Subject ST6 — Finance and Investment
Specialist Technical B

EXAMINERS' REPORT

April 2008

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.

M A Stocker
Chairman of the Board of Examiners

June 2008

Comments

The standard of attempts at this session fell from previous years, breaking the trend of recent progress. There were signs of a number of well prepared candidates, but far too many provided sketchy and cursory answers. Bookwork elements were generally known, but candidates were noticeably imprecise on some of the elements of the Core Reading that had not been tested before, despite being well flagged in the syllabus. Numerical work improved over previous sessions, possibly because the paper included the more familiar binomial model.

The ST6 exam is challenging, and derivatives is clearly a hard subject – harder maybe than it appears on paper. Quite a number of candidates seem to have under-estimated in their preparation the level of precision involved. Nevertheless, the examination is eminently passable with adequate study. For example, a significant percentage of the total marks for the examination are given simply for repeating bookwork. Candidates need to take advantage of this feature, since by providing complete and concise answers to bookwork questions, together with reasonable attempts at the application questions, they can comfortably achieve the pass mark. Far too often the examiners found one answer after another would start promisingly but then fade away, thereby making it hard for the candidate to garner enough marks.

> ***Please note that the model solutions provided are indicative, i.e. adequate to achieve full marks but without covering every possible correct response. Several points made by candidates were equally valid, and these also achieved the allocated marks.***

Q1 *This question proved manageable for most candidates and good marks were scored. Precise argumentation was required in part (iii), rather than vague concepts, but a number of alternative presentations were offered and accepted.*

Q2 *This was a straightforward bookwork question. It is hard to understand why several candidates left this to the end before attempting it as it was one of the easier questions. A commonsense analysis was required for part (iii), focusing on the dynamic nature of futures hedges and the cashflow implications of margining.*

Q3 *This was an example of a question asking candidates to do in practice what the theory states, and surprisingly almost all struggled at that point. Few obtained the Binomial Representation Theorem ratio in part (ii)(b) as per the solutions.*

It is expected that candidates will understand the BRT formula and commit it to memory, but it is also important to see how it works on an actual discrete binomial tree. If nothing else, working out a practical example provides insight into how the theory operates. Ironically, in part (iii)several candidates were able to state clearly how BRT is applied, without being able to derive the required relationship in part (ii).

Q4 *This was a long but straightforward application of fixed income theory. It was generally answered well when carried out methodically. Several candidates lost their way by not applying the results in a disciplined manner, but it was pleasing to see that most candidates understood how duration measures work and how swaps and bonds are related.*

Q5 *This question was avoided by most people, but was not as hard as it seemed. Admittedly, part (i) was daunting, particularly (i)(b), but if candidates were struggling they could have omitted this and tackled the bookwork in part (ii), assuming they knew something about LIBOR Market Models from the Core Reading.*

Part (iii) was challenging in that it was unfamiliar and more thought was required. Insight about the implicit bought floors ("coupons can never decrease") being potentially worth less than the sold caps ("the coupon can … only increase by a maximum of 0.50%") was missed by most. Questions of this sort are not complicated – they are simply looking for basic applications of derivatives theory to unusual situations.

Part (iv) was better answered. It is important to understand the natural limitations of the various pricing models.

Q6 *Part (i) of this question asked the candidate to derive an integral form for the probability distribution, without evaluating the integral.*

For part (ii), the range over which the option has value means that the Delta moves suddenly from 0 to 1 (or 1 to 0) at the boundary of the range, resulting in a spike in Gamma. Candidates who performed best in this part started with the payoff/price of the option, then derived the Delta behaviour as being the change in payoff/price, then derived the Gamma behaviour as being the change in Delta. This is a reliable technique for graphing all option strategies.

