Deloitte.

ORIC INTERNATIONAL

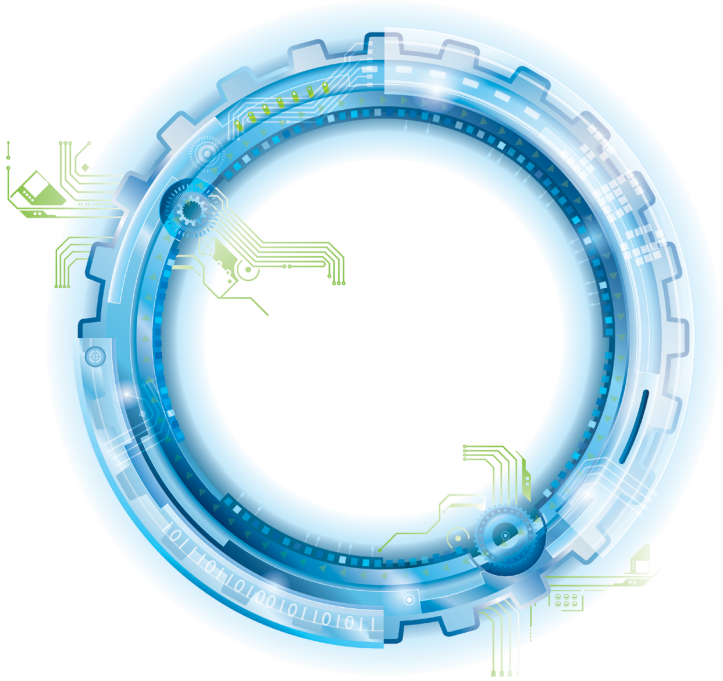

Utilising data to enhance operational risk scenario and capital assessments for investment management firms

White paper – executive summary | March 2026

Contents

Executive summary 03

1.1. Introduction 03

1.2. Headline findings 05

1.3. Using our findings within your scenario workshops 09

1.4. Key technical findings 14

1.5. Conclusion 16

Expert contributors and credentials 17

Contacts and acknowledgements 18

Endnotes 19

This executive summary provides an overview of our key findings and outcomes. A comprehensive, 77-page technical version of the white paper, offering a detailed exploration of these topics, is available via email request. In this document, references to the white paper may refer either to this executive summary or to the full technical version, depending on context. Where we intend to refer to both collectively we use the term white paper.

Please note: As a core principle of data consortia is that the benefits of shared data are reserved for contributing members, certain data elements have been redacted from this version of the report. Organisations interested in accessing the full, unredacted analysis may contact ORIC for further information.

Click to request the full paper

This is an interactive document. Click below to navigate.

1 Executive summary

1.1. Introduction

Operational risk scenario analysis plays a critical role in strengthening risk governance and informing capital assessments across investment management firms. However, many firms continue to face challenges in consistently integrating internal loss event data (ILD), external loss event data (ELD), and expert judgement within their scenario and capital frameworks. This can lead to scenario outcomes that are less comparable over time, less defensible to senior stakeholders, more difficult to assess consistently in a supervisory context, and less useful in informing capital decisions and driving actionable risk management decisions.

A key driver of this inconsistency is the absence of a clearly articulated and widely adopted, end-to-end implementation architecture for scenario-based operational risk capital in the investment management sector. Critical design decisions, including scenario risk paradigm selection, severity quantile definition, integration of ILD and ELD, and the treatment of distributional assumptions, are often embedded within practice without explicit comparison of alternatives or systematic testing of their implications. As these choices can materially influence capital outcomes, variation in implementation can lead to materially different capital assessments across firms facing similar underlying risk exposures.

In this white paper, Deloitte and ORIC International (ORIC) have built on a long-standing relationship and shared engagement with investment management firms. With many investment management firms engaging with both ORIC's data consortium and Deloitte's operational risk programmes, we have repeatedly observed common challenges in how loss data is interpreted and translated into scenario design, calibration and capital outcomes. In that context, Deloitte and ORIC partnered to apply a structured and disciplined analytical approach to ORIC's investment management loss dataset, testing hypotheses against observed loss experience and jointly interpreting the findings. The insights and techniques in this white paper are shared in the same spirit that underpins ORIC's practitioner-led network: openness, shared learning and data-driven insight with the aim of supporting continuous improvement in scenario analysis and operational risk decision-making across the sector.

This white paper presents the findings of that analysis, based on over 2,000 material operational loss events from 2005 to 2025 drawn from ORIC's investment management loss dataset. The paper has two objectives. First, it examines the operational risk loss landscape of the investment management sector as reflected in this dataset, highlighting the scenarios, risk drivers and structural patterns that most consistently shape observed loss experience. Second, it sets out a structured implementation framework for scenario-based capital modelling. The analysis is grounded in observed loss experience where empirical evidence is informative; however, scenario-based capital frameworks also involve structural design choices that cannot be resolved by data alone. We therefore examine alternative approaches explicitly, assess their implications for calibration stability and capital impact, and provide a transparent, repeatable and governance-aligned method for integrating ILD, ELD and expert judgement into scenario analysis and capital assessment.

This executive summary focuses primarily on the headline findings and key implications. The full technical version provides the detailed step-by-step framework and practical application.

While the analysis is grounded in investment management-specific loss data and written with investment management firms' scenario and capital processes in mind, many of the concepts, analytical techniques, and governance considerations are likely to be relevant to other regulated sectors. In particular, banks, undertaking Internal Capital Adequacy Assessment Process (ICAAP) submissions, and insurers, performing Own Risk and Solvency Assessment (ORSA) exercises, may find elements of the approach informative, where similar challenges arise in combining loss data, expert judgement, and capital modelling.

