under a real-world probability distribution, and hence are mathematically intractable.

However, the Cameron-Martin-Girsanov theorem tells us that we can convert such a process into a Martingale (i.e. zero drift) under a risk-neutral probability measure \mathbf{Q} .

The Binomial Representation Theorem demonstrates that for any contingent claim X based on a Martingale distribution, the process $V = E_{\mathbf{Q}}(X | \mathbf{F})$, which represents the present value of the claim, is also a Martingale ...

... and that we can find a previsible tool ϕ from which to construct it from the underlying distribution ...

... so we can find the value of V and hence price the derivative.

QUESTION 2

Syllabus section: (a) – (d), (f)

Core reading: 1 – 4, 6

(i)

Value at expiry = $\max\{F - K, 0\}$.

Hence value today = $e^{-rt} \int_K^\infty (x - K) f(x) dx$

where r is the risk-free rate, $f(x)$ is the p.d.f. for the futures price F .

(ii)

If $f(x) = \frac{x^{\lambda-1} \exp(-x/\beta)}{\beta^\lambda \Gamma(\lambda)}$, then value of European Call option

$$= e^{-rt} \left[\int_0^\infty \frac{x^\lambda \exp(-x/\beta)}{\beta^\lambda \Gamma(\lambda)} - \int_0^K \frac{x^\lambda \exp(-x/\beta)}{\beta^\lambda \Gamma(\lambda)} \right]$$

$$- Ke^{-rt} \left[\int_0^\infty \frac{x^{\lambda-1} \exp(-x/\beta)}{\beta^\lambda \Gamma(\lambda)} dx - \int_0^K \frac{x^{\lambda-1} \exp(-x/\beta)}{\beta^\lambda \Gamma(\lambda)} dx \right]$$

but since $\int_0^\infty f(x) dx = 1$, $\int_0^\infty \frac{x^{\lambda-1} \exp(-x/\beta)}{\beta^\lambda} dx = \Gamma(\lambda)$

so $\int_0^\infty \frac{x^\lambda \exp(-x/\beta)}{\beta^\lambda} dx = \beta \int_0^\infty \frac{x^\lambda \exp(-x/\beta)}{\beta^{\lambda+1}} dx = \beta \Gamma(\lambda + 1)$

$$\begin{aligned}
 &= e^{-rt} \left[\frac{\beta \Gamma(\lambda + 1)}{\Gamma(\lambda)} - \int_0^K \frac{x^\lambda \exp(-x/\beta)}{\beta^\lambda \Gamma(\lambda)} \right] \\
 &\quad - Ke^{-rt} \left[1 - \frac{1}{\beta} \int_0^K \frac{x^{\lambda-1} \exp(-x/\beta)}{\beta^{\lambda-1} \Gamma(\lambda-1)} \frac{\Gamma(\lambda-1)}{\Gamma(\lambda)} dx \right] \\
 &= e^{-rt} \left[\frac{\beta \Gamma(\lambda + 1)}{\Gamma(\lambda)} - P(\lambda, \beta, K) \right] - Ke^{-rt} \left[1 - P(\lambda - 1, \beta, K) \frac{\Gamma(\lambda-1)}{\beta \Gamma(\lambda)} \right] \\
 &= e^{-rt} [\beta \lambda - P(\lambda, \beta, K)] - Ke^{-rt} \left[1 - \frac{P(\lambda - 1, \beta, K)}{\beta(\lambda-1)} \right]
 \end{aligned}$$

[Unfortunately, the solution presented in the examination paper omitted the second bracketed term. As a result, candidates who had shown the right approach were given appropriate marks whether or not they had reached the correct answer above or the printed answer in the paper. N.B. The second bracketed term is also equivalent to $-Ke^{-rt} \Pr[F(t) > K]$.]

(iii)

Advantages of using gamma

- it appears to describe actual option prices better than the lognormal
- this could be useful to provide an alternative valuation for arbitrage purposes
- the gamma distribution has a fatter tail (higher kurtosis) – this should give better performance on large moves
- easier to manage hedging if parameters are stable
- if hedging is more stable, transactions costs should be reduced

Disadvantages of using gamma

- no clear arbitrage free hedging strategy exists, whereas it does for the lognormal distribution
- use of the gamma distribution has no economic basis
- the stability of the β parameter may be temporary
- it would be unusual – models would be non-standard and few colleagues would understand them
- as it is not a market standard approach, calibration and fitting would be harder
- there is no parameter represented that corresponds to ‘volatility’

QUESTION 3

Syllabus section: (k)

Core reading: 14

(i)

(a)

The “market price of risk” (MPR) is a measure of the trade-off that an investor makes between risk and return in the real world.