Q7 *This question included a straightforward binomial tree, and most candidates answered it well. A few came a little unstuck on valuing the lookback option at different nodes, but most managed to get the required result. Numerical accuracy was generally high. As with earlier questions, a disciplined approach in setting out the solution yields rewards.*

There were a few complaints about the initial price of the commodity being S_0 rather than the more obvious C_0. Admittedly this looks like a simple typographical error in the exam paper, but the better candidates quickly overcame their aesthetic umbrage and concentrated on more important considerations.

Q8 *Like Q5, this was quite a tough question, particularly part (i). However, a very large hint was given and several candidates managed to pick up the logic of the Margrabe approach that was represented here.*

Many candidates who omitted part (i) managed to obtain good marks for part (ii), which required an application of part (i) to convertible bonds. Incidentally, it is worth taking trouble to prepare answers to follow-up questions such as (ii)(b), asking how well a particular theoretical solution works in practice – these crop up quite frequently, and are typical of the sort of knowledge expected of actuaries.

QUESTION 1

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

Core reading: Units 8 & 9

(i)

(a)

A stochastic process $W = (W_t : t \geq 0)$ is a **P**-Brownian motion if and only if:

- W_t is continuous [*the process is said to be "unshifted" if $W_0 = 0$*]
- the value of W_t is distributed, under probability measure ***P***, as a Normal random variable with mean zero and variance t
- the increment $W_t - W_s$ $(t > s)$ is distributed, under ***P***, as a Normal random variable with mean zero and variance $t - s$
- the increment $W_t - W_s$ $(t > s)$ is independent of the filtration (path or history) F_s up to time s

(b)

A process X is a **P**-Martingale if:

$$\mathrm{E}_{\mathbf{P}}(X_t | F_s) = X_s \quad \text{for } t > s$$

where $\mathrm{E}_{\mathbf{P}}$ is expectation under probability measure **P**, and F_s is the history of the process up to time s.

(ii)

Firstly, we need to show that Brownian motion is a Martingale.

$$\mathrm{E}_{\mathbf{P}}(W_t | F_s) = \mathrm{E}_{\mathbf{P}}(W_s | F_s) + \mathrm{E}_{\mathbf{P}}(W_t - W_s | F_s) = W_s + 0$$

since $W_t - W_s$ is independent of F_s, so W_t is a **P**-Martingale.

Then $\mathrm{E}_{\mathbf{P}}(X_t | F_s) = \mathrm{E}_{\mathbf{P}}(W_t | F_s) + \alpha t = W_s + \alpha t = X_s + \alpha(t - s)$

since W_t is a **P**-Martingale.

This is true for all t, so X_t is not a Martingale unless $\alpha = 0$.

(iii)

We are told that all tradable assets are Martingales.

Suppose we have any portfolio, Π, of tradable assets, with a <u>zero value</u> today. (Π is our arbitrage portfolio.)

Suppose also that there is a non-zero probability that it will have a <u>positive value</u> at some time T in the future.

Since Π is a collection of Martingales, the expectation today of its value at time T must be its value today, i.e. zero.

For this expectation to apply, there must be a non-zero probability that Π will have a negative value at T.

Hence there can be no guaranteed arbitrage in this market.

QUESTION 2

Syllabus section: (b)

Core reading: Unit 2

(i)

(a)

The forward price of an asset is the fixed price at which a dealer agrees at the outset of a contract to buy or sell the asset at a future time.

The value of a forward contract, then as market values (e.g. the underlying asset price) change, may become positive or negative.

(b)

A forward contract is an OTC agreement between two parties with a bespoke forward date and size.

A futures contract is traded on an exchange with a standard forward date and size.

The forward contract settles on the forward date by replicating the economic effect of a real transaction on that date.

The futures contract settles daily according to a margined formula of difference.

There is counterparty credit risk on the forward contract but not the future.

(ii)

Let the variable S be the current price in Euros of 1 unit of the foreign currency (dollars), i.e. the current exchange rate, expressed in Euros per dollar.

Let F be the forward price agreed to in the contract, f the current value of the forward contract, T the term of the forward contract (= 0.5 years).