This white paper is designed for three primary audiences:

C-suite leaders, boards, and heads of operational risk, who can use the findings to benchmark scenario selection, capital drivers, and risk governance against broader industry patterns (see section 1.2: Headline findings).

Operational risk practitioners facilitating scenario workshops, who can apply scenario-level insights (including typical loss frequency and severity behaviour) to improve consistency, reduce bias, and strengthen challenge and discussion (see section 1.3: Using our findings within your scenario workshops).

Operational risk data and modelling professionals, who can use the techniques and empirical observations to inform data processing, scenario mapping, model calibration, validation, and ongoing governance (see section 1.4: Key technical findings).

Click the box above to go directly to that section.

Key features of this white paper include:

1. practical guidance to support the consistent categorisation of complex operational loss events into relevant scenarios;
2. discussion of the analytical and judgement-based decisions which firms face when interpreting loss experience - including frequency, severity magnitude and severity volatility – and the implications these decisions can have for capital outcomes; and
3. benchmarking insights intended to complement, rather than replace, firm-specific analysis and expert judgement.

Differences in business models, control environments, risk appetites, and materiality thresholds mean that no external dataset can provide a one-size-fits-all solution. Accordingly, this white paper presents a structured and evidence-based implementation framework to support more consistent, transparent and defensible scenario definition and calibration, together with stronger governance and validation practices in operational risk assessment.

The framework is intended to complement, not replace, firm-specific analysis, internal data, and expert judgement. Where helpful, the paper introduces practical working concepts and terminology used consistently throughout to support clearer discussion and more repeatable scenario workshops and calibration exercises.

The paper does not propose new academic statistical methods. Instead, it focuses on strengthening implementation discipline, benchmarking interpretation, and decision transparency when applying operational loss data and scenario analysis in practice.

We hope that by engaging with these insights, investment management firms can strengthen operational risk practices through more consistent scenario definition and calibration, clearer benchmarking and challenge, and more robust validation and backtesting, helping to align internal approaches with evolving sector standards.

Caroline Coombe
Chief Executive, ORIC International

Damian Hales
Partner, Deloitte

1.2. Headline findings

This section is primarily written for **C-suite leaders, boards, and heads of operational risk,** and is intended as a stand-alone overview, setting out the key insights, implications, and conclusions from the analysis that follows.

Leveraging sector-wide data from over 2,000 material operational loss events reported by investment management firms between 2005 and 2025, this joint analysis by Deloitte and the ORIC consortium sets out a unique perspective on operational risk in the sector. Each loss event has been rigorously categorised against a risk-sensitive scenario, enabling an in-depth view of sector risk frequencies, severities, and capital relevance. The white paper provides both substantive findings on investment management operational risk profiles and a practical example of how firms can process and integrate internal and external loss data into scenario analysis and capital modelling.

While we believe the methods and analysis we have laid out in this white paper to be robust and rigorous, such that many firms will find the resulting insights to be useful and relevant to them, there are inevitably limitations with any such analysis, which are summarised in the full technical version. The methodologies described in this white paper have been developed and refined through extensive practitioner input and challenge as summarised in section 2.

This analysis yields three headline findings with associated recommendations:

1. Under-recognised 'core' scenarios

- **Finding –** our analysis, together with recent ORIC scenario benchmarking[1], indicates that many investment management firms may be underestimating the significance of certain high-frequency, operationally 'routine' scenarios, such as fund administration errors and extended IT/system failures. These scenarios comprise a substantial proportion of observed events and exhibit combinations of frequency, severity magnitude, and severity volatility that suggest a potentially material impact on operational risk capital. In practice, such scenarios are sometimes treated as 'business-as-usual' rather than as core capital drivers.
- **Evidence –** see Table 1 for scenario-level reference points and illustrative capital contribution ranges, which highlight the potential for these 'routine' scenarios to act as core capital drivers.
- **Recommendation –** firms should assess whether these 'core' scenarios are appropriately reflected within their scenario sets, calibrations and capital assessments.

2. Limitations in capturing high-sensitivity risks

- **Finding –** the analysis also points to clear limitations in loss event databases when it comes to capturing rarer, higher-profile or high-sensitivity loss events such as cyber attacks, regulatory fines, and some forms of conduct risk. Although such risks events are widely recognised for their potential severity, from a capital perspective, they are weakly represented in reported risk data. This is due to both their low frequency and the sensitivity-linked challenges associated with recording and sharing such events.
- **Evidence –** see Table 2 for additional high-sensitivity scenarios that are weakly represented in the loss dataset and therefore require structured expert judgement overlays.
- **Recommendation –** firms should not rely solely on reported loss event data when assessing higher-sensitivity scenarios (e.g. cyber attacks, regulatory fines and certain conduct risks), as these risks are often under-represented in loss databases. In the near term, firms should apply structured expert judgement overlays (supported by scenario analysis, control assessments, and external intelligence) to support robust capital assessments. Over time, improving internal event capture, including consideration of additional sanitised event reporting processes, as well as participation in appropriately anonymised industry data sharing, will be important in strengthening the evidence base for these risks.

1. ORIC Capital Benchmark Report for investment management firms, 2025

3. Enabling robust expert assessments

- **Finding –** for more commonly reported types of operational loss events, the white paper demonstrates that external data, such as that used here, can play a valuable role in supporting and calibrating expert assessments. When processed at an appropriate scenario level, the consortium data can offer useful benchmarks for relative frequency, severity magnitude and severity volatility.
- **Evidence –** see section 1.3 for a worked example showing how external loss data can be used to support scenario calibration and workshop challenge.
- **Recommendation –** firms should use external loss data at an appropriate scenario level to support and calibrate expert judgement, reduce bias, and strengthen challenge within scenario workshops. Where combined with firm-specific context (business model, control environment, and scenario approach), scenario-level benchmarking can improve consistency and defensibility of capital outcomes while maintaining the flexibility needed for firm-specific calibration. In the next section, a detailed example is provided to illustrate precisely how this can be achieved.

