

Subject ST9 — Enterprise Risk Management Specialist Technical

April 2010 Examinations

EXAMINERS' REPORT

Introduction

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

R D Muckart
Chairman of the Board of Examiners

July 2010

- 1** It is good practice for the roles of chief executive and chairman to be held by different people. This is not the case here. This is to avoid too great a concentration of power at the head of a firm. As a temporary measure the board could meet regularly without the CEO and/or appoint an alternate lead director.

It is good practice for the majority of board members to be independent from the firm. This is not the case, due to past employment with the firm. Independence is important as it ensures that the directors act in the interests of shareholders rather than themselves or other members of the board.

It might also be useful for the board to have input and challenge from those who have experience of other industries and who may therefore be better able to see “the bigger picture”.

A significant proportion of directors' remuneration should be in the form of company stock, in order to align the interests of the directors and shareholders. Holding stock would not be a conflict of interests.

Ideally there should be a process in place to at least annually assess the performance of all directors, including the chief executive. No formal process exists here. The performance assessment should include risk based objectives. The performance assessment of the directors including the CEO should be undertaken by a committee of the board called the remuneration committee.

The CEO should not sit on the remuneration committee.

- 2** (i) The broad principle is that an employee will act to benefit his own well being rather than the financial well being of shareholders. A both ways principle statement is fine “each party not acting in concert, each in their own interests”

A specific example would be the use of the internet at work for personal purposes.

- (ii) Agency risk can be limited by:

- increased monitoring of employees
- limiting the opportunities for exploitation of opportunities; and
- using remuneration and benefits to incentivise employees to be productive in a way which is aligned to the shareholders' objectives
- setting known penalties for breaching the rules

In this case, internet sites can be blocked, and/or employees advised that connections are being monitored and personal use could result in dismissal.

- (iii) Any interventions should not breach the employee's legal rights.

The cost of any intervention or prohibition should not exceed the cost of the employee's actions before intervention or prohibition.

Any intervention should not make the employee's position untenable.

The intervention should not be unduly complicated. It will be difficult if not impossible to remove the risk altogether.

The intervention should not overly stifle initiative.

- 3**
- (i)
- (a) Risk appetite: this is the degree of risk that an organisation or individual is willing to accept in order to achieve objectives, both in terms of levels and types of risk. It can be interpreted as reflecting the setting of targets and limits across the organisation as a whole, plus the breakdown of these high level statements into more detailed risk tolerances.
- (b) Risk profile: this is a complete description of the risk exposures of an organisation, including risks that might emerge in the future and that will affect the current business of the organisation.
- (c) Risk limits: this is a group of guidelines that set limits on acceptable actions that might be taken today. If risk limits are adhered to then each individual unit of the business should be deemed to be working within its permitted risk tolerances. Risk limits can be regarded as a component of risk capacity.
- (d) Risk capacity: this is the volume of risk that an organisation can hold as measured by some consistent measure, such as Economic Capital. If there is spare capacity then it might be possible to take positive actions that add economic value to the organisation without breaching existing risk tolerances or risk limits.
- (ii) The Board's expression of its risk appetite need not be complex. Instead, it might be expressed as a short and clear set of statements related to one or more measures of risk.

The Board may wish to express its appetite with reference to metrics such as one or more of:

- the company's solvency level
- its credit rating
- its earnings and ability to pay dividends

In choosing one of the above metrics, the Board might be focusing on the interests of specific stakeholders. Reference to the solvency level will be relevant to one group of stakeholders (e.g. insurance policyholders and regulators), while earnings and dividends will be relevant to others. The Board might state different objectives each with different stakeholders in mind.

In discussions, the Board might initially state that it is unacceptable that quantity X (solvency level, credit rating etc.) falls below level Y over a

specified period (e.g. one year). Normally such breaches cannot be completely prevented, and so the initial statement often needs to be translated into a more probabilistic statement such as:

- The solvency level, X, should stay above the threshold Y with 99.5% probability over the next three years.
- The probability that the company's credit rating is reduced from the current AAA to A, or worse, in the next twelve months should be no more than 1%.
- Earnings volatility over the next year should be no more than Z%.
- The company is prepared to lose \$A with a probability of no more than 0.5% over the next 12 months and \$B with a probability of 0.1% over the next 5 years.