Typically, in the real world, the drift μ_i on asset i as seen by an investor is:

$$\mu_i = r + \lambda \sigma_i$$

where r is the risk-free rate, and σ_i is the volatility (standard deviation) of the asset's return.

Another way of writing this is: $\frac{\mu_i - r}{\sigma_i} = \lambda$ for all assets i , i.e. the risk-adjusted return above the risk-free rate is constant.

Clearly, $\lambda = 0$ in the risk-neutral world.

[*Note: other approaches were acceptable, e.g. a more general description of MPR.*]

(b)

EITHER

By algebra:

Consider two non income-producing assets $V_1 = V_1(r, t)$ and $V_2 = V_2(r, t)$, which depend only on r and t .

Using Ito's lemma we can express the processes they follow as:

$$dV_1 = \mu_1 V_1 dt + \sigma_1 V_1 dz \quad \text{and} \quad dV_2 = \mu_2 V_2 dt + \sigma_2 V_2 dz$$

for some values of μ_1 and σ_1 etc. (which will also possibly depend on r and t).

Now, considering the portfolio $\Pi = (\sigma_2 V_2) V_1 - (\sigma_1 V_1) V_2$, we have

$$d\Pi = (\mu_1 \sigma_2 V_1 V_2 - \mu_2 \sigma_1 V_1 V_2) dt$$

which is riskless, so must earn the risk-free rate over that period, i.e. $r\Pi dt$.

Hence, substituting for Π and equating the dt terms:

$$r(\sigma_2 V_2 V_1 - \sigma_1 V_1 V_2) = (\mu_1 \sigma_2 V_1 V_2 - \mu_2 \sigma_1 V_1 V_2)$$

i.e. $r\sigma_2 - r\sigma_1 = \mu_1\sigma_2 - \mu_2\sigma_1$

i.e. $\frac{\mu_1 - r}{\sigma_1} = \frac{\mu_2 - r}{\sigma_2} = \lambda$, say

and hence λ can depend on r and t , but cannot be dependent on the process for any particular asset.

OR

By argument:

Assume the short rate r follows a stochastic process under probability measure \mathbf{P} .

Changes between versions of the real world, and the risk-neutral world, only involve a change of probability measure.

But the CMG theorem states that a change of measure for r , from say \mathbf{P} to \mathbf{Q} , will result solely in the introduction of a different drift term. This term is dependent only on the change of measure.

For a given volatility, under the real world \mathbf{Q} any asset based on r must experience the same additional drift term, independent of the asset itself. The drift term is essentially the market price of risk, which therefore is the same for every asset being valued.

(ii)

(a)

Both methods integrate the PDE. This is obvious for FD as it is a technique for solving PDEs. It is less obvious for the binomial tree.

The FD method performs the integral directly by approximating continuous time and state variable values with small discrete steps ...

... by setting up a discrete rectangular grid of price changes and time steps spanning all possible outcomes of the stock evolution to final expiry.

The two types of approach are: 'implicit FD', which approximates a difference by taking values at the nearest previous time step ...

... and 'explicit FD', which does the same but for the next time step.

The explicit method is functionally the same as a trinomial tree.

The binomial tree method also evaluates the integral, but more subtly, using an expectation over the risk-neutral probability measure.

The definition of the binomial creates the risk-neutral probabilities of up and down moves – say this measure is \mathbf{P} .

Then, under \mathbf{P} , the expectation $\mathbf{E}_{\mathbf{P}}[V_t | F_s]$ is a Martingale, so equals V_s .

The binomial tree approximates this expectation directly using \mathbf{P} , hence provides an approximation of the expectation integral over each time step.

(b)

Both methods are suitable for valuing derivatives, and can also easily obtain Deltas, Gammas and Vegas by perturbation.

FD is computationally intense, slow, fiddly, prone to instabilities ...

... one problem in the implicit FD method being having to provide an entire rectangular grid of values from the lowest value of r to the highest, even though only the central ones are relevant ...

... however, lots of literature (especially from Physics) exists on how to make this method efficient (e.g. “adaptive mesh”) and solve different boundary conditions.

If the option is path dependent, then it is not possible simply to start the process at the final step.

Binomial trees are fast and efficient, intuitive, stable ...

... however, they don't converge very quickly ...

... and implementation becomes much more complex if the tree does not recombine, as tends to occur for certain processes.

