

INSTITUTE AND FACULTY OF ACTUARIES

EXAMINERS' REPORT

April 2011 examinations

Subject ST9 — Enterprise Risk Management Specialist Technical

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

July 2011

- 1** (i) The current approach is simple to understand and helps to maintain the independence of the risk management function.

However, the current approach may engender a confrontational relationship between the front office and risk management function.

The current approach is focussed on the extremes, being profit maximisation and loss minimisation. This is likely to result in the bank taking a combination of large risks and virtually no risks. This is highly unlikely to produce an optimal return for a given level of risk. In other words this approach is virtually certain to be detrimental to the performance of the bank overall.

Effective enterprise risk management requires a holistic approach taking account of the bank's objectives, return requirements and risk limits. This approach is not holistic. It seeks to maximise profits without regard to risk and alternatively to remove all risks without regard for the cost to profit of doing so.

- (ii) Instead of using an "offense and defence" approach to organising the front office and risk management function, the bank could switch to a "policy and policing" or "partnership" approach.

The policy and policing approach involves the risk management function establishing policies that set out the parameters within which the front office must operate. The front office's activities are then monitored by the risk management function for compliance with the policy.

The advantage of this approach versus the offense and defence approach is that it is less confrontational instead mirroring the roles of policeman and citizen.

However, there are numerous disadvantages with the approach. These include the following:

- The risk management function is not engaged in the day to day operations, since the front office are free to proceed as they wish provided they remain compliant with the policies.
- As a result, the risk management function may become less knowledgeable on the front office's activities as these evolve.
- Consequently, the risk policies may become out of date with new policies not established in a timely fashion.
- Similarly, the front office may become less knowledgeable on the latest risk management techniques.

- The front office may be tempted to hide an infringement of the risk policies, deliberate or otherwise, particularly where there is uncertainty as to whether a risk policy has been breached.
- The separation of the policy and policing may cause delays in the identification of and quantification of new risks as they emerge in the front office.

The partnership approach involves the front office and risk management function working collaboratively to resolve risk management issues with shared objectives.

There are numerous advantages with the approach. These include the following:

- The advantage of this approach versus the offense and defence and policy and policing approaches is that the front office and risk management function work in partnership with shared performance objectives.
- Application of risk management is more timely executed as soon as risks crystallise and in the day to day decision making process – e.g. in loan and mortgage pricing.
- Working collaboratively should result in the front office having a better appreciation of the latest risk management techniques, whilst the risk management function will be more knowledgeable on the front office's activities.

The main disadvantage is the loss of independence of the risk management function.

This may be remedied through blending the approaches described above as these are not mutually exclusive.

[It is axiomatic that an advantage of one approach would be a disadvantage in another approach. To gain full marks the advantages and disadvantages need be described with reference to the given approach.]

2 (i) Unbundling risk origination, retention and transfer...

By unbundling companies can identify their core competencies and decide where in the value chain they wish to compete

Providing a risk aggregation function across the company...

Helps clarify which risks are managed through the corporate and strategic management process...and which are managed centrally for the group

Setting risk limits and asset allocation targets...to provide complimentary targets for balancing risk and return in the business

Influencing transfer pricing, capital allocation and investment decisions...

Effectively creating an internal capital market to allocate capital and decisions accordingly.

May permit greater analysis of how risks interact with each other / reduce concentration risk

Keeping up to date with emerging risks / avoiding crisis management

Greater respect from regulators/credit rating agencies/analysts

More effective risk mitigation/transfer strategies

- (ii) Risk Capacity – An international conglomerate will have different types of businesses, of different sizes, operating in different countries with different languages, legislation and potentially regulation. This will create communication problems, difficulties in consistently identifying and quantifying risks, reporting delays and inaccuracies.

