INSTITUTE AND FACULTY OF ACTUARIES

EXAMINERS' REPORT

September 2013 examinations

Subject ST6 – Finance and Investment Specialist Technical B

Introduction

The Examiners' Report is written by the Principal Examiner with the aim of helping candidates, both those who are sitting the examination for the first time and using past papers as a revision aid and also those who have previously failed the subject.

The Examiners are charged by Council with examining the published syllabus. The Examiners have access to the Core Reading, which is designed to interpret the syllabus, and will generally base questions around it but are not required to examine the content of Core Reading specifically or exclusively.

For numerical questions the Examiners' preferred approach to the solution is reproduced in this report; other valid approaches are given appropriate credit. For essay-style questions, particularly the open-ended questions in the later subjects, the report may contain more points than the Examiners will expect from a solution that scores full marks.

The report is written based on the legislative and regulatory context pertaining to the date that the examination was set. Candidates should take into account the possibility that circumstances may have changed if using these reports for revision.

D C Bowie
Chairman of the Board of Examiners

January 2014

General comments on Subject ST6

This subject deals with mathematical techniques for the valuation and risk management of derivatives, together with some aspects of their practical application.

Different numerical answers may be obtained to those shown in these solutions, depending on whether figures were obtained from tables or from calculators. Allowance was made for these minor differences. However, candidates may have been penalised for using excessive rounding or showing insufficient working.

Derivative theory is an interesting but very exacting subject which needs to be tackled in a precise and disciplined manner. Questions that appear unfamiliar can, with some clear thinking at the start of the exam, be found to be straightforward examples of well known techniques. Thorough preparation of the course material is essential, however. This subject also requires quite a varied approach in writing solutions. Algebraic and numerical content needs to be supplemented with well reasoned but brief arguments where the question demands. Candidates should always try to provide several distinct relevant points for the discursive questions, not lengthy paragraphs developing a single idea.

Specific comments on the September 2013 paper

The overall standard of entry for this session was comparable with the April 2013 entry, but lower than that of 2012. Slight variations will naturally occur from session to session, but it was surprising to see continuation of an historically low percentage of passes when the difficulty of the paper itself was very similar to that of recent years. There is a very long history of questions available from past papers, which provides clear templates from which students can assess the types of question to expect and how to shape their responses.

Some questions contained a substantial amount of bookwork, such as Questions 4, 6 and 8, requiring mostly repetition of the course reading material. It is disappointing to see so many candidates achieving less than half marks in such questions. As in past years, several marks were lost in these questions simply by not providing enough depth to the responses. The remainder of the paper applied the basic principles to specific situations. There were no individually difficult such questions in this paper, but again many marks were needlessly lost. The intention with this type of question is simply to illustrate via examples the derivative theory in the course, so if a topic seems unfamiliar at first sight it is worth considering which part of the course might apply. All the questions adhere to the syllabus very carefully.

As for previous sessions, the solutions below have been partly written with future candidates in mind. As well as outlining a correct answer, they also often add an explanation relating to the course material from a practical perspective. These comments (in italics) are annotated where they are additional and not required to achieve the marks set. A study of these solutions will be beneficial to candidates preparing for future ST6 papers.

1 (i) Using Ito, $d(\ln S_t) = (\mu - \frac{1}{2}\sigma^2)dt + \sigma\, dz$, i.e. $\ln S_t$ is normally distributed.

The change in $\ln S_t$ during the first two years has the probability distribution:

$N(1.5 \times 2 - ½ \times 3^2 \times 2, 3^2 \times 2) = N(-6, 18)$

The change in $\ln S_t$ during the next two years has the probability distribution:

$N(2 \times 2 - ½ \times 4^2 \times 2, 4^2 \times 2) = N(-12, 32)$

The total change is a sum of two normal distributions, hence the overall distribution is normal with parameters:

$N(-6 - 12, 18 + 32) = N(-18, 50)$

Since $\ln S_0 = \ln 10 = 2.303$, $\ln S_4 = N(\ln 10 - 18, 50) = N(-15.697, 50)$.

Hence S_4 is lognormal with mean –15.697 and variance 50.

(ii) (a) $F_t(T) = S_t e^{r(T-t)}$

(b) $\frac{\partial F}{\partial S} = e^{r(T-t)}$, $\frac{\partial^2 F}{\partial S^2} = 0$, $\frac{\partial F}{\partial t} = -rSe^{r(T-t)}$

Hence using Ito:

$$dF = [e^{r(T-t)}\mu S - rSe^{r(T-t)}]dt + e^{r(T-t)}\sigma S dz$$

and so, substituting the formula for F, $dF_t = (\mu - r)F_t dt + \sigma F_t dz$

[*Hence the distribution of F_t is lognormal.*]

Part (i) of this question relied on the fact that the logarithm of the process is normally distributed, and hence that the sequence of changes in ln S_t have additive normal parameters. Ito is used to find the drift (mean) and the rest is simple numerical calculation. Some candidates incorrectly tried to simplify the question by ignoring the different parameters in the first and second two-year periods.

Part (ii) should have been familiar bookwork, using Ito to derive a stochastic differential equation for F_t and hence find that F_t is lognormally distributed. Generally this part was well answered.

