



Bank Risk Identification

The Complete Methodology

RORY ROBERTS, FRM

Former Global Head of Risk Identification | 20+
Years in Banking & Insurance Risk

Bank Risk Identification

The Complete Methodology

Rory Roberts, FRM

Contents

- 1** Why Banks Fail at Risk Identification

- 2** The Foundations: Standards and Frameworks

- 3** Governance: Who Owns What

- 4** The Risk Taxonomy

- 5** Setting the Context: External, Internal, and Risk Culture

- 6** Top-Down Identification: Workshops, SWIFT, and Delphi

- 7** Bottom-Up Identification: Templates, RCSA, and Specialist Processes

- 8** Reconciliation and the Enterprise Portfolio View

- 9** Assessment — Scoring, Multi-Dimensional Impact, and Data Quality

- 10** Risk Interaction: Bow-Ties, Matrices, and Concentration

- 11** Documentation: The Living Risk Inventory

- 12** Integration: Capital Planning, Strategy, and the Board

- 13** The Ongoing Cycle: Refresh, Events, and Audit

14 Technology: AI, ML, and Data Analytics

15 The Regulatory Landscape

16 Lessons from Bank Failures

Why Banks Fail at Risk Identification

The Question No One Can Answer

Ask any bank how it identifies its risks.

Not how it *measures* them — most banks can answer that question in exhaustive detail. Not how it *manages* them — there are frameworks, committees, and reporting lines for that. Not how it *reports* them — regulatory reporting is well-resourced and closely supervised.

How it *identifies* them. How it determines, in the first place, what risks the institution is exposed to.

The answer, in almost every case, is the same uncomfortable silence. Someone mentions the risk register. Someone else points to the ICAAP process. A third person describes a workshop held a year or two ago. But no one can describe a coherent, end-to-end process for how the institution identifies, catalogues, and validates its risk exposures. There is no methodology. No documented process. Often no named owner. No audit trail showing how the risk register was built or why certain risks appear on it and others do not.

This is not a claim about any single institution. It is an industry-wide condition. When CROs, risk leads, and consultants speak candidly, the story is almost always the same. Most banks have a risk register. Very few have a risk identification *process*. The register exists — a spreadsheet, periodically updated, dutifully presented to the Board. But the process that should produce it — the structured, repeatable methodology for determining what the institution is actually exposed to — is absent.

This book exists because, after two decades of risk practice, I became convinced of two things. First, that risk identification is the most important and least understood step in the entire risk management cycle. And second, that there was no comprehensive methodology available anywhere — no book, no standard, no framework — that told a practitioner exactly how to build one.

This is that methodology.

The Gap in the Risk Management Cycle

Every risk management framework — ISO 31000, COSO ERM, the BCBS Corporate Governance Principles — begins with the same step: *identify the risks*. It is the foundation upon which everything else is built. Assessment, mitigation, monitoring, reporting, capital allocation — all of it depends on having first identified what you are exposed to.

And yet, if you look at where banks invest their time, talent, and technology, risk identification barely registers. The measurement functions are well-resourced: credit risk modelling teams with dozens of quants, market risk systems running millions of simulations per day, operational risk databases painstakingly cataloguing loss events. Capital calculation engines are sophisticated. Reporting frameworks are elaborate. Stress testing programmes consume thousands of person-hours per year.

But the question that precedes all of this — *What are the risks?* — is typically addressed in a two-day workshop once a year, run by someone who has never been trained in facilitation techniques, using a brainstorming methodology that produces groupthink rather than genuine risk intelligence. The output is a spreadsheet. The spreadsheet becomes the risk register. The risk register is updated annually — which in practice means it is updated when someone remembers to do it, or when a regulator asks.

This is the gap. Banks have invested billions in measuring and managing risks, but almost nothing in *finding* them in the first place. It is as if a hospital had the world's best surgeons and diagnostic equipment but no process for examining patients. The treatments are excellent. The diagnosis is guesswork.

What the Evidence Shows

To understand how serious this gap is, I studied every major bank failure I could find — from the savings and loan crisis of the 1980s through the Global Financial Crisis, the post-crisis conduct era, and into the 2020s. I built a database of 179 bank loss events, covering institutions across 35 countries and six decades, with aggregate losses exceeding \$2.3 trillion.¹

For each entry, I asked the same question: *Was the risk identifiable before the loss materialised?*

The answer, in every single case, was yes.

The risks that destroyed Barings, LTCM, Enron, Northern Rock, Lehman Brothers, and dozens of other institutions were not unknowable. They were not black swans. They were identifiable — often glaringly so — using information that was available at the time. In many cases, the risks had been explicitly flagged by internal audit, by individual employees, by external analysts, or by the institutions' own risk functions. The problem was not that the risks could not be identified. The problem was that the *process* for identifying them was absent, inadequate, or actively undermined.

The database reveals ten recurring failure modes — patterns that appear again and again across different institutions, different decades, and different geographies. Understanding these failure modes is the starting point for building a methodology that prevents them.