The above headline findings are grounded in scenario-tagged loss experience and are explored further through the risk factor framework, below. Table 1 and Table 2 compare scenarios in terms of frequency, severity magnitude and severity volatility, and provide illustrative ranges for their potential contribution to operational risk capital. They also include ORIC member benchmark allocations, enabling readers to compare between data-derived scenario signals with observed industry calibration patterns. Actual capital contributions may reasonably fall outside the illustrative ranges depending on firms' business models, control environments and risk appetite.

Risk Factors

The tables below summarise key insights from the dataset using the three core risk factors commonly applied in operational risk scenario and capital frameworks:[2]

frequency - how many material events are expected to occur per year for a given scenario;

severity magnitude – the typical scale of higher-severity losses associated with the scenario, measured at a defined upper quantile of the severity distribution; and

severity volatility – the degree of dispersion of loss sizes when events occur, measured as the ratio of higher-severity to median loss and used as an indicator of tail variability.

Each loss event in the dataset was mapped to a relevant scenario based on its causal description and supporting information. This scenario tagging enables the results to be used as structured, scenario-level benchmarks when calibrating comparable scenarios within firm frameworks.

Based on the scenario-tagged data, seven scenarios, collectively representing over 90% of material events, were analysed. For these scenarios, we present relative classifications across the three risk factors alongside illustrative scenario capital contribution percents for a typical investment management firm.

These capital contribution ranges are intentionally broad. They reflect variation across firms in business model, control environment, operating scale, materiality thresholds, and scenario design choices, as well as the fact that the analysis focuses on relative scenario contribution rather than absolute capital levels. They are intended to support sense-checking and structured challenge, not to prescribe target allocations.

Firms using loss distribution approach (LDA) or similar stochastic models[3,4] are likely to derive the greatest value from the risk factor classifications, which map naturally into frequency, severity magnitude, and severity volatility/dispersion parameters. Firms using alternative or simplified approaches may instead find the illustrative percentage contribution ranges more directly informative as benchmarking context.

2. This framework uses a structured severity approach designed to support consistency and comparability across scenarios. While alternative modelling techniques exist, our focus is on aligning severity magnitude and upper-loss dispersion with the largest observed events, providing a robust and transparent benchmark for scenario calibration.
3. Originally formalised in the Basel II advanced measurement approach framework, though now widely adapted across sectors in a simplified form.
4. McNeil, A.J., Frey, R., Embrechts, P. (2015) – Quantitative Risk Management: Concepts, Techniques and Tools (2nd ed.), Princeton University Press

Table 1: Risk factor classification of key operational risk scenarios for investment management firms[5,6]

Scenario Name	Frequency	Severity magnitude	Severity volatility	Illustrative scenario capital contribution %	ORIC scenario benchmarking % of capital
01. Fund admin error	High	High	High	10%–40%	8%
02. Extended IT/system failures	High	High	Medium	5%–25%	7%
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]

In addition, five lower-frequency scenarios that are typically considered material from a capital perspective were assessed using structured expert judgement, recognising their limited representation in loss data. Together, the data-driven and expert-driven scenario classifications provide a coherent view of how operational risk manifests across the investment management sector and how different types of scenarios contribute to capital risk in different ways.

Table 2: Risk factor classification of additional operational risk scenarios

Additional scenario name	Frequency	Severity magnitude	Severity volatility	Illustrative scenario capital contribution %	ORIC scenario benchmarking % of capital
08. Cyber attack	Low	High	High	3%–15%	6%
09. Regulatory fine	Low	High	High	3%–15%	14%
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]

5. Where scenario classifications were ambiguous, a broader High/Medium/Low range has been applied to ensure robustness. See the full technical version, for a worked examples showing how each of these risk factors maps to the High/Medium/Low boundary ranges for the severity quantiles and materiality thresholds used in this white paper.
6. Firms may wish to tailor this table to reflect their specific characteristics or firm type. An illustrative example of how to do so is provided in the full technical version.

The white paper sets out four practical ways in which firms can leverage the insights summarised above:

1. As external benchmarks

- Firms may use the results to sense-check their existing scenarios and capital outcomes, i.e. by comparing risk factor rankings and the illustrative percentage contribution of each scenario to total capital to Table 1 and Table 2 above. Please note that, before benchmarking, firms should consider adjusting scenario weights to reflect their own business model, operational risk profile and risk appetite.
- An illustrative example of how scenario weights may be adjusted to reflect firm-specific business models, operational risk profiles and risk appetite is provided in the full technical version.

2. To inform expert judgement

- To provide valuable context for participants in scenario workshops, helping to ground expert assessments for firms using capital models.
- Section 1.3 provides a worked example illustrating how data can support frequency and severity calibration within scenario workshops.
- For firms using the scenario summation approach, the full technical version highlights the practical challenges of integrating structured data where capital is derived directly from expert-estimated extreme loss outcomes.

3. As a processing example

- To demonstrate how internal loss data can be systematically processed for use in scenario analysis to support and calibrate expert judgement, reduce bias, and strengthen challenge.
- Further detail on the data processing methodology underpinning this example is set out in the full technical version.

4. As a framework for utilising external data

- To demonstrate how a structured approach to harnessing external data sources, including further benefits from engaging with industry consortia such as ORIC International, can enhance the overall operational risk capital assessment process.
- Section 1.3, below, provides a worked example for use in scenario workshops, while the full technical version sets out the broader framework, including scenario classification, updating of external insights for firm-specific dynamics and backtesting.