Furthermore, the group's overall risk capacity will likely be difficult to set. Actual capital will of course be known but the economic capital will depend on the accumulation of risks across company, industry, country, currency etc.. The diversification benefits in the group economic capital will then make it difficult and in part subjective to allocate the capital and hence risk capacity to individual companies.

Risk Profile – The risk exposures will be changing over time. It is important that all of the companies take a consistent, diligent and timely approach to the maintenance of the risk profile. This will be difficult as the approach to risk management often varies by size of company, industry, country.

Risk Tolerance – To be effective, each company should set its own risk tolerance (statements). These statements may well be incompatible at the group level. The group will need to overlay a central set of risk tolerance statements to manage these inconsistencies.

Risk Limits – Risk limits are relatively straightforward to set for a standalone company. In this case each company's risk limits need to be set having regard to the group's diversification benefits. This is both complicated and partly subjective.

Risk Appetite – Risk appetite will be set at the group level and will essentially state the variation in profit that the group is willing to accept in return for its target profit. In this case it will be more difficult to monitor actual to budget profit across the group due to currency fluctuations, reporting delays and changing risk profiles across the group companies.

[This question was generally well answered. A small number of candidates showed that they were slightly confused as to the differences between risk tolerance, risk

limits and risk appetite. Put simply, risk appetite is the risk that someone wants to take (i.e. target), risk tolerance is typically a set of risks that one recognises will need to be taken at least to some degree in order to maintain the targeted risk appetite and risk limits are maximum acceptable levels of identified risks.]

- 3** (i) A QQ (quantile-quantile) plot is a plot of the quantiles of two distributions against each other. In this case there is sample data. The number of quantiles might be set to be equal to the number of data points or something similar. The quantiles are non-decreasing and there is an equal probability of being in any given quantile. Normally the sample data quantiles on the y-axis are plotted against the theoretical reference distribution quantiles on the x-axis. Of course, one first needs to decide on an appropriate theoretical distribution.

For a time series the time dependency of the sample data is removed by measuring the variation in the observed to modelled trend. This variation is then plotted against modelled variability.

- (ii) Always non-decreasing when viewed left to right.

If the two distributions are identical then the QQ plot will follow the $y = x$ line.

If the two distributions agree after linear transformation then the line will be straight but not $y = x$.

If the QQ plot is relatively flatter than $y = x$ line then the distribution plotted on the horizontal is more dispersed. Conversely, if the QQ plot is steeper than the $y = x$ line then the distribution plotted on the vertical is more dispersed.

An S shaped QQ plot indicates that one distribution is more skew than another and/or fatter tailed.

The fit is often good for the bulk of the distribution and less good at the extremes.

Interpreting a completely messy and seemingly awful fit might be due to the poor choice of the underlying theoretical distribution.

- (iii) Histograms with superimposed fitted density functions
Empirical cumulative distribution functions (CDFs) with superimposed fitted CDFs
Plotting the residuals i.e. actual minus fitted.

- (iv) Two data points represents 10% of the total sample from which the model is being built. It is difficult to argue that 10% of the data are outliers.

Removing the two most extreme values will likely reduce the volatility estimates and quite possibly overly so. The model may underestimate future volatility.

The model may underestimate both the frequency and severity of the relatively more extreme outcomes.

The parameter error will increase if the outliers are removed.

The data for the outliers should be checked to ensure they are accurate.

[Most candidates were not able to demonstrate that they could interpret a QQ plot in detail. Several candidates failed to be concerned that removing 10% of the data would create several problems with the fitting exercise.]

- 4** (i) Corporate bonds are not re-rated very often. The rating may be out of date. This could be due to company specific changes over time or due to systemic changes such as economic changes.

The rating agencies may not publish estimated probabilities of default and losses given default very often, meaning that they too may be out of date.

Also, different rating agencies will have different views meaning that the fund manager will need to choose one rating agency over another in some cases.

Whilst some bonds' probability of default may be independent of some others it is unlikely that all of the bonds are independent from one another. Empirical data shows that companies' returns and share prices are often positively correlated and particularly so in adverse times when they would be more likely to default on their debt.

