INSTITUTE AND FACULTY OF ACTUARIES

EXAMINERS' REPORT

September 2010 examinations

Subject ST6 — Finance and Investment Specialist Technical B

Introduction

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

T J Birse
Chairman of the Board of Examiners

January 2010

Comments

Question 1

This question consisted of some bookwork on Ito's Formula, together with two simple applications requiring partial differentiation similar to that learned in CT8. Most candidates were able to state the bookwork in part (i) and successfully apply it to the later two parts. However, it was disappointing to see that several had difficulty coping with the partial derivatives when both t and W_t were involved.

Question 2

This question focused on fixed income theory, which should be familiar from CT1, but it was one of three questions in the paper that was generally not well answered. Confusion arose around semi-annual compounding and the definition of a floating rate note (FRN).
For part (ii), the question specifically asked about market risk although several candidates strayed into other types of risk. This question invited a genuine use of principles of market risk beyond the bookwork, and it was pleasing to see quite a few candidates making pertinent comments. Such applications may seem unexpected, but usually there is a simple concept behind them and a commonsense response will obtain good marks.

Question 3

This question was familiar bookwork on a core part of the syllabus, which most candidates recognised and responded to successfully. This type of question gives a chance to gain a high percentage of the available marks – the guideline in the paper showed that roughly a tenth of the total marks for the exam were allotted to this one question, so it is important to be organised in response. When more than a handful of marks are allocated to bookwork, candidates should write out the relevant material relatively fully.

Question 4

This question on an alternative "real world" valuation of a range option was not well answered. The Poisson distribution seemed unfamiliar to many, despite being encountered in CT6 and many insurance models. The valuation of the option in part (ii) was based on expectations using the Poisson probability distribution.

Part (iii) asked about the management aspects of this derivative, a common type of question. Topics such as modelling assumptions, data, liquidity, controls, credit risk etc. are typical of those encountered in many real life situations. Candidates would do well to prepare a check-list of such topics from previous exam solutions, then discuss the relevant ones in their response.

Question 5

Graphical questions such as this are a regular feature of ST6. Generally, expiry patterns were well known, but "on the day of the transaction" lines appeared not be so well understood.

In part (iii), few candidates were able to discuss points in detail beyond the consequences to P&L of upward/downward movements in the FTSE. This limited the analysis to a discussion of the effects of delta. To achieve a better response, candidates should also consider the other sensitivities – at least gamma and vega, and possibly rho or theta.

Question 6

This question brought together a couple of topics from the risk management section of the syllabus.

Part (i) asked for definitions of market and credit risk, and was well answered. In part (ii), most candidates showed they were aware of the impact on these risks of the hedge proposed. However, fewer could convincingly explain why the credit risk profile resembled a straddle – many simply drew the payoff of a straddle as if that explained everything. Part (iii) on collateralised debt obligations was generally tackled well – the topic is covered in some detail in the course reading material.

Question 7

This question dealt with an economic scenario generator (ESG), used by insurance companies to model contracts with derivative characteristics, often with dependencies on yield curves. In this context, most ESGs are based on an implementation of a reasonably simple yield curve model using Monte Carlo (MC) simulation. Hence this question tested both the practical use of the single factor Hull-White model (HW-1) and aspects of MC techniques.

Very few candidates made a serious attempt at the question. Since there were nine questions in all, it is possible that several saw this as a question they would only tackle if they had sufficient time. This is understandable, but nevertheless some easy marks were foregone, since if the question had been framed instead as follows: (a) describe HW-1 model, and (b) outline how MC simulations can be used to value a derivative, candidates would have presented the answer much more readily.

Part (i) was really asking for a description of how the HW-1 model can be solved. In parts (ii) and (iii), covering the MC implementation, a few common mistakes occurred. Several candidates forgot to sum over the scenarios – the MC method produces a set of sample outcomes that form a distribution like any statistical sample. Also, there appeared to be confusion between the definition of a swap and a swaption – the latter being an option with a swap as the underlying instrument.

Part (iv) asked about confidence intervals in statistics, harking back to CT6, and part (v) was looking for reference to the difficulty of replicating the full range of possible swaptions in such a simple yield curve model as HW-1. For part (v), many instead gave generic comments about the strengths and weaknesses of interest rate models.

Question 8

This was quite a fun question, applying derivative theory to a bakery hedging its bread production. It was well answered by most.

In part (i), almost all candidates knew enough of the differences between futures and forwards to score good marks. Similarly, many candidates did well in the bookwork in part (iii), though marks were lost by not clearly defining notation or missing out some key steps in the derivation. In part (ii), however, although candidates seemed aware of basis risk, they tended to give examples of how it might arise rather than giving a clear definition.

