Operational Risk Modelling

Past, present, and future

Adél Drew, FIA
Tatiana Egoshina, FIA
Jack Huffer

Milliman

The traditional view of risk is based on an idea that there is a cause, which may lead to a risk event with a resultant effect. This linear approach is not congruent with the operational risks faced by businesses in a modern environment. Instead, risks are becoming increasingly complex and therefore difficult to model meaningfully.

Considering this traditional approach most risk managers were taught, it's not surprising that frequency-severity modelling is the standard approach for modelling operational risks. While this was a good starting point for getting us to where we are today, given the complex nature of the modern business environment, we pose the question: Will it be enough to advance our understanding in operational risk management?

In this paper we look at the common industry practices regarding operational risk modelling and explore the problems with frequency-severity models. We also look at an alternative approach in the form of causal modelling and how this can overcome some of the issues encountered with more traditional modelling techniques. Lastly, we discuss development areas in the field of operational risk modelling for consideration.

Industry Overview

PUBLIC SFCR DATA

Under Solvency II (SII) firms are required to publish Solvency and Financial Condition Reports (SFCR) which makes public information about the company to policyholders, shareholders, and other stakeholders. These reports are required to contain specific information on:

I. Business and performance
II. Systems of governance
III. Risk profile
IV. Valuation for solvency purposes
V. Capital management

The capital management section of these reports provides useful Solvency Capital Requirement (SCR) information, enabling several observations to be made regarding industry practices related to operational risk modelling.

We have performed an analysis on 2017 to 2020 reported results from around 200 UK insurers using SFCR data.[1] Almost half are non-life, with about one-third Life and the remainder composite.

- Only around 20% of all firms analysed use a full- or partial-internal model ((P)IM), and that number has remained fairly stable since the implementation of SII in 2016 (Figure 1).
- The split varies significantly by insurer type, with nearly half of composite insurers making use of an internal model, while for non-life insurers, this use is only around 10%.
- Operational risk SCR makes up less than 20% of the total SCR for around 80% of firms (Figure 2).

[1] Solvency II Wire, 2021, https://www.solvencyiiwire.com/category/sfcr-analysis/.

FIGURE 1: SCR MODEL TYPE BY INSURER CATEGORY (FROM PUBLISHED SFCR RESULTS FOR UK INSURERS [2020])

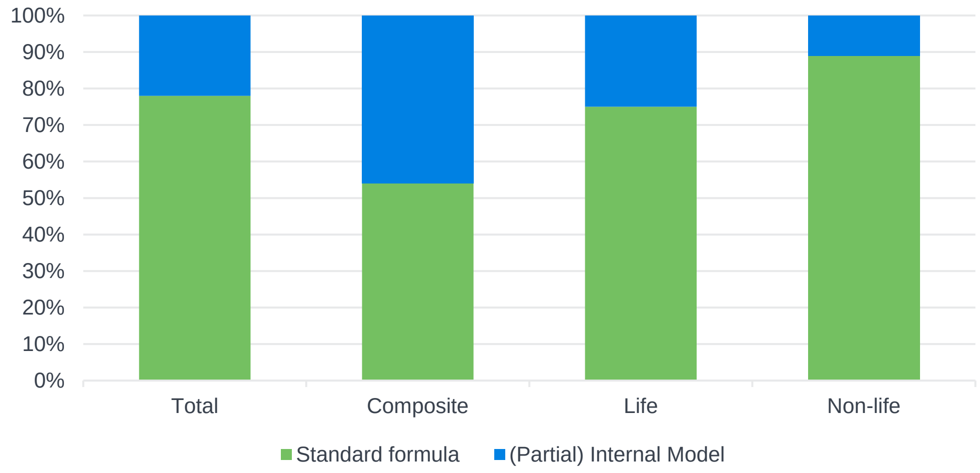

FIGURE 2: OPERATIONAL RISK SCR AS % OF TOTAL SCR BY COUNT (FROM PUBLISHED SFCR RESULTS FOR UK INSURERS INCLUDING STANDARD FORMULA AND P(IM)S [2020])

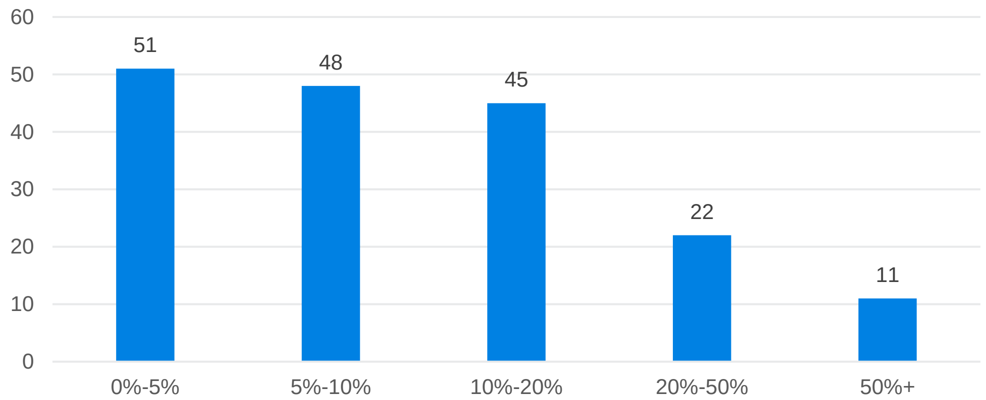

The simplistic nature of the standard formula for operational risk under SII can lead to excessive capital requirements not fully reflecting the operational risk profile of companies. From SFCR data this seems to have more of an impact on composite insurers than non-life insurers.

BENCHMARKING DATA

Results from the 2020 ORIC capital benchmarking survey[2] provides us with additional information not available from public SFCR data alone. The annual survey explores the end-to-end operational capital modelling processes conducted by their insurance and investment management member firms. This year's survey gathered responses from 24 participants across various locations in UK, Ireland, Holland, the United States, and Australia.

The majority of survey participants indicated that they still use a standard formula to calculate their operational risk capital, as well as an internal model to calculate operational risk capital for non-reporting purposes.

A reason for not using these internal models for capital reporting purposes could be that additional resources are required to get an internal model approved by the regulator, along with managing the modelling process and validating results on an ongoing basis. Internal model approval and model validation are areas we have supported a number of clients on, further evidencing the resource strain.

[2] ORIC International. 2020. 'Annual Capital Benchmarking Survey Summary Report'; https://828ff78c-7206-4ab0-bccc-4ed48e15602c.filesusr.com/ugd/44340f_2f07eaf9a5f545d9ba0c1af08a8edd64.pdf?index=true.

The ORIC survey supports this hypothesis, as the number of full-time employees (FTEs) employed by firms who are involved in the modelling processes are significantly more on average for firms with (P)IMs compared to those using a standard formula (Figure 3). Most firms using a standard formula calculation employed between one and four FTEs to perform their modelling processes. Compare this to firms who use internal models where around half employed more than four FTEs in this area. Some firms in this latter group reported employing more than 13 FTEs to perform these tasks.