The independence assumption for the losses given default is also unlikely to be true and particularly in times of large numbers of default.

The rating is a single best estimate and doesn't include variation.

- (ii) As stated in the question, a corporate bond is likely to default when the share price falls to zero or perhaps near to zero.

The more volatile a company's share price the greater the chance it has of reducing to zero. The probability of default of a corporate bond is largely driven by both its current price and the volatility in the share price.

The share price model would seek to forecast future share prices based on the current share price, past observed volatility and the dependence structure between the various companies' share prices.

The volatility estimates and dependence structures are key assumptions. Market volatility changes over time. Sometimes these changes can be both large and rapid. Hence the share price model will need to be stress tested for changing volatility and dependence.

The dependence structure is best estimated using a copula as empirical data suggests that share price probability distributions are not normally distributed.

[Many candidates handled the question well. Those that did not failed to develop a model that estimated the probability of a company's share price falling to near zero in a given time horizon.]

- 5** (i) The strong ERM classification requires that DEF can evidence that it has strong capabilities to consistently identify, measure, and manage risk exposures and losses within the enterprise's predetermined tolerance guidelines.

DEF will be able to evidence that it has processes in place seeking to optimise risk-adjusted returns.

Risk and risk management will usually be important considerations in DEF's corporate decision making.

- (ii) Each subsidiary will have to meet local regulatory capital requirements.

This makes it preferable to first calculate capital for each subsidiary before aggregating it for the group.

It is unlikely that capital can be transferred between entities quickly in response to a risk event.

This lack of fungibility must also be incorporated in the aggregation calculation.

Nevertheless, the group is still likely to be able to benefit from diversification. The overall diversification credit will be capable of being estimated.

However, it will be difficult, partially subjective and sometimes not possible to pass this diversification credit down to the individual companies.

This is because each regulator is likely to require that the capital in the regulated entity be estimated on a stand alone basis.

In order to aggregate the various subsidiaries' capital requirements into a group economic capital model it will be necessary to allow for fungibility and diversification credits.

It will also be necessary to allow for currency fluctuations.

This complexity is likely to mean that a simulation model will be needed.

It is difficult to ensure consistent calibration (to the specified confidence level) between entities since, for example, they will have different risk volatilities from country to country.

Also some territories will be better than others with respect to data collection and modelling expertise.

- (iii) One option would be to recalculate the capital required at group level. Most likely a stochastic model would be used.

This could be done by reflecting the dependencies in the diversification structure.

These reputational group risks effect the correlations between entities (in fact, these concentrations can be thought of as the opposite of diversification benefits).

It would be necessary to model both the frequency and severity distributions.

The results may then be allocated to each subsidiary and capital held at local level.

Alternatively the capital required could be calculated directly at subsidiary level.

These reputational group risks could be considered as a special case of operational risk and included in the operational risk capital.

Alternatively, at group level scenario analysis could be carried out to estimate the adequacy of capital held in each subsidiary

And the capital could be topped up as necessary.

[There were a number of different and acceptable answers to parts (ii) and (iii).]

- 6** (i) The risk report should be in keeping with the content and depth of the financial statements. In other words it should produce information at a group level.

It should include:

- A risk appetite statement
- A statement describing the risk management process of the company noting any key changes to the process over the reporting period. The statement should include the methods for identifying, monitoring, quantifying, transferring, mitigating and reporting risks.
- A table of risk categorisations.
- Top 5 risks ordered by reasonably foreseeable economic impact. The table should include a description of the risk, its relative chance of occurring and the estimated range of net cost to the group should it occur.