Parts (iv) and (v) were generally well answered, with many candidates aware of the key concepts. There were quite a large number of marks allocated for these two parts, for which the responses did not always provide the necessary level of detail. Generally, for questions of this type it is a good rule to aim to provide twice as many <u>distinct</u> points as there are marks.

Question 9

This question looked at implementing a binomial model to value an American Put option. It had some similarities with question 7, but on the whole was better answered.

In part (i), some candidates listed general advantages of a binomial tree rather than why it was particularly appropriate for the particular derivative. For part (ii), the model implementation steps, many candidates scored well but in quite a few other cases answers were light on detail. It is a good idea to map out the steps of an implementation on a piece of side paper, then write it up in full when all important aspects have been considered.

In part (iii), only a minority were aware of the control variate technique. This has an important application in derivative valuation – basically, it uses an analytical solution of a known simpler derivative to improve the accuracy of the numerical valuation of a more complex but similar derivative.

For part (iv), estimating the implied volatility as an input to the calculation, perhaps looked unfamiliar, as many candidates seemed to misunderstand what was intended by the two approaches. Manager A wants to use market implied volatility to check the price directly, whereas Manager B wants to assess implied volatility against a recent time series (historic volatility). The question was then looking for a discussion of implied and historic volatilities, which should have been familiar territory.

Syllabus section: (h)(i)-(iii)
Core reading: 8, 9

1 (i) (a) Ito's formula (or lemma) forms the basic extension of differential calculus to variables which are stochastic in nature.

It is used to derive stochastic differential equations (SDEs) for valuing derivatives whose payoffs depend on the evolution of a stochastic process.

Using suitable boundary conditions, these SDEs can then be solved analytically or numerically to give derivative prices and sensitivities.

(b) A martingale X_t is a stochastic process under a probability measure **P** for which all expected future values of X, conditional on a known history up to time t, are equal to its current value at time t, i.e.
$E_{\mathbf{P}}[X_u | F_t] = X_t$ for all $u > t$.

[*Another way of defining a martingale is as a driftless stochastic process, as below.*] [*The 1 mark above can be given for doing this.*]

The martingale must also satisfy the boundedness condition
$E_{\mathbf{P}}[|X_t|] < \infty$ for all t.

(ii) $$\frac{\partial Y_t}{\partial t} = -1, \quad \frac{\partial Y_t}{\partial W_t} = 2W_t, \quad \frac{\partial^2 Y_t}{\partial W_t^2} = 2$$

Using Ito's Lemma gives:

$$\begin{aligned} dY_t &= [-1 + \tfrac{1}{2}.2]dt + 2W_t dW_t \\ &= 2W_t dW_t \end{aligned}$$

which is driftless (no dt term), hence a martingale.

[*Another way to prove the martingale property is to note that, using the variance and mean of Brownian motion, the expectation of $W_t^2 = t$, hence process Y_t has no drift.*]

(iii) $$\frac{\partial Z_t}{\partial t} = -6W_t^2 + 2at, \quad \frac{\partial Z_t}{\partial W_t} = 4W_t^3 - 12W_t t, \quad \frac{\partial^2 Z_t}{\partial W_t^2} = 12W_t^2 - 12t$$

Using Ito's Lemma gives:

$$dZ_t = [-6W_t^2 + 2at + \tfrac{1}{2}(12W_t^2 - 12t)]dt + [4W_t^3 - 12W_t t]dW_t$$
$$= (2at - 6t)dt + (4W_t^3 - 12W_t t)dW_t$$

which is driftless if and only if $a = 3$.

Syllabus section: (e) + (j)
Core reading: 5, 13

2 (i) (a) The general valuation formula for a constant margin 5-year semi-annual FRN is:

$$\text{Value} = L_{\frac{1}{2}}v^1 + L_1 v^2 + L_{\frac{3}{2}}v^3 + \ldots + L_5 v^{10} + \frac{m}{2}\sum_{i=1}^{10} v^i + 100v^{10}$$

where L_i is the LIBOR rate for period i to $i + 1$ and $v = \left(1 + \frac{y\%}{2}\right)^{-1}$ is the semi-annual discount factor.