FIGURE 3: NUMBER OF FTES EMPLOYED BY MODEL TYPE (N=22)
ORIC ANNUAL CAPITAL BENCHMARKING SURVEY SUMMARY REPORT, (JANUARY 2020)

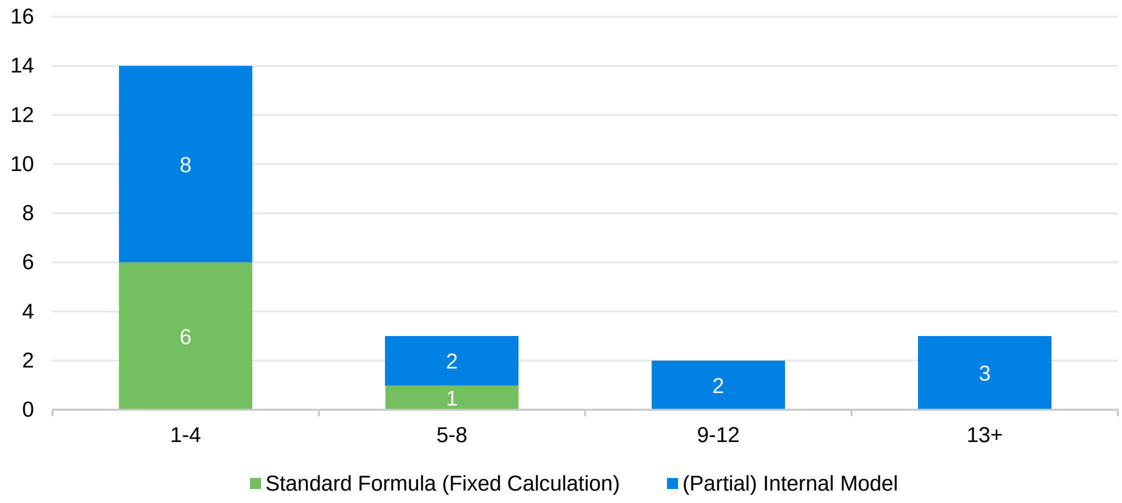

The ORIC survey also gives us an idea of the typical operational risk modelling techniques used by those firms with a (P)IM. Two distinct methodologies were employed by all participants of the survey: a scenario-based approach or a hybrid approach (Figure 4). Both approaches rely on frequency and severity modelling.

While a scenario-based approach is parameterised by expert judgement, hybrid approaches use data-driven probability distribution functions to model the frequency of loss. Scenarios would still be used to model the severity. Parameterisation of loss data distributions is discussed in more detail in the scenario calibration section.

FIGURE 4: OPERATIONAL RISK MODELLING APPROACH EMPLOYED BY FIRMS (N=20)
ORIC ANNUAL CAPITAL BENCHMARKING SURVEY SUMMARY REPORT (JANUARY 2020)

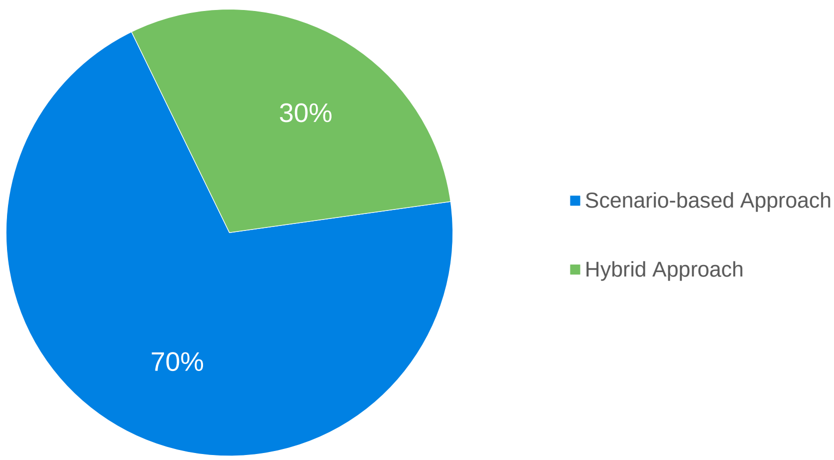

With regards to the modelling of risk loss distributions, and compared to the 2019 ORIC survey, there is evidence to suggest that firms are starting to use more advanced techniques with causal modelling being included within the 'other' category in Figure 5 below.

FIGURE 5: APPROACH EMPLOYED BY FIRMS IN OPERATIONAL RISK LOSS DISTRIBUTIONS (N=20)
ORIC ANNUAL CAPITAL BENCHMARKING SURVEY SUMMARY REPORT (JANUARY 2020)

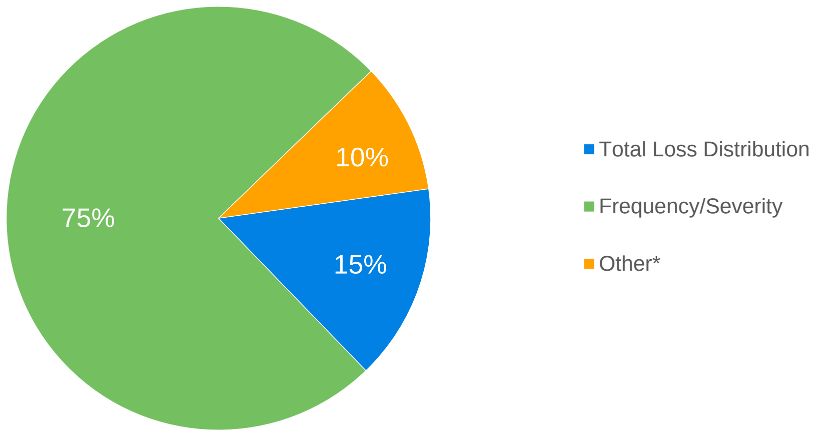

*Other includes causal modelling and deterministic scenarios.

Typical Modelling Process and Issues Faced

A TYPICAL OPERATIONAL RISK MODELLING PROCESS

The diagram below (Figure 6) shows a typical process used by firms in quantifying their operational risk capital—some steps are part of a wider risk management framework, and some are specific to the operational risk quantification. Each is discussed in turn below.

FIGURE 6: A TYPICAL OPERATIONAL RISK MODELLING PROCESS

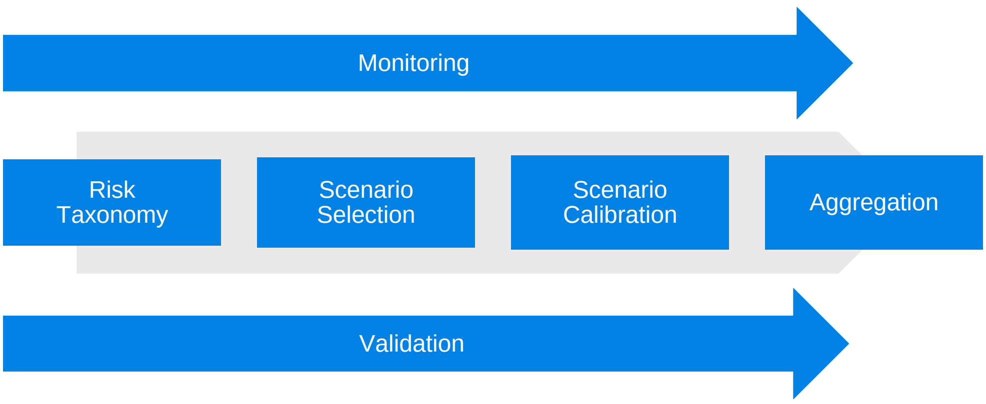

The risk taxonomy[3] informs the scenario selection process. Scenarios are analysed and then used to calibrate the distribution of losses under each scenario. The losses get aggregated to derive the total loss for operational risk, which is then aggregated with other risks.