Overall, these approaches demonstrate that disciplined scenario design, meticulous loss-data handling, and transparent integration of empirical data with expert judgement can significantly improve the credibility, consistency, and regulatory robustness of operational risk scenario and capital assessments. When applied effectively, such techniques support more resilient Internal Capital Adequacy and Risk Assessment (ICARA) or Overall Risk Assessment (ORA)[7] outcomes and strengthen governance and decision-making surrounding operational risk within investment management firms.

To quote the British mathematician, John Edensor Littlewood:

> The paradoxical feature of the laws of probability is that they make unlikely events happen unexpectedly often.
>
> **– John Edensor Littlewood**

7. As noted in CP25/42, the FCA has updated the terminology for the internal capital adequacy and risk assessment (ICARA) process to the 'Overall Risk Assessment' process.

1.3. Using our findings within your scenario workshops

This section is primarily written for **operational risk practitioners facilitating scenario workshops,** who can apply the scenario-level insights (including typical loss frequency and severity behaviour) to improve consistency, reduce bias, and strengthen challenge and discussion.

Across the investment management sector, firms use a range of calibration approaches within their operational risk scenario and capital frameworks. Some firms rely heavily on structured scenario workshops as a primary calibration input. Others use internal and external loss data directly to parameterise their models, with expert workshops serving mainly as a challenge, validation, and overlay mechanism. Many firms operate a hybrid of the two approaches.

The guidance in this section will be directly applicable to firms whose frameworks place significant weight on expert scenario workshops, including those using scenario-based loss distribution approach (LDA) Monte-Carlo models where experts provide key calibration inputs for frequency and severity dynamics consistent with established quantitative risk management practices[8, 9]. However, the same scenario-level benchmarks and risk factor mappings are also relevant to more data-driven firms, where internal and external loss data flow directly into model calibration, and scenario workshops are used to interpret, challenge and refine model outputs, rather than to generate parameters from first principles. In both cases, disciplined expert input can be strengthened through structured use of internal and external data.

Disciplined expert input

While both internal and external data can provide a foundation for credible scenario design and calibration, expert input remains essential, particularly for risks where data is sparse or events are genuinely rare. However, expert judgement is also where psychological biases, anchoring effects and narrative drift most frequently enter the capital framework. For this reason, expert input should be structured, benchmarked, and evidence-informed, rather than unconstrained.

One effective approach is to anchor experts to risk factor ranges implied by internal and external data. When used effectively, ILD and ELD allow firms to move from asking:

"What do we think our risk profile is?"

to

"Our ILD suggests our firms' risk profile is Y, while ELD suggests the typical firm's risk profile is X. Where does our firm's unique business strategy, risk profile and control environment mean we are genuinely different, and why? What impact would we expect this difference to have on our experience of the frequency and severity of loss events?"

Note: while there is debate about the potential for anchoring effects when presenting data ranges to experts, Table 1, above, and the discrepancy between empirical and ORIC scenario benchmarking results, demonstrates that scenario assessments not grounded in empirical internal and external data are more vulnerable to behavioural and narrative biases. Data-informed framing is therefore recommended as a bias-control mechanism within expert elicitation.

8. McNeil, A.J., Frey, R., Embrechts, P. (2015) – Quantitative Risk Management: Concepts, Techniques and Tools (2nd ed.), Princeton University Press
9. Cruz, M., Peters, G., Shevchenko, P. (2015) – Fundamental Aspects of operational risk and Insurance Analytics, Wiley.

This section sets out how the frequency and severity risk factor High/Medium/Low rankings from Table 1 and Table 2 in section 1.2, above, can be applied within operational risk scenario workshops. It is intended as practical guidance for firms using LDA or similar models, illustrating how the empirical insights can be translated into structured, bias-mitigating inputs to support independent expert frequency and severity calibration.

Firms adopting a scenario summation approach, in which capital is derived directly from expert-estimated extreme loss outcomes, should consider how structured data can be used to benchmark, scale and plausibility-test those tail estimates; further methodological detail is set out in the full technical version.

Assumptions:

1. The firm is using an LDA model, with frequency, severity magnitude and severity volatility (the risk factors) assessed independently.
2. The firm possesses sufficient relevant internal and external loss data to inform all three risk factors for the chosen scenario. In practice, it is often the case that adequate data is only available for a subset of risk factors or, occasionally, for none at all. In such instances, practitioners should adapt accordingly, omitting or adjusting those components where the supporting data is lacking.

There are three key components that firms should bring into each scenario workshop:

1. **A cleaned chronology** of historic scenario-tagged loss events, sorted by date and magnitude. Where data volumes are low, e.g. <10 events, it may be possible to go through each in the workshops, whereas, where data volumes are higher, we suggest firms only show causal descriptions for the larger events and show summary statistics and appropriate visuals/charts to capture the insights of smaller events.
2. **Side-by-side internal/external loss data implications** for frequency and severity, benchmarked to the chosen scenario approach (e.g. materiality thresholds, severity quantiles).
3. **Visualisation (optional but recommended)** such as time-series charts of events and losses, and severity probability distributions demonstrating the most frequently observed loss amounts historically.

These components must be aligned with your firm's scenario framework. In this example, we use:

- a materiality threshold of £25k;
- higher severity loss measured as the worst-of-four-events at the 87.5th quantile[10]; and
- lower severity loss measured as the typical events at the 50th quantile.

While it is generally straightforward to address frequency and severity magnitude (defined by the higher severity loss) risk factors, the severity volatility risk factor is often more challenging, as it depends on the ratio between the higher and lower severity losses. Our recommended approach is to establish the higher severity loss first, then the expected ratio, before deriving the expected level of typical or the 50th quantile lower severity loss. Data-driven analysis is particularly valuable here and can be used to challenge both the higher and lower severity loss simultaneously.