It is important to bear in mind that Boards may express their risk appetite using a combination of statements linked to several metrics.

As a variant, the Board of an insurance company might express the objective to maximise the economic value of a company for shareholders subject to one or more constraints that focus on the policyholders or the regulator, such as keeping the solvency level above a certain threshold with a given probability.

In order to establish the Board's risk appetite statement in the company, the senior risk managers (possibly involving further discussions with the Board, or the risk committee of the Board) would then have the task of translating the higher level statements of risk appetite into a more detailed set of risk tolerances and risk limits across the enterprise. This needs to be carried out in a holistic way to take advantage of synergies and to avoid unanticipated concentrations of risk.

A statement of risk tolerances needs to cover the company's attitude to all risks, both quantifiable and non-quantifiable. Quantifiable categories of risks might have risk tolerance expressed in probabilistic terms, as described above. There should be similar statements related to each category of risk, and a full set of statements might be made for each combination of category of risk and business unit.

Non-quantifiable risks also require a clear statement of what is acceptable or not acceptable. For example, the Board has no tolerance for reputational damage due to risk crystallisation.

Ultimately, these statements of risk tolerance must be expressed in a way that can be easily understood and implemented by all staff within the organisation. In many cases this might manifest itself through statements of risk limits.

- (iii) In order to assess alternative strategies all of the chosen risk metrics, ideally including the balance sheet and profit and loss accounts, will need to be projected over the longer term. The insurance company will need to compare these before and after allowing for the anticipated volume of protection and unit linked business, and will need to ensure that no existing metrics are breached.

However, systems do need to be in place that will allow the company to adjust its risk tolerances and limits if these might inadvertently thwart beneficial alternative strategies.

Further, processes need to be in place regarding how to choose between alternative strategies if more than one falls within the acceptable risk appetite, e.g. maximising risk-adjusted economic value, or the one which closest matches a desired risk profile

The insurance company should also conduct scenario analysis for the status quo and alternative strategies to understand how the switch exposes the insurance company to new risks or increases/decreases exposure to current risks.

- (iv) The additional risks will depend on the exact policy features for the unit-linked business.

Market risks – typically equity price, property price and credit spread risks – relating to the fact that typically charges are deducted as a percentage of the funds under management. Any reduction in the prices of the underlying assets will reduce the funds under management and hence the fund management charge received.

Should the prices of the underlying assets not increase in line with expectations then it may not cover the costs incurred by the insurance company or take much longer for the insurance company to cover the costs associated with writing unit-linked business.

Poor market performance relative to competitors also introduces the risk of lower new business sales and higher lapses.

Operational risks related to the establishment of new processes within the insurance company (particularly unit pricing) and the training of staff to execute these processes.

Related to these operational risks, compliance risks may arise relating to the sale of a new type of product.

Additional product sellers in the market place or competitive action by existing sellers may put pressure on the volume of business sold and the profit margin. Where the volumes of business are less than anticipated then the expected expense allocation will be in error leading to expense risk, particularly the recovery of the product development and launch expenses.

- 4** (i) Pearson's correlation coefficient and the resulting correlation matrix can be calculated for any distributions of any type.

However, it is a linear concept based on moments meaning that it only fully defines the dependency structure when the cumulative distribution is elliptic as with, for example, the multivariate normal distribution and hence when the marginal distributions (post any transforms) are elliptically (e.g. normally) distributed.

The pricing actuary is most likely planning to use the correlation coefficients for forecasting and is therefore assuming that the underlying claim frequency and severity distributions (or total claims cost distributions) post any transforms are normally distributed.

