

**Subject ST6 — Finance and Investment
Specialist Technical B**

EXAMINERS' REPORT

April 2009

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 2009

QUESTION 1

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

Core reading: Units 8 & 9

(i)

(a) **P** and **Q** are equivalent if and only if they operate on the same sample space and every positive probability outcome for **P** has an equivalent positive probability outcome under **Q**. This is clearly the case because each distribution has positive probabilities.

[This could also be stated in symbols.]

(b) Starting at $t = 2$, take each branch of the tree and discount under **P** and **Q**.

Top branch: under **P**, the expected value = $180 \times 0.5 + 100 \times 0.5 = 140$.

Top branch: under **Q**, the expected value = $180 \times 0.25 + 100 \times 0.75 = 120$.

The value of the tree at $t = 1$ for that branch is 120, so **Q** gives that value not **P**.

Similarly, for the lower branch, under **P** the value is 55, under **Q** it is 60 – same result.

Finally, for the initial branch, under **P** the value is 90, under **Q** it is 80 – same result.

Hence **Q** is a martingale measure since $E_{\mathbf{Q}}(X_j | F_i) = X_i$ for all $i < j$, but this does not apply under **P**, so **P** is not a martingale. (F_i is the filtration i.e. history up to time i .)

(ii)

(a) Under **P** we have:

$$\mu_0 = E_{\mathbf{P}}((S_1 - S_0) | F_0) = (120 - 80) \times 0.5 + (60 - 80) \times 0.5 = 10$$

and

$$\sigma_0^2 = E_{\mathbf{P}}((S_1 - S_0)^2 | F_0) - \mu_0^2 = (120 - 80)^2 \times 0.5 + (60 - 80)^2 \times 0.5 - 100 = 900$$

so $\sigma_0 = 30$.

[Note that in continuous time changes of measure do not change the volatility.

*Under measure **Q**, $\sigma_0 = 28.3$ on the discrete formulation. This could be used in the calculation below instead of 30.]*

(b) For the first time step, $\frac{dQ}{dP}\Big|_{t=1}^{up} = \frac{\frac{1}{3}}{\frac{1}{2}} = \frac{2}{3}$ and $\frac{dQ}{dP}\Big|_{t=1}^{down} = \frac{\frac{2}{3}}{\frac{1}{2}} = \frac{4}{3}$

For the second time step, $\frac{dQ}{dP}\Big|_{t=2}^{up,up} = \frac{\frac{1}{4}}{\frac{1}{2}} \frac{dQ}{dP}\Big|_{t=1}^{up} = \frac{1}{2} \frac{2}{3} = \frac{1}{3}$, $\frac{dQ}{dP}\Big|_{t=2}^{up,down} = \frac{\frac{3}{4}}{\frac{1}{2}} \frac{dQ}{dP}\Big|_{t=1}^{up} = \frac{3}{2} \frac{2}{3} = 1$,

$$\frac{dQ}{dP}\Big|_{t=2}^{down,up} = \frac{\frac{3}{5}}{\frac{1}{2}} \frac{dQ}{dP}\Big|_{t=1}^{down} = \frac{6}{5} \frac{4}{3} = \frac{8}{5} \text{ and } \frac{dQ}{dP}\Big|_{t=2}^{down,down} = \frac{\frac{2}{5}}{\frac{1}{2}} \frac{dQ}{dP}\Big|_{t=1}^{down} = \frac{4}{5} \frac{4}{3} = \frac{16}{15}.$$

(iii)

(a) The CMG theorem states that, if W_t is a \mathbf{P} -Brownian motion and γ_t is a pre-visible bounded* process, then there exists a measure \mathbf{Q} equivalent to \mathbf{P} such that

$$\frac{dQ}{dP} = \exp\left(-\int_0^T \gamma_t dW_t - \frac{1}{2} \int_0^T \gamma_t^2 dt\right)$$

and $\tilde{W}_t = W_t + \int_0^t \gamma_s ds$ is a \mathbf{Q} -Brownian motion.

[* It is not really necessary to say how the process should be bounded, though it is strictly correct to do so.]