Let r_{EU} be the continuously compounded risk free rate in Euros and let r_{US} be the risk free rate in dollars (for the period to time T).

The two portfolios that enable us to price a forward contract on a foreign currency are:

A: An amount $e^{-0.5r_{US}}$ of US dollars; and

B: A long forward contract to buy 1 dollar, plus an amount $Fe^{-0.5r_{EU}}$ of Euros.

Both of these portfolios will become worth the same as one unit of the foreign currency (i.e. one dollar) at time *T*, so, to be arbitrage free, must be equal at time 0.

Hence in Euros:

$$f + Fe^{-0.5r_{EU}} = Se^{-0.5r_{US}}$$

Forward contracts are entered into at zero cost, so $f = 0$.

Hence price of forward contract $F = Se^{0.5(r_{EU} - r_{US})}$.

(iii)

To hedge $ cashflows in Euros, the manufacturer would need to buy dollar-Euro futures (i.e. sell dollars, buy Euros).

Divide the total amount by the contract size to find how many contracts.

Problems in using futures:

- must generally roll contracts every 3-months => additional (small) trading costs
- futures basis will not be aligned to the forward market => might enter/exit the contracts at disadvantageous basis
- variation margin needs to be paid in cash and might be large if $-Euro rate moves substantially => could cause cashflow problems
- possibly not enough liquidity in the futures market to cover large positions when required

QUESTION 3

Syllabus section: (h)

Core reading: Units 8 & 9

(i)

(a) We will say that a process S_i, $0 \le i \le T$ is <u>previsible</u> if S_i depends only on the filtration (history) F_{i-1}, i.e. up to the previous time step. Once its value is known at the previous time step, there is no stochastic uncertainty about its value at the next time step. (This is not true of most stochastic processes.)

[*Credit is also given to candidates for the equivalent continuous version.*]

(b) The Binomial Representation Theorem (BRT) states that, if X is a **P**-Martingale, i.e. $E_{\mathbf{P}}(X_i | F_j) = X_j$ where $i > j$ (for filtration F_j up to time step j), and Y another **P**-Martingale, then there exists a previsible process ϕ such that:

$$Y_{i+1} - Y_i = \phi_i (X_{i+1} - X_i) \quad \text{or} \quad \Delta Y_i = \phi_i \Delta X_i$$

where **P** is the risk neutral probability measure.

(ii)

(a) Set up the table of possible values for S:

$n = 0$	$n = 1$	$n = 2$	$n = 3$
			55
		50	
	45		47
40		42	
	37		39
		34	
			31

A discrete random walk is a Martingale under the risk neutral probability measure.

Hence $40 = 45p + 37(1 - p)$, whereby $p = 0.375$ for all time steps.

(b) Now filling in the table of values for X, starting at n = 3 for the payoff, then at each previous node back to n = 0 take expectations under **P**.

$n = 0$	$n = 1$	$n = 2$	$n = 3$
			1.416198
		1.065859	
	0.695943		0.855655
0.407959		0.473994	
	0.235169		0.244998
		0.091874	
			0

We need to verify for n = 1 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,

(45 – 40) / (0.695942 – 0.407959) = 17.362

and (37 – 40) / (0.235169 – 0.407959) = 17.362

so the ratio process at n = 0 is previsible (same whether the stock goes up or down).

(iii)

Find the probability measure **P** that makes the stock price S a **P**-Martingale.

The second tree shows process $X_i = \mathrm{E}_{\mathbf{P}}(X | F_i)$ for filtrations F_i to time i, which is also a **P**-Martingale.

The Binomial Representation Theorem states that the changes in S_i and X_i between i and i + 1 are proportional with a (hedge) ratio that is known at time i.

Each node going through the second tree represents a previsible point at which the hedge ratio is known ...

... so if the claim X is an option payoff, the option value can be constructed from known (dynamic) positions in the stock S.

Thus the option can be replicated by taking delta positions in the stock.