The appropriateness of the model and its calibration will need to be monitored; the Risk Management Function (RMF) should validate the modelling approach and the calibration.

The next sections will look into each step of the process in more detail, the questions companies need to answer, and the issues to solve.

[3] A risk taxonomy is a comprehensive, common, and stable set of risk categories an organization is exposed to. It may be referred to as a 'risk universe' as well.

RISK TAXONOMY

Operational risk modelling starts with defining a taxonomy, as it helps answer the question 'What is being modelled?' Operational risk is a complex, multi-faceted risk which involves many departments within an organization. Therefore loss modelling is usually undertaken at a more granular level, or level 2 in the taxonomy—that is, the potential losses are segmented into homogeneous categories rather than modelling operational risk (level 1) as a whole. There is no set standard taxonomy used by the insurance industry, although many base their taxonomy on Basel II or external data providers' categories such as ORIC[4] or ORX.[5] An operational risk profile typically differs from company to company and what can be important for one organisation may be of little focus and attention to another. The exact taxonomy used may depend on how the business model is structured or who the risk owners are for various risks. Figure 7 below is an indicative example of an operational risk taxonomy with level 2 and 3 shown:

FIGURE 7: EXAMPLE OF OPERATIONAL RISK TAXONOMY

LEVEL 2	LEVEL 3
Financial Crime	• Internal Fraud • External Fraud • Money Laundering
Information Security	• Technology • Cyber • Data Protection
Data and Processes	• Model Risk • Data Management • Process Risks
Legal and Compliance	• Legal • Regulatory Compliance • Statutory Reporting and Tax
People	• Health and Safety • Conduct
Third-party and Outsourcing	

The aim of the taxonomy is to segment operational events into homogenous, mutually exclusive groups; taxonomies can employ different levels of granularity, e.g., level 2 risks can be split further into more granular risks. Increased granularity may help with scenario generation and its further calibration as it might be easier to put scenarios into context. On the other hand, a very granular taxonomy with a lot of risks increases the difficulty in calibration and aggregation of scenarios. It is important to strike a balance here.

To make matters more complex, operational risks almost always fall into more than one category simultaneously, irrespective of the granularity of the taxonomy. Consider for example data protection—this may be viewed as an information security risk but also a compliance risk as it relates to the General Data Protection Regulation (GDPR). Additionally, operational risks can overlap with other non-operational risks in what is called a 'boundary risk.' Boundary risks are the occurrence of market, credit, underwriting, or liquidity losses that are triggered, or increased, by the occurrence of an operational risk event. For example, a data entry error in the accounting system might generate an increase or decrease of the credit position of a counterparty; this effect is originated from an operational event but has also an impact related to credit risk.

Therefore, the taxonomy employed needs to be precisely defined to avoid double counting within the operational risk modelling framework and within the overall risk capital calculation.

[4] Operational Risk Consortium Limited International (ORIC), https://www.oricinternational.com.

[5] Operational Riskdata eXchange Association (ORX), https://managingrisktogether.orx.org/.

SCENARIO SELECTION

Taxonomy and risk assessment inform the scenario analysis and selection phase, i.e., the risks to be discussed and investigated at the workshops with subject matter experts (SMEs). This in turn informs the scenario calibration.

The scenario selection phase provides a high-level evaluation of the risk exposure, typically using a heat map approach. This process is usually reliant on expert opinion to identify and evaluate major potential risk events and to assess their potential outcomes, ensuring a forward-looking view. However, in practice, developing truly forward-looking views can be challenging and scenarios may tend towards events that seem plausible based on the past. Many experts assume measuring risk in the past is a good way to predict the future; however this does not work in a complex world where the risks are always evolving and emerging.

The aim of the scenario selection process is twofold:

1. To provide an input for the capital estimation for operational risk, as outlined in the scenario calibration section
2. To fulfil risk management purposes, as controls and mitigations are discussed

Scenario analysis provides the context for the calibration discussions specific to the insurer's risk environment. Additionally, the aim should be to support effective risk management through:

- Improving the understanding of risk factors among risk owners with the sharing of experiences and best practices
- Assessing the control framework and detecting possible future failures

The assessment process is typically undertaken by risk owners facilitated by the RMF and involves considering a range of large to catastrophic scenarios that the organisation could suffer. Examples of scenarios include:

- Failure to deliver a significant change programme due to operational errors
- Failure to recruit and retain staff of appropriate quality and in sufficient numbers
- Large fraud event
- Cyber attack

Scenario selection, however, relies heavily on expert judgement within the traditional risk modelling framework. Even the most advanced frequency and severity model will only be as good as the scenarios selected in determining a truly reflective set of risk scenarios.

SCENARIO CALIBRATION

As discussed earlier, frequency and severity modelling is the most commonly used approach for the calibration of operational risks. Under this approach, the number of operational events, and the severity of those events, are assumed to follow particular probability distributions. The total aggregate loss for the risk is then estimated as follows:

$$L = \sum_{i=1}^{N} X_i$$

Where L denotes the total loss;

N denotes the number of operational events, following an assumed frequency distribution; and

X_i denotes the severity of the impact of the operational risk being modelled.

The most common choices for modelling frequency distributions are Poisson and Negative Binomial, as these models provide the best fit for operational risk events.

A Poisson distribution is commonly used as the basis for statistical models that describe event incidence; it is the most popular choice for modelling frequency distributions for the following practical reasons:

a) The Poisson distribution requires only one parameter to be estimated, making the parameterisation process and understanding of the parameters easier for the business experts.

b) Under regular circumstances it appears reasonable to assume operational risk losses to be a Poisson process.

Under hybrid approaches discussed in the earlier section, frequency distribution is calibrated using data; under scenario-based approaches it is calibrated with expert judgement as described below.

The choice of the distribution for modelling severity, however, is also subjective and is driven by pragmatic considerations.

Distributions assumed for severity are normally:

- LogNormal
- Weibull
- Generalized Pareto Distribution
- Log-logistic
- Similar heavy-tailed distributions

The LogNormal distribution is a popular choice, for practical reasons, as it lends itself well to the truncated form, is simple to use and explain, and it is relatively heavy tailed. Truncation is used as small-size events are of little interest for the purposes of capital estimation. These small events can also be allowed for in best estimate assumptions, while the upper limit of the truncation can be kept to the 'observable' part of the distribution. Calibration can be performed with the use of external or internal data, but is most often done with the use of expert judgement.

Combining the frequency and severity of the impact distributions is normally performed using Monte Carlo simulations, as closed form analytical solutions for the total distributions are not available.

The distribution of operational risk losses is highly skewed due to a combination of:

- The low probability for an event occurring
- The skewed distribution of losses given an event has occurred

In order to reduce simulation error to an acceptable level, it is necessary to run a high number of simulations for each scenario—the number of simulations is chosen so that the result is stable, i.e., Monte Carlo simulations converge.

Calibration using internal data

It may be the case that for certain risks within a firm's operational risk taxonomy there is some internal data with which to fit a distribution. This approach has the advantage that the distribution will more closely reflect the specific organisation's profile. Using internal data also ensures firms have better control over this data; however, the data still potentially needs to be adjusted for the considerations outlined below.

As with all data, there is the usual actuarial conundrum of the trade-off between relevance and credibility—there should be enough data to be credible, therefore data collection should go back as far as possible. On the other hand, the data needs to be relevant and reflect the risks the company is exposed to now. However, risk modelling should be forward-looking so the data should ideally reflect the ever-changing landscape of operational risk the organisation is exposed to and is expected to be exposed to in the future. Unfortunately, historical loss data cannot account for the new and innovative risks corporations are expected to face, in particular in the areas of cyber risk and risks related to climate change.