(b) From CMG, with $\gamma = \frac{\mu_0}{\sigma_0}$, the required Radon-Nikodym derivative at time 1 is:

$$\begin{aligned} \frac{dQ}{dP} &= \exp\left(-\int_0^1 \frac{\mu_0}{\sigma_0} dW_s - \frac{1}{2} \int_0^1 \frac{\mu_0^2}{\sigma_0^2} ds\right) = \exp\left(-\frac{\mu_0}{\sigma_0} W_1 - \frac{\mu_0^2}{2\sigma_0^2}\right) \\ &= \exp\left(-\frac{1}{3} W_1 - \frac{1}{2} \left(\frac{1}{3}\right)^2\right) \end{aligned}$$

using the values for μ_0 and σ_0 in part (ii)(b).

QUESTION 2

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

Core reading: Units 8 & 9

(i) Since the moment generating function for a $X \sim N(\mu, \sigma^2)$ Normal distribution is $M(\theta) = E[\exp(\theta X)] = \exp(\mu\theta + \frac{1}{2}\sigma^2\theta^2)$, the result follows immediately.

OR using integrals:

$$\begin{aligned} E[\exp(\theta X)] &= \frac{1}{\sqrt{2\pi v^2}} \int_{-\infty}^{\infty} \exp(\theta x) \exp\left(-\frac{1}{2}\frac{x^2}{v^2}\right) dx \\ &= \frac{1}{\sqrt{2\pi v^2}} \int_{-\infty}^{\infty} \exp\left(-\frac{1}{2}\left\{\frac{x}{v} - \theta v\right\}^2 + \frac{1}{2}\theta^2 v^2\right) dx \\ &= \frac{1}{\sqrt{2\pi v^2}} \exp\left(\frac{1}{2}\theta^2 v^2\right) \int_{-\infty}^{\infty} \exp\left(-\frac{1}{2}\left\{\frac{y}{v}\right\}^2\right) dy \quad ** \\ &= \exp\left(\frac{1}{2}\theta^2 v^2\right) \end{aligned}$$

[** substituting $y = x - \theta v^2$]

(ii) The process $B = \{B_t : t \geq 0\}$ is a Brownian Motion under probability measure \mathbf{P} , if:

- B_t is continuous,
- $B_0 = 0$
- B_t is distributed as $N(0, t)$ under \mathbf{P} . (*)
- $B_t - B_s$ for $t > s > 0$ is distributed as Normal $N(0, t - s)$ under \mathbf{P} , and is independent of any filtration (or history) of the process up to time s .

[*Note: (*) is usually given as a separate condition, although it can be derived from the other criteria. Candidates were not penalised for omitting it.*]

(iii)

This uses the result from (i), i.e. for $B_t \sim N(0, t)$ and $B_t - B_s \sim N(0, t - s)$ then $E\{\exp(\theta B_t)\} = \exp(\frac{1}{2}\theta^2 t)$ and $E\{\exp(\theta(B_t - B_s))\} = \exp(\frac{1}{2}\theta^2(t - s))$.

Remember that since $t - s > 0$, $B_t - B_s$ is independent of B_s etc.

Therefore, $E\{\exp(B_s + B_t + B_u)\} = E\{\exp[(B_u - B_t) + 2(B_t - B_s) + 3B_s]\}$
 $= \exp\frac{1}{2}[(u - t) + 4(t - s) + 9s] = \exp\frac{1}{2}(u + 3t + 5s)$.

(iv)

Let B_t be a Brownian Motion process and let S_t be the stock price process.

Brownian motion wanders over time with an element of randomness which is suitable for modelling the uncertainty in the behaviour of equity market price changes ...

... but we need to scale for the volatility of the relevant stock or index by adding a volatility parameter σ such that $S_t = \sigma B_t$.

Brownian motion has zero drift (mean), whereas the equity value of a company (or index of companies) grows or falls at some rate dependent on the company's performance and other economic factors, e.g. the rate of inflation ...

... but we can adjust for this by adding a drift parameter μ such that $S_t = \sigma B_t + \mu t$ (or in differential form $dS_t = \sigma dB_t + \mu dt$).

If the stock price process followed a simple Brownian Motion there is a non zero probability that the stock price can take on negative values ...

... so we can adjust our model by using exponential growth, e.g. $S_t = \exp(\sigma B_t + \mu t)$ (or in differential form $dS_t = \sigma S dB_t + \mu S dt$).

QUESTION 3

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

Core reading: Units 10 - 13

(i)

The cost of guarantees and smoothing can be approximated by two straight lines connecting three points on the scatterplot:

- Cost of £10m at FTSE = 0
- Cost of zero at FTSE = 3,000
- Cost of -£5m if FTSE = 10,000 (i.e. a benefit)

The simple "stock + option" hedge will be constructed by:

1. Using a forward deal to make the gradient of the graph zero on one side of 3,000
2. Using an option to hedge the residual risk on the other side of 3,000

There are two possible hedges, one involving buying a Call and one involving buying a Put – either is acceptable.