2 (i) An over-the-counter (OTC) derivatives trade between two counterparties involves credit exposure between the two parties who depend on each other for mutual delivery of the terms of the trade.

After the 2008 "credit crunch", regulators became concerned that large numbers of these OTC trades could cause multiple failures of financial firms, and hence represented a systemic risk to the market.

An independent separately capitalised and regulated central counterparty (CCP) stands in between the two sides of an OTC trade to remove the immediate credit exposure.

By charging initial and variation margin, the CCP is able to protect both itself and other counterparties from a default of one of its members (since it can close out one side of the trade for no loss in the event of a default).
This system operates in the same way as clearing of futures at a futures exchanges.

By reducing the possibility of loss on default, the CCP system creates a more orderly market and reduced risk to the financial system.

It also provides transparency to the marketplace, giving regulators a reliable source of derivative transaction statistics not previously available for OTC trades.

(ii) To perform the following calculation, a normal distribution for price changes is assumed.

From normal distribution $N(0, 1)$, abscissa for 99% probability = 2.326

Daily volatility = 24% / $\sqrt{252}$ = 24% / 16 = 1.5% approx.

Hence VaR = 2.326 × 1.5% = 3.48% on an index priced at 100%.

So ten-day VaR = 3.48% × $\sqrt{10}$ = 11%.

(iii) (a) PCH is taking a portfolio approach, so is including the offsetting effects of correlation …

… so if there are two identical but uncorrelated assets, the margin is only $\sqrt{2}$ × the margin of one of them instead of twice for ICH.

So the initial margin requirement for the bank using ICH is higher than PCH.

The decision may depend on whether the bank has readily available funds or assets for posting as margin.

But in most cases, more margin will cost the bank more.

One mitigant is that ICH is likely to be a safer institution as it has higher margin from all its members …

… but the bank may consider all CCPs as too important to fail.

(b) A weakness of VaR is that it aims to calculate losses in rare situations, yet is calculated from distributional information which is by its nature sparse around the tail.

In addition, certain VaR methods suppose distributional assumptions, such as lognormality of prices, which do not tend to apply at the tail.

Also, VaR is usually backward looking, in that it is likely to have been calibrated using historic data over an appropriate period (probably only a single year of history).

The past is not a good guide to the future, as in a stressed period there could be a sudden surge in volatility …

… which could make the margins inadequate, and hence expose the CCP to loss if a default occurs.

In addition, PCH relies on correlation effects which can change quickly, and in stressed conditions asset price falls become more correlated.

(iv) Potential problems include:

- Non-standard contracts and definitions
- Lack of price transparency for calculating variation margin
- Price data is subjective and infrequently updated, with possible basis effects (e.g. using an index as a proxy for actual properties)
- Estimating volatility for calculating initial margin
- Inability to close out / replace trades in the event of default (market liquidity poor, especially if a major market-maker was in trouble)
- Return on investment in setting up these derivatives on the CCP's systems due to
low volumes

[*Other points could be allowed here.*]

This question picked up on some recently added course reading material on the use of central counterparties (CCPs) for derivatives, linking that with the VaR methodologies two such CCPs might use. This topic is very relevant to the current regulatory environment. As with so many ST6 questions, this combined bookwork with an application to a specific less *familiar example. The key to success in this question was to apply risk-related knowledge carefully and appropriately.*

Part (i) focused on the bookwork element and allowed a range of points to be made for the two marks on offer. Part (ii) required a simple manipulation of the normal distribution which underlies the basic assumption for VaR, including the well known $\sqrt{n}$ *multiplier to convert between 1-day and n-day VaR. These parts were well answered.*

Responses to parts (iii) and (iv) were less confident. Here candidates need to apply wider knowledge of derivative theory gained from the entire course. Part (iii) in particular required quite a few distinct observations – often only one or two were given. Key points were the impact of changes in future volatility vs the past volatility on which the margins are based, and the instability in correlations in PCH's case. For part (iv), a lack of liquidity and absence of standard contracts hamper central clearing for property derivatives.

3 (i) (a) Graph showing option values:

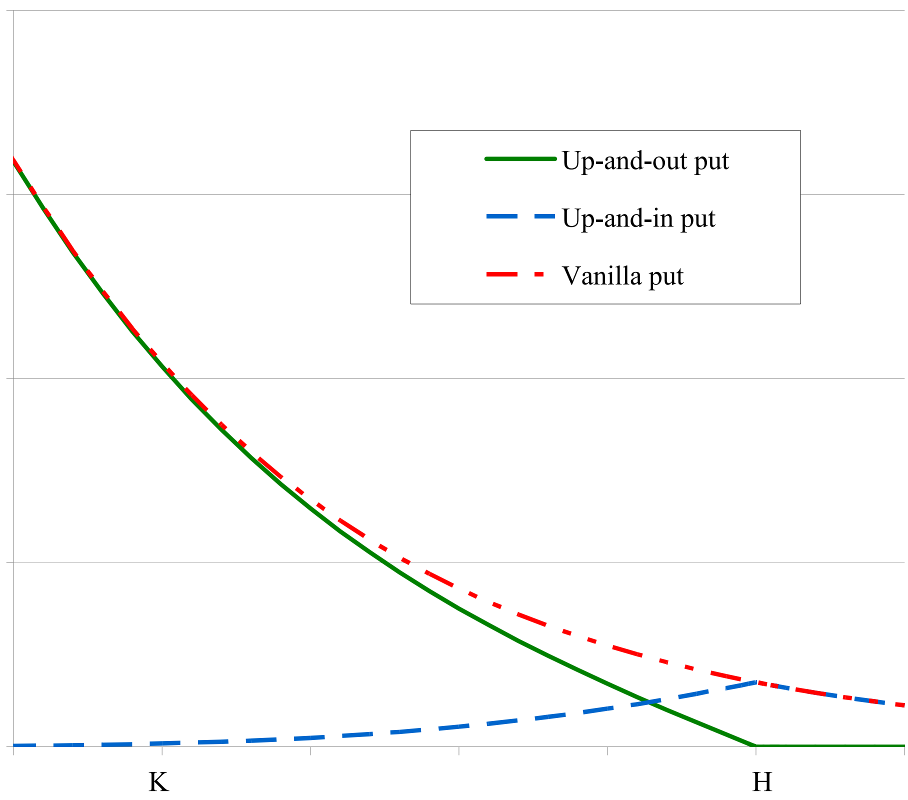

[*This chart and the one below have been created by computer for clarity. The graphs were not required to be so precisely sketched, but the key features and relationships should be clearly shown. Note: An implied asymptote to the up-and-out put line crosses the x-axis at the strike K.*]

(b) Graph showing the deltas:

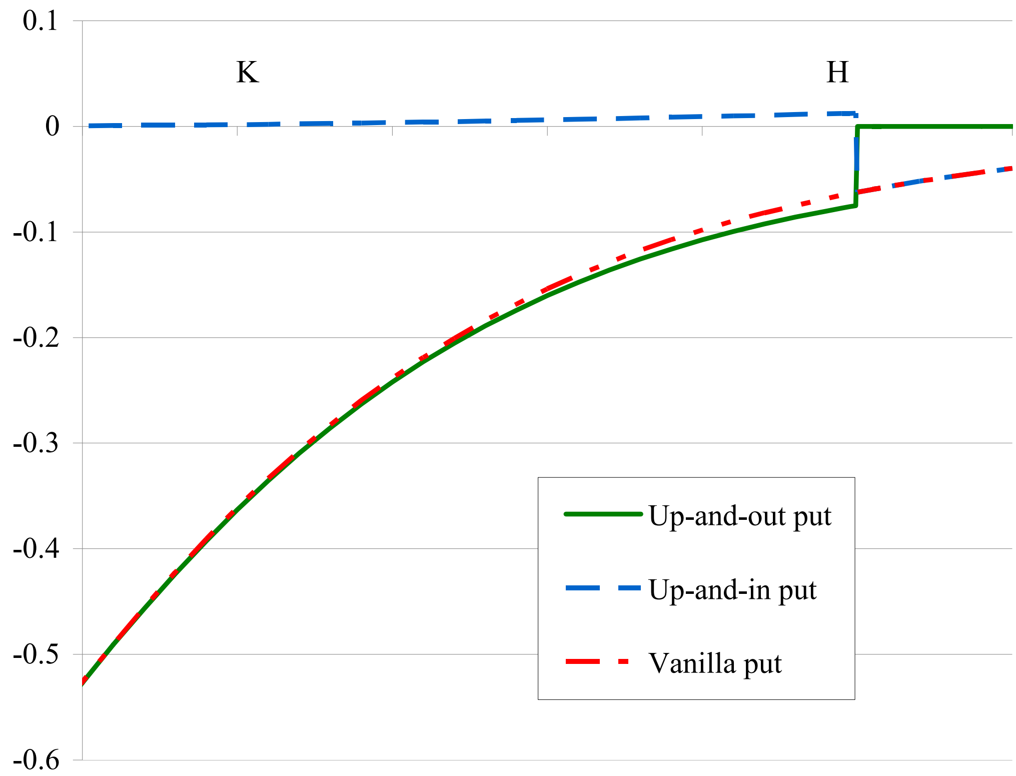

[*The delta of the up-and-in put is slightly positive up to the barrier H.*]

(ii) It is conceivable that for an x-and-out option where the barrier has not yet been breached, the vega could be negative.

This is because an increase in the volatility of the underlying asset would increase the (risk-neutral) probability of the barrier being breached and of the option becoming worthless.

This could happen

- Close to maturity
- When the share price is close to the barrier

[*Other points could be allowed here.*]

(iii) (a) <u>Formulae</u>

- If formulae exist for valuing up-and-in options, they exist for up-and-outs given that a normal option can be expressed as the sum of an up-and-in and an up-and-out.

- The formulae assume continuous monitoring of barrier breaches: because continuous monitoring is unlikely to apply to options written in the real world, this will overvalue x-and-in options and under-value x-and-out options.

- This method leads to fast pricing, which could be useful even if not 100% accurate.

(b) <u>Trees</u>

- Trees can be used, with the added complexity of having to record at each node in the tree the proportion of scenarios at that node that have already breached the barrier or alternatively using non-recombining trees.