10. Note: this approach to higher severity quantile is a type of frequency-aligned-worst-case (FAWC) risk capture, which focuses expert assessment upon imaginable event quantiles without involving a time dimension (e.g. 1-in-X-years). The worst-of-four-events is selected as, even at the commonly applied lower bound frequency of 1-in-10-years, it will not take more than 40 years for four events to occur. We also note that some readers may expect to find the worst-of-four-events associated with the 75th quantile; however, we have used 87.5th as the midpoint of the last quartile of the probability range.

Figure 1 below provides a worked example of how firms might use internal and external data within their scenario workshops.

Figure 1: Example use of internal and external loss data in operational risk scenario workshops

Scenario Event History

ID	Date	Loss (£k)	Description
1	2005	100	Trader fat finger error
2	2009	50	Trader bought instead of sold
3	2015	75	Small trade error with large market mvmt.
4	2022	300	Large trade error with large market mvmt.
5	2024	30	Large trade error with small market mvmt.

Key data questions for the experts

- Should any events be excluded?
- Any missing events or near misses?
- Is the data representative of our expected future risk profile?

Scenario Empirical Loss Distribution

ID	Date	Loss (£k)	1 in X event	Empirical CDF
4	2022	300	1 in 5.0	90%
1	2005	100	1 in 2.5	70%
3	2015	75	1 in 1.7	50%
2	2009	50	1 in 1.3	30%
5	2024	30	1 in 1.1	10%

Empirical loss observations

- Median loss: £75k
- 90th quantile loss: £300k
- The 87.5th quantile (worst-of-four-events) can be approximated from the empirical data (suggesting a value below £300k), or estimated using a fitted lognormal distribution, which implies ~£260k

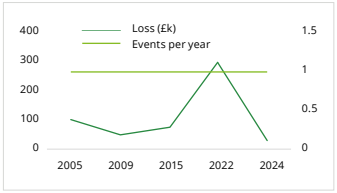

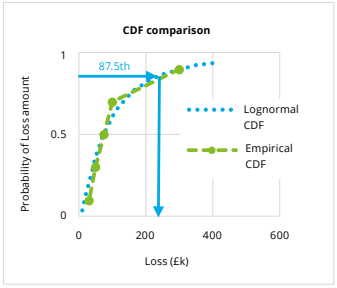

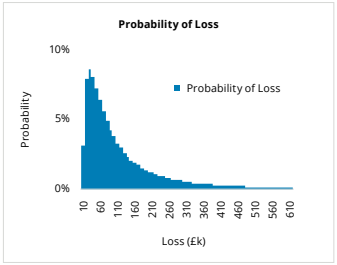

Insights from the data

	ILD	ELD
1.Frequency	With five in-scope events in the last 20 years, the data implied frequency is **0.25,** i.e. one event every four years.	The scenario is ranked as low frequency, i.e. for most firms it's rarer than **once every two years (0.5)**.
Workshop Question: Given ILD & ELD agree this is a low frequency scenario, are we happy that this is likely to continue? If so, what frequency less than 1-in-2 years should we use?		
2. Higher-severity-loss (1-in-4, i.e. 87.5th)	The worst-of-four-events (87.5th quantile) sat **between £100k and £300k (lognormal fitting suggests £260k)**.	The scenario is ranked as high severity magnitude, i.e. for most firms, the 1-in-4 would be **> £600k**.
Workshop Question: Our ILD has suggested a milder experience than external firms. Does that align with our understanding of our risk profile, such that we can use a relatively low 1-in-4?		
3. Typical event	The median in-scope event was £75k, implying a 1-in-4 typical **between 1.3 and 4**.	The scenario is ranked as low severity volatility, i.e. for most firms its 1-in-4 typical would be **< 4**.
Workshop Question: Given ILD & ELD agree this is a low severity volatility scenario, are we happy that this likely to continue and thus to use a relatively high typical loss to capture this?		

When internal and external loss data are present in a disciplined format within scenario workshops, as illustrated in Figure 1, expert input is anchored within a clear, empirical context, producing more stable frequency assessments, more credible estimates of extreme/upper-tail losses, and fewer inconsistencies.

Note that, as supporting data volumes are sometimes relatively low, there is a risk that simply extracting specific quantiles (e.g. worst-event-out-of-10) may, by chance, be unrepresentative of the wider loss profile – in our testing (see full technical version for methodological detail), this dynamic was observed in approximately 25% of cases. In such cases, fitting a parametric distribution (e.g. Lognormal) can be used to smooth sampling noise and reduce the reliance on individual loss observations, supporting a more stable basis for risk factor classification; see full technical version for a description of how we have used such methods in a limited manner for the 25% of cases where we identified this dynamic.[11]

To support practitioners, this white paper provides lookup tables (Table 3 to Table 6) mapping the High/Medium/Low bands from Table 1 and Table 2 in section 1.2 to each risk factor based on external loss data relevant for scenario workshops, as illustrated in Figure 1. Cells highlighted in green have been used in Figure 1 based on the illustrative firm's workshop approach (£25k materiality threshold, 87.5th higher severity quantile, and 50th lower severity quantile).