Existing data may come from a time when the company was smaller in size, or the risk management framework was less mature, and therefore its exposure to operational risk events may be different; both lesser and greater depending on the risk category. It may be possible to adjust past data for such factors as companies evolve, but it is not a clear-cut assumption. Exposure to the severity of operational risk impacts is typically non-linear with regard to the size of the company, i.e., if the company is now twice the size it was five years ago, it does not necessarily mean that its exposure is double as well.

Other factors could also impact an organisation's risk profile over time; for example, introduction of new products, distribution to new markets or using different channels, changing regulations, changes in actuarial and administration systems, and even external factors (consider the pandemic). Therefore, unpicking exposure impacts related to various business change impacts, then adjusting historical data is often not a plausible exercise.

It is also important to consider the purpose of the operational risk modelling and quantification exercise, as impacts may differ. For example, for model risk, if the aim of quantification is SII reporting then errors affecting the SII balance sheet should be taken into account. If we are looking to quantify operational risk for economic capital purposes, then model errors affecting the economic value of the company should be considered (and those are not easy to estimate as, for example, pricing errors can emerge over a long period of time).

These considerations will also apply to external data and expert judgement-driven calibrations.

Calibration using external data

For some types of operational risk, there might be publicly available information on the losses incurred by various organisations, for example regulatory fines; there may be anonymised information available for other risks, e.g., cyber risk. Therefore, it may be possible to use this data to fit the severity distribution.

As with internal data, one of the main questions is whether it would reflect the risk profile of the organisation. Additional considerations for external data are the question of how the data was collected, what assumptions were made on collection, bias in reporting of data, cleanliness of data, and data errors—all of which would affect the distribution fitting.

Expert judgement calibration

The calibration of scenarios can occur at the same time as scenarios are discussed with the SMEs as outlined in the scenario selection section. During workshops, the SMEs will form an opinion about the capital losses which could arise from each scenario. These workshops are normally facilitated by the risk function, who also might send preliminary information to SMEs before the workshops, such as typical losses incurred by the industry, observed frequency of events and emerging trends, and further reference materials.

During workshops scenarios will be evaluated; alongside the operational risk and controls assessment, the following quantitative characteristics of the risk will be discussed:

- Frequency: Possible frequency of occurrence, within a one-year time horizon; this will be used to inform the parameters of the frequency distribution.
- Severity: Scenarios can be discussed in the context of ‘typical impact’—material loss in ordinary conditions, and ‘worst case’—severe economic loss derived from an extreme but plausible scenario. The ‘worst case’ will need to be associated with the extremity of such loss, e.g., 1-in-150. These estimates will be used to calibrate the parameters of the chosen distribution.

Calibration of scenarios in this way has an advantage in that it allows for the specific risk profile and exposure of the company, it is done in the context of controls and mitigation actions, and can be made forward-looking.

As with all expert judgement calibrations, however, it is not without its own issues—quality relies heavily on SME input which is subjective by nature. Depending on the organisation’s risk culture, information offered up may also be tainted by a desire to make an SME’s risk area appear better managed than it is. This bias is exacerbated when workshops are conducted with only one or two SMEs per risk category. Facilitation of the workshops by an experienced risk professional may help the discussions; however, this again is dependent on the risk function’s authority within the organisation and specialist skills.

The choice of distribution is again subjective as it is not informed by data. Calibration requires SMEs to formulate not just estimates of the losses, but also to make a judgement on what percentile of the distribution these losses represent—not an easy task!

Below we cover three case studies that highlight just how much of an impact these subjective, expert judgement choices can have on the outcomes.

Case Study 1: Comparison of different distributions

As discussed earlier within this paper, the operational risk severity distributions are often modelled using a LogNormal, Weibull, Generalized Pareto distribution, or Log-logistic distribution.

Figure 8 shows a simulation of the total loss distribution under each of these four severity distributions. The severity distributions were all calibrated using a typical case loss (median) of £1.1m and a 1-in-100 worst case loss of £70m; parameters that would be informed by expert judgement.

For this example, a Poisson distribution is used to model the frequency of operational events. Individual losses, X_i, and frequency, N, are assumed to be independent. We discuss the implications of this independence assumption in Case Study 3.

Figure 9 shows, in each case, the capital required to cover a 1-in-200 loss. We can see that the choice of individual loss distribution can have a large impact on the outcome of the capital model.

FIGURE 8: TOTAL LOSS DISTRIBUTION (£) UNDER FOUR DIFFERENT SEVERITY DISTRIBUTIONS

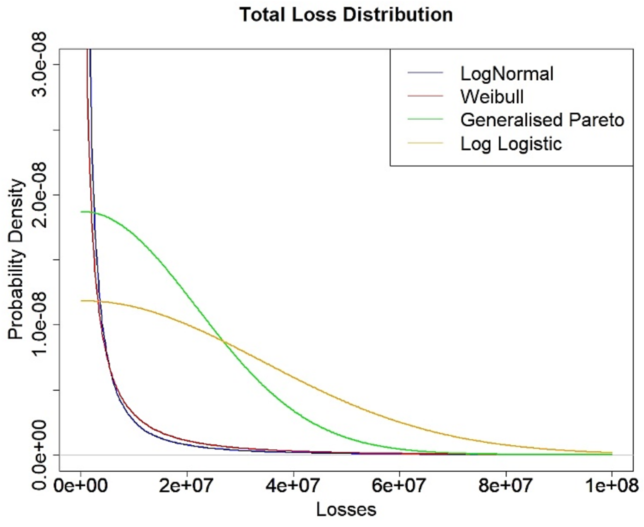

FIGURE 9: CAPITAL REQUIRED (1 IN 200 LOSS) UNDER FOUR DIFFERENT SEVERITY DISTRIBUTIONS

INDIVIDUAL LOSS DISTRIBUTION	CAPITAL REQUIRED (£M)
LogNormal	43.9
Weibull	50.2
Pareto	38.6
Log Logistic	37.7

Case Study 2: 1-in-100 vs 1-in-150 worst case calibration

When deciding on a worst plausible case for the severity of an extreme loss event it will likely be hard to gauge exactly which percentile of the individual loss distribution this represents.

Figures 10 and 11 show the impact of moving from a 1-in-100 to a 1-in-150 worst case calibration. The same parameters as in Case Study 1 and a LogNormal individual loss distribution were used.

Assuming a 1-in-150 rather than 1-in-100 calibration results in a lighter tailed total loss distribution and hence a materially lower 1-in-200 capital requirement.