- It makes things easier if the tree is built so that there are nodes along the barrier.

- Otherwise, the barrier lies between nodes and the nodes on either side of it represent "inner and outer barriers". One approach to pricing with trees could be to interpolate between prices that assume the inner or outer barrier to be the true barrier.

- The timestep should be set equal to the barrier monitoring frequency to minimise run times while picking up all breaches.

(c) <u>Monte Carlo</u>

- This is the most straightforward technique, suitable for path dependence.

- The timestep should be set equal to the barrier monitoring frequency, as with Trees.

- Pricing is slower and more complex to set up, but is more flexible and accurate.

[*Other points could be allowed here.*]

There is nearly always an option strategy question in an ST6 paper offering at least six marks for a sketched graph, based on the inter-relationship of option value, delta (first derivative) and gamma (second derivative) for a portfolio of options. This question focused on simple knock-in/out barrier options, whose theoretical pricing is outside the scope of the course but whose price on a chart can be constructed fairly easily given the relationship with vanilla puts and calls. Candidates who identified that "vanilla = up-and-in + up-and-out", for both value and delta, tended to produce clearer and more accurate answers.

Part (i) involved sketching option values and deltas, and part (ii) briefly touched on how increased volatility would affect barrier option pricing. A small point on drawing sketches is worth emphasising. Although the graph lines will be drawn freehand, and there is considerable leeway in the marking scheme regarding precision for these, there seems little excuse for not using a ruler for the axes!

Part (iii) asked the candidate to apply knowledge of trees, formulae and Monte Carlo methods to barrier options. Several different points were allowed that were accurately applied to the specifics of valuing barrier options as asked in the question.

4 (i) (A) LONGEVITY RISK

For longevity hedges, the main risk factor is mortality risk, which is the risk that mortality rates are lower than expected, i.e. people live longer.

Pension funds are sensitive to this risk because pensions will be more expensive if people live longer. Hence pension funds wish to buy protection.

Longevity risk hedges typically fall into the following categories of product:

A mortality swap of term T is an OTC contract that swaps a floating notional amount times the specified survivor index (based on actual mortality rates for a reference population) each period $t = 1, 2, ..., T$ for regular fixed payments, which are agreed when the contract is arranged.

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

A longevity bond of maturity T is a tradable security that has coupons linked to the specified survivor index (based on a reference population) at times t = 1, 2, ... , 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 specified survivor index at
time t.

For example, the principal repaid at t might be reduced only if the survivor index at t is above a certain threshold, fixed at outset.

Hedging is effected by issuing a strip of these bonds.

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 $S(t)$ is the specified survivor index and $K(t)$ is an agreed fixed strike amount.

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

One possible source of residual risk is the basis between the population used to calculate the derivative and the actual pensioner population being hedged.

(B) INFLATION RISK

The main inflation risk for a pension fund is that the inflation index, e.g. RPI, evolves differently from that anticipated and hence makes the pension scheme under funded …

… either with respect to the pensioners in terms of their pensions in payment (if inflation-linked) …

… or the active staff / deferred pensioners in terms of their accrued benefits.

Inflation risk is typically hedged by the following types of contract:

RPI (or CPI) Swaps

These are swaps that pay out an index-linked coupon that tracks RPI [*or CPI – many pension schemes now link to CPI*] against a floating payment on the other side of the swap.

LPI Swaps

These are the same as RPI swaps but limit the index-linking to a maximum, usually 5% in the UK to mirror pension scheme liabilities.

Often there is a floor of zero as well.

LPI Bonds

These are structured bonds created by adding an LPI swap to a principal amount.

Inflation options

Caps and floors on inflation itself could be offered so pension funds could hedge inflation directly without a specific swap structure.

These can be created by extracting the options embedded in LPI swaps.

[*Not all the above points need to be made for full marks, and other points may be valid. As ever, additional risks are involved, such as operational, liquidity and reputational, but these are not the main thrust of the question.*]

(ii) Probabilities implied from option prices and tradable instruments are risk-neutral ...

… because these are obtained from expectations in a martingale analysis that assumes discounting at the risk-free rate.

Methods that use risk-neutral probabilities are called arbitrage free …

… and are very suitable where other tradable instruments need to be priced based on these probabilities, because they will create a replicable price for such instruments that has no inherent arbitrage opportunities.

Hence these are used when the product is hedged with other products.

By contrast, probability distributions obtained from historical data are "real-world" …

… in that they are the actual observed outcomes in a risky world.

Methods that use real-world probabilities are sometimes called actuarial …

… and are very suitable when the absolute financial implications of the underlying risk need to estimated.

Hence these are used when outright risk is being taken ...

… examples being insurance contracts or extreme stress tests.

Real-world probabilities are intuitive and readily accessible …

… as there is always some form of historical data available to price a contract, or some means of empirical estimation of probabilities …

On the other hand, risk-neutral probabilities are opaque and often counter-intuitive …

… and there may not enough liquid tradable instruments to provide a risk-neutral price, then real-world techniques must be used.