[The calculations below allow for no difference between the current value of the FTSE and the 5-year forward price of the FTSE. This is not necessarily deductible from the question but was assumed by all candidates who attempted a solution.]

EITHER

Hedge 1

To zeroise the gradient on the left of the graph, we need a forward deal that pays £10m more at a FTSE of zero than at a FTSE of 3,000.

Can do this by forward selling $£10m / 3,000 = £3,333.3$ per unit of FTSE with exercise date in five years' time.

To flatten the profile for $FTSE > 3,000$, we need to buy a 5 year FTSE Call option with strike 3,000.

At a FTSE level of 10,000, the option needs to pay off $£3,333.3 \times 7,000$ (negative payoff from forward deal) less £5m (negative cost of guarantees and smoothing) = £18.333m. So we need to buy a Call option on $£18.333m / 7,000 = £2,619.0$ per unit of FTSE.

OR

Hedge 2

To zeroise the gradient on the right of the graph, we need a forward deal that pays £5m less at a FTSE of 10,000 than at a FTSE of 3,000.

Can do this by forward selling $£5m / 7,000 = £714.3$ per unit of FTSE with exercise date in five years' time.

To flatten the profile for $FTSE < 3,000$, we need to buy a 5 year FTSE Put option with strike 3,000.

At a FTSE level of 0, the option needs to pay off £10m (negative cost of guarantees and smoothing) less $£714.3 \times 3,000$ (payoff from forward deal) = £7.857m.

So we need to buy a Put option on $£7.857m / 3,000 = £2,619.0$ per unit of FTSE.

(ii)

[The sketch is not shown here, but should follow the description given.]

Sketch will show points clustered around a horizontal line for a FTSE ranging from 0 to 10,000, with a number of outlying points as in the original scatter-plot.

The level around which the points are clustered needs to be labelled on the y-axis. This level is $£10m / 3 = £3.333m$ for Hedge 1 and $£5m / 7 = £0.714m$ for Hedge 2.

(iii)

The cost of guarantees and smoothing is more likely to be dependent on total equity returns than on the performance of a capital only index.

The life insurance company might therefore find that the hedge is less accurate than it thought if:

- dividends on the FTSE vary from those that were assumed within the stochastic analysis
- reinvestment of the dividends is not at the rate/value assumed

This would not have been so problematic were the analysis to have been based on a total return index.

[Other common difficulties in modelling using an index, e.g. basis risk, roll risk and model risk, are not all relevant to the question.]

QUESTION 4

Syllabus section: (l) & (m)

Core reading: Units 15, 16

(i)

(a) Features of n th-to-default basket swap

This is a leveraged investment on the default risk of a reference portfolio (basket) where the buyer of protection is protected against the n th default in the basket by means of a special credit default swap (CDS), and the seller is exposed to it.

A CDS is a swap that pays a premium based on a notional principal in return for insuring against any credit losses on that principal according to the attachment points specified in the contract.

If and when the n th default occurs, the swap buyer is compensated for the amount of loss by the seller, and the swap is terminated.

As an example, if there are 10 names, a 4th-to-default basket swap is exposed only to losses of the 4th out of the 10 names to default in order of timing of default.

These swaps are similar to CDO tranches, with $n = 1$ being equivalent to the equity tranche etc.

A key difference, though, is that the return profile of the protection seller in the n th-to-default basket for higher values of n is independent of the severity of the first loss.

Usually these are structured and priced as bespoke over-the-counter deals. This means there is considerable flexibility in e.g. choosing the basket, maturity etc.

(b) Features of CDO

At its simplest, a CDO is a collection of credit-risky bonds.

The CDO is usually structured (tranching) so that the first losses go to the lowest (equity) tranche, then further losses go to the next (mezzanine) tranches, with the top (senior) tranche only receiving losses once all the lower tranches have been exhausted.

The margins on the different tranches are priced to reflect the different inherent credit risks. The tranches may then be rated according to their perceived credit worthiness.

Credit exposure can be taken in physical form by setting up a Special Purpose Vehicle (SPV) that invests in those bonds and passes the cashflows to the investors of the different tranches. The cashflow “waterfall” effectively determines the return on the bond.