We are given that the L_i are constant for all maturities, and equal to $\frac{r}{2}$, so:

$$\text{Value of FRN} = \frac{(r+m)}{2}\sum_{i=1}^{10} v^i + 100v^{10} = \frac{(r+m)}{2}\left(\frac{1-v^{10}}{y\%/2}\right) + 100v^{10}$$

[*In simplifying the above equation, candidates may choose to identify the term in brackets as a semi-annual annuity – this is a valid alternative if carefully defined.*]

(b) If $m = 1$ and $r = 5$, then coupon payments are 3% each half-year.

$y = 7\%$ and so $v^{10} = (1.035)^{-10} = 0.708919$

Hence value of FRN = 3.(1 – 0.708919)/0.035 + 70.8919 = 95.842

(c) To obtain the 5 year FRN's duration, we multiply each cashflow at time t by t and sum, then divide by the value of the FRN:

$$\text{Duration} = \left[\frac{(r+m)}{2}\sum_{i=1}^{10}\frac{i}{2}v^i + 500v^{10}\right] / \text{Value of FRN}$$

[Alternatively, the values of r and m from part (b) can be substituted.]

(d) The majority of the contribution to the duration value will come from the final payment, which has duration 5 × 103 × 0.708919 / 95.842 = 3.81 years.

The coupons will contribute further, suggesting the duration lies between 4 and 4.5 years. [*In fact it is 4.38 years.*]

[*Other similar explanations could be provided.*]

(ii) At coupon dates, floating LIBOR payments show zero market risk because they will always be set at the prevailing level of rates, assumed to be risk-free.

[*Between coupon dates a small residual market risk is created.*]

However, the margins on the FRN are like fixed rate swap payments, so these will show up as having interest rate risk.

If the margins are modest, as they usually are, this will not be a large risk compared with an outright fixed rate bond position ...

... but it gives real market risk – for example, if rates fall generally, then this set of fixed payments will be worth more (and similarly less if rates rise).

Syllabus section: (h)(i)-(iii)
Core reading:8, 9

3 (i) <u>The Cameron-Martin-Girsanov Theorem (CMG)</u>

If W_t is a **P**-Brownian motion and γ_t is an F-previsible process (i.e. a variable known at time t based on past history filtration F)

satisfying the boundedness condition $E_{\mathbf{P}}\left[\exp\left(\frac{1}{2}\int_0^T \gamma_t^2 dt\right)\right] < \infty$,

then there exists a measure **Q** such that:

- **Q** is equivalent to **P**

[*equivalence means they map to the same set of events, i.e. events are impossible under* **Q** *if and only if they are impossible under* **P**]

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

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

[$\frac{d\mathbf{Q}}{d\mathbf{P}}$ *is the* <u>*Radon-Nikodym derivative*</u>*, which can be defined as the ratio that the joint probability density function of W under measure* **Q** *bears to the probability function of W under measure* **P**.]

(ii) <u>Martingale Representation Theorem</u>

The definition of a martingale is given in Q1.

If M_t is a **Q**-martingale process ...
... whose volatility σ_t satisfies the additional condition that it is (with probability one) always non-zero.

Then if N_t is any other **Q**-martingale,
there exists an F-previsible process ϕ
such that N_t can be written as $N_t = N_0 + \int_0^t \varphi_s dM_s$.
[*Alternatively, this condition can be written as* $dN_t = \phi_t dM_t$. *We also need boundedness, i.e.* $\int_0^T \varphi^2\sigma^2 dt < \infty$ *with probability one*.]

(iii) <u>Significance of Cameron Martin Girsanov (CMG) theorem and Martingale Representation Theorem (MRT) for valuing derivatives</u>

Derivatives can be valued by finding a portfolio of cash and stock which replicates the derivative. The value of the derivative must be equal to the value of the replicating portfolio otherwise arbitrage opportunities would exist.

The approach for finding the replicating portfolio involves finding a probability measure **Q** that makes the underlying discounted stock price process $Z_t = B_t^{-1} S_t$ a martingale, where S_t is the stock price and B_t is the risk-free zero coupon bond price at time t. This means that the expectation of all future values of Z_t is equal to its value at time 0, i.e. S_0.

Converting a stochastic process into a martingale involves finding measure **Q** under which it is driftless. The CMG theorem confirms the existence of such a measure. It gives us a powerful tool for controlling the drift of most stochastic processes that are encountered in practice.

If the process Z above is a martingale, then the process $E_t = \mathrm{E}_{\mathbf{Q}}[B_T^{-1} X \mid F_t]$, which represents the expected value of a discounted claim $B_T^{-1} X$ on S at time T (conditional on the history up to time t), is also a martingale.