- Top 5 risks ordered by likelihood of occurrence over the next 12 months (for risks which are deemed to have a large potential economic impact). The table should include a description of the risk, its relative chance of occurring and the estimated range of net cost to the group should it occur.
- Commentary on any changes to the risk lists since the last report.
- A statement comparing large risk exposures to stated risk tolerances
- A statement describing the risk modelling (quantification) process including the validation process.
- A statement describing the main operational risks and the systems in place to mitigate or transfer them. Corporate governance, HR rules, Compliance Dept, management structure and oversight, internal reporting to mitigate risks and insurance and outsourcing to transfer them. This statement could reference the risk mitigation strategies in place for the top 5 risks mentioned if any.
- A statement of risk crystallisations / risk events which happened over the past period.

(ii) Information From Risk Reports

- Frequency and severity estimates for the major risks.
- The risk description in the risk list can be compared across companies (same industry, same size).
- The description of mitigation strategies for operational risks will help to decide on the company's exposure to operational risks.

Information from the Financial Statements

- The net profit
- Any components in the net profit which arise from risk costs. These should be added back to estimate the return before risk costs.
- Average shareholders' equity during the period (adjusted for rights issues, bonus issues etc..)

Would also need to use the average risk free rate of return (say 3 month government debt return) for the purpose of calculating the risk adjusted ROE

The risk reports will be different between companies and over time.

Adjust all of the risk reports to cover the same time horizon (probably one year)

A perusal of all of the risk reports may give insight into typical risks of companies, by industry, by market cap, overall.

The risk report will provide frequency and severity information for the largest risks. Use this information to estimate a risk cost distribution for the company. Judgmentally, gross up this risk cost distribution. The gross up factor would depend on the extent of the risk list in the given company's risk report and the comparable risk lists of other companies in the same industry.

Overlay a (judgmental) risk factor to allow for operational risk.

The percentile chosen of the risk cost distribution will depend on the purpose of the risk adjusted return table. For example, use the mean risk cost to estimate each company's expected risk adjusted ROE in the upcoming 12 months. Alternatively use the 75th percentile to estimate the sensitivity of each company's net profit to a changing risk outcome.

Risk adjusted ROE = (Pre-dividend net of tax profits adjusted for risk based losses minus mean risk cost minus risk free rate of return) divided by shareholders' equity during the period adjusted for any rights issues or bonus issues.

[This question was intended to be difficult. Many candidates noted that the risk report should include some sort of quantification of the company's key risks at the time of each report. A smaller number of candidates realised that these key risks would likely change over time and hence could be compared over time for a given company and across different companies in the same or different industries.]

- 7 (i) The generator function for this Gumbel copula is $(-\ln F(x))^2$.

We are interested here in $x = -0.6$, for which $F(x) = 0.1$.

So the generated result for a single loss is $(-\ln(0.1))^2 = 5.3019$.

Multiply this by 10 to allow for the number of bonds in the portfolio, giving 53.019.

Apply the pseudo inverse of the generator function (square root, take the negative, take the exponential) to obtain 0.000688 or 0.0688%.

- (ii) The correlation coefficients indicate the overall level of dependence.

A copula describes the shape of this relationship.

Both are important in describing the relationship between variables.

- (iii) The Gumbel copula has upper tail dependence. This is appropriate if you think that extreme positive returns are likely to be more closely related than extreme negative ones. For a high-yield bond portfolio, the opposite is more

likely to be the case. This suggests that a copula such as the Clayton copula is more likely to be appropriate. However, as single-parameter Archimedean copulas, the Gumbel and Clayton copulas can be parameterised only with a single variable. This means that there is an implicit assumption that the shape and degree of correlation between each bond is assumed to be identical, which might not be the case. A wider range of relationships could be described by an elliptical copula. In particular, a t-copula could also be used to describe data with fat tails.

[This question required candidates to be relatively familiar with the calculation of a copula.]

- 8** (i) A CDO is formed by setting up an investment entity known as a special purpose vehicle (SPV). This is used to purchase a portfolio of bonds, mortgages, credit derivatives or, in this case, mortality-related cash flows.