Alternatively, exposure can be taken in synthetic form via a CDS with various attachment and detachment points. There is then the question of whether to be funded or unfunded, i.e. whether to pay a principal amount into a special risk-free fund (effectively to mirror the cashflow CDO) or just to hold the derivative.

(ii)

Risk appetite

The investor wants to avoid first loss so need more senior tranches.

Probably also avoid currency risk \Rightarrow unfunded exposure. Easy to do with either type.

Use of *n*th-to-default basket swap

Maybe sell protection on a third-to-default through to 10th-to-default.

This is like a senior (or super-senior) tranche of a CDO.

Don't want to be too senior or won't get any real exposure at all.

Good for the investor because he is worried about certain names but doesn't know which ... and does not need to worry about the severity of the first two losses ... but possibly may find the liquidity of this bespoke derivative poor and hence less cost-effective.

[Note: an alternative strategy is to buy the bonds and buy protection on the first and second to default.]

Use of senior CDO

Choose senior tranche – this does not take first losses, so in line with risk appetite.

Can probably adjust attachment point to reflect credit concerns.

Not quite so effective as *n*th-to-default for the particular view on a couple of names ...
... but if the market is active for those bonds, the CDO may be more keenly priced.

(iii)

Correlation effect

Increase in correlation spreads the losses further up the CDO tranches towards senior.

This has the impact of reducing the value of the senior tranches and hence (assuming no change in default probability) increasing the value of the junior tranches.

QUESTION 5

Syllabus section: (e) & (j)

Core reading: Units 5, 13

(i)

(a) LIBOR stands for "London InterBank Offer Rate".

It is the rate at which large banks active in a particular money market charge each other for borrowing (hence the word "offer") in a currency for a particular term.

LIBOR rates apply to the short-term international interbank market for large loans with maturities from 1 day (overnight) to 12 months and beyond.

It is not restricted to Sterling – there are LIBOR rates for almost every currency.

LIBOR is the most common reference base for pricing interest rate sensitive instruments, e.g. swaps, and loans to larger companies and institutions.

It is officially fixed in London once a day by a defined group of large banks, but the rate changes throughout the day.

[Other points on the practical operation of LIBOR could gain marks.]

(b) LIBOR rates are different from Government bond yields because bank counterparties are not credit risk free (in the way Government debt is, at least in its own country).

The spread between Government and LIBOR rates is usually positive (LIBOR is higher), but occasionally supply and demand effects will distort this relationship.

This does not prevent the LIBOR rate from being the standard rate for swaps since the interbank market is the place where most are transacted.

(ii)

We have the following relationship that applies to par yield g_n at term n (whether Gilt or swap):

$$100 = g_n \sum_{i=1}^n d_i + 100d_n \quad (*)$$

where d_i are the discount factors (zero coupon bond prices, in decimals) at year i .

For Gilts at year 25, $d_{25} = (1.0295)^{-25} = 0.48344$.

For swaps at year 25, $d_{25} = (1.0305)^{-25} = 0.47185$

so, from (*) rearranged, $\sum_{i=1}^{25} d_i = \frac{100(1-d_{25})}{g_{25}} = 15.09000$.

Putting these values in the table given, adding the cumulative column for swaps (sum of the discount rate column – not required for the Gilts as we have all the values):

<i>Term</i>	<i>Gilt par yield</i>	<i>Swap par rate</i>	<i>Gilt discount rate</i>	<i>Swap discount rate</i>	<i>Swap cumulative disc rate</i>
25	3.20	3.50	0.48344	0.47185	15.09000
26	3.15	3.40	0.47650	0.47093	15.56093
27	3.10	3.30	0.47023	0.47095	16.03188
28	3.05	3.20	0.46461	0.47188	16.50376
29	3.00	3.10	0.45960	0.47370	16.97746
30	2.95	3.00	0.45518		

[We actually only need the cumulative discount rate at $t = 29$.]

For the final value, i.e. the swap d_{30} , re-arrange (*) again to:

$$d_n = \frac{100 - g_n \sum_{i=1}^{n-1} d_i}{100 + g_n}$$

which gives swap $d_{30} = (100 - 3 * 16.97746) / 103 = 0.47638$

Then:

$$(a) \text{ Gilt zero coupon rate at 30 year} = 100 * ((d_{30})^{-1/30} - 1) = 2.658\%.$$

Similarly, swap zero coupon rate at 30 year = 2.503%.

$$(b) \text{ Gilt forward rate from 25 to 30 years} = 100 * \left[\left(\frac{d_{25}}{d_{30}} \right)^{1/5} - 1 \right] = 1.212\%$$

Similarly, swap forward rate from 25 to 30 years = -0.191%.