In the context of interest rates, additional branching patterns may be required depending on where the interest rate lies currently compared with the maximum and minimum anticipated rates.

For either model, one factor is not ideal for an option on any asset where correlation might affect the outcome – but each model can be extended to multiple factors.

QUESTION 4

Syllabus section: (l) + (m)

Core reading: 15, 16

(i)

[Pension funds experience longevity risk on their savings contracts. To some extent this is offset by liabilities from term assurance contracts. The insurance company is concerned specifically about high survival rates.]

The main approach in every case is to express survival probabilities from the particular cohort of lives into a *survivor index* calculated up to some horizon T ,

$$\text{i.e. } S(t) = p(0, x).p(1, x)....p(t-1, x) \quad t = 1, 2, \dots, T$$

where $p(k, x)$ is the one-period survival probability for group of lives x from time k to $k + 1$, calculated with reference to a specified relevant population.

The requisite derivative instruments are swaps, bonds and options (caps). Swaps and bonds hedge this risk outright, and hence there is no surplus gain to be generated if survival rates are lower than anticipated, whilst options allow the company to continue to take one-sided gain in such a case.

Swaps

A mortality swap of term T is an OTC contract that swaps a notional amount times the survivor index each period $t = 1, 2, \dots, T$ for regular fixed or LIBOR-linked floating payments.

By receiving or paying mortality payments on this swap in the appropriate nominal amount, the company will receive payments that match its outflows on longevity risk.

Longevity Bonds

A longevity bond of maturity T is a tradable security that has coupons linked to the survivor index at times $t = 1, 2, \dots, T$.

By holding this bond in the appropriate nominal amount, the company will receive payments that match its outflows on longevity risk.

Principal-at-risk bonds

Instead of using the coupons to hedge the longevity risk, bonds can be constructed where the final principal at time t reflects the survivor index at time t .

Hedging is effected by issuing a strip of these bonds.

Survivor caps

Survivor caps are OTC contracts like interest-rate caps. A survivor caplet will pay at time t the maximum of $S(t) - K(t)$ and 0, where $K(t)$ is an agreed fixed strike amount.

A survivor cap of term T is simply a collection of survivor caplets with payment dates $t = 1, 2, \dots, T$.

(ii)

In an OTC derivative contract, both parties undertake to pay each other cashflows at certain times. Cashflows due on the same date are usually offset, i.e. only the net balance passes between the parties in a single payment. Counterparty risk arises from the value of future cashflows being at risk to default, and the amount of exposure is often linked to (and hence varies with) underlying market variables such as interest rates, foreign exchange rates, equity prices etc.

If a particular firm has a counterparty that defaults, future payments will not be made, leading potentially to loss for the firm. However, there can only be a loss on a derivative contract if a counterparty defaults and the contract with them has positive value to the firm (i.e. net positive cashflows payable to the firm) at the time of default.

Mitigants are:

- Break clauses – allowing a party to break the agreements say every 5 or 10 years in the event of a serious credit downgrade of the counterparty; also, could allow the contract to be re-negotiated at market rates periodically.
- Netting – ensure the derivative is subject to a legal agreement to net.*
- Collateralisation – asking for temporary cash collateral to cover the mark-to-market of the trade, thereby ensuring that a positive mark-to-market value of the OTC trade can be recovered if the counterparty defaults since one side will always hold money to cover their potential loss; probably link thresholds to the credit rating of the counterparty.
- Diversification – using a range of banks to avoid concentration on any one.
- Design of contract – avoiding getting the market value too large.
- Rating quality – policy guidelines on the lowest allowable rating for a counterparty bank, for example S&P single A (or equivalent).
- Monitoring exposure – careful controls on the extent of counterparty risk.
- Credit insurance – using credit derivatives (e.g. credit default swaps), or seeking a parental or third party guarantee.

[This is a very comprehensive list – candidates would not be expected to include all of these.]

[**Note: the signing of bilateral ISDA legal contracts (ISDA master agreements or confirmations) ensures that all derivatives are amalgamated together for offset, thereby avoiding a defaulting institution calling cash in on some transactions but not paying out others. Most banks and major counterparties sign ISDA agreements when trading derivatives, and these would often require daily margining (i.e. collateralisation) via a Credit Support Annex.*]

(iii)

(a)

LPI is the same as RPI (Retail Price Indexation) inflation except that it is capped at a certain level (say 5%) and usually floored at zero.

Hence a company with an LPI-linked liability has basic RPI risk but has effectively bought a cap on RPI and sold a 0% floor.