[*Other points could be allowed here.*]

Part (i) of this question asked the candidate to summarise course material relating to derivatives based on longevity and inflation. The marks were reserved for derivative-based products only, including structured bonds, but there are many examples, as the model solution shows. Candidates should have had little difficulty in getting many of the marks, though disappointingly few achieved more than half. Answers needed to provide an explanation of why the pension fund would find the products useful – just one brief statement for each was all that was required.

Part (ii) asked for a discussion of the merits of risk-neutral (arbitrage-free) valuations vs actuarial (real-world) valuations. Arbitrage-free valuations must be used for consistent market pricing where tradable derivatives are involved, for example in hedging. Real-world pricing is more common in insurance contracts, where there are probably no tradable derivatives and valuations are based on historical experience. This part was not well answered, with few candidates making proper balanced comparisons.

5 (i) Value of the put using Black-Scholes is $K_P e^{-rt} N(-d_2) - SN(-d_1)$

where S = £50m (50% of £100m), K_P = £40m

$$d_1 = \frac{\ln(S/K_P) + (r + \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}, \; d_2 = d_1 - \sigma\sqrt{T}$$

$d_1 = [\ln(50\text{m} / 40\text{m}) + (0.02 + 0.25^2/2) \times 1] / (0.25 \times \sqrt{1}) = 1.09757$

$N(-d_1) = 0.13620$

$d_2 = 1.0976 - 0.25 \times \sqrt{1} = 0.84757$

$N(-d_2) = 0.19834$

So value of put = £40m × $e^{-0.02}$ × 0.1983 – £50m × 0.1362 ≈ £966,500

(ii) Delta of the put option is $-S\,N(-d_1)$ = –50m × 0.13620 = –£6,810,000

Delta of the (shorted) call is $-S\,N(d_1)$

where $d_1 = \dfrac{\ln(S/K_C) + (r + \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}$ with K_C = £62m

$= [\ln(50\text{m}/62\text{m}) + (0.02 + 0.2^2/2) \times 1] / (0.2 \times \sqrt{1}) = -0.87556$

$N(d_1) = 0.19063$

So delta of (shorted) call is –£50m × 0.1906 = –£9,532,000

So effective equity investment in fund is £50m – £6.810m – £9.532m = £33.658m

and effective equity exposure = 33.7%

[*In the above, percentage of fund could be used instead of absolute amounts.*]

(iii) (a) If equity prices increase, effective equity exposure (EEE) reduces further …

… tending towards 0% as the index increases.

This is due to the call becoming increasingly in the money …

… making the combined equity/derivative portfolio look more like a bond paying £64m at time $t = 1$.

(b) If implied volatilities increase, EEE reduces …

… due to the increase in time value of the options.

(c) If a shorter term is used, EEE increases …

… due to the reduced time value of the options.

(d) If the collar is widened, EEE increases …
… because the option values decrease (further out of the money) hence their delta offset reduces.

[*Graphical explanations could be used instead.*]

(iv) Rho of the (shorted) call is:

$-£62\text{m} \times 1 \times e^{-0.02} \times N(d_2)$

where $d_2 = d_1 - \sigma\sqrt{T} = -0.87556 - 0.2 = -1.0756$

hence $N(d_2)$ is 0.14106

and so rho is $-£62\text{m} \times 1 \times e^{-0.02} \times 0.14106 = -£8{,}572{,}700$.

Rho of the put is:

$-£40\text{m} \times 1 \times e^{-0.02} \times N(-d_2) = -£40\text{m} \times 1 \times e^{-0.02} \times 0.19834 = -£7{,}776{,}500$.

Total rho = –£8,572,500 – £7,776,500 = –£16,349,000.

The cash part of the fund is £50m, and the duration of a zero-coupon bond of maturity t
is $-t$

so effective bond duration = £16,349,000 / £50m = 0.327 years ≈ 4 months

(v) Volatility risk – vega: Changes in the value of the fund's assets as a result of changes in market implied volatilities. An increase in volatilities will increase the values of the put and the shorted call.

While these are offsetting, this is still a risk that needs to be understood, especially if there are skew effects between strikes.

Volatility risk – gamma: Non-linear changes in the value of the fund's assets as a result of market movements. A change in equity prices will alter the effective equity exposure.

Counterparty risk: the risk of the derivative counterparty defaulting on its obligations at a point in time when the fund has positive credit exposure to them.

[*"Model risk" is another possible additional risk, but not considered a main change to the risk profile in view of the vanilla nature of the options.*]

This question introduced a fund invested 50% in equities, with 50 % in cash, and considered the effect of adding a derivative hedge to both equity and interest rate exposure. The long put and short call options add interest rate exposure that was not present in the original investment strategy.

Part (i) asked for a simple Black-Scholes numerical valuation of the put, and was answered well by all. Part (ii) then looked at the net delta of the portfolio (fund + put – call). This is conceptually straightforward, as the long put and short call both reduce the equity exposure from its initial 50%. The use of the term "effective equity exposure" seemed to confuse, but the question was clearly speaking about some sort of residual delta. Part (iii) then asked for an assessment of the impact on the delta of various changes in parameters, which was well answered by those who completed part (ii).