FIGURE 10: TOTAL LOSS DISTRIBUTION (£) UNDER A 1-IN-100 AND 1-IN-150 WORST CASE INDIVIDUAL LOSS CALIBRATION

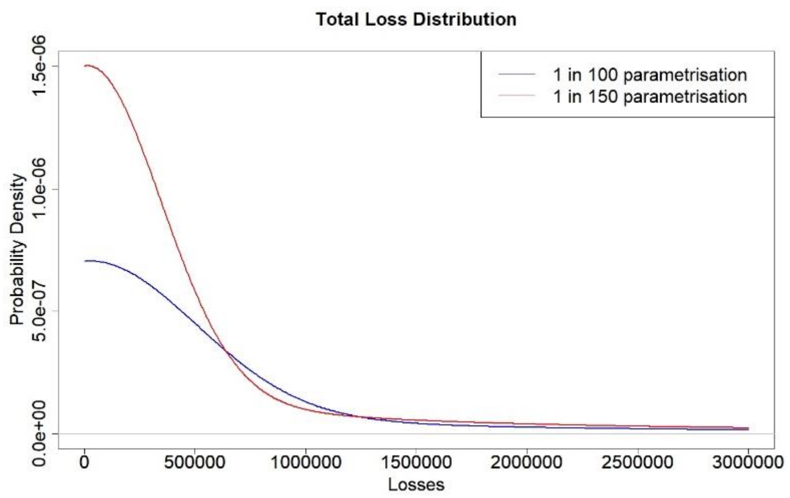

FIGURE 11: CAPITAL REQUIRED (1 IN 200 LOSS) UNDER A 1-IN-100 AND 1-IN-150 WORST CASE INDIVIDUAL LOSS CALIBRATION

WORST CASE CALIBRATION USED	CAPITAL REQUIRED (£M)	CHANGE
1-in-100	43.9	
1-in-150	35.3	-20%

Case Study 3: Frequency and severity dependency

Frequency and severity are typically assumed to be independent within frequency and severity modelling approaches. However, this is not necessarily the case as increased frequency of events might lead to higher (or lower) severity for certain operational risks.

Consider, for example, cyber risk where risk events often start with 'probing' (non-extensive small loss attacks to explore vulnerabilities) followed by an extensive attack with high loss.[6]

A sufficient level of data would be required to accurately capture a dependency structure in this case.

Different techniques exist to allow for dependence between frequency and severity. One approach is to use copulas.

Where dependence does exist, methods to calibrate an appropriate copula to use are more complicated in this situation as a result of the frequency distribution being discrete (rather than continuous).[7] The standard continuous case estimate of the level of dependence to use in a copula when applied to discrete marginal distributions would be biased due to the lack of continuity of the discrete marginal distribution function.[8]

For this case study we have used a Gaussian copula to induce a dependence structure between frequency and severity. The severity of losses was assumed to be conditionally independent when given the frequency.

The Monte-Carlo simulations are trickier to perform in this case compared to Case Study 1 (where for each simulated N, the corresponding number of independent LogNormal variables were simulated).[9]

A zero-truncated Poisson distribution was used as the marginal frequency distribution applied to the Gaussian copula when determining the dependency structure. The marginal severity distributions remain as LogNormals, parameterised in the same way as in Case Study 1.

The use of a zero-truncated Poisson ensures the individual losses remain observable, i.e., there should be no individual losses simulated when the number of loss events N is zero. The possibility of no loss events occurring is then taken into account when simulating the aggregate loss.

Through the copula approach the LogNormal severity distributions were sampled simultaneously with the frequency, such that the individual severity distributions are independent when given the frequency, but overall the severity and frequency are dependent.

Figures 12 and 13 show the impact of assuming an independent frequency and severity when they are in fact dependent. It is possible to observe operational events where either a positive or negative correlation between the frequency and severity exists. In both cases the dependent total loss distribution is much flatter.

Where positive dependence exists, the distribution has a longer tail, and, as a result, the 1-in-200 capital required has increased by almost 80%. Where negative dependence exists, the distribution has a shorter tail. In this case the 1-in-200 capital required has decreased by 12.5%.

This shows that capturing the correct dependency structure can ultimately be more important than the choice of the frequency and severity distributions.

[6] Chris Hamer, Chris Beck, Blake Fleisher. 'Know your cyber blind spots,' Milliman white paper.

[7] Channouf, N & L'Ecuyer, P. 2009. 'FITTING A NORMAL COPULA FOR A MULTIVARIATE DISTRIBUTION WITH BOTH DISCRETE AND CONTINUOUS MARGINALS,' *Proceedings of the 2009 Winter Simulation Conference*; https://www.informs-sim.org/wsc09papers/033.pdf, pp. 1-2.

[8] Trivedi, P & Zimmer, D. 2017. 'A Note on Identification of Bivariate Copulas for Discrete Count Data,' *Econometrics 2017, 5(1), 10;* https://doi.org/10.3390/econometrics5010010, pp. 1-3.

[9] More detail about the approach taken to simulate a Gaussian Copula can be found here: Haugh, M. Spring 2016. 'An Introduction to Copulas,' *IEOR E4602: Quantitative Risk Management;* http://www.columbia.edu/~mh2078/QRM/Copulas.pdf, p. 5.

FIGURE 12: COMPARISON OF THE TOTAL LOSS DISTRIBUTIONS (£) IN THE CASES WHERE FREQUENCY AND SEVERITY ARE INDEPENDENT AND DEPENDENT

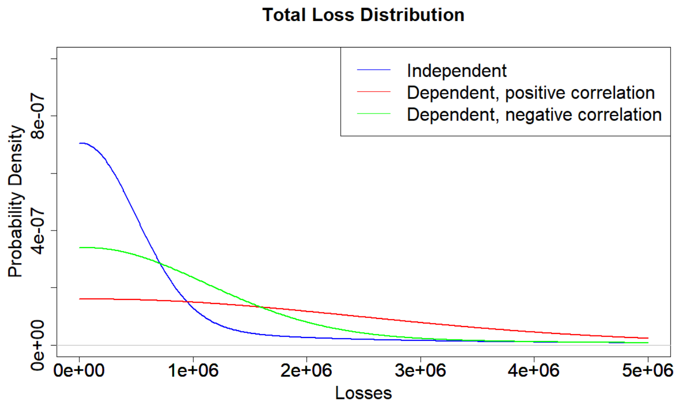

FIGURE 13: CAPITAL REQUIRED (1 IN 200 LOSS) IN THE CASES WHERE FREQUENCY AND SEVERITY ARE INDEPENDENT AND DEPENDENT

LEVEL OF DEPENDENCE	CAPITAL REQUIRED (£M)	CHANGE
Independent	43.9	base
Dependent (Kendall's rank correlation coefficient between the severity and the frequency equal to c. 40%)	78.3	+78.4%
Dependent (Kendall's rank correlation coefficient between the severity and the frequency equal to c. - 40%)	38.4	-12.5%

AGGREGATION

The standalone capital requirements from different scenarios need to be aggregated to obtain the overall operational risk capital requirement, typically allowing for diversification between scenarios. A common approach is to model dependencies between loss distributions using a Gaussian copula parameterized by a correlation matrix,[10]; t-copula is also used, though parameterisation becomes more complicated in this case.

The calibration of the correlations is normally based on expert judgement, due to lack of historical data. The SMEs are asked to think about the relationship between operational risk scenarios—how likely it is that one scenario could cause another scenario, or for scenarios to be impacted by a common driver, e.g., market fall or systems failure.

Expert judgement on correlation should not be expressed with spurious accuracy, and is typically recorded qualitatively as 'non-correlated,' 'low,' 'medium,' 'high,' and 'very high.' at best, which are then translated into correlation parameters in 0.25 steps.

Following the correlation matrix calibration exercise, the matrix (or matrices, if multiple step aggregation is used) needs to be adjusted to be positive semi-definite; the resulting adjusted matrix needs to be inspected and compared to the matrix prior to adjustments. Any large adjustments could indicate that the initial matrix had some significant internal inconsistencies—those would require further investigations with SMEs.