The money used to purchase these cash flows comes from external investors. The money from the investors is used by the SPV to make payments to the life insurance company, which then pays any claims on the term assurance policies it has written. This means that the greater the number of deaths, the greater the level of payments from the SPV to the insurance company and the smaller the funds left for the investors.

These investors can purchase different classes of share in the SPV, each of which receives returns from the SPV. The riskiest tranche of shares – known as the equity tranche – suffers the full impact if the mortality rate is higher than expected. In other words, for the first term assurance claims, only holders of the equity tranche suffer a reduction in their income stream. However, to compensate for this increased risk these investors have the highest expected returns relative to their initial investment. At the other end of the scale the safest tranche of shares does not suffer the impact of any term assurance claims until all of the funds allocated to lower tranches have been exhausted through prior claims. The high level of security means that investors in this tranche have the lowest expected return. For this reason, it is known as the senior or super-senior tranche. In the middle, with a moderate level of both risk and return, is the mezzanine tranche.

The aggregate loss at which the payments on a particular tranche start to reduce are often defined by attachment points. The returns for investors in a particular tranche can therefore be defined as follows:

- if the loss for the portfolio as a whole is less than the attachment point for this tranche, then the investor will receive the maximum possible investment;
- if the loss is greater than the attachment point for the next most senior tranche – which can also be regarded as the detachment point for the investor's own tranche – then the investor suffers a total loss; and

- if the loss is between these two points, then the return to the investor is the fund value less the detachment point.

Other methods to share loss include coupon reduction or default and different approaches to allocating partial capital losses of the various tranches.

The return received for each tranche is in return for an initial investment. The total investment over all tranches must equal the total initial value of the fund, but the greater the investment required for investment in a particular tranche the lower the potential return for that tranche. The return for each tranche therefore depends on both the attachment points and the initial investment required from investors in each tranche.

The attachment points and levels of investment are determined using quantitative models that are frequently agreed with credit rating agencies. This means that the tranches themselves get credit ratings.

- (ii) Both will provide protection against extreme losses.

The CDO is described as only covering losses up to a certain limit; the reinsurance may not have a limit. The ultimate payout will depend on the specific terms and conditions of the contracts.

This will reduce the amount of capital required to write business if the protection reduces the amount of capital required under Solvency II – in other words, if it reduces the SCR and/or the ORSA.

However, depending on the premiums paid, the protection may reduce the expected level of profit.

The extent to which this is true will depend on whether the cost of the reinsurance or the profit foregone with the catastrophe bond exceeds the additional profit that can be generated from the increase in the volume of business written.

With the reinsurance, there is credit risk from the possibility of reinsurer failure. This can be managed by ensuring that the reinsurer has a sufficiently high credit rating. This risk does not exist with the catastrophe bond as the capital is set aside in an SPV for the term of the bond.

However, in this instance there is instead the risk that too high an expected return might need to be embedded in the bond to make it attractive to investors. Such a bond might not be particularly liquid, so a higher premium might be needed as compensation. Investors might also be concerned about asymmetry of information, since the insurer would know more about the lives being placed in the bond. This is also potentially true with reinsurance, but the reinsurer is more likely to have seen past claim data so will be more confident on the pricing.

The CDO structure could make the product more attractive to a wider range of investors.

Market capacity is more likely to be known and stable for the reinsurance deal, but the market for catastrophe bonds may be less well developed and there may or may not be appetite for this particular issue.

There may not be a sufficiently great economy of scale to enter into the CDO arrangement.

There may be different levels of administrative burden. The reinsurance deal is likely to be relatively straightforward to enter into and broker, but will require regular information feeds to/from the reinsurer thereafter. The CDO is likely to take more effort initially with the design and getting it to market, but there may be less ongoing administrative burden.

There may be potential tax differences in the jurisdiction involved.

