INSTITUTE OF ACTUARIES OF INDIA

Subject SP9 - Enterprise Risk Management

November 2025 Examination

INDICATIVE SOLUTION

Introduction:

The indicative solution has been written by the Examiners with the aim of helping candidates. The solutions given are only indicative. It is realized that there could be other points as valid answers and examiners have given credit for any alternative approach or interpretation which they consider to be reasonable

Solution 1:

i) Risks arising from business:

Since JSS's main business is term insurance products, the main category of risk is likely to be insurance risk, specifically mortality risk.

An adverse mortality experience could be a result of:

- a generally higher than expected population mortality rates (level risk),
- uncertainty due to having a finite pool of policies (volatility risk), or
- a one-off event such as a pandemic (catastrophe risk).

Other relevant and potentially important insurance risks include:

- persistency risk (particularly relevant early on where new business strain is a factor)
- expense risk

Risks arising from investments:

Since JSS's main business is term insurance products, their main investments are likely to be in bonds (either government or corporate) and deposits.

As such, the main category of risks is likely to be interest rate risk and credit/counterparty risk.

Interest rate risk may arise in various ways such as:

- Drop in interest rate leading to depressed bond prices / lower returns on floating rate bonds / deposits
- Future premiums received invested at lower-than-expected interest rates

Credit / counterparty rate risk may arise in various ways such as:

- Default by issuer of (mostly corporate) bonds
- Counterparties such as reinsurers failing to meet obligations [4 marks]

ii) Legal risk is the risk arising from the understanding of and adherence to legislation, including changes in accepted interpretation.

The three ways that it might arise are as follows:

1. There may be new legislation in response to political or social pressures, which could impact on JSS either positively or negatively.
   Example: A decrease in the statutory capital requirement for life insurers may usher a multitude of new entrants, thereby squeezing margins and affecting profitability.
2. Provisions in important contracts could give rise to significant problems if certain circumstances were to occur.
   Example: If there are clauses which might make reinsurance contracts voidable at the option of the reinsurer under certain circumstances and should those circumstances occur, JSS might be left footing a claims bill much larger than expected.
3. Adverse court / legal judgements can affect JSS significantly.
   Example: Ambiguous wordings may be interpreted in a way that favours the insured. OR Courts may choose interpretations that effectively extend coverage to include risks not purported to be covered. [4 marks]

iii) The following are some of the ways to manage / control operational risk:

- Outsourcing non-core operations such as IT or call centre
  - Tackles: People / process / system risk
  - Residual risk: Possible failure of third-party to deliver / reduced control
- Background verification (BGV) of prospective employees
  - Tackles: People risk
  - Residual risk: Possibly delayed onboarding of new joiners
- Horizon scanning
  - Tackles: Regulatory / legal / event risk
  - Residual risk: Occasional over-reaction to noise (rather than signal)
- Well-tested and regularly updated business continuity plan (BCP)
  - Tackles: Reputational / External event risk
  - Residual risk: Costly, if not proportionate
- Structured approach to change management
  - Tackles: Technology / process / system risk
  - Residual risk: Testing of edge cases / interactions missed out
- Strong relationships with key stakeholders
  - Tackles: Reputational / People / regulatory risk
  - Residual risk: Agency risk
- Employee training, checks and balances
  - Tackles: People / compliance risk
  - Residual risk: Reduced to a checkbox exercise
- Independent and strong validation and audit functions
  - Tackles: Crime / people / process / model risk
  - Residual risk: Slower pace of things moving
- Having a sound and well-integrated ERM framework
  - Tackles: All risks, especially reputational
  - Residual risk: Minimal [10 marks]

iv) It is a good idea to set up a risk sub-committee for various reasons such as:

- Ensuring adequate focus on risk management, which may not be possible to achieve in full Board meetings with multiple agenda items
- To provide for an increased role to independent / non-executive directors (and possibly external experts), who may have a greater representation on this sub-committee
- To develop specific Board members as experts / specialists in risk management

This charter should include:

- Clearly laid out roles and responsibilities of the sub-committee
- Its membership (including split of independent and non-independent directors)
- Frequency of meetings
- Criteria for performance assessment
- Resources available (including internal deparatments as well as external consultants, if any) [5 marks]

v) The main reason for getting JSS rated by a CRA would be to (hopefully) obtain a good credit rating on the strength of which JSS may be able to borrow more cheaply, thereby reducing its cost of capital.
However, the rating may also serve additional purposes such as:

- Providing investors with a broad indication of the creditworthiness of the company
- Inspiring trust among other stakeholders such as regulators, counterparties, reinsurers, etc. [2 marks]

vi) The five main areas that are evaluated by S&P to assess the ERM capability of a company like JSS are:

- Risk management culture
- Risk control
- Extreme event management (as part of risk control)
- Risk models and capital models
- Strategic risk management

On the basis of these areas, S&P will grade JSS's ERM capability as falling in one of the categories: weak, adequate, strong or excellent.
This grading then feeds into the overall credit rating of JSS through the Management & Governance component…
… depending on the:

- Complexity of the risks faced by JSS
- Amount of available capital and ease of access to it (i.e. capability to absorb risks) [4 marks]

vii) Risk appetite is the degree of risk that JSS an organization is willing to accept in order to achieve its objectives.

Examples for JSS:

1. JSS is prepared to lose a maximum of $500m (in terms of earnings) with a probability of no more than 0.5% over the next 12 months.
2. JSS's solvency ratio should stay above 175% with a 99.5% probability over the next three years.

Risk tolerance refers to a set of statements (more detailed than risk appetite, and typically broken down by risk categories or business units) as to how much risk a company wants to take.

Examples for JSS:

1. JSS will tolerate a probability of no more than 0.5% that losses attributable to interest rate risk should exceed $100m over the next 12 months.
2. No individual with a criminal record shall be employed with JSS – whether as an employee, whether full-time or part-time or a contractor.

Risk limits translate the risk tolerance statements into operational limits or restrictions that can be used by the business on a day-to-day basis.

Examples for JSS:

1. JSS may not have more than 2.5% of its tangible assets in corporate bonds of a single issuer.
2. No more than 4 senior executives (defined as CXOs and heads of department) of JSS may fly together on an aircraft / ship. [7 marks]

viii) An ARIMA model has three orders (p, d, q) corresponding to the AR (auto-regressive), I (integrated) and MA (moving average) components of the model.

Firstly, we check the given time series data X(t) for stationarity. Stationarity essentially means the statistical properties of the process stay the same over time. If the time series is stationary, then d = 0 and an ARMA model would suffice.
If the time-series is not stationary, we difference the time series repeatedly until we get a stationary time series. The number of times we had to difference the time series until we get a stationary time series is the order d. [We then consider this stationary, differenced time series for the rest of the discussion.]

Next, we look at the correlograms of the stationary time series, which are plots of the:

- Sample autocorrelation function (sample ACF) – correlation between the process values at lag (h), i.e. corr(X(t),X(t-h))
- Sample partial autocorrelation function (sample PACF) – conditional correlation of X(t) with X(t-h) given the times series values X(t-1), X(t-2), … X(t-h+1)

The way to use the Sample ACF and Sample PACF is as follows:

- If Sample ACF tails off but Sample PACF cuts off after lag p, then q = 0 and AR(p) model is appropriate.
- If Sample PACF tails off but Sample ACF cuts off after lag q, then p = 0 and MA(q) model is appropriate.
- If both Sample ACF and Sample PACF tail off but with kinks at q and p respectively, then ARMA(p,q) model is appropriate. [5 marks]

ix) To test whether a time series model such as ARIMA is a good fit, we can look at the residuals. These residuals are expected to behave like white noise, i.e. close to 0 in value and uncorrelated with each other (random, without any pattern).

Some of the tests to check if the residuals are white noise include:

- A visual inspection of a plot of residuals against time [to check if they are close to zero and without patterns]
- A plot of the sample ACF of the residuals [expected to be close to zero]
- The turning point test where the null hypothesis is that the graph of residuals is patternless
- The Ljung-Box portmanteau test where the null hypothesis is that there is no correlation between residuals [2 marks]

x) The likelihood ratio (LR) test is a test under the hypothesis testing framework to compare nested models only. Two models are said to be nested if the second contains all the independent variables of the first and some additional ones.
The LR test checks whether the additional variables significantly improve the explanatory power of the model – the null hypothesis being that this is not the case.
Unlike LR test, the information criteria (IC) such as Akaike information criterion (AIC) and Bayesian information criterion (BIC) can be used to compare alternative models, even if not nested.
However, they only rank models and do not quantify the statistical significance of any difference between the alternatives.
In most cases, the BIC penalizes the introduction of another independent variable more strictly than AIC, and as such it tends to result in less complex models being selected.
[3 marks]

xi) For the model in question, let the likelihood be L and the number of parameters be N. Also, we are given the total number of observations T = 2981.