The MRT then leads us to the construction of the replicating portfolio, i.e. the appropriate volumes ϕ_t of stock to hold. Since E and Z are both martingales, the MRT tells us that there is a previsible process ϕ_t such that $dE_t = \phi_t dZ_t$ under measure **Q**.

This is important because if we are holding a volume ϕ_t of stock, then changes in the value of our stock and cash portfolio will match changes in the derivative's expected value, i.e. it is self-financing.

To complete the replication, the volume of cash needed is then $\psi_t = E_t - \phi_t Z_t$. The portfolio consisting of ϕ_t units of stock and ψ_t units of cash will always have value equal to the value of the derivative since its current value is equal to the value of the derivative and changes in the value of the components of the portfolio exactly match changes in the value of the derivative.

Hence we obtain the risk-neutral pricing formula for the value V_t of the derivative: $V_t = B_t \mathrm{E}_{\mathbf{Q}}[B_T^{-1} X \mid F_t] = \exp(-r(T-t))\mathrm{E}_{\mathbf{Q}}[X \mid F_t]$.

Syllabus section: (a) – (d), (f), (m)
Core reading: 1 – 4, 6, 16

4 (i) The Poisson distribution can be used to model the number of events occurring over a given time interval, if these events take place at a known average rate λ independently of the last event occurring.

If there are n identical observations in the time interval, each with probability p of an event occurring during that observation, λ is given by $n\,p$ ($n = 60$ in this case).

[*The Poisson distribution is the limiting case of the Binomial distribution when n is large and p small, which is likely to be the case here. It actually works best if n is at least 20 and p is less than or equal to 0.05.*]

The value of p to be used can be estimated from the five years of historical data.

(ii) (a) Expected number of occurrences in three months
$\lambda = 60\,p = 60 \times 0.05 = 3$

(b) We have that $p(k) = \dfrac{\lambda^k e^{-\lambda}}{k!}$ $\quad k = 0, 1, 2 \ldots$

The "strike" number of occurrences is 3, so the value of the option is the sum of the cases $k \geq 3$.

Hence value of option = $X(1 - e^{-\lambda} - \lambda e^{-\lambda} - \dfrac{\lambda^2}{2} e^{-\lambda}) = X(1 - 8.5e^{-3}) =$ 0.577 X.

(c) The valuation is <u>not</u> a "risk neutral price" because:

- The probability estimate is not derived from observable market prices

- There is no mechanism for hedging the outcome perfectly (even an approximate hedge would be expensive due to transaction costs)

(iii) (a) There is unlikely to be a transparent mechanism for determining price from the market independently ...
... hence a "directors' valuation" needed ...
... taking into account historical experience and a justification for any future variation from history.

Theoretically, should use a risk-adjusted discount rate but unlikely to make a big difference.

(b) Assurance would be needed that the managers of the fund have acted within their proper authority in buying the option.

Some assurance around the validity of the assumptions may also be needed; this could be partially mitigated by robust stress testing.

Assuming the purchase is permitted by fund mandate, the instrument is illiquid and lacks transparency, so is very hard to include properly in risk reports.

Is there a clear rationale for buying the option? It seems it would have been better to have purchased straddles which at least have a clear market valuation.

Supervision of the trade post execution may be weak because it falls outside usual governance, and this could give rise to reputational risk. Stronger monitoring procedures will be required.

The option is hard to value and probably illiquid, so this may add to operational risk.

If the trade is not collateralised, there will be counterparty credit risk.

[*Other valid points could be made.*]