The operational risk capital requirement will then need to be aggregated with other risks, following the methodology adopted by the firm.

[10] ORIC 2019 Capital Benchmarking Survey: 89% of firms use a copula approach, and of those 88% of firms use a Gaussian copula.

PRACTICAL CONSIDERATIONS

Apart from modelling decisions there are a number of practical considerations firms might take into account in their approach to modelling of operational risk.

Frequency of model calibration

As can be seen from the earlier sections, calibration of a model for operational risk quantification is a lengthy and elaborate process, so there is a decision to be made on how often the process should be performed. As the operational risk environment is rapidly changing, calibration exercises cannot be performed infrequently, or alternatively, too regularly to present a heavy resource burden. The typical approach is to calibrate scenarios and aggregation annually, unless there is a trigger requiring a re-calibration in the interim; methodology is also reviewed annually but often remains unchanged year on year.

Solvency II (partial) internal model approval

In order to include an operational risk component into a (P)IM, firms will need to engage with the regulator through a formal application process.

In any application for approval, insurance and reinsurance undertakings are required to submit, as a minimum, documentary evidence that the (P)IM fulfils the requirements set out in Articles 120 to 125 of SII Directive.[11]

These requirements involve:

- **Use Test:** Firms need to demonstrate that the model is widely used and plays an important role in their system of governance, in particular, risk management and economic and solvency assessment processes.
- **Statistical Quality Standards**: Several standards are outlined, including, but not limited to, requirements regarding probability distribution methods used, data (which should be accurate, complete, and appropriate), and allowance for diversification.
- **Calibration Standards**: This includes requirements to use a Value at Risk measure derived from the probability distribution as a metric for the SCR directly; supervisory authorities may require firms to run their internal model on relevant benchmark portfolios and to use assumptions based on external, rather than internal, data in order to verify the calibration of the internal model and to check that its specification is in line with generally accepted market practice.
- **Profit and loss attribution**: Firms shall demonstrate how the categorisation of risk chosen in the internal model explains the causes and sources of profits and losses. The categorisation of risk and attribution of profits and losses shall reflect the risk profile of the firm.
- **Validation standards**: Firms shall have a regular cycle of model validation which includes monitoring the performance of the internal model, reviewing the ongoing appropriateness of its specification, and testing its results against experience.
- **Documentation standards**: Firms shall document the design and operational details of their internal model.

Firms also need to demonstrate ongoing compliance with the requirements outlined above.

Given the nature of operational risk, and how its modelling and quantification differs from insurance and market risks, the above requirements may not be that straightforward to comply with. For example, due to the heavy reliance on expert judgement in calibration and in some key modelling decisions, proof of compliance with statistical quality standards and calibration standards will be different, and more onerous, than other risks.

Model documentation will need to be updated to reflect inclusion of operational risk in the (P)IM, as well as model governance and model validation processes. For model validation it is likely that new validation tools will need to be developed. Validation of expert judgement is not straightforward, and SMEs may be needed to help validate the calibration. In this case it may be difficult to maintain the independence of calibration and validation.

The model for quantification of operational risk may alternatively be used in Pillar 2 SII reporting, or Economic Capital used for decision-making, rather than Pillar I reporting, which would ease some of the requirements outlined above.

[11] Directive 2009/138/EC of the European Parliament and of the Council of 25 November 2009 on the taking-up and pursuit of the business of Insurance and Reinsurance (Solvency II), 30/06/2021, Article 120-125, https://eur-lex.europa.eu/eli/dir/2009/138/2021-06-30.

An Alternative to Frequency and Severity Models

All operational risk events have causes. Often these causes occur in multiple levels which are not easy to identify, let alone model. Consider for example the potential for a data security breach. On simplistic terms this breach could be because of process failures or human error. Stepping further back, the root cause, or driver, could be an inexperienced data security team or even a pandemic.

FIGURE 14: EXAMPLE OF A RISK CHAIN OF EVENTS

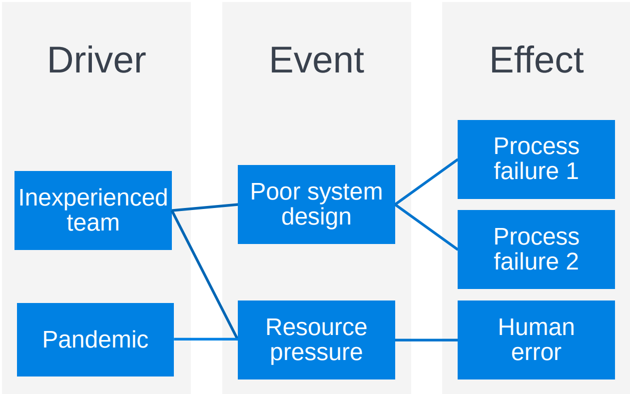

In linking drivers to outcomes through various causal steps, what emerges is a chain of events. Add in more detail, or additional interconnected risks, and what you end up with is a very complex picture that no frequency-severity model can replicate.

Generally, insurers apply their modelling at the event or effect stage, usually grouped according to a risk taxonomy, in an attempt to predict outcomes. The modelling process then involves breaking down the complex system into smaller pieces according to a chosen taxonomy, making simplifying assumptions, and hoping that by solving these simpler parts and adding them back together to get the whole, this will give the 'correct' answer. As described earlier, even defining the operational risk taxonomy precisely enough to enable the correct risks to appear in the correct buckets without double counting as a start is tricky enough.

Additionally, this approach misses out on the effect of the chain of events and where the problem comes from. By understanding the drivers that may result in events across several categories of the risk taxonomy, and which of these are the most material, it is possible to understand what needs to be monitored in order to spot issues before they crystallise.

CAUSAL MODEL

A causal model conditions loss outcomes upon the states of underlying business drivers that cause those events to occur. By directly incorporating the business drivers into the modelling framework, it is possible to see how loss outcomes are driven by the states in which business drivers exist. This approach overcomes some of the issues with frequency and severity modelling described earlier in that:

- Expert judgement is incorporated in a narrative of the business and not by asking experts to pinpoint exact loss frequencies or maximum loss amounts.
- A causal structure and a conditional probability map are used to build up a holistic, joint probability distribution of all losses, thereby removing the need for separate aggregation and correlation assumptions which is another area heavily reliant on expert judgement.
- Different perspectives can be integrated into one narrative of the complex system across individuals, risk categories and over time; as the process is repeated these perspectives are integrated allowing the model to reinforce the most important causal links and reduce individual bias.
- Scenarios are driven by the risk profile of the business, and therefore readily available and interconnected; scenario analysis is not anchored by past analysis or what is current in the insurance zeitgeist.
- The ability to dynamically model capital is based on observed states of important business drivers rather than a static structure that is slow to respond to change within business and regulatory governance structures.

Expert judgement

Within traditional operational risk modelling approaches, experts in particular risk categories are consulted in order to parameterise the operational risk models. These discussions usually focus on qualitative enquiries followed by a handful of quantitative questions. It is an approach that relies on your RMF asking the right questions, in the right way, so as not to elicit defensive responses and therefore be able to make a call on the maturity of risk management processes within each risk category.

The range of quantitative questions may go along the lines of:

- In worst case scenarios, what is the maximum loss that you would expect from risk x?
- Do you consider risk y to be a 1-in-100 year event or 1-in-200?
- How many risk events do you expect over the next 10 years?

