

INSTITUTE AND FACULTY OF ACTUARIES

EXAMINERS' REPORT

April 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

July 2013

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.

Comments on the April 2013 paper

Derivative theory is an interesting but very exacting subject which needs to be tackled in a precise and clear 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. 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.

The overall standard of entry for this session was slightly lower than last year's, with a reduced percentage of passes. Slight variations will naturally occur from session to session, but it was surprising to see this decline when the difficulty of the paper itself was very similar to that of recent papers, especially given the long history of questions available from past papers.

Some questions contained what should be familiar bookwork, such as Questions 4, 7 and 8. It was disappointing that candidates lost several marks in these questions simply by not providing enough depth to their responses. Question 4 was a case in point here: a pure bookwork question on a familiar topic where very few achieved over half marks. The remainder of the paper applied the basic principles to less familiar situations. There were some challenges, particularly Question 5, which introduced a new approach to deriving implied volatility. But candidates who pressed ahead and carefully wrote down a few relevant points still achieved a good mark in this question, even if they were unable to derive a full answer.

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.

QUESTION 1

Syllabus section: (e) & (j)

Core reading: 5, 13

(i)

If the annually compounded rate at time t is r_t , then $B(t) = (1 + r_t)^{-t}$ is the value of a zero coupon bond with term t .

The forward swap rate S_F (payable semi-annually) satisfies

$$B(5) = \frac{1}{2}[B(5.5) + B(6) + B(6.5) + B(7) + B(7.5) + B(8)] \times S_F + B(8)$$

$$\text{Hence } 0.9501 = \frac{1}{2} [0.9411 + 0.9309 + 0.9195 + 0.9079 + 0.8950 + 0.8821] \times S_F + 0.8821$$

so $S_F = (0.9501 - 0.8821) / 2.7383 = 0.0024833$ or 2.48% rounded, as required.

(ii)

Value of the payer swaption P of term $T = 5$, according to Black's formula is

$$P = LA[S_F N(d_1) - S_K N(d_2)]$$

where $L = \text{principal} = \text{£}1,000,000$

and A is the value of an annuity from $t = 5.5$ to $t = 8$ (which was calculated in part (i))

$S_F = \text{current forward swap rate} = 2.48\%$ [rounded amount]; $S_K = \text{strike swap rate} = 2\%$

$$d_1 = \frac{\ln\left(\frac{S_F}{S_K}\right) + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}; \quad d_2 = d_1 - \sigma\sqrt{T}$$

$$A = \frac{1}{2}[B(5.5) + B(6) + B(6.5) + B(7) + B(7.5) + B(8)] = 2.73826$$

$$d_1 = [\ln(2.4833\% / 2\%) + \frac{1}{2} 0.2^2 \times 5] / (0.2 \times \sqrt{5}) = 0.70461$$

$$N(d_1) = 0.75947$$

$$d_2 = 0.70758 - 0.2 \times \sqrt{5} = 0.25740$$

$$N(d_2) = 0.60156$$

$$\text{So } P = \text{£}1,000,000 \times 2.73826 \times (0.0248 \times 0.75947 - 0.02 \times 0.60156) = \text{£}18,630$$

[Note: slightly different values could be obtained due to rounding.]

(iii)

Put-call parity says:

Value of payer swaption = Value of receiver swaption + Value of swap to fixed payer

Value of swap as fixed payer is

$$\text{£}1\text{m} * [\text{B}(5) - A * 2\% - \text{B}(8)]$$

$$= \text{£}1,000,000 \times [0.95005 - 2.73826 \times 2\% - 0.88213] = \text{£}13,155$$

So value of receiver swaption would be $\text{£}18,630 - \text{£}13,155 = \text{£}5,475$

[Note: slightly different values could be obtained due to rounding.]

Alternative solution to part (iii):

An alternative Put-call parity formula, analogous to that used for equity options, is:

Value of payer swaption + Value of swap fixed payments at strike rate S_K

$$= \text{Value of receiver swaption} + \text{Value of swap fixed payments at forward rate } S_F$$

which leads to:

$$\text{value of receiver swaption} = \text{£}18,630 + \text{£}1,000,000 \times 2.73826 \times [0.02 - 0.0248] = \text{£}5,486$$

(iv)

The trader is ignoring correlation effects between the five year discount rate and the forward swap rates.

The trader does not understand that there is a difference between forward (interest/swap) rates at which he should be prepared to enter deals at zero cost and future (interest/swap) rates that he should be prepared to exchange for realised rates in a deal that is worth zero today.

In any probability measure (including the risk-neutral world) the expectation of a bond yield will be higher than the yield that can be back-solved from the expected bond price. The two rates are linked by a convexity adjustment.

This translates into the future swap rate being higher than the forward swap rate of 2.4833%.

This means that the trader is quoting too low a price for the derivative...

... so someone can arbitrage against him by buying at the quoted price and hedging (maybe by shorting the derivative to another bank that prices it properly)

This question was a straightforward yield curve application, seen many times in previous ST6 papers.