Syllabus section: (g) + (i)
Core reading: 7, 12

5 (i) P&L chart

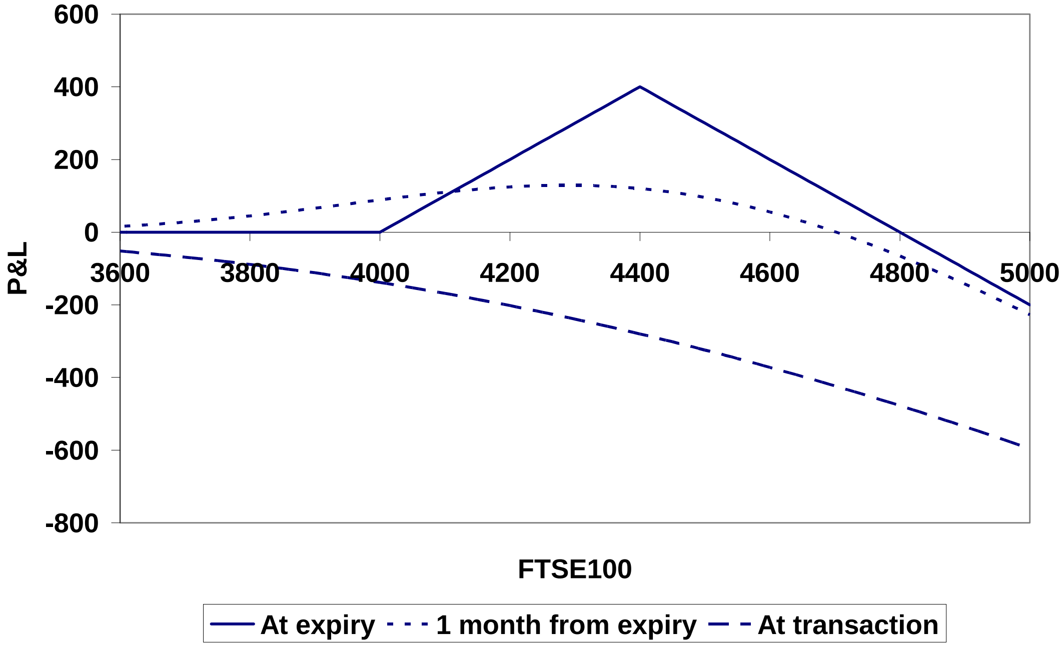

(ii) Gamma chart

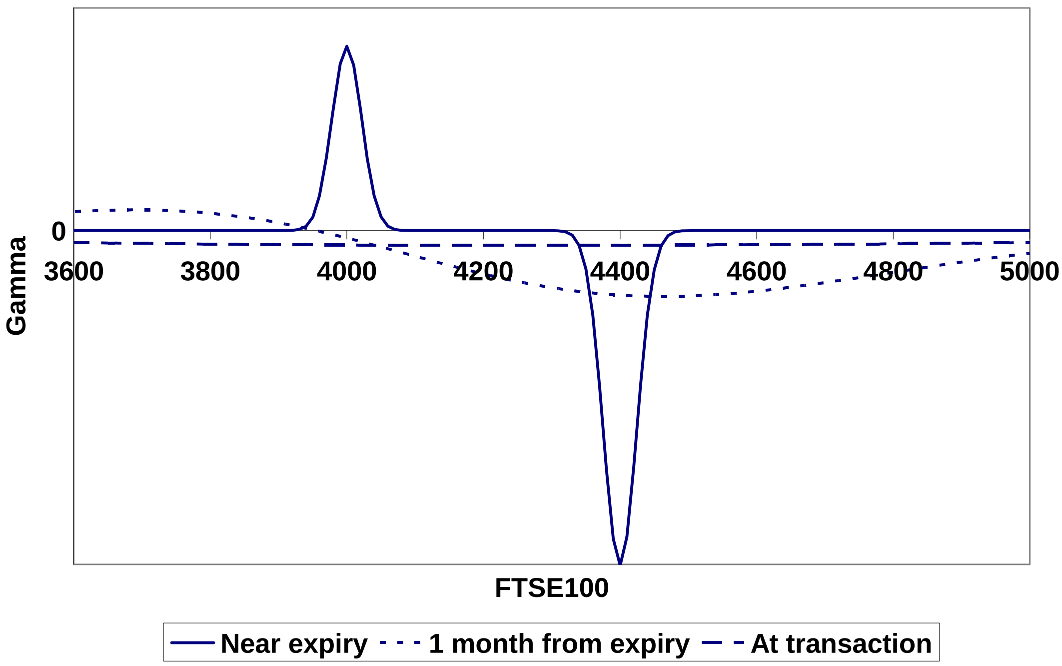

[Note: the two dotted lines are very close to zero; the spikes in Gamma very close to expiry would likely be even higher than those shown in the sketch.]

(iii) The transaction will give a gain if FTSE 100 lies within the range 4,000 to 4,800 at the expiry date.

The transaction has no impact if FTSE 100 is below 4,000, and is very negative if FTSE rises strongly.

The transaction is mostly negative delta and gamma, especially as the market rises.

Thus any sharp upward move in the market will increase the strain ...
... although this might be relieved by time decay in due course.

Writing gamma in a volatile market is not attractive unless premiums are very high, which is not the case here.

A "two-for-one" strategy rarely works well in a very volatile market.

The transaction is not at all well matched to a long term expectation that the market will rise.

Syllabus section: (l) + (m)
Core reading: 15, 16

6 (i) Market risk of a derivative is the risk that its value will change adversely ...
... due to movements in underlying variables such as interest rates, FX rates etc and other variables such as implied volatility.