- (ii) Directly from the data

- Provided the datasets are sufficiently large to produce credible estimates this should be the best approach as it reflects its own dependency structure.
- This data will be more recent than the industry data.
- This is likely to be the most time consuming approach.
- The actuary will need to consider the accuracy and completeness of the data.
- The actuary will need to consider the likely homogeneity of the data both over time
- and between the risks allocated to each line of business.
- The claims data will require transforming to take account of past inflation rates.
- If it is not possible to estimate the past levels of claims inflation from the data then the transformation will involve adopting external data.
- Results can be very sensitive to extreme outcomes.
- This is a common problem for heavy tailed distributions if the actuary has wrongly assumed that the claims cost distributions are normally distributed.
- The randomness associated with small data sets could result in the actuary adopting inappropriate correlation coefficients and as a consequence over or under pricing.

Using estimates published by other actuaries for various lines of business using various data sources over the years

- This information may not be produced regularly or recently.
- Nor is it subject to any particular standardisation. It is possible that the various studies may produce results which are not comparable to each other.
- It is even more likely that the various results might not be directly useful to the particular risks written by the insurer.
- Gathering the studies and interpreting them should be relatively straightforward.
- The adopted correlations can be reviewed and updated as more studies are released in the future.
- The studies are likely to use relatively larger datasets than the insurer's own datasets. The larger datasets may reduce the chance of randomness over-influencing the results.
- The studies do not offer any direct insight into the dependency structure of the insurer's own data. To the extent that the insurer's own experience is different from the published data then it would lead to over or under pricing.

Using scenario testing for a range of possible correlation coefficients

- Assuming that the claims distributions are sufficiently symmetrical (or can be transformed to be so) then subjecting the underlying claims distributions to a range of possible correlation matrices might be quite a useful way of estimating the range of likely capital needed for the underwriting side of the business.
- This approach might give some insight into potential claim costs under extreme conditions.
- Clearly, the ultimate choice from the range of results will be subjective and, to the extent that the choice is wrong, would lead to over or under pricing.
- The method does not offer any direct insight into the dependency structure of the insurer's own data.
- The method should be relatively straight-forward to employ.

$$(iii) \quad \rho_{\tau}(X, Y) = E(\text{sign}((X - \tilde{X})(Y - \tilde{Y})))$$

where (\tilde{X}, \tilde{Y}) is an independent copy of (X, Y) .

- Kendall's tau is a rank correlation measure.
- The numerical difference between the random variables (rvs) is discarded and only the sign of the difference is analysed.
- Kendall's tau works by generating pseudo distributions of the rvs being analysed which have the same distribution but are independent from one another.
- Kendall's tau is equal to the mean of the sign of the calculations of the product of the differences between the actual rvs and the pseudo rvs.
- Kendall's tau is a symmetric dependence measure
- between $[-1, 1]$.
- Independent rvs have a tau of zero
- but a tau of zero does not ensure independence.
- When tau equals 1, X and Y are comonotonic.
- When tau equals -1 , X and Y are countermonotonic.
- For continuous marginal distributions Kendall's tau depends only on the unique copula of the distributions
- and is thus invariant under strictly increasing transformations.

$$(iv) \quad \rho_{\tau}(X, Y) = \binom{40}{2}^{-1} \sum_{1 \leq t < s \leq 40} \text{sign}((X_t - X_s)(Y_t - Y_s))$$

$$(v) \quad \text{Assumed model } F(X, Y) = C_{\theta}(F(X), F(Y))$$

where θ is the single parameter to be estimated.

Let r^{τ} = sample value for Kendall's tau.

Solve $r^{\tau} = f(\hat{\theta})$ for $\hat{\theta}$ to estimate copula parameter θ .

For example Gumbel's copula is calibrated by taking

$$\hat{\theta} = (1 - r^{\tau})^{-1}$$

5 (i) ABC's credit risk exposure falls into three categories:

- The risk that the counterparty for a particular bond defaults on (all or some of) its payments resulting in a shortfall of cash or assets
- The risk that the counterparty for a particular bond is downgraded triggering a widening of spreads and a fall in the value of the asset relative to its liabilities
- Concentration risk: The risk should be considered across all the bond holdings with a particular counterparty rather than on an individual security basis.

The policy described above sets relative risk limits for default risk. It does not directly mitigate against any of these three risks at the point of sale which is when the risks are crystalising.

In fact it may actually make things worse as it would require ABC to sell bonds just after a downgrade when the market in that bond is already depressed triggering further falls.