QUESTION 4

Syllabus section: (e), (j), (l) & (m)

Core reading: Units 5, 13, 15, 16

(i)

(a) Price of a 5-year 10% annual coupon bond at current rates:

$$P = 10\sum_{i=1}^{5} d_i + 100 d_5 = 10(4.30784) + 100(0.77796) = 120.875.$$

where d_i are the values of the zero coupon bonds of length i (per nominal of 1 unit).

(b) The value of a 5-year annual fixed-floating swap with 10% fixed coupon is the difference between a 10% 5-year bond and an FRN. But the latter is valued at 100.

Hence swap value = 120.875 – 100 = 20.875.

[*This is a receiver swap - the payer swap is of opposite sign.*]

(ii)

Working in percentage of par:

(a) Modified duration $M = -\frac{1}{P}\frac{\partial P}{\partial r}$ where r is the level of rates (or yields).

Absolute yield sensitivity $\zeta = \frac{P_r - P_{r+\Delta r}}{\Delta r}$,

so $\zeta = PM$ if Δr is small and the entire yield curve shifts in parallel.

(b) First we need to calculate the absolute yield sensitivity of the bond and swap.

Using the same technique as in (i), we get:

Bond price (perturbed) = 10 (4.296) + 100 (0.77428) = 120.388.

Hence ζ_{bond} = (120.875 – 120.388) / 0.10 = 4.87.

Similarly, the swap values (perturbed) = 120.388 – 100 = 20.388.

Hence ζ_{swap} = (20.875 – 20.388) / 0.10 = 4.87.

[*This is also clear from the fact that, in the original and perturbed states, the swap value is the bond price less 100.*]

(iii)

(a) The cash flows of the reverse floater are:

Years 1 to 5:	$10 - f_i$ i = 1 to 5	(coupons)	(A)
+ Year 5:	100	(redemption)	(B)

But the coupons on an annual fixed-floating swap are the same as in (A) ...

... and the 5-year zero coupon bond pays 100 at the end of year 5 as in (B).

Hence the decomposition required.

(b) Price and absolute yield sensitivity are additive.

So the price of the reverse floater = swap + 5 year zero = 20.875 + 77.796 = 98.671.

Also, the absolute yield sensitivities = swap + 5 year zero

= 4.873 + (77.796 – 77.428) / 0.10 = 4.873 + 3.680 = 8.553

so Modified Duration of reverse floater = 8.553 * 100 / 98.671 = 8.668

which is roughly twice the Modified Duration of the par yield bond (4.299).

(iv)

There are two possible alternative answers, both equally valid.

<u>Bond world</u>

In a bond world, the portfolio is risk managed by ensuring that modified durations match.

The total duration is the sum of the individual bonds, weighted by nominal amount.

Longer bonds will have greater convexity.

Swaps are treated as a long (or short) fixed bond position and an equal short (or long) floating bond position, and their durations obtained separately.

The reverse floater is decomposed into a swap and zero coupon bond using the analysis in (iii) above.

<u>OR</u>

<u>Swap world</u>

In a swap world, the risk management is performed by sensitivity analysis to the underlying yield curve.

The bonds are mapped to the relevant yield curve points by taking their price sensitivity to each of the components of the curve in turn, i.e. 1 year zero, 2 year zero, etc.

Swaps are decomposed into fixed and floating legs ...

... and the inverse floater is decomposed as in (iii).

QUESTION 5

Syllabus section: (k), (l) & (m)

Core reading: Units 14 - 16

(i)

(a) P and F are related as:

$$\frac{P(t,t_k)}{P(t,t_{k+1})} = 1 + \delta_k F_k(t) \text{ where } \delta_k = t_{k+1} - t_k \text{ (annual compounding)}$$

<u>OR</u> $\dfrac{P(t,t_k)}{P(t,t_{k+1})} = \exp(\delta_k F_k(t))$ where $\delta_k = t_{k+1} - t_k$ (continuous compounding)

(b) The Equivalent Martingale Measure (EMM) result states that, for two processes f and g that have the same single underlying source of uncertainty, $\dfrac{f}{g}$ is a Martingale, i.e. $f_t = g_t \mathrm{E}\left(\dfrac{f_T}{g_T}\right)$ for $T > t$, taking expected values in a world in which g is forward risk neutral.