To mitigate the risk, the company can find an LPI-linked hedge – take out an LPI swap, or buy an LPI bond.

Although inflation is capped under LPI, in a low inflation environment LPI and RPI are similar, hence an RPI hedge might be effective, as well as giving upside ...

... but care must be taken with an RPI hedge (or LPI hedge without the floor), as there is also a risk that inflation falls below zero, i.e. turns to deflation, because in that situation the 0% floor will kick in and the hedge will be mismatched.

Alternative hedges are: transfers of liabilities, securitisations, buyouts, reinsurance.

There can be a second order cross-impact between LPI and longevity: inflating payments increases the exposure to longevity (by decreasing the effective discount rate), or if longevity is under-estimated then any LPI hedge will be insufficient.

(b)

Those with risk from inflation increasing are ones whose liabilities are inflation-linked ...

... such as pension funds running Defined Benefit (DB) schemes, companies with fixed price contracts and costs that rise with inflation (e.g. energy companies), issuers of index-linked bonds or annuities.

Those with risk from inflation decreasing are ones whose assets are inflation-linked ...

... such as housing trusts, hospitals and schools with income linked to government contracts, companies with products or services that rise in line with inflation (e.g. shops, estate agents etc), pensioners in DB schemes.

Of these, those firms with inflation-linked income are very likely to sell inflation, and the pension funds are the main buyers.

[The above points are the main ones that should be made. Partial recognition was given for mentioning other related risks, such as commodity price risk for producers of raw materials, and manufacturers using them.]

QUESTION 5

Syllabus section: (h(iv)-(ix), (i))

Core reading: 10 – 12

(i)

Delta

A small Delta loss has been made, so the book was generally well hedged.

Markets must have moved more often against the institution, given its Delta mismatch, than in favour of it.

If delta hedging is working properly, Delta profits/losses will be immaterial over long periods, so the Delta loss is not of concern.

Gamma

Gamma gives a loss in each period: this is to be expected given that the option has negative Gamma and the futures have zero Gamma.

Unsurprisingly, the Gamma losses were biggest in the volatile period.

Gamma losses have been exacerbated as the hedge was being moved backwards and forwards during the volatile period, effectively selling equities when cheap and buying equities when expensive.

Theta

Theta profits represent how the time value of the option would have decayed had it been held.

There are Theta profits in each period – these are to be expected given that the option has negative Theta and the futures have zero Theta.

Theta and Gamma must be considered together since they are opposite sides of the same coin (appearing on opposite sides of the Black-Scholes partial differential equation) ...

... since by delta hedging, the institution is giving up the Theta losses that buying an option as a hedge would have incurred, and taking on Gamma losses instead.

Rho

There are big Rho losses, which is surprising.

It looks as if interest rates fell (particularly in the Volatile period), making the value of the option increase but not impacting on futures payoffs.

It seems that the institution was not hedged against changes in interest rates. Questions need to be asked about why it wasn't.

Vega

Vega profits/losses must be zero as it was assumed that volatility did not change.

General lessons learned

The strategy of choosing to delta hedge rather than buying options appears to have lost money.

If part of the rationale behind delta hedging was that implied volatilities were higher than expected realised market volatilities, then one would expect Theta profits to more than offset Gamma losses ...

... but it's the other way round here, as a result of realised market volatility being so high during the middle period.

If the view is still that implied volatility is higher than future realised volatility, then the institution should continue to delta hedge and put down the experience to date as "bad luck".

Perhaps volatility should have been increased in the delta hedging calculations during the volatile period to provide more accurate hedging.

Using a single volatility input is a weakness in the analysis – to get a clearer picture, one should see how the numbers change when the options are valued each day using the actual market implied volatilities.

(ii)

The difference in P&L from -54 to -52 is almost certainly due to beginning and end valuation differences in the options.

Delta

No exposure to report – again the book is well hedged.

Gamma

Gamma losses are smaller, because with higher implied volatility the option becomes less sensitive to underlying price movements.

Vega

Vega shows big losses, because market implied volatilities rose significantly during the volatile period and Vega was not hedged.

It was not necessarily a mistake to leave Vega unhedged – if the institution wanted to hedge Vega as well as delta, it might as well have bought the option!

Vega losses will largely revert when implied volatilities return to normal levels ...

... which looks as if it hasn't happened yet, given there's still a Vega loss there.

Theta

Theta profits are bigger, because time value is now bigger (with bigger volatility), so disappears faster.

Rho

Rho losses now look slightly worse, because the time value of the option will have been much bigger (due to higher volatility), so more sensitive to interest rate changes.