Part (i) asked the candidate to derive a given forward-starting swap rate. The working out of this rate is fairly mechanical, so ability to calculate discount rates quickly and accurately was needed. The examiners were conscious that a few candidates, who were clearly heading on the wrong track, suddenly seemed to arrive at the correct answer. For example, some candidates used continuously compounded rates even though the question explicitly stated annual compounding. Others calculated the three-year forward rate in year 5 and converted it into a semi-annual rate to give the required swap rate, which is at best an approximation. For part (ii), a swaption valuation, most candidates quoted the correct Black formula, which was a good start. The evaluation of d_1 and d_2 was generally done well; however, there was less success in the calculation of the annuity part. The Black formula can look simple, but it is always worth going through specific worked examples for caps, floors, swaptions etc. Part (iii) asked for an application of put-call parity to swaptions. Since the swaption is exercisable into a swap, the value of that swap is the difference between payer and receiver swaptions. Intermediate rounding seemed to affect the answer considerably, so a fair amount of leeway was allowed in the final answer. Very few candidates attempted part (iv), which was trying to elicit a discussion of convexity adjustment for futures vs forwards due to different settlement timings.

QUESTION 2

Syllabus section: (k)

Core reading: 14

(i)

Let r be the short rate (which can be stochastic).

The money market account is worth 1 at time 0 and earns $r\Delta t$ interest over the time interval Δt .

Hence the money market account follows the zero volatility process g

where $dg = rgdt$

i.e. $g_t = e^{\int_0^t r_s ds}$.

For any other process f which is based on the same underlying uncertainty, the Equivalent Martingale Measure tells us that the ratio f/g is a martingale in a world where the market price of risk is equal to the volatility of g , i.e. zero:

Hence
$$\frac{f_t}{g_t} = \hat{E} \left(\frac{f_T}{g_T} \right)$$

where \hat{E} is the expectation (from t to T) in the risk-neutral world.

Hence

$$\begin{aligned} f_t &= g_t \hat{E}(e^{-\int_0^T r_s ds} f_T) \\ &= \hat{E}(e^{-\int_t^T r_s ds} f_T) \\ &= \hat{E}(e^{-\bar{r}(T-t)} f_T) \end{aligned}$$

where \bar{r} is the average short rate from t to T , as required.

(ii)

Let $P(t, T)$ be the price at time t of a zero-coupon bond that pays 1 at time T .

Then from part (i), $P(t, T) = \hat{E}(e^{-\bar{r}(T-t)})$

Let $R(t, T)$ be the continuously compounded interest rate at time t for a term of $T - t$.

Then $P(t, T) = e^{-R(t, T)(T-t)}$

or, equivalently, $R(t, T) = -\frac{1}{T-t} \ln P(t, T) = -\frac{1}{T-t} \ln \hat{E}(e^{-\bar{r}(T-t)})$

so the entire term structure (R for all possible values of t and T) is determined by r .

(iii)

Let $s_t = e^{at} r_t$.

Then $ds_t = d(e^{at} r_t) = e^{at} dr_t + e^{at} ar_t dt = e^{at} (abdt + \sigma dW_t)$

Hence $s_t = s_0 + ab \int_0^t e^{as} ds + \sigma \int_0^t e^{as} dW_s$

i.e. $s_t = r_0 + b(e^{at} - 1) + \sigma \int_0^t e^{as} dW_s$

Hence $r_t = b + (r_0 - b)e^{-at} + \sigma \int_0^t e^{a(s-t)} dW_s$

It follows that r_t follows a Normal distribution with

$$\text{mean} = b + (r_0 - b)e^{-at}$$

$$\text{variance} = \sigma^2 \int_{s=0}^t e^{2a(s-t)} ds = \sigma^2 \frac{(1 - e^{-2at})}{2a}$$

This question took up a part of the syllabus covered in Hull's book relating to the use of the Equivalent Martingale Measure. It asked the candidate to show that using the money market account as numeraire allows all zero-coupon bonds to be expressed in terms of the evolution of the short rate r . Part (i), the use of the EMM, was better answered than part (ii), which generalised the result.

Part (iii) asked candidates to derive the underlying distribution implied by a particular stochastic process for r . This had some simple algebra that most found relatively straightforward.

QUESTION 3

Syllabus section: (a)–(d), (f)

Core reading: 1–4, 6

(i)

(a)

Cross hedging is the practice of hedging exposure to the price of one asset (in this case rocket fuel) using futures with a different underlying asset (in this case crude oil).

(b)

The hedge ratio is the ratio of the position taken in futures to the size of the exposure.

(c)

The minimum variance hedge ratio is the hedge ratio that minimises the variance of the hedger's net position.

(d)

Tailing the hedge is a small adjustment to make to the hedge to allow for how futures (with daily settlements and $\delta \neq 1$ rather than forward deals with $\delta = 1$) are being used as hedging instruments.

(ii)

The formula for the minimum variance hedge ratio is $h = \frac{\rho\sigma_S}{\sigma_F}$

In this case, $h = 0.75 \times 4.2 / 2.5 = 1.26$

Number of futures to buy = $1.26 * 100000 / 1000 = 126$ before tailing the hedge

The hedge can be tailed by multiplying by spot price / futures price to get

$$126 \times 101.2 / 106.3 = 119.95, \text{ or } 120 \text{ futures}$$

(iii)

The t -year futures price satisfies $F_t \leq Se^{(r+u)t}$, where r is the risk-free rate and u represents the rate of storage costs ...

... otherwise investors could make arbitrage profits by borrowing cash, buying oil and shorting oil futures.

However, in certain circumstances where physical storage capacity (or indeed borrowing capability) is limited, prices can remain distorted despite this arbitrage occurring.