The test statistics are:

- AIC = 2N – 2ln(L)
- BIC = Nln(T) – 2ln(L) = N*ln(2981) – 2ln(L) = 8N – 2ln(L).

But, BIC – AIC = 60. So, 60 = (8N-ln(L)) – (2N-ln(L)) = 6N. So, N = 60/6 = 10.
The model has 10 independent parameters. [2 marks]

xii) Choice of threshold:
The empirical version of the mean excess loss function over threshold u, theoretically defined as E(X-u | X>u), is useful in fitting a GPD for the excess distribution over a threshold.
It is best to visualize the empirical mean excess loss function by plotting the mean excess on y-axis against the various thresholds on x-axis. [This is called the mean excess plot.]
The lowest threshold above which the empirical mean excess loss function e(u) is linear in u is a typical choice.
It is usual for the chosen threshold to be in the range of 90^{th} to 95^{th} percentile of the complete distribution.
The choice of the threshold also often reflects the context, e.g. it could represent the retention limit over which a reinsurer becomes liable to cover claims.
Trade-off:
There is a trade-off embedded in the quality of approximation of the GPD v/s level of bias in the fit. The former is helped by choosing a higher threshold, while the latter is better off with a lower threshold.
Fitting of parameters:
A GPD is then fit to the data above the chosen threshold. A GPD has two parameters – scale (beta / β) and shape (gamma / γ).
These parameters are determined using methods such as:

- Maximum likelihood estimate (MLE), or
- Method of moments.

The shape parameter (gamma / γ) is also linked to the slope of the linear part of the mean excess plot. [4 marks]

xiii) Model risk is the risk arising from the use of inappropriate or inaccurate model when assessing or managing risks.
The specific choice of model structure for M1 is more subjective, flexible and open-ended, and highly susceptible to model risk…
… while the same for M2 is driven by a statistical result from the Extreme Value Theory regarding the asymptotic distribution for the excess loss distribution over thresholds.
As such, prima facie, one may regard M2 to be less prone to model risk than M1.

However, if the threshold in M2 is chosen very low, the asymptotic behaviour might not kick in and the model risk may be substantial too.

As regards M2, the higher the chosen threshold, the lower will be the model risk.

But correspondingly, the parameter risk will be higher owing to being based on much less data. [3 marks]

xiv) A. C(u,v) = uv is a valid copula and corresponds to the independence / product copula.
B. C(u,v) = u+v-1 is not a valid copula as it can take values less than 0, e.g. C(0.25, 0.25) = -0.5.
C. C(u,v) = uv / (u+v-uv) is a valid copula of the Clayton type.
D. C(u,v) = min{u^2, v^2} is not a valid copula as it doesn't satisfy the property that C(u,1) = u for 0<u<1, because C(u,1) = u^2 < u whenever 0 < u < 1.

[4 marks]

Side note: B can be tweaked slightly to make a valid copula as C(u,v) = max{0, u+v-1} which is, in fact, the counter-monotonicity copula.

xv) Given the 1-day Value-at-Risk is estimated as a multiple of σ [the 1-day portfolio return volatility], it is likely that the 1-day portfolio returns are assumed to be normally distributed with mean 0 and variance σ^2.
Mean return may be assumed 0 given the very short time horizon and that it errs on the side of prudence.
In order to derive a weekly Value-at-Risk (which is effectively a 5-day VaR), we need the following additional assumptions:

- Daily returns for different days are identically distributed (zero mean and same variance σ^2)
- Daily returns are uncorrelated with each other.

With these, we can conclude that the 5-day return distribution follows a Normal distribution with mean 0 and variance $5\sigma^2$.