Table 3: Annual frequency risk class lookup

	Materiality Threshold				
Frequency	**£10K**	**£25K**	**£50K**	**£75K**	**£100K**
High		> 1			
Medium		Everything in between			
Low		< 0.5			

Using Table 3: *for example, a firm which uses a £25k materiality threshold, and had identified their cyber scenario as being Low for frequency in Table 1 or Table 2, would interpret the industry norm for cyber frequency to lie below 0.5 events per year.*

Table 4: Severity magnitude risk class lookup at 87.5th quantile

	Materiality Threshold				
Severity magnitude	**£10K**	**£25K**	**£50K**	**£75K**	**£100K**
High		> £600K			
Medium		Everything in between			
Low		< £300K			

Table 5: Severity magnitude risk class lookup at 95th quantile

	Materiality Threshold				
Severity magnitude	**£10K**	**£25K**	**£50K**	**£75K**	**£100K**
High					
Medium		Everything in between			
Low					

Using Table 4 and Table 5: *for example, a firm which uses a £25k materiality threshold and an 87.5th higher severity quantile, and had identified their Cyber scenario as being high for Severity Magnitude in Table 1 or Table 2 above, would interpret the industry norm for Cyber Severity Magnitude to lie above £600k per event.*

11. We note that many in the industry prefer to minimise complexity in data processing and, as such, prefer use of empirical rather than parametrically fitted quantiles. We empathise with this highly defensible position, but also believe our findings here present evidence for the benefits of the alternative approach, as described further in the full technical version.

Table 6: Severity volatility risk class lookup

Severity volatility	87.5th/50th	95th/50th
High	> 9	
Medium	Everything in between	
Low	< 4	

Using Table 6: *for example, a firm which uses a £25k materiality threshold and an 87.5th higher severity quantile, and had identified their cyber scenario as being low for severity volatility in Table 1 or Table 2 above, would interpret the industry norm for cyber severity volatility to lie below 4 (i.e. higher severity loss divided by median loss below 4).*

Given the wide range of materiality thresholds and severity quantile definitions observed across the sector, it is not practical to provide an exhaustive set of combinations; however, various simple methods exist to extrapolate from the results shown here to other combinations. Firms with access to external loss event data, including ORIC members, may tailor these to match their scenario approach and are encouraged to replicate this analysis for their own frameworks. For further advice, please contact the authors.

1.4. Key technical findings

This section is directed primarily at **operational risk data and modelling professionals** and summarises some of the more technical findings arising from the research, analysis and discussion that has arisen during the process of analysing ORIC's specific investment management dataset, while simultaneously acting as a summary of the full technical version underpinning this executive summary's findings. In addition to the primary findings of the executive summary, several other important insights are introduced below.

1. Scenario tagging at scale

One of the major challenges in conducting analysis of this nature lies in tagging thousands of events based on causal descriptions. Based on considerable experience in conducting this task for various firms, we recommend the following:

a. Establish a small, operational risk experienced, and collaborative team (our 'core analytical team' as described in section 3) with dedicated time linked to number of events requiring tagging.
b. Establish clear scenario definitions, addressing edge cases and refining them through regular team discussions.
c. Conduct training and interactive sessions, where each team member tags a sample, presents back, and engages in structured discussion to uncover any subtle discrepancies in interpretation.
d. Allow multiple tags and 'unclear' options, to be revisited and resolved later.
e. Use four-eyes checks to challenge tags and refer disagreements for independent review.
f. Assign random samples of events to each team member and monitor consistency in the number of scenario tags per person to identify cases warranting further review.
g. Document all tagging judgements to support auditability, challenge and facilitate future review.

2. Scenario approach risk paradigm matters

For firms using loss distribution approach models or similar, the choice of risk paradigm to apply within their scenario approach greatly affects operational risk capital assessments.

a. The **nightmare capture (NC) method** asks specialists to estimate robust 'worst-case' losses, prioritising rare and plausible tail risks, while explicitly assuming that experts cannot reliably internalise or apply frequency information when reasoning about extreme events, and therefore should not be expected to do so. While inviting strong stress testing, this approach can lead to experts implicitly estimating events beyond working human lifetime or even capital severities, which may meet regulatory scepticism if not anchored in suitable supporting data/evidence.
b. By contrast, the **frequency-aligned-worst-case (FAWC) method** ties severe loss estimates to observed event frequencies, anchoring probabilities within plausible human lifetimes and enhancing regulatory acceptance. However, firms need to be careful not to implement approaches which overestimate experts' ability to reliably internalise or apply frequency information when reasoning about extreme events, and to be aware that such approaches may have more potential to understate the risk of extreme outliers. The optimal framework depends on data richness, regulatory expectations, and firm-specific risk sensitivities.

3. Distinguishing severity

Differentiating severity volatilities across scenarios remains a key challenge during scenario workshops. External data can play a valuable role here, helping to anchor estimations of relative dispersion rather than absolute severity.

4. Materiality and methodological choices

Decisions regarding materiality thresholds, severity quantiles and the risk paradigm used in workshop questions can have highly material and often unintentionally distortive impact when loss history is linked to expert assessments. To mitigate this risk, expert judgement is elicited through structured, workshops, with challenge from multiple perspectives, clear documentation of assumptions, and governance review to reduce individual bias and support consistency.

5. Smoothing and extrapolation via parametric distributions

Consistent with established quantitative risk management practices,[12] some of the methodological issues above can be addressed by applying parametric distributions to smooth, gap-fill and extrapolate from observed loss histories. However, given the common focus of capital models on simulated single large loss events, ensuring plausible and intuitive alignment of these parametric distributions with the largest historical loss events is critical, whether achieved through expert-judgement, or statistical techniques.

12. McNeil, A.J., Frey, R., Embrechts, P. (2015) – Quantitative Risk Management: Concepts, Techniques and Tools (2nd ed.), Princeton University Press

6. Suitability of the Lognormal distribution

Consistent with established quantitative risk management practices,[13, 14] the Lognormal distribution, in particular, is found to have desirable 'tail-decay' properties that make it a practical modelling option for smoothing, gap-filling and extrapolating loss severities in this context. However, it should be noted that it is not intended to prescribe a preferred distribution; firms should test, validate and justify the choice of severity distribution considering their own data, scenarios and modelling objectives.