It is not possible to construct a no arbitrage argument from the opposite side of the arbitrage to prove that $F_t = Se^{(r+u)t}$ because:

- there are not enough investors holding crude oil as an investment (and ready to take arbitrage profits by selling the oil and buying futures)
- there are benefits in holding the oil physically, meaning that people holding the oil physically benefit from the “convenience yield”

So futures prices can only give us a lower bound on crude oil storage costs of

$$u \geq \frac{1}{t} \ln(F_t/S) - r.$$

[The above solution could use fixed storage costs, U say, or be expressed in terms of a convenience yield, or even be stated verbally without algebra. These are all valid approaches. However, since the question said “Discuss in general terms ...”, credit was not given for references to the specific ESA example in part (ii).]

This question was devoted to hedging with futures. It began in part (i) with some definitions, which were answered well except for (i)(d), where few seemed to be aware of the “tailing” of a futures hedge to allow for daily settlement, a concept recently introduced into Hull’s book. This also affected part (ii) where the hedge was mostly calculated correctly, but very few then went on to “tail” it.

Part (iii) ask for a general discussion about the relation between storage costs and futures prices for crude oil. Key points to make were the basic equation of forward pricing including storage costs, and some insight into how arbitrage could enforce that equation (or not). Some candidates referred only to the example in part (ii), which was not asked for.

QUESTION 4

Syllabus section: (h)(iv)–(ix), (i)

Core reading: 10–12

(i)

(a)

The binomial tree method performs an expectation using the risk-neutral probability measure.

The binomial tree creates a pair of risk-neutral probabilities of up and down moves – say this measure is \mathbf{P} with up probability p , down probability $1 - p$.

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

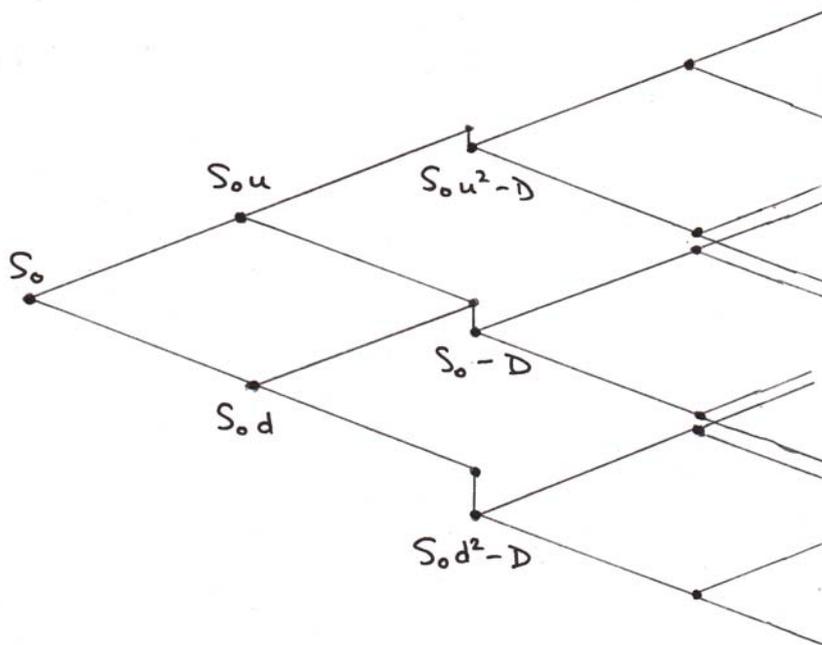
The value of a node is $\exp(-(r - q)\Delta t) \cdot [pX_{\text{up}} + (1 - p)X_{\text{down}}]$.

The binomial tree approximates this expectation in discrete time using \mathbf{P} , so must provide an approximation of the integral that creates the continuous time expectation ...

... which converges to the true value when n is large.

[The assumption here is of a continuous dividend rate q ; other approaches are acceptable.]

(b)



Example shown for starting value S_0 , up-step u and down-step d .

[The fixed dividend D has more effect on the lower values, hence the tree becomes non-recombining after $t = 2$. This should be shown clearly. As an alternative, the initial price S_0 may be adjusted to deduct the discounted dividends, i.e. $S_0^* = S_0 - De^{-2r}$; then the tree is recombining.]

(ii)

(a)

The FD method approximates the solution of the PDE by setting up a discrete rectangular grid of price changes ΔS and time steps Δt spanning all possible outcomes of the stock evolution over the time 0 to final expiry T .

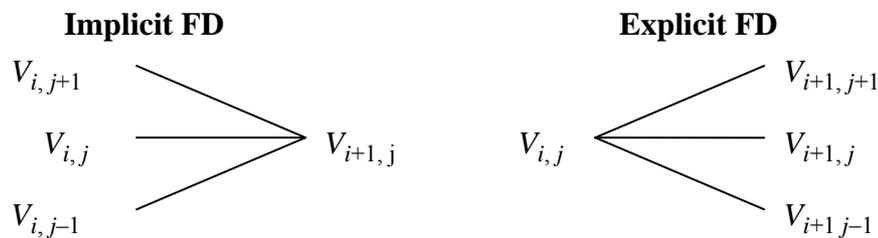
These steps must be small enough to make the approximation accurate, but not so small that the number of steps is computationally intense (leading to rounding errors).

Let the steps in time go from $i = 1$ to n , and in stock price from $j = 1$ to m , where $T = n\Delta t$, $S_{\max} = m\Delta S$ (S_{\max} = the highest price in the grid).