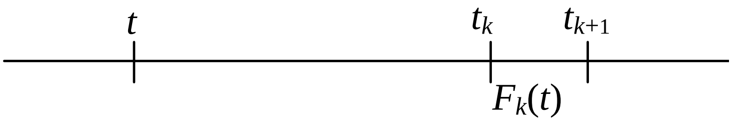

Let $f_t = F_k(t)$ and $g_t = P(t, t_{k+1})$, and set $T = t_{k+1}$.

Then we note that $g_T = P(T, T) = 1$, hence:

$$F_k(t) = P(t, t_{k+1})\,\mathrm{E}(F_k(t_{k+1})).$$

(ii)

The Libor Market Model (LMM) works in a <u>rolling</u> forward risk neutral world that sets the risk neutral measure to $P(t, t_{k+1})$ for each k = 1, 2 etc. (t is the current time.)

We know that $F_k(t)$ is a Martingale under the measure that makes $P(t, t_{k+1})$ risk neutral for each k. So, if the volatilities of $F_k(t)$ are $\zeta_k(t)$, we have:

$$dF_k(t) = \zeta_k(t) F_k(t) dz \quad \text{for each } k,$$

for some Wiener process dz (same for all k).

Converting to the rolling forward risk neutral measure introduces a drift $\mu_k(t)$ into the process for $F_k(t)$ that depends only on the volatilities up to k, i.e. $\zeta_i(t)$, i = 1, 2, …, k.

To calibrate the model to observed caplet volatilities, the LMM decomposes these volatility into a continuous sets of 'instantaneous' (or, in the discrete notation we are using, step-by-step) volatilities that correspond to the forward rate volatility.

The LMM is very suitable for solving using a Monte Carlo approach which follows the forward rates as they evolve.

[*The above presentation follows the notation used in Hull section 29.2.*]

(iii)

The ratchet coupon can never fall, but if rates rise sharply it will pay a coupon less than the market rate, at least for a while (it may catch up). Thus the floater has good downside protection on coupons in exchange for less upside.

The option effect here is that the holder of the ratchet floater has bought a series of floorlets and sold a series of caplets based on rising strikes. Each floorlet strike is at the same level as the previous coupon, and the corresponding caplet strike is 0.50% above it.

With a flat yield curve, the strikes on the caplets are further away from the current LIBOR level than the strikes on the floorlets, so the floorlets are more valuable.

However, if there is a rising yield curve, the expected forward rates will be increasing and each caplet strike may be nearer the forward rate than the corresponding floorlet strike (or at least, they might be overall). This would mean that the cap on the increase could be worth more than the ratchet effect. [*This is not an obvious result at first glance.*]

(iv)

The two key features to model for these types of complex interest rate option are:

(1) Path dependency – for the 'ratchet', each coupon depends on the previous one.

(2) Yield curve slope – the value depends on the slope as well as the level of the yield curve. This implies two correlated factors.

Point (1) suggests a Monte Carlo (MC) approach would be better than a binomial tree, hence favouring the LMM which needs a MC.

Effect (2) suggests a two-factor model using principal components, along the lines of level and slope of the yield curve.

Even a two-factor Black model would not work, as it cannot be calibrated for such a decomposition (it is more suitable for two correlated assets, not risk factors).

It is true that the LMM is complex, and time-consuming difficult to implement.

It requires a specialist team of quant modellers and programmers to achieve the calibration and MC calculator.

However, the MC method is very suitable for extending to multiple factors.

But the danger of using an inappropriate model is that you do not know how inaccurate the pricing will be ...