Credit risk is the risk that a counterparty to a transaction defaults wholly or in part on its obligation to pay amounts as they fall due.

For a derivative, there will only be an actual credit loss if both the counterparty defaults <u>and</u> the transaction has a positive value to the original party.

(ii) (a) Market risk is neutralised in the absence of counterparty defaults.

Credit risk now exists to two counterparties, since the first transaction had credit risk to the customer and now there is also credit risk to the bank selling the hedge.

These credit risks are affected by market movements but in opposite directions, i.e. an FX rate change will increase one of the two credit risks, not both simultaneously.

(b) & (c)

Suppose a forward contract struck at X provides a payoff based on the value of S at time T.

The value of a long forward contract is $e^{-rT}(S_T - X)$.

The credit exposure on the sold customer contract is
$e^{-rT}\max(X - S_T, 0)$ ...
... which behaves like a Put on the asset price with strike price X.

Also, the credit exposure on the bought hedging contract is $e^{-rT}\max(S_T - X, 0)$ ...
... which behaves like a Call on the asset price with strike price X.

The total credit exposure, therefore, behaves like a straddle with strike price X.

[*Credit was also given for showing (b) graphically, provided that candidates clearly labelled the component parts of the straddle.*]

(iii) (a) Default correlation is the correlation between two or more entities that reflects their propensity to default around the same time (i.e. it is correlation of defaults).

Possible reasons:

- Companies in the same sector are subject to similar economic factors.

- Knock-on impacts of one company on another, e.g. motor manufactures and their suppliers.

- Major downturns (e.g. Great Depression) can cause mass bankruptcies.

The impact this has on CDOs is to increase the incidence of loss feeding into the higher rated (senior) tranches.

This means that, if default probabilities remain unchanged, rising default correlation will cause the more senior tranches to fall in value and corresponding junior tranches to rise in value.

(b) The CDOs of corporate bonds will be affected by general economic conditions and specific risks at the individual companies involved.

There will be an increase in correlation but companies generally will still default more or less independently.

With mortgages, there is much more likelihood of a systemic reaction – the vast majority of mortgages will either perform well together or badly, hence correlation impacts will be greater.

[*In the recent "credit crunch", US mortgage holders walked away from their debts in droves, causing an almost 100% correlation of defaults in CDOs of RMBS.*]

Syllabus section: (j), (k)
Core reading: 13, 14

7 (i) Using the tree defined in the question, and assuming that the SDE will apply to Δt and ΔR (the increment in R for a timestep Δt), we have:

$$\Delta R = -aR\Delta t + \sigma\Delta z\ .$$

This process has mean $-aR\Delta t$ and variance $\sigma^2\Delta t$.

The probabilities have to be chosen to match this mean and variance over the time interval Δt at every node, as well as add up to 1 (since they are probabilities).

If p_u, p_m and p_d are the three probabilities (up, middle and down), then $p_u + p_m + p_d = 1$. (*1)

For the jth row of the tree, $R = j\ \Delta R$, where $j = -j_{\max}, -j_{\max} + 1,\ \ldots, j_{\max} - 1, j_{\max}$.

From the mean:

$$p_u\Delta R - p_d\Delta R = -aj\Delta R\Delta t \qquad (*2)$$

and from the identity variance = $E(X^2) - (E(X))^2$:

$$p_u(\Delta R)^2 + p_d(\Delta R)^2 - a^2 j^2(\Delta R)^2(\Delta t)^2 = \sigma^2\Delta t \qquad (*3)$$

These equations (*1), (*2) and (*3) provide the means of calculating the probabilities.

(ii) Five-year zero coupon bond price = $\dfrac{1}{1{,}000}\displaystyle\sum_{n=1}^{1{,}000}\dfrac{1}{C(5,n)}$.

(iii) Five-year swap rate at time 3 = $S(n) = \dfrac{1 - B(3,8,n)}{\displaystyle\sum_{t=1}^{5} B(3,3+t,n)}$.

Swaption payoff at time 3 in nth run = $\max[5\% - S(n), 0].\sum_{t=1}^{5} B(3, 3+t, n)$.

This would be deflated back by dividing by $C(3, n)$, so:

$$\text{Swaption value} = £1\text{m}.\frac{1}{1{,}000}\sum_{n=1}^{1{,}000}\frac{\text{payoff}}{C(3,n)}$$

$$= £1\text{m}.\frac{1}{1{,}000}\sum_{n=1}^{1{,}000}\frac{\max[5\% - S(n), 0].\sum_{t=1}^{5} B(3, 3+t, n)}{C(3,n)}$$

$$= £1\text{m}.\frac{1}{1{,}000}\sum_{n=1}^{1{,}000}\frac{\max[5\%.\sum_{t=1}^{5} B(3{,}3+t,n) - 1 + B(3{,}8,n),0]}{\text{C}(3,\text{n})}$$

[*There are other ways of expressing this result – for example, using* $\frac{B(3,t,n)}{C(3,n)} = B(0,t,n)$.]