The approach is to approximate the differential terms $\frac{\partial V}{\partial S}$, $\frac{\partial V}{\partial t}$ etc with values from neighbouring nodes.

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 approaches can be summarised in a diagram, or alternatively (*not shown*) a grid:



The implicit method is stable and robust, and always converges ...

... but can only be solved implicitly, hence the name – there are many methods (such as the Hopscotch method) which efficiently solve the resulting matrix of relationships at each time step.

The explicit method is easy to compute, functionally the same as the trinomial tree ...

... but can introduce instabilities if the “pseudo-probabilities” created by the three branches are invalid anywhere (i.e. < 0 or > 1).

(b)

Initial and boundary conditions

In the example given, the option is a 2-year American call, so $T = n\Delta t = 2$.

Initial conditions occur at the option expiry:

$$V_{n,j} = \max(j\Delta S - K, 0) \text{ for } j = 1, 2, 3, \dots m$$

Boundary conditions for the implicit method only occur at $S = 0$ and $S = S_{\max}$:

$$V_{i,0} = 0, \quad V_{i,m} = S_{\max} - K \text{ for } i = 1 \text{ to } n$$

The American feature creates a “free” boundary – at each node, compare node value (option value) with the early exercise value. If early exercise is optimal, replace the option value with the early exercise value.

This question was basic bookwork on first the binomial tree method, in part (i), then the finite different method, in part (ii). A variety of answers were acceptable for part (i)(a): the essential point is that the binomial tree approximates to the process of obtaining the expected value of the claim (asset) X, with the approximation improving as the number of

steps increases. Part (i)(b), the tree with dividend added, was well attempted. There are two different ways of allowing for dividends at time $t = 2$: by adjusting the tree at the nodes $t = 2$, or deducting the known dividend from the stock price at $t = 0$. All candidates knew the two finite difference methods, but relatively few managed to write enough about them to garner the full six marks.

QUESTION 5

Syllabus section: (h)(iv)–(ix), (i)

Core reading: 10–12

(i)

(a)

$$P(X, t) = e^{-rt} E[\max(X - S(t), 0)]$$

$$= e^{-rt} \left[X \int_0^X f(S, t) dS - \int_0^X S(t) f(S, t) dS \right]$$

(b)

$$\frac{\partial P}{\partial X} = e^{-rt} \left[\int_0^X f(S, t) dS + Xf(X, t) - Xf(X, t) \right] = e^{-rt} \left[\int_0^X f(S, t) dS \right]$$

$$\frac{\partial^2 P}{\partial X^2} = e^{-rt} f(X, t)$$

$$\text{so } f(X, t) = e^{rt} \frac{\partial^2 P}{\partial X^2}$$

(ii)

(a)

A different pdf will need to be constructed for each time horizon t .

For each t , start by collecting recent put option trading prices $P(X, t)$ for a range of different strikes X .

The trader could then extrapolate between the prices (or fit a formula) to get an expression for $P(X, t)$, then differentiate this twice and derive the pdfs ...

... or she could convert the prices to implied volatilities, extrapolate (or fit a formula to) the implied volatilities, substitute the volatility expression into Black-Scholes, differentiate twice and derive the pdfs.

Of these two methods, the second is probably better because the volatility smile is likely to be more straightforward to extrapolate (or fit a formula to) than a set of option prices.

(b)

For an option trader, one of the most important features of a pricing tool is that it does not allow the market to arbitrage against her.

[Hence the importance of calibrating the model to traded option prices in (ii)(a).]

Related to this, her pricing tool also needs to be able to react quickly to recalibrate itself as options are traded throughout the day. The recalibration of the model (i.e. derivation of new pdfs) after every option exercise is a more complicated exercise than just calculating implied volatilities for each option trade and maintaining a record of the volatility matrix ...

... but this isn't necessarily a "deal breaker" as traders may well be able to automate the calculations to be performed almost instantly ...

... although the number of strikes available for the analysis might be limited.

Some thought would need to be put into whether the model was appropriate for pricing options whose strikes lay outside the range of strikes traded in the market. With the implied volatility matrix a trader would have a better understanding of how she was extrapolating prices outside the more liquid areas of the market.

Overall, provided practicalities can be dealt with, this approach could be used by the trader to price options ...

... but it is not immediately clear that this pricing approach would be better than simply maintaining a volatility matrix.

Once advantage is that there is no assumption of lognormality required ...

... but the model is less intuitive as it is hard to interpret in economic terms.

[Not all the above points needed to be made for full marks. Also, other comments or explanations may be valid.]

(iii)

The insurance company will probably need to calculate the market consistent values of its guarantees for financial reporting purposes.

[No specific analysis of insurance companies is needed here.]

The guarantees are likely to be valued via Monte Carlo techniques.

In valuing the guarantees via Monte Carlo, it will be easier for the insurer to use a single pdf that applies to all different strikes, and $f_i(S)$ does this.

Unlike the option trader, the insurer does have time available to derive $f_t(S)$.

In conclusion, the risk-neutral pdfs could be useful to the insurer.

Maybe the insurer would also like to feed the analysis into its solvency calculations, although this is unlikely to be effective since these rely on estimating tail events.

[Not all the above points needed to be made for full marks. Also, other comments or explanations may be valid.]

This question was the hardest in the paper. It proposed a novel (and, as it turns out, not particularly good) way to assess implied volatility in options.