7. Limited impact of smaller events

Consistent with established quantitative risk management practices,[15] our testing suggests that expanding the analysis to include smaller events that are often excluded due to the data-collection thresholds ('truncation bias') would not significantly alter outcomes. This suggests that smaller loss events are less relevant to capital outcomes suggesting that modelling effort is often better spent on upper-tail calibration rather than extending coverage of low-severity losses.

8. Backtesting

Although not a principle finding from our data analysis, our statistical experts observe a general lack in robust or methodologically sound backtesting. It is worth noting that the data processing methods presented in this white paper can, and should, be leveraged to support improved backtesting practices. We recommend viewing backtesting not primarily as a formal statistical pass/fail exercise, but as a structured, judgement-based validation tool to (i) assess model conservatism, (ii) evaluate broad alignment between historical loss experience and scenario-driven outcomes, and (iii) require explicit explanation where forward-looking assumptions differ from observed history. Formal statistical backtesting tests have recognised limitations for tail-focused and scenario-driven risk models.[16]

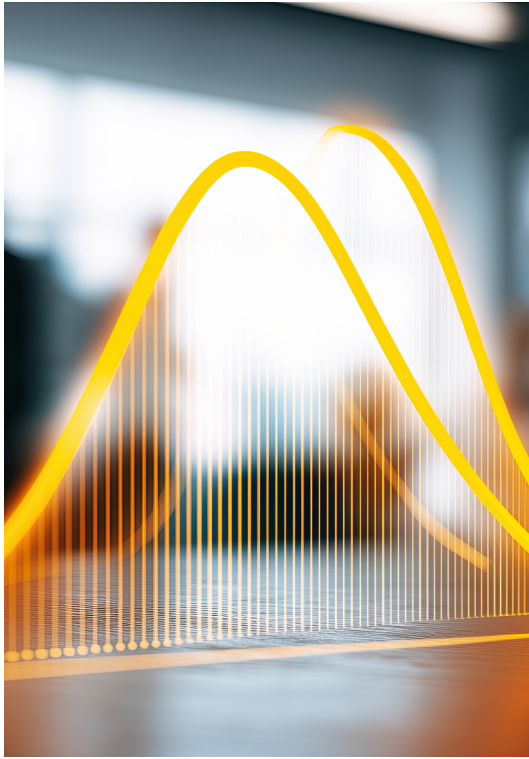

13. McNeil, A.J., Frey, R., Embrechts, P. (2015) – Quantitative Risk Management: Concepts, Techniques and Tools (2nd ed.), Princeton University Press
14. Cruz, M., Peters, G., Shevchenko, P. (2015) – Fundamental Aspects of operational risk and Insurance Analytics, Wiley.
15. Cruz, M., Peters, G., Shevchenko, P. (2015) – Fundamental Aspects of operational risk and Insurance Analytics, Wiley.
16. Alexander, C. (2008) — "Statistical Models of Operational Loss", in Market Risk Analysis, Volume III, Wiley.

1.5. Conclusion

Operational risk scenario analysis lies at the intersection of data, expert judgement, and governance. While the importance of each element is well established, many firms continue to face practical challenges in integrating them to achieve capital outcomes that are stable, credible, and decision-useful. This white paper aims to bridge that gap by explicitly rooting scenario analysis in observed loss experience, while preserving the central role of expert judgement.

Drawing on a unique investment management specific dataset of more than 2,000 material operational loss events, we have demonstrated how internal and external loss data can be systematically processed, interpreted, and translated into scenario-relevant insights. In doing so, we have highlighted both the value of data in anchoring scenario discussions and the limitations that arise when data volumes are low, definitions are unclear, or methodological choices are not carefully considered.

A key theme throughout this white paper is that data should not replace expert judgement but rather frame and discipline it. When loss experience is presented in a structured and transparent manner, our experience is that scenario workshops produce more coherent frequency assessments, more realistic views of tail severity, and fewer internal inconsistencies across risk factors. Equally, where data is sparse or noisy, we have shown how methodological choices – including smoothing, extrapolation, and distributional assumptions – can materially influence outcomes and therefore require explicit governance and challenge.

The full technical version can be requested separately to this executive summary and is designed to support practitioners and quantitative specialists wishing to engage more rigorously in these issues. By detailing the analytical choices, testing methodologies, and statistical properties underlying our findings, these additional sections are provided to promote transparency, reproducibility, and technical defensibility, rather than to prescribe a single 'correct' methodology.

Ultimately, the onus for sound scenario assessment and capital adequacy lies with firm management. Nevertheless, our analysis supports the view that integrating structured expert elicitation techniques with high-quality internal loss data and credible external benchmarks can materially enhance the credibility, resilience, and regulatory acceptability of operational risk frameworks, especially in light of evolving ICARA / ORA requirements.

As supervisory expectations around ICARA and ORA processes continue to evolve, approaches that improve transparency in the linkage between loss experience, expert judgement and capital outputs may support more constructive supervisory engagement and greater comparability across firms.

We hope that the analysis and practical examples set out in this white paper provide a useful reference point for firms seeking to enhance their scenario analysis practices, whether at board level, within scenario workshops, or in the technical design and validation of capital models.

ORIC International and the Deloitte operational risk quantification team

Please feel free to reach out to any of the authors of this report to understand more (contacts provided below).

2 Expert contributors and credentials

This white paper draws on the combined experience of senior practitioners in operational risk, capital modelling, and regulatory engagement. Contributors include current and former risk leaders from investment management firms, banking and insurance institutions, quantitative modelling specialists, and practitioners with direct experience of supervisory reviews and model validation.

For contacts and a full list of contributing experts, please see section 3, below.

Collectively, contributors have:

- designed, implemented or reviewed LDA and scenario-based operational risk capital frameworks across firms of varying size and complexity;
- supported numerous firms with all aspects of model implementation, embedding and validation, including SREP review preparation and remediation;
- led scenario workshops and loss data frameworks used in ICAAP, ICARA and ORSA contexts; and
- worked with both internal model builds and pragmatic, simplified approaches for mid-sized firms.