Considering what is asked of the experts, it is not surprising that these quantitative measures may be glossed over, or heavy reliance is placed on what was answered the last time this exercise was performed. It is very rarely the case that the modelling parameters are changed as this would require justification. Other than in the most obvious scenarios (think pandemic or increased industry focus on cybercrime) there may not be much to go on to substantiate a change.

Within the causal modelling approach experts are asked to describe the factors which could influence the behaviours and outcomes of the chosen risk scenario, as one would for more traditional modelling approaches. However, for causal modelling it is a narrative discussion that allows the experts to explore ‘what if’ scenarios without the pressure of placing unknown values on unknowable events.

Questions are along the lines of:

- What controls would need to fail before event x could happen?
- How many events must happen at once for event y to occur?
- What are your main concerns?

Cognitive mapping

These questions help to develop a systems structure of a business and its operations through the use of a cognitive map (Figure 15). This is the first step in building a causal model which is essentially a visual representation of the narrative of the problem.

The mapping involves breaking down the story of the complex business system into its constituent elements. Distinct concepts are connected to produce a non-linear, robust, and transparent framework to represent the underlying complexity of the system in a visual format. Concepts are linked across risks, and to higher order effects, which helps to identify how the various parts of the story interact.

FIGURE 15: EXAMPLE OF A COGNITIVE MAP

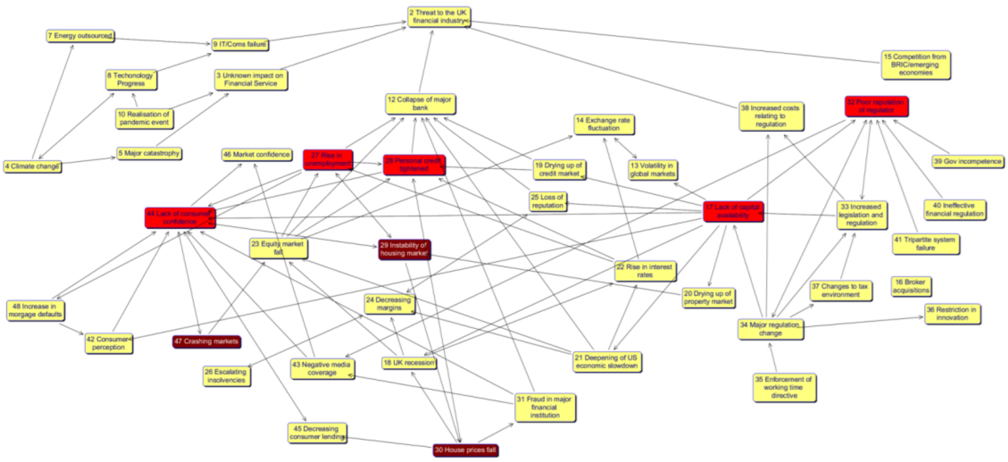

By applying mathematical techniques from graph theory, key nodes are identified (shown in bright red in Figure 15). Key nodes are the parts of the story that explain most of the variation and are most important to the story being told. These nodes directly, or indirectly, link to lots of concepts or outcomes across the cognitive map.

Next, it is possible to identify which beginnings of the story most often lead to these key nodes (shown in dark red in Figure 15). These are the key business drivers that are central to the functioning of an operational activity. These drivers are important concepts as identifying evidence of their occurrence indicates which important part of the story could happen next.

Lastly, a cognitive map also allows identification of the gaps in the story. If a story line jumps unexplained from an early stage in the story to a later one, this may be an indication that parts of the narrative are missing. Gaps may be due to a lack of understanding about how a particular scenario might emerge or that information is missing on how different factors interact to produce a particular outcome. These gaps provide the visual prompt to conduct further investigations in order to complete the narrative.

FIGURE 16: EXAMPLE OF INSIGHTS FROM A COGNITIVE MAP

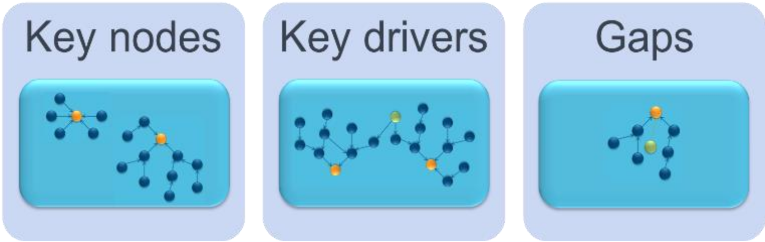

Cognitive mapping therefore provides a holistic, dynamic description of the complex business system with limited cognitive bias. It describes where one thing leads to another, and another. It foresees forces working to keep something in balance or alternatively pushing things out of control. Understanding these dynamics is critical to grasping how risks may spread through the system, but also potentially how they may be controlled or avoided.

Bayesian network

This cognitive map is then converted into a minimally complex version. The challenge is to maintain the most important structures and dynamics of the cognitive map, without destroying the information value captured.

Collapsing onto key nodes provides a robust framework for developing a structural model as a network of interconnected elements. It is upon this subset of concepts that a practical quantitative modelling framework can initially be built where nodes represent a variable, and the edges represent the conditional dependencies between the variables.

Nodes are assigned probability distributions by considering the states that they can be in within the designed framework. Drivers are typically calibrated as discrete probabilities with a smaller number of states, each of which is assigned a probability, which drive the conditional distributions of intermediate nodes. Outcomes will typically be continuous probability distributions of operational losses. By repeating the process, the Bayesian framework links operational outcomes to business drivers all the way through to individual risk factors. The result is a conditional set of outcomes which helps to infer the posterior distribution, meaning that the distribution will update with new information and give an updated probability.

Nodes are calibrated with the use of expert judgement and data, where it is available. The calibration is performed in the context of the narrative described above, i.e., in the context of specific risk environment and controls that exist. The estimates experts need to come up with are split in ‘tangible’ pieces—‘If we are in this state, what would be the probability of this situation be?’—rather than expert judgement of the complex system as a whole.

Consider the sample illustration in Figure 17 where the outcome is a distribution of reserving errors (part of the scenario ‘Model risk’) as a function of accuracy of asset and liability valuations. The drivers of the accuracy of the valuation in this example are accuracy of the asset and policyholder data, appropriateness of assumptions, and adequacy of actuarial and accounting models. Each of the drivers has been assigned discrete probability states which directs the conditional probability distribution of the ultimate outcome. Note, some of the drivers interact with each other as well.

Multiple risk elements can then be brought together to assess total risk without the need for separately defining correlations and/or copulas. The interrelationships between different elements of an operational risk framework are accounted for directly through the causal relationship structure and common drivers.