Thus, the weekly VaR can be estimated as $\sqrt{5}*\sigma*2.5758 = 5.7597\sigma$. [3 marks]

xvi) A. Higher expected return on asset (relative to liability) will make ruin less likely and reduce required capital.
B. Higher volatility of asset (relative to liability) will make ruin more likely and increase required capital.
C. Higher positive correlation between asset and liability values will make ruin less likely and reduce required capital.
D. Higher diversification in asset types constituting A(t) will lead to lower volatility, make ruin less likely and reduce required capital.
E. Higher risk tolerance α will makes things less prudent and reduce required capital.
F. Higher comfort ratio k will make things prudent and increase required capital.

[3 marks]

xvii) Some of the challenges / limitations in the stated approach are as follows:

- Difficulty in obtaining consistent valuations for both assets and liabilities (all must be marked-to-market or marked-to-model)
- Difficulty in formulating assumptions about future (e.g. how assets will be invested)
- Large number of parameters required (increases drastically with number of variables)
- Difficulty in estimating parameters such as correlations and variances easily in a robust manner
- Consideration only at discrete points, while ruin may occur in between those points

- Difficulty in incorporating hard-to-quantify risks (e.g. operational risks, reputational risks) within this framework
- Need to select an appropriate risk measure (ruin may not work always)

[3 marks]

xviii) A practical, bottom-up process that can serve as an alternative for calculating capital requirements can be as follows:

1. For each source of risk, generate stand-alone distribution of change in the value of the company due to that risk.
2. While allowing for any diversifications as deemed appropriate, combine these distributions.
3. For this combined distribution, calculate the total capital at the desired level of significance / for the selected risk metric.
4. Attribute capital to each activity (by region / business unit or otherwise) based on the amount of risk generated by the same.

An additional important consideration is the correlation between risk factors. This is a necessary allowance to make, whether implicitly or explicitly, and at a micro or at a macro level. [3 marks]

[71 Marks]

Solution 2:

i) While it is true that ERM is conducted centrally (as against the traditional risk management that operates in silos), that is not the only distinguishing feature. The main distinguishing element of ERM is a holistic approach.
This entails:

- Consideration of all risks from all sources across the whole enterprise
- Led by the Board (top-down), coordinated by a risk management function (RMF) led by a Chief Risk Officer (CRO) and responsibility of all personnel
- A risk management failure in another company (presumably in a similar sector)
- A regulatory requirement
- Consistent application of risk management techniques

Additionally, there is also a thrust on value creation in ERM by:

- Integrating risk management / measurement into business processes and strategic decision making
- Considering upside risks and not just downside risks [2 marks]

ii) Apart from a previous risk management failure at the company, the following are the typical triggers causing the company to initiate ERM:

- A near-miss, even if there wasn't a full-blown failure, within the company
- A high-profile disaster in another similar organisation
- Demand / criticism from a regulator / auditor
- Concern expressed by other stakeholders [2 marks]

iii) Even if all investments of a company are denominated in its domestic currency, there may still be exposure to foreign exchange risk by way of:

- Uncertainty in revenue if sales are partly to another country with a different currency (e.g. GCCs serving onshore entities)
- Uncertainty in costs if raw materials are sourced from another country with a different currency (e.g. industries relying on oil imports)
- Consolidated accounts, if other group entities rolling into that company are domiciled in another country with a different *currency* [2 marks]

iv) A risk register contains much more than just a list of risks. In particular, it includes:

- A consistent labelling / numbering system fot easy identification of risks
- A categorization of all the identified risks
- A clear, understandable description to enable a consistent shared understanding
- An estimate of:
  - Likelihood of risk occurring
  - Its impact
  - Applicable timeframe
  - Its relationship (e.g. positive or negative correlation) with other risks
- Risk response action, along with any residual risks
- Individuals responsible for monitoring / managing risks [3 marks]

v) The CORREL function in Excel simply computes the linear correlation – Pearson's rho.

Pearson's rho is a measure of linear dependence between the variables.

Unfortunately, it has a number of limitations such as:

- Value not necessarily being preserved under a strictly increasing transformation
- Being undefined in case of heavy-tailed distributions (where variance is infinite)
- Being zero even when these is a strong relationship between variables (simply because there is no *linear* relationship)
- Strong assumptions in terms of marginal distributions (being jointly elliptical)

It may be appropriate to explore rank correlation metrics instead such as Spearman's rho and Kendall's tau, which overcome many of the above-stated limitations.