[*Full marks can be given for this part by providing the above answer without all the interim steps.*]

(iv) We are estimating prices as average of 1,000 observations of $X(n)$ = "payoff*deflator in run n".

By central limit theorem, this average is approximately normally distributed whatever the distribution of X.

An x% confidence interval for the underlying expectation of X (i.e. the value that the ESG would place on the payoffs) is:

Average ± Std Dev * $N^{-1}(1 - \frac{(100-x)}{200}) / \sqrt{1{,}000}$

where N^{-1} is the inverse Normal function
and the Average and Std Dev values are derived from the sample.

If the market value falls outside this confidence interval obtained, then this would be a statistically significant deviation.

[*Another valid method of estimating the percentiles would be to use the distribution obtained from the 1,000 sample paths.*]

(v) We assume the model has been calibrated properly.

It will value bonds accurately ...

... because given a and σ, $\theta(t)$ can be chosen so that the model replicates the starting yield curve.

It will not replicate all swaption prices ...
... because we have only two parameters (a, σ) to adjust and countless swaption prices to calibrate to.

One factor Hull-White allows negative interest rates, so this could in some cases give rise to inaccurate swaption prices.

Syllabus section: (b), (c), (h)(iv)-(ix), (i)
Core reading: 2, 3, 10 – 12

8 (i) Futures are traded on an exchange; forwards are private transactions between two parties.

Futures are standardised; forwards can come in many forms

Forwards usually have a single delivery date; futures have a range of delivery dates.

Forwards are settled at the end; futures are settled daily.

Physical delivery normally takes place with forwards; futures are usually closed out before maturity.

Forwards will carry some credit risk; futures have virtually no credit risk.

(ii) Basis risk is the risk to the bakery arising from differences between how wheat futures prices and cost of producing bread behave.

(iii) Minimum variance hedge ratio is the value of
[number of wheat futures] / [number of loaves ordered]
that minimises the variance of the cost of producing the bread after allowing for the hedge.

Define notation:

ΔB = change in cost of producing a loaf of bread over lifetime of hedge
ΔF = change in wheat futures price over lifetime of hedge
σ_B = standard deviation of ΔB
σ_F = standard deviation of ΔF
ρ = correlation between ΔB and ΔF
h = hedge ratio = N_F/N_B = number of futures / number of loaves

Cost at the end of the hedge of producing bread after allowing for hedge payoff is:

[cost of producing a loaf today] * $N_B + N_B\ (\Delta B - h\Delta F)$

Variance of this is minimised when var($\Delta B - h\Delta F$) is minimised
i.e. when $f(h) = \sigma_B^2 + h^2\sigma_F^2 - 2h\rho\sigma_B\sigma_F$ is minimised.

To find minimum, set $f'(h) = 2h\sigma_F^2 - 2\rho\sigma_B\sigma_F = 0$
which gives $h = \rho\sigma_B/\sigma_F$.

Check it is a minimum by seeing that $f''(h) = 2\sigma_F^2 > 0$

(iv) (a) If wheat futures fall, then the bakery will be hit by margin cash calls on its long portfolio of wheat futures.

This could cause a cashflow problem because margin has to be paid in cash at once ...
... but the offsetting asset that arises on the balance sheet (the greater margin from selling bread at fixed future prices but with lower input costs) is intangible and can't be converted into cash immediately.

The bakery could possibly borrow the cash from a bank to tide it over, but there would be a natural limit to how much it could fund this way at any reasonable cost.

In an extreme case, the bakery could go bankrupt due to the hedge, even though it has ostensibly acted sensibly and covered its commodity input costs.
[*For discussion of a practical example on this topic, candidates could read about the Metallgesellschaft case in Hull.*]

(b) The bakery could avoid the cashflow risks by taking out forward contacts without entering into any margining agreement.

However, this depends on finding a (bilateral) counterparty willing to take credit risk to the bakery for the life of the contract.

This should be possible since the bakery is large – generally at minimum the bank that provides it with overdraft and business banking facilities would offer such terms.

Even if the forwards were margined, the agreement would provide materiality thresholds and periodic payments, hence would be less onerous.