... and the risk sensitivities will probably be even more inaccurate.

QUESTION 6

Syllabus section: (g),(i) & (j)

Core reading: Units 7 & 12

(i)

(a) We are given that $dS = \mu S dt + \sigma S dz$, so by Ito on ln S (i.e. $\log_e S$):

$$d(\ln S) = \left[\frac{\partial \ln S}{\partial t} + \mu S \frac{\partial \ln S}{\partial S} + \tfrac{1}{2}\sigma^2 S^2 \frac{\partial^2 \ln S}{\partial S^2} \right] dt + \sigma S \frac{\partial \ln S}{\partial S} dz$$

$$= (\mu - \tfrac{1}{2}\sigma^2) dt + \sigma dz$$

which means that ln S has a Normal distribution with mean $\mu' t = (\mu - \tfrac{1}{2}\sigma^2)t$ and variance $\sigma^2 t$.

Hence $\dfrac{\ln S - \mu' t}{\sigma\sqrt{t}}$ ~ N(0, 1), so the probability density function is thus:

$$f(s) = \frac{1}{\sqrt{2\pi\sigma^2 t}} \frac{1}{s} \exp\left\{ -\tfrac{1}{2} \frac{(\ln s - \mu' t)^2}{\sigma^2 t} \right\}.$$

(b)

The risk neutral measure is the measure that makes the return on the asset equal to the risk-free rate, so $\mu = r$.

(ii)

(a) Value of the option = $Ke^{-rt} \int_L^U f(s)ds = Ke^{-rt} \left[\int_0^U f(s)ds - \int_0^L f(s)ds \right]$

and substitute for the density function as above.

[*Not required to evaluate the integrals.*]

(b)

Chart of Gamma profile (dotted line = now, solid line = near expiry):

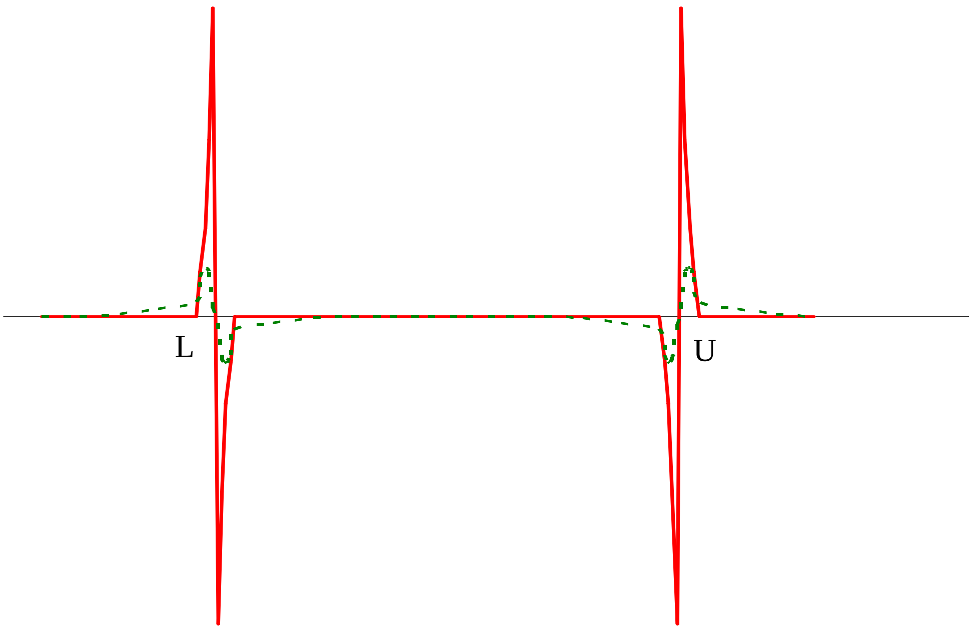

(c) Problems with hedging the option:

- Payoff changes suddenly at *U* and *L*, so Gamma is large (potentially infinite) there
- This will make delta hedging difficult if the stock is close to either the up or down range boundary towards expiry
- Spikes in Delta and Gamma do not apply to vanilla calls and puts
- Only options with some barrier effect have similar gamma profiles, so could perhaps hedge with those, but they are not as liquid as vanilla options

QUESTION 7

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

Core reading: Units 8 – 10

(i)

Setting up the commodity tree using *u* for up move and *d* for down move:

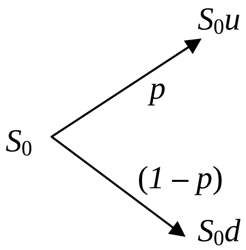

where *p* is the up probability, (1 – *p*) the down probability.

Then $E(C_t) = S_0\left[pu + (1-p)d\right]$, and

$$Var(C_t) = S_0{}^2\left[pu^2 + (1-p)d^2 - \left(pu + (1-p)d\right)^2\right]$$
$$= S_0{}^2\left[p(1-p)u^2 + p(1-p)d^2 - 2p(1-p)\right]$$
$$= S_0{}^2\left[p(1-p)(u-d)^2\right]$$

since $u.d = 1$.

Equating moments:

$$S_0e^{rt} = S_0\left[pu + (1-p)d\right] \qquad \text{(A)}$$

and $\sigma^2 S_0{}^2 t = S_0{}^2\left[p(1-p)(u-d)^2\right]$ (B)

The solution to equation {A} is:

$$p = \frac{e^{rt} - d}{u - d}$$

Substituting into equation {B}gives:

$$\sigma^2 t = -(e^{rt} - u)(e^{rt} - d) = (u + d)e^{rt} - (1 + e^{2rt}) \text{ when } d = 1/u$$

Multiplying through by *u* gives:

$$u^2e^{rt} - u(1 + e^{2rt} + \sigma^2 t) + e^{rt} = 0$$

This is a quadratic in *u* which can be solved in the usual way.

(ii)

(a) σ = 0.15 and t = 0.25, so u = exp(0.075) = 1.077884, d = 1 / u = 0.927744

The tree looks like this:

***t* = 0**	**0.25**	**0.5**	**0.75**	
			93.924	*Node A*
		87.138		
	80.841		80.841	*Node B*
75		75		
	69.581		69.581	*Node C*
		64.553		
			59.889	*Node D*

(b)

Since r = 0, we have $p = (1 - d) / (u - d) = 0.48126$

To value the vanilla option, set the final column (t = 0.75) to be the values above less the strike of 75, and then 'discount' back along the tree using p and $(1 - p)$ to t = 0.

***t* = 0**	**0.25**	**0.5**	**0.75**	
			18.924	*Node A*
		12.138		
	7.300		5.841	*Node B*
4.215		2.811		
	1.353		0	*Node C*
		0		
			0	*Node D*

Hence the tree value of the vanilla call is 4.215.

[*Alternatively, a variant of the method used below in (c) is possible using individual paths. The answer will be identical.*]

(c)

The lookback call pays the difference between the minimum value and the final value.

Notate the paths by U for up and D for down, in order.

The lookback payoffs are, for each successful path (i.e. with a non-zero result):

UUU = (93.924 – 75) = 18.924 @ *Node A*

UDU = (80.841 – 75) = 5.841 @ *Node B*

UUD = (80.841 – 75) = 5.841 @ *Node B*

DUU = (80.841 – 69.581) = 11.261 @ *Node B*

DDU = (69.581 – 64.553) = 5.028 @ *Node C*

with the remaining paths not generating a lookback profit since the minimum value is at the final node.

The probabilities of arriving at each node are:

Node A = p^3 = 0.11147

Node B = $p^2(1 - p)$ = 0.12015

Node C = $p(1 - p)^2$ = 0.12950

Hence the tree value of the lookback option is:

= (0.11147 x 18.924) + (0.12015 x [5.841 + 5.841 + 11.261]) + (0.12950 x 5.028)

= 5.517.