Part (iv) was more tricky, but the essential point was simple, namely to consider the rho of the option and then express this in bond terms. Part (v) was short and relatively familiar bookwork relating to the risks in using derivatives, though it was not in fact handled well.

6 (i) (a) $\tilde{S}_t = S_t \exp(\delta t)$.

(b) Define $\tilde{Z}_t = B_t^{-1}\tilde{S}_t$.

The Cameron-Martin-Girsanov theorem tells us that there exists a measure **Q** such that under **Q** the process $\tilde{Z}_t$ is a Martingale.

Consider the discounted expected claim process $E_t = \mathrm{E}_{\mathbf{Q}}[B_T^{-1}X \mid F_t]$.

This too is a **Q**-measure Martingale, since for any $s < t$:

$$\mathrm{E}_{\mathbf{Q}}[E_t \mid F_s] = \mathrm{E}_{\mathbf{Q}}[\mathrm{E}_{\mathbf{Q}}[B_T^{-1}X \mid F_t] \mid F_s] = \mathrm{E}_{\mathbf{Q}}[B_T^{-1}X \mid F_s] = E_s .$$

Since both $\tilde{Z}_t$ and E_t are **Q**-Martingales, the Martingale Representation theorem states that there is a pre-visible process $\tilde{\varphi}_t$ such that $dE_t = \tilde{\varphi}_t d\tilde{Z}_t$.

The replication strategy then consists of holding a portfolio of $\varphi_t = e^{\delta t}\tilde{\varphi}_t$ of stock and $\psi_t = E_t - \tilde{\varphi}_t\tilde{Z}_t$ of riskless bonds.

Then the value of the portfolio V_t at time t is given by
$V_t = \tilde{\varphi}_t \tilde{S}_t + \psi_t B_t = \varphi_t S_t + \psi_t B_t$.

Then $dV_t = \varphi_t dS_t + \psi_t dB_t$, hence the portfolio is self-financing.

As the portfolio replicates the claim, the arbitrage-free condition requires that the value of the claim equals the value of the replicating strategy.

(ii) (a) Lognormality is a reasonable assumption because price changes are proportional to the level of the underlying asset price …

… and, of the random fluctuations which occur around that, a large enough sample size will tend towards making the proportionate changes (ΔP / P) normally distributed.

Another feature of lognormality is that prices cannot become negative.

Lognormality may not be precise, though, because unexpected shocks can occur – these tend to be more severe than the normal distribution allows, i.e. "fatter tails".

Also, markets can show skew in that one direction is favoured more than others – for example, the panic factor when equity indices fall means movements on the downside can be more extreme than those on the upside, yet the normal distribution is symmetrical.

(b) The actual distribution of the asset price changes can vary from the lognormal in a number of ways, as shown in the various diagrams below.

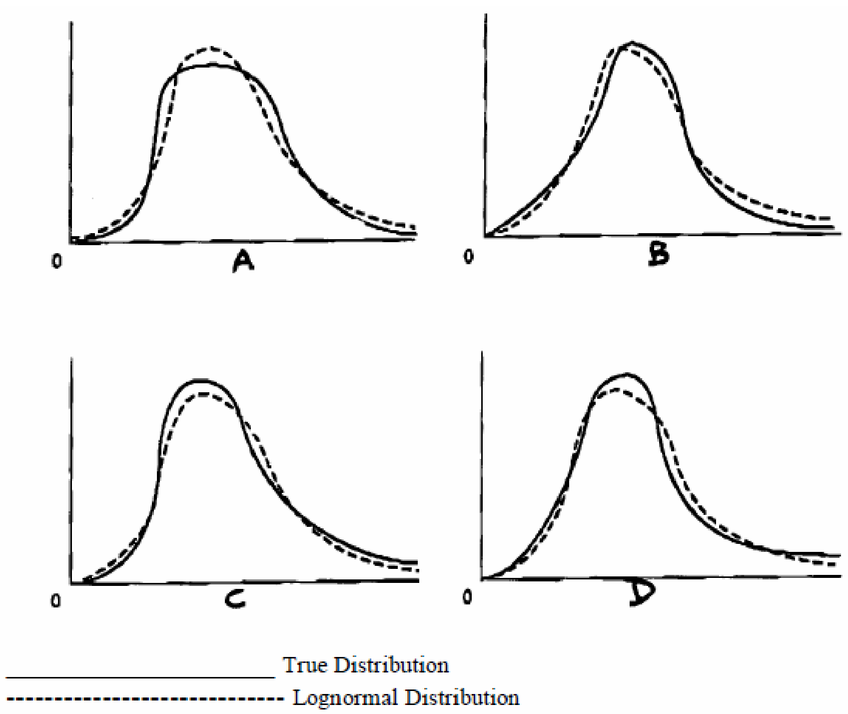

If the distribution favours higher values of the asset, this will increase the value of calls, if lower values, then puts.

The pricing bias can be seen most clearly when pricing an out-of-the-money option, whose value is very dependent on the probability distribution in the relevant tail …

… but in- and at-the-money options are also affected.

In charts A and B, the RH tails are thinner than those of the lognormal distribution, so calls will be overpriced by Black-Scholes.