(v) Futures will bear basis risk; with forwards, basis risk could conceivably be reduced if the OTC provider is willing and if the cost of bread can be suitably defined.

Futures bring the cashflow risk – see part (iv).
With forwards, the bank might still want collateral posted if wheat / bread production prices fall, which would still leave the bakery with problems....
...or the bank might be happy bearing the credit risk if it feels that the fixed price contracts provide it with enough security.....

If wheat /bread production prices increase, the bakery will have credit exposure to the bank....
...which could be controlled by having the bank post collateral with the bakery.

Whether futures or forwards are used, the bakery will want to close the hedge out before maturity for operational simplicity.

The term of forwards can be longer than that for futures, meaning "rollover risk" can be avoided.....
...in fact, the maturity dates of the futures could be fitted to the bakery's order profile.

Dealing in futures will be cheaper, with no OTC derivative provider seeking to make a profit….
...and also they are likely to be more liquid when entering or exiting....
...and if the bakery wants especially non standard forwards (e.g. based on costs of making bread), it will be seeking bigger profit margins than usual.

Correlation between interest rates and wheat prices would alter the balance between futures and forwards (the margining requirements of futures would become less attractive).

Syllabus section: (h)(iv)-(ix), (i)
Core reading: 10 – 12

9 (i) No formula exists for valuing an American Put option so an algebraic approach cannot be used.

A tree integrates backwards from the payoff, so allows us easily to model the decision of whether to exercise the option before the expiry date because the option value at the exercise date is always known.

A Monte Carlo (MC) approach projects stochastic variables forward to capture possible time dependent outcomes, so is less easy to adapt to American options ...

... because it does not inherently produce the option value at the time – to do this in MC, some tree-like valuation of European option value (or a further iteration of the MC engine) would be needed.

(ii) Construct the Tree

In what follows below, $\Delta t = T / 30$.

Define $a = \exp(r\Delta t)$.

At each step, $F(t)$ at time t can move to $F(t + \Delta t)$ which can be either $uF(t)$ or $dF(t)$
where $u = \exp(\sigma\sqrt{\Delta t}\,)$ and $d = 1 / u$.

Probability of an upward move $p_{\text{up}} = (a - d) / (u - d)$.

Probability of a downward move $p_{\text{down}} = 1 - p_{\text{up}}$. These probabilities applying at every node.

Filling Option Values in the Tree

At time T, value of put is the payoff

$P = \max(X - F(T), 0)$ at each terminal node.

Then work backwards through the tree, starting from penultimate set of nodes.

At each point at time t where $F(t)$ could move to $uF(t)$ or $dF(t)$, price of a European put at that node would be

$$p(t) = \exp(-r\Delta t)\,[p_{\text{up}}.uF(t) + p_{\text{down}}.dF(t)].$$

But we are valuing an American option, so we need to model it as if the option were exercised if the payoff exceeded $p(t)$, so instead we value the option at this point as

$$P(t) = \max\{\,\exp(-r\Delta t)\,[p_{\text{up}}.uF(t) + p_{\text{down}}.dF(t)]\,,\, X - F(t)\,\}.$$

(iii) A typical control variate procedure would be:

- Value the American option using the tree
- Value the corresponding European option using the tree
- Value the European option using Black-Scholes formula
- Use as the value of the American option

Value of American option using tree

less Value of European option using tree
plus Value of European option using Black-Scholes

(iv) (a) Manager A

We could either:

- Calculate implied volatilities of similar options whose prices are published in the daily press OR
- Have daily reports from a bank with implied volatilities for a range of liquid options OR
- Use the VIX volatility index

and

- have a simple rule to convert these to an appropriate volatility for the option being valued

(b) Manager B

We could either:

- Keep a record of the daily implied volatilities
- Maintain a scatter-plot of implied volatility vs equity index
- Watch out for volatilities that look out of line with recent experience
- Check whether strange looking volatilities are a result of market movements by
  - Querying with option writer OR
  - Keeping daily record of one of the three volatility sources in (iv)(a) to compare against

(c) Comparison

The fundamental difference between the approaches is that one is a reasonableness check and one is the calculation of an alternative price.

The alternative price calculation is unlikely to exactly match the value quoted by the option provider, so judgement will still be required on whether the differences are of concern.

Translating prices into implied volatilities is useful for comparing prices for consistency, which suggests that a sense check on implied volatility will be performed in any case.

Comparing implied volatilities is also useful going forward to help understanding of how the implied volatility of the option in question is related to other readily available volatilities.
[*Other valid points could be made.*]

END OF MARKING SCHEDULE