ORIC International's dataset

ORIC International operates a global operational risk data consortium focused on the insurance, investment and banking sectors. Its structured loss database is built from standardised member submissions and provides a cross-firm evidence base for operational risk scenario development, severity benchmarking and capital model calibration.

The dataset spans multiple operational risk categories and event types and is supported by consistent taxonomy standards and validation processes. It is widely used by participating firms and independent advisors to support model development, expert judgement calibration, and independent challenge within regulatory capital and risk assessment frameworks.

3 Contacts and acknowledgements

Deloitte

Ian Wilson
Modelling Partner
Deloitte United Kingdom
iwilson@deloitte.co.uk

Damian Hales
Modelling Partner
Deloitte United Kingdom
dhales@deloitte.co.uk

Lead Author
William Shields
Director
Deloitte United Kingdom
williamshields@deloitte.co.uk

ORIC International

Caroline Coombe
Chief Executive
caroline.coombe@oricinternational.com

Chris Watson
Head of Risk Intelligence
chris.watson@oricinternational.com

A special thank you to Dikshita Baid (Deloitte Senior Manager), Harry Danias (Deloitte Manager), and Mallika Arora (Deloitte Analyst) who, together with William Shields, formed the 'core analytical team' in the development of this white paper.

The white paper team would also like to thank the following for their expert input, guidance, review and challenge contributions:

- **Adrian Whitaker –** ORIC Board Member and Aviva Group Senior Operational Risk Manager
- **Thanos Alevizos –** ORIC Member and FNZ Head of Risk and Analytics
- **Ciaran Hosty –** ORIC Member and Aviva Investors Senior Emerging and Capital Risk Manager
- **Brian Thornhill –** Deloitte Director and Prudential Lead
- **Geyer Bisschoff –** Deloitte Director and Technical Specialist
- **Harrison Bye –** Deloitte Director and Capital Clarity Product Lead
- **Dan Oakey –** Deloitte Director and Quality Control Lead
- **Justin Le Blanc –** Deloitte Director and Model Risk Management Lead
- **Devang Shah –** Deloitte Senior Manager and Model Risk Management Specialist
- **Muskan Aggarwal –** Deloitte Analyst and Model Risk Management Analyst

This white paper - executive summary is an abridged version of our full technical version. Further details of which can be made available to full ORIC and Deloitte Capital Clarity members, and to other interested parties by email request.

Click to request the full paper

4 Endnotes

1. Deloitte LLP and ORIC International (2026). Utilising data to enhance operational risk scenario and capital assessments for investment management firms: Full technical white paper.
2. Financial Conduct Authority (2021). CP21/7: A new UK prudential regime for MiFID investment firms.
3. Financial Conduct Authority. FCA Handbook – MiFIDPRU 4 & 7 – ICARA Process.
4. Basel Committee on Banking Supervision (2006). International Convergence of Capital Measurement and Capital Standards: A Revised Framework. Bank for International Settlements.
5. Basel Committee on Banking Supervision (2011). Operational Risk – Supervisory Guidelines for the Advanced Measurement Approaches. Bank for International Settlements.
6. ORIC's Annual Capital Benchmarking Survey Summary Report, 2025
7. McNeil, A.J., Frey, R., Embrechts, P. (2015) – Quantitative Risk Management: Concepts, Techniques and Tools (2nd ed.), Princeton University Press.
8. Cruz, M., Peters, G., Shevchenko, P. (2015) – Fundamental Aspects of operational risk and Insurance Analytics, Wiley.
9. Alexander, C. (2008) – "Statistical Models of Operational Loss", in Market Risk Analysis, Volume III, Wiley.
10. David, H. A., & Nagajara, H. N. (2003) – Order Statistics (3rd ed.). Wiley.
11. Embrechts, P., Klüppelberg, C., & Mikosch, T. (1997) – Modelling Extremal Events: for Insurance and Finance. Berlin: Springer

Deloitte.

Deloitte LLP Statement of Liability

This publication has been written in general terms and we recommend that you obtain professional advice before acting or refraining from action on any of the contents of this publication. Deloitte LLP accepts no liability for any loss occasioned to any person acting or refraining from action as a result of any material in this publication.

Deloitte LLP is a limited liability partnership registered in England and Wales with registered number OC303675 and its registered office at 1 New Street Square, London EC4A 3HQ, United Kingdom.

Deloitte LLP is the United Kingdom affiliate of Deloitte NSE LLP, a member firm of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee ("DTTL"). DTTL and each of its member firms are legally separate and independent entities. DTTL and Deloitte NSE LLP do not provide services to clients. Please click here to learn more about our global network of member firms.

ORIC INTERNATIONAL

ORIC International Statement of Liability

ORIC International shall not have liability to any third party with respect to this report or any actions taken or decisions made as a consequence of the results, advice or recommendations set forth herein.

The opinions expressed herein are valid only for the purpose stated herein and as of the date hereof. Information furnished by others, upon which all or a portion of this report is based, is believed to be reliable but has not been verified. No warranty is given as to the accuracy of such information. Public information and industry and statistical data are from sources ORIC International deem to be reliable; however, ORIC International make no representation as to the accuracy or completeness of such information and has accepted the information without further verification. No responsibility is taken for changes in market conditions or laws or regulations and no obligation is assumed to revise this report to reflect changes, events or conditions, which occur subsequent to the date hereof.

Operational Risk Consortium Limited. Registered in England and Wales. CRN 05510364.

Registered Office: 2nd Floor, 168 Shoreditch High Street, London, E1 6RA. VAT No. 882 3901 09