Part (i) was really just a differentiation. Part (ii) looked at the trader's perspective, but seemed to baffle many candidates due to the unfamiliar situation. One possible way to tackle this part would be to ask: how does the usual implied volatility calculation work, and what are its difficulties? Then map this across to the new method and highlight the relevant point in terms of whether they generate improvements or problems.

Part (iii) looked at the new method as applied to financial reporting, where market consistency is important. There is no real right or wrong answer here, since the method is just a theoretical concept, and a wide range of views were acceptable.

QUESTION 6

Syllabus section: (g), (i)

Core reading: 7, 12

(i)

The value of the option at time T_1 is $\max(C, P)$ where C = call value, P = put value.

Using put-call parity, this can be expressed as:

$$\max(C, C + Ke^{-r(T_2-T_1)} - S) = C + \max(0, Ke^{-r(T_2-T_1)} - S)$$

which shows the chooser option can be separated into a call with strike K and maturity T_2 , and a put with strike $Ke^{-r(T_2-T_1)}$ and maturity T_1 .

[A solution that does not use algebra could also be acceptable.]

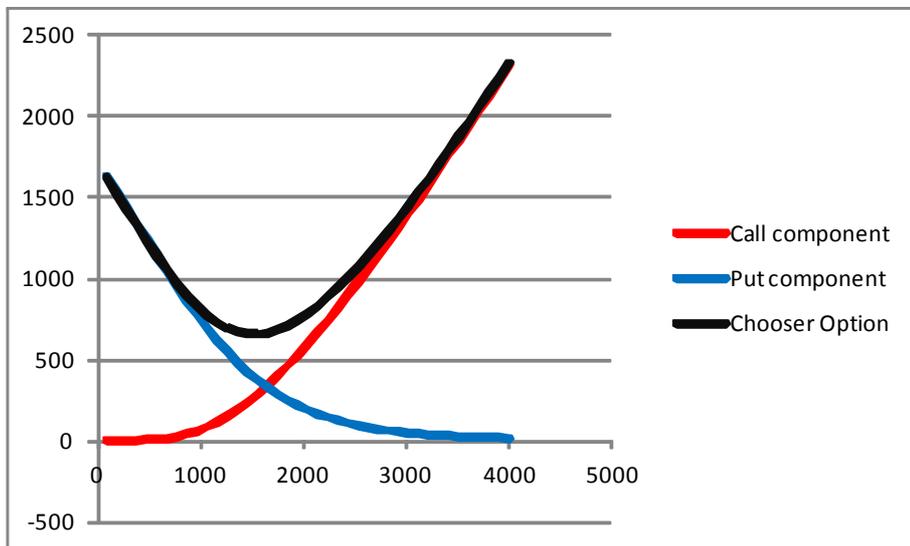
(ii)

[Note: The graphs below have been produced on a computer. It is not necessary to have numerical labels on the axes for the price and gamma graphs (except for showing the range -1 to 1 on the delta graph). Key features of the graphs are described below.]

(a)

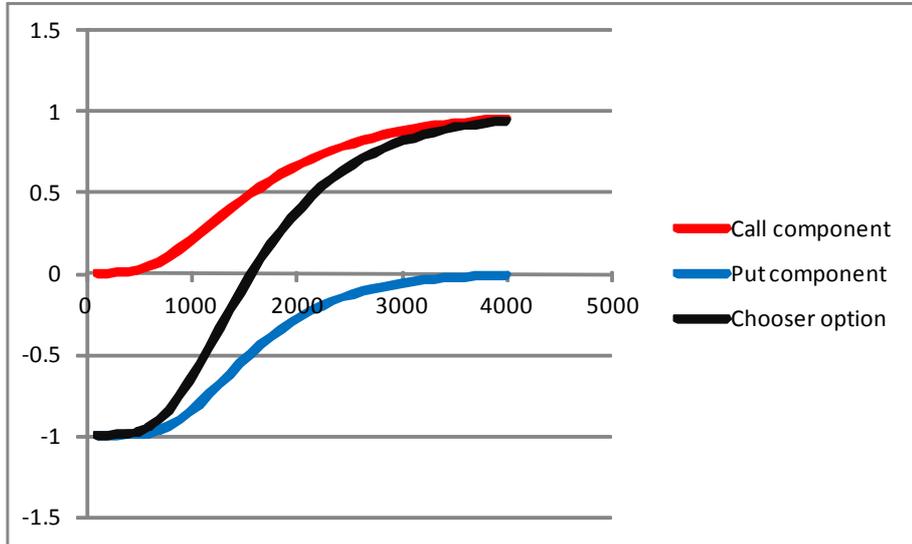
Graph of values when option first written

The put and call lines are straightforward, and asymptotically approach straight lines that both cross the x -axis at Ke^{-rT_2} . A full solution shows these straight lines and labels the point where they cross the x -axis. The chooser option is the sum of the two and tends towards the put at low S and the call at high S .