There may be differences relating to the flexibility regarding adaptation to writing new business, so that the company can continue to manage its mortality risk on an open book, rather than just on the existing portfolio. For example, with the reinsurance it may be possible to keep the treaty open to new business.

The reinsurer may provide extra benefits such as data, pricing advice and underwriting expertise.

[This question was intended to have difficult components. Part (i) was generally well answered. Many candidates' answers to part (ii) appeared to be a little rushed. They were quite superficial. For example, many candidates failed to reflect on the likely differences in terms and conditions between the two contracts.]

- 9** A number of options are available to Scilly Sandals plc to help ease the severe liquidity pressures.

Whatever the company does to address the short term issue it should endeavour at the same time to reduce the chance of the liquidity strain recurring in the future.

A cash buffer may have been set aside on the balance sheet so as to help the company manage through liquidity pressures. However, the cash buffer may have been utilised already, since the liquidity pressures now facing the company are described as severe. Therefore, further more severe actions to ease the liquidity pressures may now need to be considered.

Actions relating to the company's working capital may be undertaken so as to maximise short term cashflow.

- The company could sell surplus stock, although the value achieved for the surplus stock may depend on the season due to the seasonal nature of the products. In

addition, there may be longer term pricing and brand implications where aggressive pricing is required to shift the surplus stock. Finally, the sale of surplus stock may take some time, particularly where this has not been undertaken previously so that established channels for the sale of surplus stock are not readily available.

- The company could monetise receivables, although there will be cost implications due to the third party's interest charges and other costs. As with the sale of surplus stock, the monetising of receivables may take some time, particularly where this has not been undertaken previously so that established channels for monetising of receivables are not readily available. It might be simplest for the company to first chase debtors to see if any will pay.
- The company may delay payments to suppliers, although this may negatively impact upon the relationship of the company with its suppliers. Worse still, the delaying of payments to suppliers may lead to alarm spreading amongst the company's suppliers regarding the company's financial health. This may in turn result in suppliers becoming unwilling to ship goods.
- The company could sell assets. This might generate cash quickly but some of the assets might need to be repurchased in the future at possibly higher cost.
- The company could reduce expenditure eg. Raw materials and/or staff costs. This action may reduce the quality of the company's goods and lower staff morale.

Actions relating to the company's capital expenditure may be undertaken so as to maximise short term cashflow.

- The company may temporarily reduce the planned maintenance expenditure – e.g. spending on refurbishing the head office could be delayed. However, reduced maintenance spending may not be possible or may take time to implement where contractual agreements have been entered into or where there are regulatory requirements.
- The company may temporarily reduce investment in new products from planned levels. However, reduced investment in new products may not be possible where the company has entered into contractual agreements – e.g. official flip flop supplier to the 2012 Olympic Games.

Actions relating to the capital markets and the company's bank credit may be undertaken so as to access further or at least preserve capital.

- The company may defer dividends or pay scrip dividends. However, such measures may send a distress signal to the capital markets with possible implications for the company's ability to access further capital. Similar comments apply where dividends on preference shares are deferred with such deferrals potentially creating a liability on the company's balance sheet.
- The company may tap capital markets for further capital utilising the debt or equity markets. However, with the severe liquidity pressures being experienced

the company may already be experiencing rolling over commercial paper let alone accessing further capital.

- The company may be able to negotiate credit facilities with its bank. However, unless there are committed lines of credit the company may find the door closed where the bank is not confident in the ability of management to negotiate the severe liquidity pressures.

Finally, where the company has a defined benefit pension scheme the company may be able to negotiate the delay of payments into the scheme with the trustees. This is most likely to be successful where the scheme is in surplus and where the trustees are confident in the ability of management to negotiate the severe liquidity pressures.

[The key point of this question was that Scilly Sandals had an immediate issue that needed immediate or at least near term solutions. Several candidates focussed on relatively long term solutions which would not have worked in time to save the company.]

END OF EXAMINERS' REPORT