If the situation so warrants, especially in the context of risk management, one can even look at tail correlations and tail dependence (e.g. by way of copula) [4 marks]

vi) Firstly, risks need not always be downside risks. There can be upside risks too. So not all uncertainly / risk is bad.

Notwithstanding the above, it may be appropriate or even profitable to retain a particular risk if:

- The risk is trivial
- Taking the risk is a component of its core business (e.g. credit risk for banks)
- The cost of any treatment (transfer / reduction etc) is greater than the exposure to that risk, i.e. it is more economical to retain
- That risk acts as a natural hedge / diversifier with its other risks
- There is no alternative (say, the risk is uninsurable and non-transferable)

[3 marks]

vii) It is not true that risk can only be retained or transferred. In fact, there are many ways of reducing risk without transfer such as the following:

- Risks can be diversified by taking on uncorrelated risks, e.g. spreading investments across asset classes
- Risks associated with random fluctuation can be reduced by increasing the portfolio size (allowing the law of large numbers to reduce the volatility), e.g. investing in mutual funds instead of individual stock
- Some risks can even be (at least partially) hedged by taking on other appropriately chosen risks, e.g. mortality risk from protection business v/s longevity risk from annuity business for a life insurer
- Closer matching of assets and liabilities (in terms of their nature, term, currency, etc) can reduce the overall risk
- Operational risks can be reduced through strong internal controls and governance.
- Underwriting
- Credit and counterparty risks can be reduced through robust due diligence practices and tightly worded agreements.
- Agency risk can be reduced through appropriate structuring of incentives (such as remuneration and bonus systems) for better alignment of interests
- Appropriate preventative measures can greatly reduce certain risks, e.g. the risk of property damage due to fire can be mitigated by having fire alarms, sprinkler systems and regular audits. [4 marks]

viii) While in theory, it may seem to keep one's portfolio neutral with respect to the three greeks (delta, gamma, vega), it may be impractical to do so at all times.
For a portfolio linked to the value of multiple (say m) underlying indices or asset prices, the number of derivatives required can be very large ($> m^2$).
The rebalancing costs are often significant, limiting how frequently one can adjust one's portfolio.
Often, suitable traded derivatives may not be available or (even if available) they may be illiquid.

In practice, it is common for organizations to give individual traders some limits on how far they can deviate from delta, gamma, and vega neutrality. Perhaps individuals can do something similar.

Practically speaking, it is easier to neutralize delta rather than gamma or vega, due to easy availability of futures contracts which are cost-effective (as against options which are more illiquid and costly). As such, it is usual to achieve delta neutrality more frequently (say daily) while adjusting for gamma and vega neutrality less frequently.

Even if one does manage to make one's portfolio neutral, it is important to remember that one loses out on the upside as well, because of this approach. [3 marks]

ix) The idea of doing away with capital allocation and holding total capital centrally does have the advantage of protecting the business units from the vagaries of different capital allocation methods.

However, it can lead to a lack of understanding on the part of each business line regarding the impact their actions have on the capital requirements of the business.

It may even lead to agency risk, whereby business units may focus on absolute returns rather than risk-adjusted returns (which shareholders are typically interested in optimizing for).

Having a sound capital allocation approach can become a strategic level to encourage the growth of the right kind of business, while limiting the growth of the wrong kind.

This can even feed into a performance metric for business units which may have implications on bonuses for the leaders of that business unit.

If the business units are situated in different tax and regulatory jurisdictions, this can even have associated cashflows and real financial impact on the organization

[3 marks]

x) Simplistic approaches to capital allocation (e.g. a pro-rata approach based on business reserves or premium income) do have some advantages such as the following:

- It ensures all the capital (no more, no less) is allocated.
- It is easy to compute
- It is easy to communicate and explain to all stakeholders (including business units themselves, auditors, tax authorities etc).

However, this is not necessarily the best approach always, because of the following reasons:

- It does not correspond to the financial principle of marginal pricing (which could be achieved by allocating according to marginal capital costs)
- Correlation and dependency structure may not have been adequately considered.
- If the basis is not risk based, large amounts of capital may be allocated to business units that generate relatively little risk, and vice versa.
- As such, it may lead to over-investment in risky business units.

[3 marks]
[29 Marks]

****