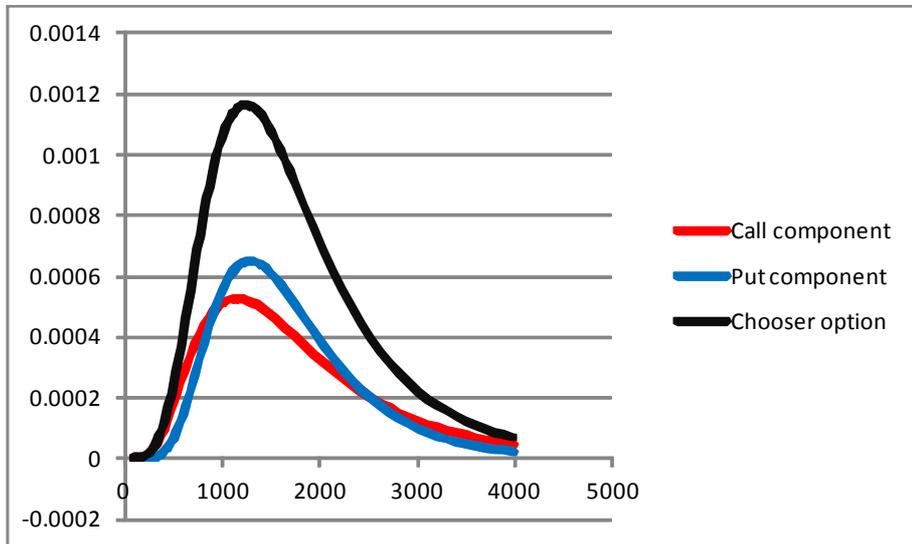
Graph of deltas when option first written

A key feature of this graph is the behaviour of delta as S tends to zero or infinity.



Graph of gammas when chooser option is first written

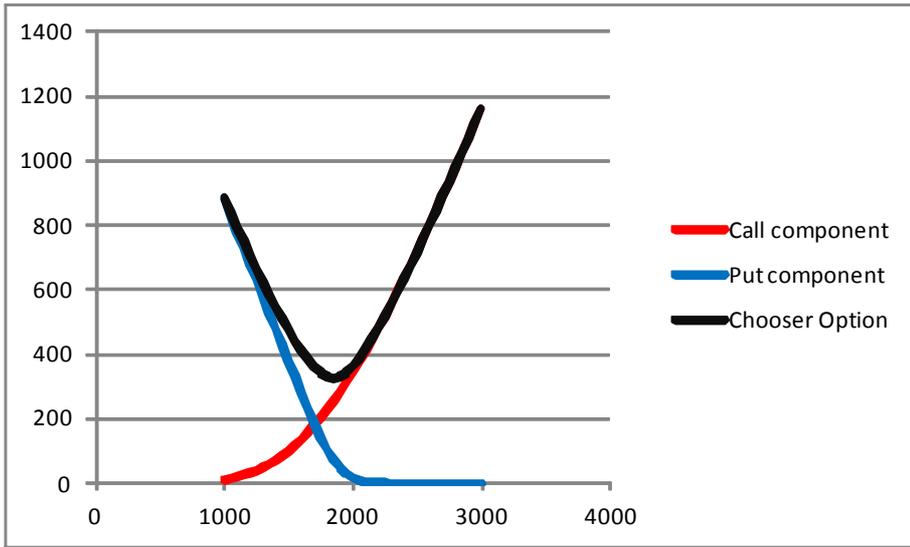
A surprise is that the peak of the put lies to the right of the peak of the call. [*Within reason, any two bell shaped curves that add together to give a third would be acceptable.*]



(b)

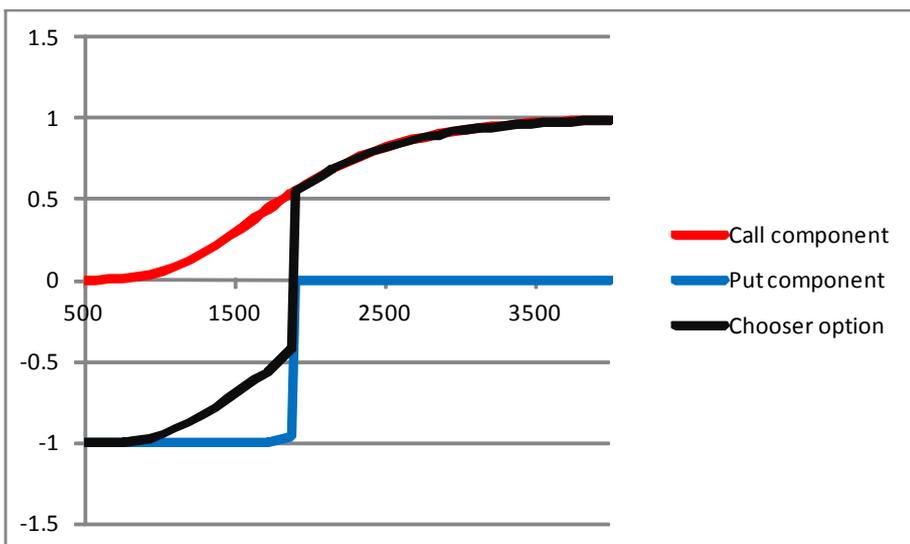
Graph of prices close to chooser date T_1

The put option line needs to look almost like two straight lines but the call option line should still be smoothly curved. As a result of this, the line for the chooser option will be closer to the call line than the put line, as seen below. Again, the asymptotes to the put and call lines will cross the x-axis at $K \exp(-rT_2)$ but with T_2 being very small, this will be very close to K .