[Other valid points could be made in both parts (i) and (ii).]

QUESTION 6

Syllabus section: (e) + (j)

Core reading: 5, 13

(i)

It is important that the family of formulae used to fit to the yield curve allow for a rich range of possible shapes, including the current shape.

Cubic splines are often used to fit the curve through all the points in a piecewise smooth fashion, although they given little additional information by doing so ...

... but there are other higher parameterisation curves that can be used.

On the choice of parameters:

- With 15 parameters and 15 swap rates, there will be a unique combination of 15 factors that fits to the 15 published swap rates.
- With more than 15 parameters, there will be an infinite number of possible fits: the range of parameters is too wide.
- With less than 15 parameters, the model will not reproduce all 15 swap rates exactly: some interpolation or smoothing will be required.

Adding parameters would mean that the model would fit better, but might have redundant parameters and very likely would be less stable.

Removing parameters would mean that the model would only fit approximately and hence may not be arbitrage free (i.e. does not reproduce the benchmark swap rates).

The resulting yield curve needs to be checked for reasonableness ...

... for an imperfect *non* arbitrage-free fit, calculate a measure of the fit (for example sum of squares of differences), possibly with weightings to reflect the swap rates that it is most important to replicate closely ...

... ideally giving the best fit for the smallest number of parameters.

A sensible constraint for the longer end ($t = 30$) is to tend to a constant value.

A sensible constraint for the short end ($t = 0$) is to have a constant gradient.

In general, a 'once per curve' modelling is used, i.e. there is no guarantee that the parameters have any economic meaning or will be stable from one period to the next.

If the parameters are not 'orthogonal' (i.e. relating to independent effects) then adding new parameters will change the previous ones.

(ii)

Avoidance of arbitrage is the main reason for using bootstrapping. The bootstrapping method exactly fits all the main observable swap rates, whereas the curve fitting method may not exactly replicate published swap rates, so this could open up the possibility of arbitrage.

Bootstrapping is easy and reliable – it always works. Curve fitting may not give a solution for some shapes of curve.

Normally, bootstrapping for swaps uses money-market instruments such as LIBOR rates and interest rate futures to improve the fit at the short-end ...

... but there are discontinuities across the boundary between swaps, futures and money-markets.

Bootstrapping will often include some assumptions in interpolating spot or forward rates between different points, whereas the curve fitting method will likely only use the points specified.

Bootstrapping will be more difficult for swaps than it is for zero coupon bonds, e.g. the one year swap can be used to determine $r(0.5)$ and $r(1)$ but an assumption needs to be made about how the two are related.

What constitutes an acceptable family of functions today might not work for tomorrow's yield curve. The family of functions may need to be repeatedly revised.

Various techniques for interpolation (between benchmark rates) and extrapolation (beyond the final rate) are also used.

(iii)

Mainly the dealer would use the sensitivity of the swap book to individual swap rates to determine a set of swap trades that would neutralise (instantaneous) interest rate risk. (Of course, they do not have to hedge unless wanting to reduce risk.)

The interest rate sensitivity of a swap portfolio can always be reduced to a portfolio of benchmark swaps using perturbation analysis, since the benchmark curve forms a spanning set.

The dealer chooses the main benchmark swaps for this analysis because they are the most liquid to trade (i.e. can be dealt in large size with the lowest transaction costs).

[Using principal component analysis, most swap books could be expressed using very few parameters – probably no more than four or five – but there is no guarantee that these sparse parameters would coincide with easily tradable instruments.]

In addition, the dealer can use benchmark sensitivities to amalgamate interest rate exposure by currency with other dealers in the bank, so that aggregated risk management and reporting can take place.

Looking at the effect of changing individual swap rates will also enable the dealer to see the impact of particular significant scenarios, e.g. yield curve rising or falling, steepening or flattening, changing shape. This is part of what is called stress testing, and aims to mirror the way certain economic scenarios could play out, particularly useful if these have been taken from past extreme situations.

(iv)

(a)

For each period, the discount factor at period i $d_i = \frac{d_{i-1}}{(1 + f_i)}$. Bootstrapping:

$$d_1 = \frac{1}{1.005} = 0.995025, \quad d_2 = \frac{0.995025}{1.0125} = 0.982741, \quad d_3 = \frac{0.982741}{1.015} = 0.968217$$

The swap fixed rate for term n , g_n , is given by $g_n = \frac{1-d_n}{\sum_{k=1}^n d_k}$, so

$$g_3 = \frac{1-d_3}{d_1+d_2+d_3} = \frac{0.031783}{2.945983} = 0.0107886, \text{ i.e. the 3-year swap rate is 1.07886\%}$$

which on €1m will be a payment of €10,788.60 per year.