(iii)

There are two surprising features:

- Gilt zero coupon yields are higher than swap rates at the long end of the curve ...
... this could be due to an excess supply of long-dated Gilts (heavy Government funding) ...
... or a feature of the slope of the yield curve caused by the more actively traded middle section (the longer end is less liquid so arbitrages can persist there).
- Swap forward rates are negative ...
... this could be due to an imbalance of supply and demand at the long end ...
... or absence of arbitrage at the long end of the curve during market disruption.

QUESTION 6

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

Core reading: Units 10 - 13

(i)

The bank is effectively writing a Put option on residential property values against the value of the accumulated non-interest bearing loans.

S_0 is the current value of a residential property.

K is the strike price of the Put option. It is equal to the initial loan accumulated for T years at rate f .

r is the continuously compounded risk-free interest rate for term T .

q represents the "yield" on the property as a way of obtaining the forward price ...
... such as the cost of renting a similar property for term T , expressed as a percentage per annum of the property value ...

... but should also include an allowance for any extra wear and tear the bank expects on these properties compared to those owned by people without equity release deals.

σ is the volatility of the residential property price.

(ii)

Property value S_0

- The bank is likely to have had the property professionally valued at the outset of the contract, so it is the change in value since then that is of interest.
- Some building societies and banks publish data on property prices, and there are commercial indices.
- Depending on how granular the data is, the bank might need to assume that price changes in a particular region are identical for all properties, independent of property size and of location within the region.
- The bank might deduct from price growth a “wear and tear” rate if it thinks that equity release clients are less careful about maintaining their homes than other homeowners.

Yield q (e.g. rental cost)

There are various ways of estimating the yield on the property.

- Data may be available on market rent rates if the bank offers buy-to-let mortgages. Similar comments apply as with S_0 regarding granularity of data.
- The bank might include within q an allowance for future wear and tear if it thinks that equity release clients are less careful about maintaining their homes than other homeowners.

Volatility σ

- There are unlikely to be suitable property derivatives in the market for volatilities to be inferred from market prices ...
- ... or analysis of historic price movements is the alternative.
- Similar considerations apply to granularity of published data as for S_0 .
- Historic price indices will understate volatility if there is price smoothing hidden within them: the bank might add an arbitrary adjustment to historical volatility to allow for this.
- The use of historic volatilities always involves two assumptions: (i) that the past is a guide to the future, and (ii) that market implied volatilities (if they existed) were an accurate reflection of underlying price volatilities. In the absence of a proper hedge, these assumptions would need to be stress-tested.

[Further valid points could be made here. For example, some candidates gave more details of how to use historic volatility data to estimate implied volatilities; others interpreted the question as asking about the assumptions behind the use of the Garman-Kohlhagen model such as time-constant volatility and interest rates.]

(iii)

The bank would purchase Put options on a relevant house price index, if available, to hedge its position ...

... however, this would be retaining considerable basis risk because the distribution of the properties by region and size is unlikely to be identical to that underlying whatever index that the option is based on.

There is also unlikely to be a liquid market in options on residential property, even over the counter.

Securitisation would, until recently, have been a possible solution (hedging out all the risks in the equity release portfolio – not just property price risk) but is unlikely to work during the current “credit crunch”.

Also (minor point) in practice T is not fixed.

[Further valid points could be made here.]

QUESTION 7

Syllabus section: (g) & (i)

Core reading: Units 7, 12

(i)

(a) For the “at-the-money” forward option Y , $K = S \exp(rT)$, hence \log_e of the ratio disappears (= 0). (r is the risk-free rate)

Then the Black-Scholes European Call price C is therefore:

$$C = S \left[N\left(\frac{1}{2} \sigma \sqrt{T}\right) - N\left(-\frac{1}{2} \sigma \sqrt{T}\right) \right]$$

(b) Expanding the $N(\cdot)$ using Taylor's theorem:

$$N(x) = N(0) + xN'(0) + \frac{x^2}{2} N''(0) + o(x^3)$$

We have $N'(0) = \frac{1}{\sqrt{2\pi}}$ (given).