The Ten Failure Modes

1. Concentration Blindness

The most common failure mode in the database. Institutions fail to identify that their exposure to a single sector, counterparty, asset class, or geography has reached a level where it poses an existential threat. The concentration is often masked by a narrative of diversification — the portfolio *looks* diversified at the instrument level, but the underlying risk factors are correlated.

The pattern: Fannie Mae and Freddie Mac, Countrywide, Northern Rock, the Icelandic banks, Washington Mutual, Anglo Irish Bank, and dozens of others all shared the same fundamental problem — extreme concentration in property-related credit risk, disguised by the apparent diversification of having many individual loans. When the housing market turned, the diversification disappeared because every loan was exposed to the same macro factor.

What was missing: An enterprise portfolio view that aggregated exposures across business lines and looked for common underlying risk factors, not just instrument-level diversification. A process that asked: *What single scenario would impair the largest number of our assets simultaneously?*

2. Model Overreliance

Institutions place excessive trust in quantitative models — Value at Risk, credit ratings, internal models — that fail to capture tail risks, correlation breakdowns, or regime changes. The model becomes a substitute for judgement rather than an input to it.

The pattern: Long-Term Capital Management's Nobel Prize-winning founders built a \$125 billion portfolio on models calibrated to normal market conditions. When the Russian crisis caused correlations to spike across all markets simultaneously, the models were useless. LTCM's value-at-risk model told them their maximum expected daily loss was \$35 million. On 21 August 1998, they lost \$553 million in a single day.²

Merrill Lynch, Citigroup, and Bear Stearns all treated AAA credit ratings on CDO tranches as definitive assessments of risk, substituting an external agency's opinion for independent risk analysis. When the ratings proved wrong, the losses were catastrophic.

What was missing: A risk identification process that treated models as one input among many, not as the final word. A culture that asked: *What does this model NOT capture?* And a governance framework that required independent challenge of model assumptions.

3. Governance Bypass

Risk identification frameworks exist on paper but are circumvented by dominant individuals, commercial pressure, or cultural inertia. Risks are identified at the working level but never reach the people who can act on them — or they reach those people and are ignored.

The pattern: At Barings, internal audit identified the segregation-of-duties violation in Nick Leeson's Singapore operation. Management did not act.³ At Standard Chartered, the compliance function identified the sanctions risk in the Iranian business. Senior management overruled the concerns because the business was commercially important. At Lehman Brothers, the Chief Risk Officer raised increasingly urgent warnings about leverage and mortgage exposure. She was first marginalised and then replaced.

What was missing: A governance structure with clear escalation paths, where risk identification findings cannot be unilaterally overruled by the business, and where the Board receives risk information that has not been filtered through the very management responsible for generating the risk.

4. Silo Thinking

Risk identification is conducted within individual business lines or risk types without aggregation across the enterprise. Each silo sees its own piece of the picture but no one assembles the whole.

The pattern: At JPMorgan, the Chief Investment Office's "London Whale" trades were classified as hedging rather than proprietary trading. Because of this classification, the CIO was exempted from the bank's trading risk limits and VaR monitoring. The organisational silo determined the level of risk oversight, and the hedge designation was never independently challenged.⁴

At Citigroup, off-balance-sheet SIV exposures were excluded from the bank's risk aggregation framework because the liquidity puts were considered remote contingencies. No one asked what would happen if *all* the contingent exposures crystallised simultaneously.

What was missing: A reconciliation process that brings together top-down and bottom-up risk views and explicitly asks: *What risks are we missing because of the way we are organised?* An enterprise portfolio view that aggregates across silos.

5. Cultural Suppression

The institutional culture discourages challenge, dissent, or escalation of risk concerns. People who identify risks are treated as obstacles to the business rather than as performing an essential function.

The pattern: At Wells Fargo, an estimated 3.5 million potentially unauthorised accounts were opened to meet cross-selling targets.⁵ Whistleblower complaints were filed — and treated as HR issues, not risk events. The sales culture was so dominant that the risk function could not challenge it. At Lehman Brothers, internal critics of the leverage strategy were marginalised. At Credit Suisse, the risk culture that allowed Archegos to accumulate \$20 billion in exposure through total return swaps⁶ was described by regulators as one where commercial considerations consistently overrode risk management.

What was missing: A risk identification process that includes anonymous reporting channels, formal protection for risk identifiers, and Board-level visibility into whether the risk culture enables or suppresses the identification of uncomfortable truths.

6. Emerging Risk Blindness

Institutions fail to identify new risk types that fall outside their existing taxonomy or historical experience. The risk is novel, and because no one has lost money on it before, no one recognises it as a risk.

The pattern: Before the PPI scandal, no UK bank had “product mis-selling” as a named risk in its risk register. Before Wirecard, “fintech fraud” was not a standard risk category. Before the GFC, “securitisation market closure” was not a scenario that banks modelled. In each case, the risk was identifiable — the warning signals existed — but the institutions’ risk taxonomies and identification processes did not have a category for it.