(b)

Reducing the three year rate to 1.25% changes only $d_3 = \frac{0.982741}{1.0125} = 0.970608$.

$$\begin{aligned} \text{Value of swap} &= g_n \left(\sum_{k=1}^n d_k \right) - (1-d_n) = 0.0107886 \cdot 2.948374 - 1 + 0.970608 \\ &= 0.00241683 \end{aligned}$$

so on €1m this will be a mark-to-market increase of €2,416.83.

[Other valid points could be made in parts (i), (ii) and (iii).]

QUESTION 7

Syllabus section: (h)(i)–(iii)

Core reading: 8, 9

(i)

(a)

Using Black-Scholes to price the Call option:

$$d_1 = (\ln(500/600) + (5\% + 0.5 \cdot 22.5\% \cdot 22.5\%)) / 22.5\% = -0.4756$$

$$d_2 = d_1 - 22.5\% = -0.7006$$

$$C = 500 * N(-0.4756) - 600 * \exp(-5\%) * N(-0.7006) = 20.60 \text{ as required}$$

Using Black-Scholes to price the Put option:

$$d_1 = (\ln(500/400) + (5\% + 0.5 * 22.5\% * 22.5\%)) / 22.5\% = 1.3265$$

$$d_2 = d_1 - 22.5\% = 1.1015$$

$$P = 400 * \exp(-5\%) * N(-1.1015) - 500 * N(-1.3265) = 5.33 \text{ as required}$$

(b)

Let the four risk-neutral probabilities be p_1, p_2, p_3, p_4 with p_1 at the bottom of the tree as shown in the diagram.

An asset that pays X_i in i th scenario is worth $[X_1 p_1 + X_2 p_2 + X_3 p_3 + X_4 p_4] \exp(-5\%)$

Bonds will be priced correctly because

$$[\exp(5\%) p_1 + \exp(5\%) p_2 + \exp(5\%) p_3 + \exp(5\%) p_4] \exp(-5\%) = 1$$

For the Call to be priced properly: $[0 p_1 + 0 p_2 + 0 p_3 + 100 p_4] \exp(-5\%) = 20.60$

$$\text{So } p_4 = 0.217$$

For the Put to be priced properly: $[100 p_1 + 0 p_2 + 0 p_3 + 0 p_4] \exp(-5\%) = 5.33$

$$\text{So } p_1 = 0.056$$

Note that p 's sum up to 1, so $p_3 = 1 - 0.056 - 0.217 - p_2 = 0.727 - p_2$

For equities to be priced properly:

$$[300*0.056 + 450p_2 + 550(0.727 - p_2) + 700*0.217] \exp(-5\%) = 500$$

So $p_2 = 0.428$ and $p_3 = 0.299$.

(ii)

[These answers reference the calculations in (i).]

Deriving the p's involves solving four simultaneous equations as before, but with different values for the Call and the Put.

The Call option has lower volatility so will be lower in price ...

... and since the Call has gone down in value, p_4 will be lower.

The Put option has higher volatility so will be higher in price ...

... and since the Put has gone up in value, p_1 will be higher.

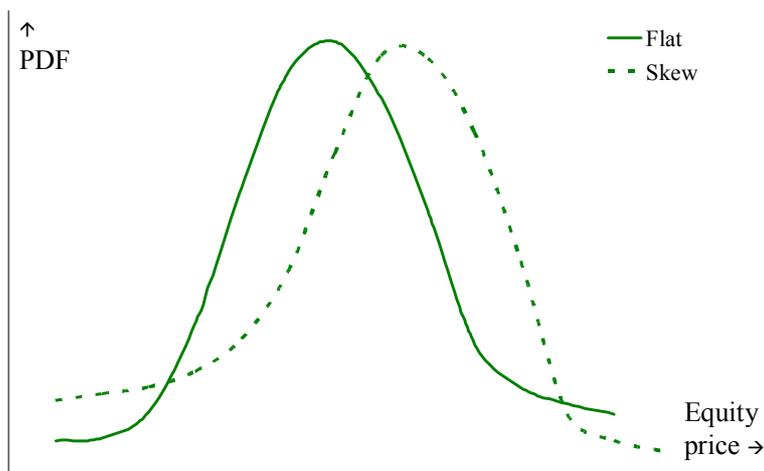
To quantify impact on p_2 and p_3 , we note that:

- $p_2 + p_3$ is constrained by the four probabilities summing to 1
- the expected one-period equity price in the risk neutral measure will be unchanged, i.e. $[300*p_1 + 450p_2 + 550p_3 + 700*p_4] \exp(-5\%) = 500$

so to leave the expected value unchanged after p_1 rises and p_4 falls we need to increase weighting on scenario 3 at the expense of scenario 2.