(iii)

The lookback option appears to be priced higher than the vanilla call.

Normally this would be the case since a lookback is more valuable, provided the strike is sensible, i.e. the vanilla option is not in-the-money at outset.

In general, the tree is too coarse to get accurate prices ...

... but one could use smaller time-steps to improve accuracy.

QUESTION 8

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

Core reading: Units 8 – 10

(i)

(a) At expiry at time T, payoff of option = $\max[S_2(T) - K.S_1(T), 0]$.

(b) We can write the payoff = $S_1 \max\left[\dfrac{S_2(T)}{S_1(T)} - K,\ 0\right]$

and hence can value $\Phi(X,t) = \dfrac{V(t)}{S_1(t)}$

using the risk neutral measure associated with S_1 (i.e. with S_1 as the numeraire).

Let the stochastic process for each stock be represented by $dS_i = \mu_i S_i dt + \sigma_i S_i dW_i$ for i = 1 and 2, where dW_i are the correlated standard Brownian motions.

Using the result we were told, under the risk neutral measure for S_1, in which S_1 is the numeraire, we must have $\mu_1 = r + \sigma_1^2$.

Then, using Ito, $d\left(\dfrac{1}{S_1}\right) = -\dfrac{r}{S_1}dt - \dfrac{\sigma_1}{S_1}dW_1$

so $d\left(\dfrac{S_2}{S_1}\right) = \dfrac{S_2}{S_1}(\mu_2 - r - \rho\sigma_1\sigma_2)dt + \dfrac{S_2}{S_1}(-\sigma_1 dW_1 + \sigma_2 dW_2)$

since $dW_1 dW_2 = \rho\ dt$.

Hence if we put $\mu_2 = r + \rho\sigma_1\sigma_2$ then $\dfrac{S_2}{S_1}$ is driftless, and hence a Martingale,

and $\dfrac{S_2}{S_1}$ has volatility $\hat{\sigma} = \sqrt{\sigma_1^2 - 2\rho\sigma_1\sigma_2 + \sigma_2^2}$.

This is the same starting point of a Black Scholes option with "stock price" $\dfrac{S_2}{S_1}$, strike K, volatility $\hat{\sigma}$

From part (a), we can see that the boundary conditions are also similar, with K as the equivalent strike price for $\dfrac{S_2}{S_1}$.

Therefore the closed-form solution is of the form:

$$\Phi(t) = \frac{S_2}{S_1}N(d_1) - KN(d_2)$$

i.e. $V(t) = S_2 N(d_1) - KS_1 N(d_2)$

where $d_1 = \dfrac{\ln\left(S_2 / KS_1\right) + \frac{1}{2}\hat{\sigma}^2 (T-t)}{\hat{\sigma}\sqrt{T-t}}$, $d_2 = d_1 - \hat{\sigma}\sqrt{T-t}$,

and N(¨) is the cumulative Normal distribution.

[*This is called the Margrabe formula.*]

(ii)

(a) A convertible bond is a bond issued by a company that can be converted into the equity of the company on certain a date or dates at a predetermined exchange ratio.

In the case outlined in (i):

- the underlying corporate bond is S_1 (non-convertible equivalent, not the convertible itself)
- the equity is S_2
- derive the volatilities from the market
- assess correlation between underlying bond and equity prices
- use the Margrabe formula from (i) for the convertible price and the Greeks.

(b) It would work but there are problems:

- the model ignores credit risk, which is a big factor in convertibles; need to allow for the possibility of default in the corporate bond
- the model would have to be adapted to allow for the equity's dividends
- correlation might be unstable or hard to measure – perhaps use benchmarks
- correlation might change if the equity performed well or badly; in stress situations, correlations become nearer 1
- the exact terms of conversion may depend on other non-financial factors
- multiple conversion dates are hard to deal with
- bonds will not necessarily follow a geometric Brownian Motion

END OF EXAMINERS' REPORT