Putting this in the formula for C , the x^2 terms cancel and we assume we can ignore terms of x^3 and higher:

$$\Rightarrow C = S \left[N'(0) \sigma \sqrt{T} + o((\sigma \sqrt{T})^3) \right] \cong \frac{S \sigma \sqrt{T}}{\sqrt{2\pi}}$$

[This result shows that an at-the-money option increases linearly with volatility and stock price, and proportional to the square root of time.]

(c) Put-call parity states that $C - P = S - K\exp(-rT)$.

However, the RHS is zero, since the strike is “at-the-money”.

Hence the Call and Put have the same price.

(ii)

(a) Hedge ratio

From (i)(b), the option prices are roughly proportional to \sqrt{T} .

From (i)(c), the prices of the Calls and Puts for a given maturity are the same.

Hence the dealer will need to sell twice as many 3-month straddles as she buys 1-year straddles.

(b) Strategy

The dealer is almost certainly taking a view that the range of trading will narrow for a while, before picking up again. This might be because she thinks that the market will “range trade” for a while, or because there is some anticipated event in the next three months that would prevent a break-out until then.

The strategy benefits (all other parameters being equal) because the time decay is based on \sqrt{T} , so is faster for the shorter-dated options.

The dealer is unlikely to profit from lower experienced volatility feeding into implied volatilities. If implied volatilities fall, it should affect the value of the longer-dated options just as much as the shorter-dated (prices are proportional to volatility).

It depends on the balance of the actual market prices of the options, but assuming that prices are near their theoretical values as in part (i), a fall in volatility would be neutral to the portfolio.

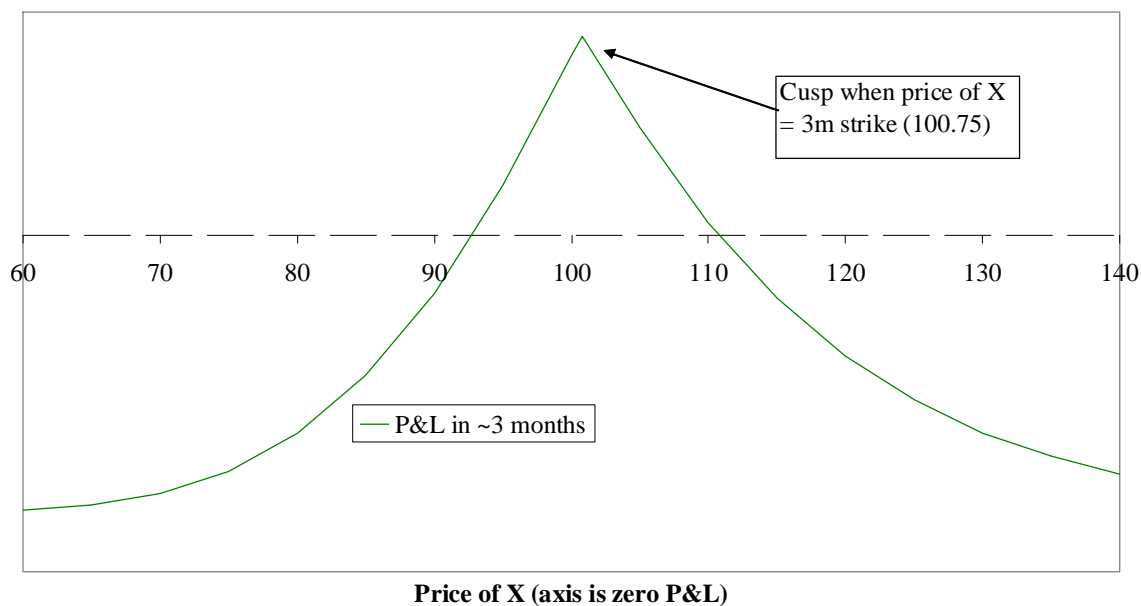
(c) Sketch

The longer option will retain its time premium whilst the short one moves towards intrinsic value. This creates “curvy” expiry P&L lines on the graph.

The shape will be cusp-like with the cusp of the P&L line at the shorter-dated strike.

An example sketch is given below using particular values for the strikes (the y-axis scale is not important):

P&L vs asset X: Sold 3 month / Bought 1 year Straddle Spread
 3m Strike = 100.75, 1yr Strike = 103.05, Risk-free Rate = 3%, Volatility = 25%



(iii)

[Note that the question says the asset's market price is not affected – hence there will not be a difference to the value of the strategy from direct movements in S .]