In charts C and D, the RH tails are fatter than those of the lognormal distribution, so calls will be underpriced by Black-Scholes.

The thicker the RH tail the more valuable a call option will be because it is more likely that the price will rise or remain at a higher level.

Similarly for the LH tails and puts (A and C over-priced, B and D under-priced).

[*There are numerous ways of answering this part of the question correctly. However, at least two specific examples must be discussed to obtain full marks.*]

Part (i) was familiar bookwork concerning replicating strategies, this time involving the use of dividends as a slight complication. Some candidates derived the formula for dividends correctly in (i)(a) but then omitted dividends in (i)(b).

Part (ii), concerning the effects of the lognormal assumption, should also have allowed candidates to score well. This required a careful assessment of firstly the derivative aspects in (ii)(a), and then an application via examples in (ii)(b) using statistical theory. Although (ii)(a) was well covered, mostly too few examples were given in (ii)(b) for a full answer.

7 (i) Let the spot rate at end of year t be i_t.

The three-year swap rate S_3 satisfies

$$1 = S_3\,[\exp(-i_1) + \exp(-2i_2) + \exp(-3i_3)] + \exp(-3i_3)$$

$$= S_3\,[0.99392 + 0.98708 + 0.97893] + 0.97893$$

Hence $S_3 = (1 - 0.97893) / 2.95993 = 0.00712$, or 0.71%

(ii) Denote by F_t the continuously compounded one period forward rate from time t to time $t + 1$.

Then $F_t = \ln(\exp[(t+1)\, i_{t+1}] / \exp[t\, i_t]) / 1 = (t+1)\, i_{t+1} - t\, i_t$

So $F_9 = 10 \times 1.90\% - 9 \times 1.73\% = 3.43\%$

The first t forwards up to F_{t-1} (i.e. including F_0) sum to $t\, i_t$, so

50-year spot rate = $(F_0 + F_1 + ... + F_{49}) / 50$

$= [9i_9 + (F_9 + F_{10} + ... + F_{49})] / 50$ since the first t forwards sum to $t\, i_t$

$= [9 \times 1.73\% + 41 \times (F_9 + F_{49}) / 2] / 50$ using sum of an arithmetic progression

$= [15.57\% + 41 \times (3.43\% + 4\%) / 2] / 50$

$= 3.3577\%$

[*Note: the neat answer above uses the formula for the sum of n forwards. This avoids the much longer route of interpolating all the forwards after t = 9.*]

(iii) Rates are rising beyond $t = 10$, so the term for 3.5% will obviously be beyond $t = 50$.

Beyond $t = 50$, the forward rates are fixed, so we can say

$$t\, i_t = 50\, i_{50} + (t - 50) \times 4\%$$

and so $i_t = 3.5\%$ when $t \times 3.5\% = 50 \times 3.3577\% + (t - 50) \times 4\%$

i.e. $t \times 0.5\% = 200\% - 167.885\% = 32.115\%$

Hence $t = 64.2$, i.e. 64 years

(iv) The cashflows at times 1, ..., 10 can be hedged directly by buying zero coupon bonds with maturity values equal to the payments that the firm is liable to at those points in time.

This hedge will work because the bonds (being priced consistently with swaps) will also behave as if valued using the official yield curve (this part of the yield curve being consistent with swap rates).

Official yields beyond 10 years are formula driven, and dependent only on the rates up to 10-years (since future forward rates are determined by the [9, 10] one-year forward rate).

Hence to hedge a cashflow at time $t > 10$ years, the company can therefore match the sensitivity of this cashflow with that of a portfolio of zero coupon bonds up to 10 years.

This can be achieved by perturbating the rate that leads to the cashflow and comparing this with the sensitivity of the other bonds up to year 10, in the same manner as a normal interest rate swap hedge.

The above will lead to a portfolio of 1, …, 10 year zero coupon bonds as a hedge.

Given that the relationship between the *t*-year official rate and the zero coupon rates is non-linear, the optimal hedge ratio will not stay fixed over time, so the hedge will need to be dynamic.

(v) In constructing a hedge, the insurer would need to know (or make assumptions about) how the ultimate forward rate will be linked to future economic conditions …

… and decide what economic indicators would drive this linkage, and how.

Some of those indicators may have a derivative market associated with them, in which case the insurer can hedge against them.

Any indicators with no derivative market will have to remain as risks on the insurer's balance sheet.

In any case, if the dependency of the ultimate forward rate on economic conditions is only an assumption, the insurer will be bearing the risk of the assumption being incorrect.

This question focused on swap rates in a particular yield curve, with the added complication of Solvency II features such as discontinuities due to regulatory specifications, which actually made this a harder question than it might have initially seemed. Although the yield curve concepts were familiar, there were some complexities involved in deriving the forward rates efficiently, so thought and planning was required, particularly for part (ii).

Part (i) was a straightforward swap valuation, and was well answered. Part (ii) required the candidate to derive all the forward rates out to 50 years. The solution given applies a shortcut to obtain the sum of the first n forwards as needed to calculate the spot rate. Once this sum is obtained, the values in parts (ii) and (iii) are easy to obtain; without it, the numerical workload increases significantly. However, a logical but perhaps incomplete methodology attracted a good proportion of the marks. Some candidates were clearly heading down a wrong path in part (ii), but then contrived to pluck the correct answer (as given in the question) apparently from nowhere.