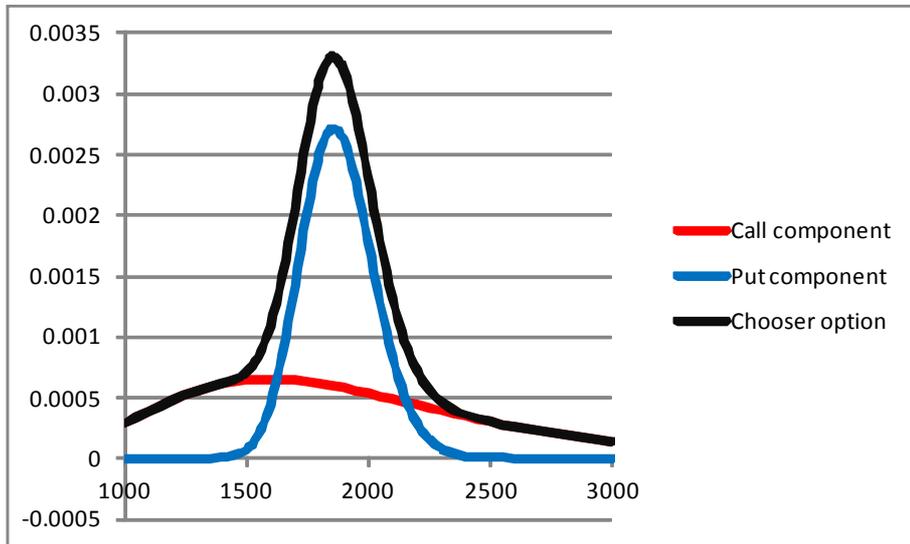
Graph of deltas close to chooser date T_1

The put option delta looks close to three straight lines whereas the call is still a more gradual move from 0 to 1. Combined effect for the chooser option is a combination of smooth bits with the sudden jump at the strike of the put component.



Graph of gammas close to chooser date T_1

The gamma of the put option (being close to expiry) is starting to get taller and narrower compared to the gamma of the call, which looks short and wide in comparison. The chooser option curve may have some kinks in it, being dominated by the put in the middle and the call at the extremes.



This was a fairly familiar graphical question based on a new type of option, the “chooser”. The question made it clear how to tackle the composition of the “chooser”, so all that was needed was to construct the graphs carefully from put and call values. The delta and gamma are just the first and second derivatives of these. This exercise was generally performed well, although there was less consensus on the gamma graphs. Two gamma humps with different peaks (strikes) actually combine to create another hump with one peak but a larger amplitude.

QUESTION 7

Syllabus section: (l) & (m)

Core reading: 15, 16

(i)

(a)

VaR is the expected loss on a risk-sensitive portfolio from an adverse market movement with a specified probability over a particular period of time.

A 99% one-day VaR means that there is a $(100 - 99)\% = 1\%$ probability of experiencing an adverse change in the portfolio over one day in excess of the calculated amount.

(b)

VaR expects that the underlying distribution is normal.

A 95% tail corresponds to an x -axis value of 1.645 and a 99% tail corresponds to an x -axis value of 2.326 ...

... so taking 95% confidence decreases the VaR by a factor of $1.645 / 2.326 = 0.71$.

VaR is proportional to the square root of the time interval ...

... so taking ten-day instead of one-day increases the VaR by a factor of $\sqrt{10} = 3.16$.

Total effect is an increase in VaR by a factor of 2.23.

[This solution provides calculations to show how VaR numbers can be linked via the normal distribution, but equally acceptable is a more general discussion of relative movements.]

(ii)

Three methods of calculating VaR

Parametric (Variance-Covariance)

This method assesses the VaR directly from the assumption of normality of price changes in all the constituent risk factors.

A linear combination of multiple normal distributions is also normal, so the loss distribution of the entire portfolio can be derived from a combination of underlying correlated values.

Correlations are estimated from past benchmark data and compiled into a large matrix.

The Parametric method has the advantage that VaR can be calculated quickly and simply provided there are not too many factors (i.e. matrix is not too large) ...

... but using only a normal distribution it cannot cope with non-linear effects ("fat tails").

It would suit a small linear portfolio of relatively few normally distributed instruments.

Historical Simulation

This method assesses the loss distribution of the portfolio based on a set of actual historical scenarios from the recent past.

It re-runs the portfolio valuation for each day and creates a distribution of scenario outcomes.

VaR is then simply the 99% percentile of these (or whatever confidence interval).

The Hist Sim method is good for non-linear instruments with complex interactions ...
... but the revaluation can be complicated and new instruments have no past history.
It would suit a reasonably large portfolio of mixed instruments.

Monte Carlo Simulation

This method involves modelling future price returns of the portfolio directly ...
... then running many hypothetical trials to obtain a distribution of portfolio losses.
As with Hist Sim, VaR is the relevant (e.g. 99%) percentile of the distribution.

This method is the most complex to apply, as future price movements and correlations have to be modelled ...

... but less tractable distributions than the normal can be used, which enables better modelling of “fat tails”.

This method is well suited to portfolios of complex options using the simulation-based models that are already set up for pricing the options.

(iii)

(a)

The main problem with VaR is that it is backward looking.

Past history does not predict future ...

... and low volatility has not helped anticipate the emerging currency risk.

VaR is supposed to model the tail of the distribution ...

... but most of the values are not in the tail, hence allowing for fat tails is difficult because by their nature they occur less frequently ...

... and there is no information in the VaR statistic to indicate the expected severity of losses outside the confidence interval.

VaR is not always well tailored to a portfolio if the risk factors used are generic ...