The approach relies on credit rating agencies rating the bonds accurately and frequently.

(ii)

INFORMATION	SOURCE
Market value of equity (S_0)	Share price \times shares issued
Value of liabilities (B)	Balance sheet
Market value of assets (V_0)	Modelled – using option pricing
Asset volatility (σ_V)	Modelled – using option pricing
Default threshold (\tilde{B})	Liabilities due over the next year, including the “close out” cost of longer term liabilities
Leverage ratio of the firm (d)	Balance Sheet
Average coupon paid (c)	Past Financial Statements
Risk free rate	Market information
Time horizon	Determined by company

(iii) Advantages:

- Reacts quickly to changes in economic prospects of the company– faster than rating agencies
- More sensitive to current macroeconomic conditions than historic default / transition probabilities
- Accommodates complex liability structures

Disadvantages:

- Only applies to firms with publicly traded stock
- Sensitive to any over- or under- valuation in equity markets

(iv)

- ABC pays fees/premiums to bank GHI
- If the defined event does not occur then no other payments are made.
- A payment is made from bank GHI to ABC if a credit default event on the company DEF corporate bond occurs.
- The amount of the payment is the difference between the original price of the DEF bond and the recovery value of that bond.
- Alternatively, the CDS may be settled physically i.e. ABC receives the full agreed value of the bonds from bank GHI and ABC gives the defaulted bond to bank GHI.
- The payment from GHI to ABC is subject also to the risk of GHI defaulting.

(v) The credit spread is made up of four parts:

- The expected probability of default and the expected loss given default
- A risk premium covering unexpected defaults
- A liquidity premium
- A premium for structural/documentation complexity

The CDS only estimates the first of these components. Deducting the cost of the CDS from the spread will not help us distinguish between the remaining parts.

Also reflected in the CDS price will be the profit loadings and credit risk premiums attached to the counterparty GHI causing a further distortion.

The profit loading includes the risk premium for covering unexpected defaults.

A CDS may not be available for a large number of assets in the portfolio.

Different banks will have different profit loadings.

6 (i)

- 99.5% is potentially misleading due to model risk and estimation error.
- The failure could have been due to a 1 in 200 year event or worse, i.e. in excess of the 99.5% confidence interval.
- The cause of the failure should be analysed carefully.
- It might have resulted from an event which was not previously considered to be a risk and so was not being modelled by either the banks or the regulator.
- For example, a court ruling imposing new and unexpected liability on the bank to cease trading a given product type and to repay all losses to customers who had lost moneys but not be entitled to collect profits from customers who had made money.
- The bankrupt bank's model could have been flawed. It could have underestimated one or more risks.
- The bankrupt bank's model variance could have been insufficient.
- The regulator should consider whether there have been other bank failures or near failures in past years.
- Regulation should be proportionate. Consumer protection and confidence in the system doesn't mean that individual banks should not be able to or allowed to fail.
- The regulator could consider augmenting the current approach with additional requirements e.g. governance, reporting requirements etc..

(ii) Politician & Regulator

The politician wants more capital, safer banking and customer protection.

The regulator wants safe banking and customer protection but also appreciates that too high a capital requirement could lead to banks boosting charges to customers and/or banks finding it difficult to attract capital and/or domestic banks being forced out of business by foreign banks with lower capital requirements.

The choice of 99.5% probability of sufficiency in the ensuing year has been made by the regulator. It is subjective and could have been higher or lower. The regulator will not want to make changes unless it is proven to be warranted.

The “five times” suggestion is a simple and relatively practical solution given the subjectivity of measuring the extreme ends of a distribution. Moments are very sensitive to extreme values and hence the extreme end of a distribution can be driven by very few extreme observations. The standard deviation is a more robust calculation than the 99.5% probability of sufficiency statistic.

Model

The suggestion is still internal model dependent and as such the model errors and parameter errors remain.

The capital models will still produce both the standard deviation and the 99.5% probability of sufficiency.

Hence, regardless of how the minimum capital is being calculated, the regulator will be able to continue to compare the minimum capital under both approaches.