Hence p_2 will be lower and p_3 will be higher.

(iii)



Features required are: axes and bell shapes for curves, fatter tail at low end for Skew, thinner tail at high end for Skew, middle moving to right for Skew.

QUESTION 8

Syllabus section: (g) + (i)

Core reading: 7, 12

(i)

The life insurance company could be hedging to protect its solvency position against the risk of equity price falls ...

... especially in uncertain times or if equity prices are falling / have fallen

Legislation may require it ...

... or there could be equity guarantees that are at risk to equity prices.

Collars could be of interest because they reduce the cost of hedging ...

... which may be important if market implied volatilities are currently high...

Also, there may not be enough free cashflow to allow the purchase of a Put.

Collars may look especially attractive at the moment if equity volatilities are not as skewed as normal (i.e. if implied volatility is not as quickly a decreasing function of strike as it normally is) ...

... or (for whatever reason) equity upside may not be especially important to the company.

(ii)

The one-year term of the option could be based on:

- A solvency test based on a one year long stress
- The longest period over which the company is prepared to sacrifice equity upside
- The period over which market uncertainty is expected to last

Insisting on the collar being zero cost eliminates one degree of freedom, leaving two choices: term, and either width of collar ($Y - X$) or strike X of floor (Put).

If the company is reviewing an OTC option, a bank is likely to have quoted a range of possible collar widths for each term ...

... whereas if an exchange-traded option, there may be liquidity considerations.

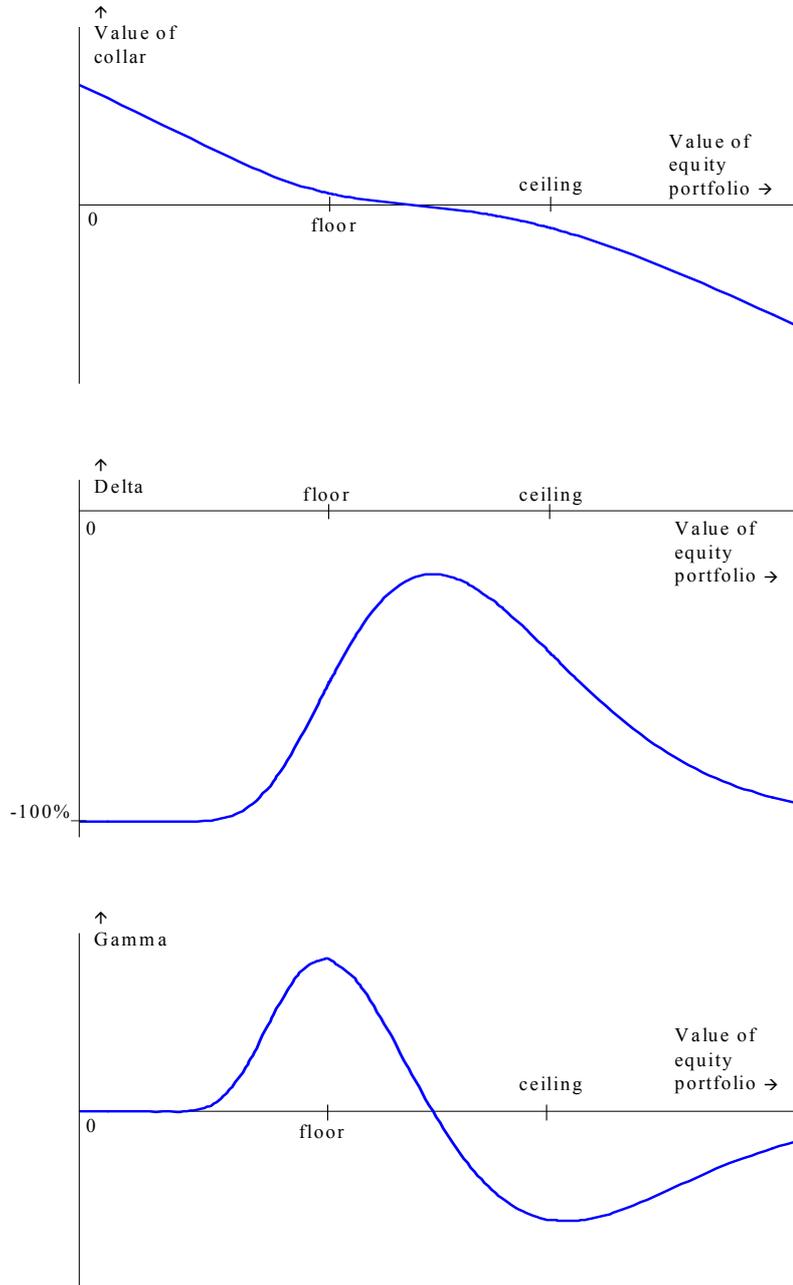
The collar width (or strike of floor) could have been set based on:

- the floor being the floor that absolutely protects solvency
- the floor being at the level that restricts the probability of ruin to the level at which the company is comfortable (i.e. within risk appetite)
- the maximum amount of upside the company is prepared to give up
- for an exchange-traded option, the availability of market prices

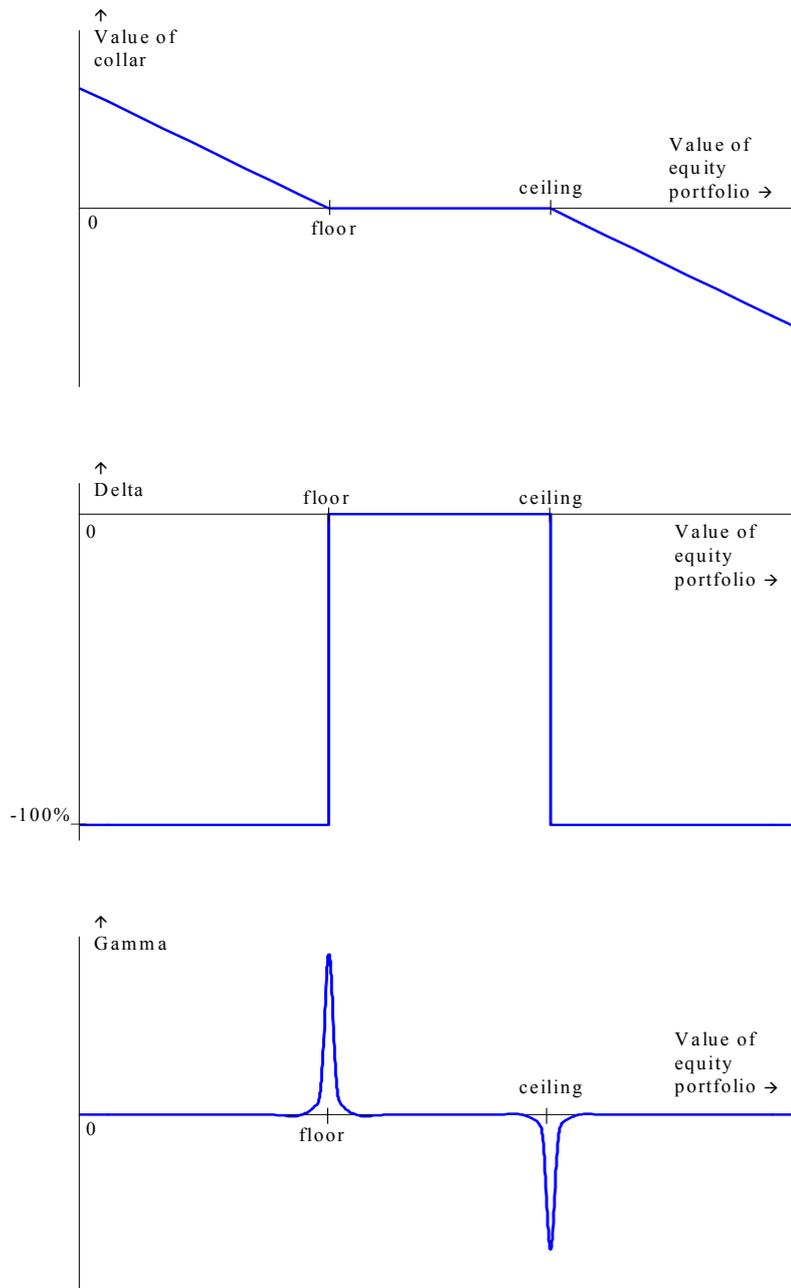
[*Other valid points could be made in parts (i) and (ii).*]

(iii)

Today



Close to maturity



END OF EXAMINERS' REPORT