... so that the price movements that drive P&L are badly matched.

It is not clear how VaR and credit default can be linked, which could affect the specific situation in the question.

(b)

Stress tests are the most useful additional risk measures ...

... as they give a forward-looking component and do not rely on past volatilities and correlations.

The manager should test the portfolio on both historical and hypothetical scenarios.

Other possibilities are to use implied rather than historic volatilities in the calculation ...

... and to use more tailored risk factors to improve fit of VaR to portfolio.

A better estimate of tail risk could be attempted, such as using Extreme Value Theory.

This question covered value-at-risk, and essentially replicated the bookwork from the course reading in parts (i), (ii) and (iii)(a). Part (i)(b) could be answered in numbers or in words, but the numerical conversion between different confidence levels (using the normal distribution) and time periods (using \sqrt{t}) has been highlighted in the solution as it has important uses.

Part (iii)(b) was a simple request for other (and better) techniques for assessing tail risk, such as stress testing and extreme value theory.

QUESTION 8

Syllabus section: (h)

Core reading: Units 8 & 9

(i)

The Binomial Representation Theorem (BRT) states that, if M is a \mathbf{Q} -Martingale and N any other \mathbf{Q} -Martingale, then there exists a previsible process ϕ_i such that:

$$N_i = N_0 + \sum_{k=1}^i \phi_k (M_k - M_{k-1})$$

or, equivalently, $\Delta N_i = \phi_i \Delta M_i$.

A process X is a \mathbf{Q} -Martingale if:

$$E_{\mathbf{Q}}(X_j | F_i) = X_i \text{ for } j > i$$

where $E_{\mathbf{Q}}$ is expectation under probability measure \mathbf{Q} , and F_i is the history of the process up to time i .

A process ϕ_i , $0 \leq i \leq T$ is previsible if ϕ_i depends only on the filtration (history) F_{i-1} , i.e. up to the previous time step.

(ii)

(a)

Set up the table of possible values for S :

$i = 0$	$i = 1$	$i = 2$	$i = 3$
			4.32
		3.6	
	3		2.52
2.5		2.1	
	1.75		1.47
		1.225	
			0.8575

A discrete random walk is a martingale under \mathbf{Q} , the risk-neutral measure.

Hence $2.5 = 3q + 1.75(1 - q)$, whereby $q = 0.6$.

The fact that the table has a common ratio means that this probability is constant.

(b)

Now filling in the table of values for $X = \ln(S)$ at $i = 3$ for the payoff, then at each previous node back to $i = 0$ take expectations under \mathbf{Q} .

$i = 0$	$i = 1$	$i = 2$	$i = 3$
			1.463255
		1.247657	
	1.032058		0.924259
0.826298		0.708660	
	0.517659		0.385262
		0.231157	
			0

Method: We need to verify that BRT gives us a relationship between the two tables such that, at any node, the change in the second table if the upper path is taken is in the same proportion to the stock tree as if the lower path were taken.

For $i = 1$, we take just the first up step (U) and the first down step (D):

$$U: (1.032058 - 0.826298) / (3 - 2.5) = 0.4115$$

$$D: (0.517659 - 0.826298) / (1.75 - 2.5) = 0.4115$$

For $i = 2$, we start by looking at the first up step and where this goes next:

$$\text{UU: } (1.247657 - 1.032058) / (3.6 - 3) = 0.3593$$

$$\text{UD: } (0.708660 - 1.032058) / (2.1 - 3) = 0.3593$$

then we look at the first down step and where this goes next:

$$\text{DU: } (0.708660 - 0.517659) / (2.1 - 1.75) = 0.5457$$

$$\text{DD: } (0.231157 - 0.517659) / (1.225 - 1.75) = 0.5457$$

In each case, the ratio of the steps up and down is constant, showing that the process φ_i is previsible (has the same value whether the stock goes up or down).

(iii)

(a)

The expectation $Y_i = E_{\mathbf{Q}}(X | F_i)$ of any process is always a martingale, by the Tower law, hence can always be used with the BRT.

(b)

The Martingale Representation Theorem states that if M is a \mathbf{Q} -Martingale with non-zero volatility, and N any other \mathbf{Q} -Martingale, then there exists a (bounded) previsible process φ such that

$$N_t = N_0 + \int_0^t \phi_s dM_s \quad \text{or} \quad dN_t = \phi_t dM_t$$

i.e. replacing differences in BRT by derivatives.

Further, φ is unique.

(c)

We know that W_t , a \mathbf{Q} -Brownian motion, is a martingale.

So any other martingale X_t can be expressed as:

$$X_t = X_0 + \int_0^t \phi_s dW_s \quad \text{or} \quad dX_t = \phi_t dW_t$$

and hence X is a process with zero drift.

The converse is harder to prove, i.e. that a process with zero drift is a martingale, but (subject to a technical condition) does apply due to the uniqueness of φ .

Hence the absence of drift is used as a confirmation of a martingale.

This question asked for a definition of the Binomial Representation Theorem in part (i), then in part (ii) applied it to a specific binomial tree to show numerically how it works. The calculations were not at all hard, but the variety of responses showed that it can be easy to learn an expression in a formula without considering how it works in practice. Also, a surprising number of candidates only looked at trees to $i = 2$ as opposed to $i = 3$. Part (iii) then asked for the extension to continuous time. This was pure bookwork and was generally well answered.

END OF EXAMINERS' REPORT