Industry

The five times standard deviation suggestion could result in a very high capital charge. For example the 99.5% probability of sufficiency under a normal distribution is circa three times the standard deviation. Assuming that the suggestion results in a dramatic increase in capital then it will inevitably impact the banks' profits and the prices that they charge to customers.

Moving to the “five times” suggestion removes some of the individual nature of the capital requirement. Some banks may spend less time focussing on the extreme end of their enterprise capital model.

This may reduce the accuracy of the banks' capital models at the extreme end. It may also result in the banks taking more risk without first testing the consequence of the risk taking in their models.

The “five times” suggestion may encourage previously less risky banks to become more risky in an effort to utilise the additional capital. Conversely it is unlikely to encourage the historically more risky banks to become less risky. As a consequence the banking industry as a whole might be expected to adopt more risky strategies if the “five times” method was introduced.

If all of the banks operate with very similar risk profiles then the suggestion may be equivalent to setting the probability of sufficiency (currently 99.5%) to be equivalent to five times the standard deviation.

The regulator should consider whether the “five times” approach would have saved the failed bank and if so was it just due to the additional capital or was it due to the 99.5% estimate being flawed in the case of the failed bank.

7

(i)

- Basis risk – the risk that the change in the value of the swaps will not match exactly the change in value of the assets.
- Credit/Counterparty risk – the risk that the investment bank will become insolvent.
- Operational risk – the risk that the payments due under the swaps will not be made, or will not be made at the correct level.
- Longevity/demographic risk – this is the risk that the pensioners will live longer than anticipated by the swap.
- Reinvestment risk – Interest rate / inflation residual risks on the roll forward of swaps if not available for a long enough term.

(ii)

Basis risk:

Can be mitigated by ensuring that a sufficiently complex swap overlay is put in place both in terms of the various swap terms and underlying payment types.

Credit/counterparty risk:

- Can be limited through the use of collateralisation.
- Can be mitigated by the right to call for additional collateral should the counterparty's credit rating fall below a stated level.
- Can be mitigated by having the swaps automatically unwind should the counterparty's credit rating fall below a stated level.
- Can be mitigated by ensuring full netting arrangements are in place between the counterparties.
- Can be mitigated by only dealing with banks with a minimum credit rating issued by one or more acceptable credit rating agencies.
- Can be mitigated by using multiple counterparties to limit the exposure to any one counterparty.
- Can be mitigated by buying credit protection from another counterparty.

Operational risk can be limited by ensuring that appropriate systems and controls are put in place. An external assessment of the adequacy of these systems could be sought.

Longevity swaps can also be added to limit demographic risk.

- (iii) The comparison might have excluded the costs of restructuring the physical portfolio (buy/sell costs).

The comparison might have used hypothetical market values (mark to model) for both the current portfolio and the new hypothetical portfolio. These prices and quantities may not be available increasing the time and cost of restructuring the physical portfolio.

It might be that the trustees wish to regularly change the interest rate and inflation rate profiles of the asset portfolio. The transaction costs of trading swaps can be lower than those for trading bonds.

Also the time taken to trade swaps may be much less than the time to buy and sell the physical assets.

It might be felt that the current asset portfolio has relatively greater chance of potential future gains from contracting credit and liquidity spreads as the portfolio duration shortens than would be the case for the hypothetical portfolio.

Trading the physical assets may crystallise taxable gains.

The bank trading the swaps may be able to supply other services at a reduced cost.

The physical assets may not be available in the requisite amounts which affect both the price and time taken to trade the physical assets.

The pension fund might not be allowed by scheme rules or regulation to hold some of the hypothetical assets.

- (iv) One alternative is to sell and repurchase all of the larger holdings in liquid bonds. The transaction costs would be relatively low. The trustees would then need fewer (relatively expensive) swaps.

Whilst the initial cost of moving the portfolio might be relatively high, if the assets are then held to maturity, the pension scheme might benefit from credit and liquidity premiums.

Repositioning the physical portfolio takes time and money and so only makes sense for the part of the portfolio that the pension scheme wants to hold for a long period of time.

The trustees could explore the potential for using other derivatives such as interest rate caps and floors. It may be possible to achieve the required profiles at a lower cost.

END OF EXAMINERS' REPORT