The interest rate fall affects the relative value of the strikes vs asset price S .

In particular, the hedge will be out of sync now, because the strikes will not be “at-the-money”.

In fact, the Puts will be “in the money” and the Calls “out of the money” ...
 ... so the impact for the options is equivalent to a large fall in the price of X .

However, the strikes of the longer-dated options will be more affected by the rate move than those of the shorter-dated ones, so given the trader is long the former and short the latter, the impact is positive for the strategy.

QUESTION 8

Syllabus section: (k)

Core reading: Unit 14

(i)

Parameters to be set are $\theta(t)$, a , σ , v ...
... and ρ (the correlation between processes z_1 and z_2).

It is intended to calibrate as far as possible to a consistent set of market variables, usually vanilla liquid instruments such as caps.

Because the equity model cascades off the yield curve model, the yield curve model needs to be calibrated first.

Yield curve model

In calibrating this model, we need to decide what represents the yield curve for risk-free zero coupon bonds. This could be derived from the dealing desk's yield curve (based on LIBOR and swaps) or a yield curve derived from government bonds.

In both cases there could be an adjustment of "x basis points" to the yield curve derived from the market.

The parameter a is the mean reversion "speed" of the Hull-White model.

Given a and σ , the function $\theta(t)$ can be expressed as a formula in terms of a , σ and the starting yield curve. Defining $\theta(t)$ in this way ensures that all zero coupon risk free bonds are priced exactly by the ESG.

So a and σ will be parameterised simultaneously, with $\theta(t)$ automatically linked to them. The idea is to choose a and σ so that the ESG replicates the implied volatilities for appropriate fixed interest options.

These fixed interest options could be caps or swaptions. In practice it is difficult to calibrate models to both simultaneously due to incompatibility of log-normal vs normal distributions reflected in the market quotes.

Whichever is chosen, there will be some compromises as the model is unlikely to be able to fit to the whole volatility matrix. It is more likely that the model will be calibrated to the cell(s) in the matrix that are of most significance to the life insurance company.

Equity model

Correlation parameter ρ will be calibrated first. Given the lack of liquid instruments from whose prices correlations can be inferred, ρ is likely to be chosen based on historical correlation data.

Some judgement needs to be exercised in setting ρ since the past is not necessarily any guide to the future.

Decisions will need to be made whether to make ρ constant or not – if constant, the calibration is easier but it may not give such a good model.

Finally, v will be set so that the ESG replicates the implied volatilities for appropriate vanilla options.

There will be some compromises as the model is unlikely to be able to fit to the whole volatility matrix: it is more likely that the model will be calibrated to the cell(s) in the matrix that are of most significance to the life insurance company.

[Further valid points could be made here.]

(ii)

The output from the ESG needs to be converted into a form suitable to be used in projecting guarantee payoffs.

For the equity model, future equity values are a direct output.

For the bond model, a yield curve needs to be calculated at each time step within each run as a function of $\theta(t)$, a , σ and r . This can be converted to bond prices and then into either bond returns or a bond price index, as required.

A model office system will calculate the guarantee payoff at the end of each run ...
... which then need to be discounted/deflated back to the present. The bond model will provide this [*since the numeraire underlying the model is the cash account*].

The value of guarantees will be calculated by performing a large number of stochastic runs (Monte Carlo simulation) ...
... then averaging the deflated payoff over all the runs.

(iii)

The life insurance company is effectively estimating the mean from a probability distribution by taking the average of n samples from the distribution (n being the number of runs).

According to the central limit theorem, the average of n samples from the distribution will be distributed as $N(M, \frac{1}{n}\Omega^2)$ where M and Ω^2 are the mean and variance of the underlying distribution.

Hence construct a confidence interval for the value of the guarantees based on a Normal distribution using μ (the sample mean) as an estimate for M and σ (the sample standard deviation) for Ω . This will be tighter for higher n .

[Making the assumption that the sample standard deviation of our n deflated payoffs is equal to the true underlying standard deviations, it can be shown that a $(100 - x)\%$ confidence interval for the value of guarantees is:

$$\mu \pm N^{-1}\left(1 - \frac{x}{200}\right) \frac{\sigma}{\sqrt{n}}$$

where $N^{-1}(\dots)$ is the inverse cumulative Normal. This exact formulation is not required so long as the method is clear.]

END OF EXAMINERS' REPORT