Part (iv) asked for a general hedging strategy, which involved using the insight that the forward rates are formulaic at and beyond 10 years, hence the hedge relationship changes there. A short part (v) then looked at the implication of varying the controlled rate on the insurer's risk profile, which focused on the insurer being able to anticipate the parameters that would govern how the rate might be set in future.

8 (i) (a) The value at time t of the forward rate spanning period T_1 to T_2 is given by:

$$f(t,T_1,T_2) = -\frac{\ln P(t,T_2) - \ln P(t,T_1)}{T_2 - T_1} \qquad (*)$$

(b) The value at time t instantaneous forward rate of maturity T is given by:

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

Alternatively, this can be defined as the limit of $f(t, T, T')$ as $T' \to T$.

(ii) Using Ito's lemma and the fact that $\frac{d}{dP}(\ln P) = \frac{1}{P}$ and $\frac{d^2}{dP^2}(\ln P) = -\frac{1}{P^2}$,

$$\begin{aligned} d(\ln P) &= [\frac{d}{dP}(\ln P)rP + \tfrac{1}{2}\frac{d^2}{dP^2}(\ln P)(vP)^2]dt + \frac{d}{dP}(\ln P)vPdz \\ &= (r - \tfrac{1}{2}v^2)dt + vdz \end{aligned}$$

so inserting this in the formula from (*) above for T_1 and T_2 gives:

$$df(t,T_1,T_2) = \tfrac{1}{2}\frac{v(t,T_2)^2 - v(t,T_1)^2}{T_2 - T_1}dt - \frac{v(t,T_2) - v(t,T_1)}{T_2 - T_1}dz$$

Putting $T_2 = T_1 + \Delta T$ and then letting $T_2 \to T_1 = T$ as the question hint suggests, the LHS of the above equation becomes the instantaneous forward rate $F(t, T)$.

Hence

$$dF(t,T) = \tfrac{1}{2}\frac{\partial\left[v(t,T)^2\right]}{\partial T}dt - \frac{\partial v(t,T)}{\partial T}dz = m(t,T)dt + s(t,T)dz$$

from the representation given in the question.

Hence $m(t,T) = \tfrac{1}{2}\frac{\partial\left[v(t,T)^2\right]}{\partial T} = v(t,T)\frac{\partial v(t,T)}{\partial T}$ and $s(t,T) = -\frac{\partial v(t,T)}{\partial T}$.

[*Equally, we could have chosen* $s(t,T) = \frac{\partial v(t,T)}{\partial T}$ *and switched to random process –z(t). This is the approach used implicitly by Hull, which gives a more intuitively correct positive standard deviation in the example in part (iii). Both methods are correct, since we are only concerned with the amplitude of the random fluctuations given z(t) is symmetric.*]

Now, integrating, $v(t,T) = \int_t^T \frac{\partial v(t,\tau)}{\partial \tau} d\tau$ since $v(t,t) = 0$

Hence $m(t,T) = s(t,T)\int_t^T s(t,\tau)d\tau$, showing the required dependency.

(iii) <u>For the HJM model in the question</u>:

From the definition in part (ii), we have $s(t,T) = -\frac{\partial v(t,T)}{\partial T} = \sigma e^{-a(T-t)}$.
Integrating:

$$v(\tau,T) = -\int_t^T \sigma e^{-a(\tau-t)} d\tau = \frac{1}{a}\left[\sigma e^{-a(\tau-t)}\right]_t^T = \frac{\sigma}{a}\left[e^{-a(T-t)} - 1\right] = -\sigma B(t,T).$$

So the SDE for P is $dP(t, T)$ = [drift term] $P\, dt - \sigma\, B(t, T)\, P\, dz$

<u>For the Hull-White model</u>:

We are told that $P(t,T) = A(t,T)e^{-rB(t,T)}$ and $dr = [...]dt + \sigma\, dz$,

so using Ito's Lemma without determining the precise nature of the drift component, the SDE for P is:

dP = [drift term] $P\, dt + \sigma \frac{\partial P}{\partial r} dz$ = [drift term] $P\, dt - \sigma\, B(t, T)\, P\, dz$

since $\frac{\partial P}{\partial r} = -B(t,T)A(t,T)e^{-rB(t,T)} = -B(t,T)P$

So for both models the volatility factor (standard deviation) of bond prices is $\sigma B(t,T)P(t,T)$ and hence the two models are intrinsically the same.

This was another bookwork based question on the HJM model of the yield curve for parts (i) and (ii), and most candidates scored well here. The derivation of the volatility parameter of the instantaneous forward rate, expressed in terms of the drift parameter, should have been familiar territory.

Part (iii) was less familiar and more tricky to obtain, but an approach comparing the volatility term of two SDEs was clearly hinted at, since stochastic processes with the same volatility component are essentially the same. Candidates should have been able to obtain some derivation here for both Hull-White and HJM, even if incomplete. As the solution shows, the drift term can be ignored in the calculations, greatly simplifying the algebra.

END OF EXAMINERS' REPORT