FIGURE 17: LINKING DRIVERS TO OUTCOMES WITHIN A CAUSAL MODEL

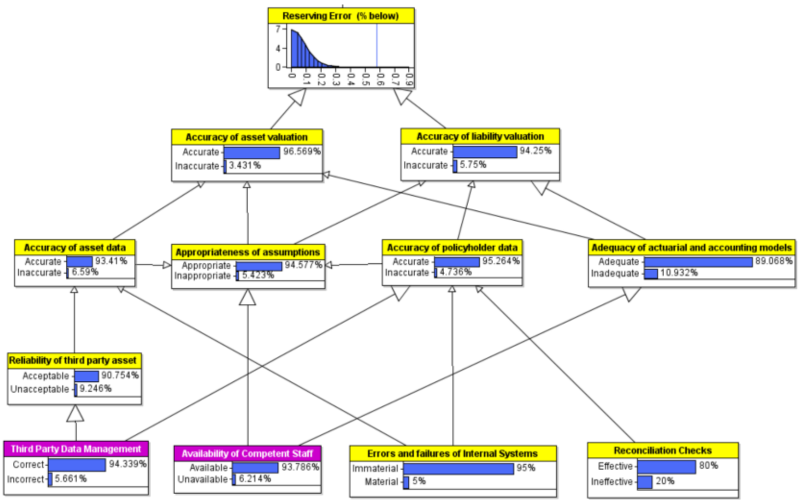

Through the use of causal modelling, what we end up with is a framework that is 'causal' in the sense that it is modelling a network of causes and effects, where relationships between causes and effects are governed by probability distributions. This approach goes beyond looking at frequencies of outcomes to understand 'how' those outcomes are produced. They are particularly suited to assessing capital requirements as a Bayesian framework can simultaneously account for the full range of operational outcomes, both positive and negative, and extremes.

The approach is not without its challenges, being relatively new and less established than frequency and severity modelling. The causal modelling concept, being unfamiliar, may take some time to get 'buy-in' from senior management. It will also require a new set of modelling expertise and modelling tools from existing processes. Regulators are getting more familiar with causal models, but generally the 'status quo' as of now is using frequency and severity modelling, as could be seen from the surveys outlined in the introduction.

Areas of Development in Operational Risk Modelling

Operational risk as a risk category has been under constant development since the days before SII when it was considered a catch-all for 'other' risks. The introduction of SII saw the enhancement of the frequency-severity models discussed earlier in this paper. However, companies are beginning to realise that these past techniques are not sufficient to model the complex nature of operational risk in today's business environments.

Along with increasing the sophistication of operational risk modelling, we are also seeing firms focusing on other areas to improve their modelling and understanding of operational risks. Below we discuss some of these.

OPERATIONAL RESILIENCE

The expression of operational risk as a capital event does not always make sense from a practical perspective. Consider low frequency, high impact events (like pandemics), particularly when capitalised with a frequency-severity approach. The recent regulatory focus on operational resilience[12] is an acknowledgment that certain operational risk events fall under Pillar 2 and not Pillar 1 of SII. Therefore, we expect that there could be a number of lessons learned by leveraging across both areas as insurers develop their operational resilience capabilities.

Firstly, operational resilience is an outcome of effective operational risk management. Therefore, firms have scope to leverage work already done on operational risk management, and quantification, to get a head-start on defining their operational resilience approaches.

[12] Operational resilience is defined by the UK's Prudential Regulation Authority (PRA) as referring to firms' ability to prevent, respond to, recover and learn from operational disruptions as set out within the PRA's Supervisory Statement SS1/21. [Bank of England Prudential Regulation Authority. 2021. 'Operational resilience: Impact tolerances for important business services,' *Supervisory Statement | SS1/21*; https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/supervisory-statement/2021/ss121-march-21.pdf?la=en&hash=C69464DA1603A288F387ADF55F2596004D8640FC, p. 1.]

Additionally, insights gained from further developing operational resilience frameworks can be used to enhance existing operational risk modelling approaches by asking whether or not the existing operational risk framework is driving the right outcomes for the business.

EXTERNAL DATA

The analysis of data is the bedrock of all capital modelling. As discussed earlier in this paper, the lack of internal data and relevance of external data presents a Catch-22 problem for modelling operational risk. Nevertheless, there is increased focus in the market on supplementing firms' internal losses with external loss data, such as ORIC[13] and ORX,[14] particularly since data quality has been an area of challenge by regulators.

External data enables insurers to capture any relevant events based on the experience of others, and help challenge expert judgement calibrations. The key to this is to ensure that external data used is adjusted so that it is proportionate to the firm, for example, scaling the losses based on the size of the firm and filtering the type of events that are relevant.

INSURANCE RECOVERIES

A continuing area of development is that of allowances for insurance recoveries when quantifying operational risk capital. Under Pillar 1 SII requirements, insurance can be used to offset the operational risk contribution to the SCR, provided that credit risk and other risks arising from the use of such mitigations are properly reflected in the capital calculations undertaken.[15]

Not only can operational risk capital modelling inform the optimal risk transfer strategy and insurance program to follow, it can also be used to potentially reduce capital requirements. While most firms make use of insurance programs to reduce their operational risks, we have not seen many clients making full use of these benefits within their operational risk capital modelling frameworks. Insurance recoveries, for example, could be considered when setting parameters for modelling purposes.

The way in which insurance recoveries can be incorporated will depend on a firm's modelling approach. However, a few factors should be considered, such as any limitations and exclusions, sum insured limits across a number of operational risk scenarios, and the counterparty default risk of the insurer.

DEPENDENCIES

One of the key challenges for operational risk modelling is the modelling of dependencies across operational risks. However, dependencies with non-operational risks are another area in which practices continue to develop. It is unlikely that all operational risks will crystallise at the same time. However, if we take away one lesson from the COVID pandemic, it is the interconnectedness of operational risks with market, credit, and insurance. We are starting to see firms pick up on this point in their capital modelling approaches to avoid the under- or over-estimation of dependencies.

The solution though is not clear cut, particularly with more traditional frequency and severity modelling approaches. Consider for example the asymmetry of operational risk correlations; while a non-operational risk event could impact a firm's operational risk, it is unlikely that operational risks would have a wider impact on market and claims experience. Correlation assumptions are typically scaled back to allow for this asymmetry.

---

[13] Operational Risk Consortium Limited International (ORIC), https://www.oricinternational.com.

[14] Operational Riskdata eXchange Association (ORX), https://managingrisktogether.orx.org/.

[15] Directive 2009/138/EC of the European Parliament and of the Council of 25 November 2009 on the taking-up and pursuit of the business of Insurance and Reinsurance (Solvency II), 30/06/2021, Article 101(5), https://eur-lex.europa.eu/eli/dir/2009/138/2021-06-30.

Conclusion

Traditional risk quantification models provide a range of outcomes depending on the assumptions made regarding model distributions and parameters, mostly based on expert judgement. Additionally, these models generally rely on historical loss data, which can be hard to come by for many operational risk events and cannot account for future risks in a fast-changing risk landscape like the current business environment.

Causal modelling bypasses many of the weaknesses and difficulties with more traditional modelling approaches for understanding from where risk can originate in your business model and how it connects to the wider picture. It does not rely solely on past events and expert judgement; instead it enables insights into future events by understanding the drivers of risk. This in turn enables effective capitalisation for operational risk driven by data insights.

Additionally, traditional approaches, such as frequency and severity modelling, do not solve the business problem on their own and need to be developed for very specific solutions. This may result in multiple models for various operational risk types that require aggregation in an attempt to produce a complete picture.

A causal model, on the other hand, produces a holistic spectrum of all losses and risk paths that can be used for multiple purposes. By understanding where, and how, risk can propagate through your organisation, more effective risk management solutions can be implemented and more insightful stress scenarios can be inferred, which is not only useful for P(IM) quantifications but for ORSA reporting as well.

Operational risks are complex and the field of modelling these risks is still maturing in both technique and scope. Areas that are the current focus for development include regulatory attention on operational resilience and how this integrates with firms' operational risk management approaches.

CONTACT

Adél Drew
Adel.Drew@milliman.com

Tatiana Egoshina
Tatiana.Egoshina@milliman.com

Jack Huffer
Jack.Huffer@milliman.com