What was missing: A structured horizon-scanning process — such as the Delphi Method — that systematically asks: *What risks could emerge in the next three to five years that are not currently in our taxonomy?*

7. Control Environment Failure

Weak internal controls allow fraud, unauthorised trading, or accounting manipulation to go undetected for extended periods. The risk identification process assumes a functioning control environment that does not actually exist.

The pattern: At Enron, shareholders lost \$74 billion in value after accounting fraud — concealed through off-balance-sheet special purpose entities — was exposed.⁷ At WorldCom, \$180 billion in shareholder value was destroyed by systematic capitalisation of operating expenses.⁸ At Wirecard, €1.9 billion in cash balances simply did not exist.⁹ In each case, the basic controls that should have prevented or detected the fraud — reconciliations, segregation of duties, independent verification — were absent or overridden.

What was missing: A risk identification process that includes assessment of the internal control environment as a precondition — recognising that if the control environment is compromised, all other risk assessments built on top of it are unreliable.

8. Information Asymmetry

Key risk information is held by the front office, individual traders, or senior management and is not shared with the risk function or the Board. The people who need the information to identify risks do not have it.

The pattern: At Bear Stearns, the collapse of two hedge funds in June 2007 was treated as an isolated subsidiary event rather than identified as an early warning signal of the same concentration risk in the parent bank's own balance sheet.¹⁰ At Lehman, the Repo 105 programme — used to temporarily remove assets from the balance sheet at reporting dates — was known to a small group of executives but not to the risk function or the Board in a way that allowed them to assess its implications.¹¹

What was missing: A risk identification process that actively seeks out information from across the institution rather than waiting for it to be volunteered. Bottom-up templates that require front-line disclosure. Top-down workshops that include people with direct operational knowledge, not just senior management.

9. Regulatory Arbitrage Masking

Complex financial structures designed to optimise capital or accounting treatment simultaneously obscure the true risk profile. The structure is legal but the risk is hidden.

The pattern: Citigroup's SIVs, Lehman's Repo 105, the entire shadow banking system's use of off-balance-sheet vehicles — all of these were structures that reduced reported risk while increasing actual risk. The regulatory capital framework said the risk was low. The economic reality was that the risk was enormous.

What was missing: A risk identification process that asks: *Where is there a material gap between our regulatory risk profile and our economic risk profile?* And a culture that treats regulatory arbitrage as a risk to be identified, not a benefit to be celebrated.

10. Complacency

Extended periods of low losses breed overconfidence and reduced vigilance. The absence of recent failures is interpreted as evidence that risks do not exist, rather than as a period during which risks are accumulating.

The pattern: Northern Rock's business model had worked for years. LTCM had a track record of exceptional returns. The US housing market had not experienced a nationwide decline in living memory. In each case, the recent track record was used as evidence that the risk was low, when in fact the absence of losses was a lagging indicator that said nothing about the risks currently being accumulated.

What was missing: A risk identification process that explicitly challenges the "it's always worked before" assumption. Scenario analysis that asks: *What would have to happen for this to go badly wrong?* — regardless of whether it has happened recently.

The Common Thread

These ten failure modes are different in their specifics but identical in their root cause: **the institution did not have a process designed to find the risks it was exposed to.** In some cases, there was no process at all. In others, the process existed but was structurally incapable of catching the risks that mattered — because it was top-down only, or annual only, or disconnected from the business, or run without a structured methodology, or undermined by a culture that did not want to hear the answers.

The losses that resulted were not inevitable. They were preventable. Every one of the 179 failures in the database could have been mitigated — not necessarily prevented entirely, but identified early enough to reduce the exposure before it became existential. The risk was there. The information was there. What was missing was the process.

What This Book Is For

This book provides that process.

It is not a textbook on risk management theory. It is not an academic survey of risk identification techniques. It is a practitioner's methodology — a complete, end-to-end process for identifying risks in a banking institution, designed to be implemented by real practitioners in real banks under real regulatory scrutiny.

The methodology is built on three foundations:

Standards. The process is grounded in ISO 31000 (Risk Management Principles and Guidelines), ISO 31010 (Risk Assessment Techniques), and the COSO Enterprise Risk Management Integrated Framework. These are the international standards that regulators reference and against which processes are assessed. Every element of the methodology traces back to a specific provision of these standards.

Regulation. The process maps to 16 regulatory frameworks across the BCBS, PRA, EBA, ECB, Fed, OCC, and FCA. Every regulatory requirement for risk identification is traceable to a specific section of the methodology. This is not an academic framework — it is a process designed to pass regulatory examination.

Evidence. The process is informed by the Industry Loss Database — the 179 bank failures, the ten failure modes, and the \$2.3 trillion in losses that demonstrate what goes wrong when risk identification fails. Every phase of the methodology exists because history showed what happens without it.

The book is organised around a six-phase process:

- **Phase 1: Foundation Setting** — establishing the external and internal context, defining risk criteria and appetite, and building the starting universe of risks to investigate
- **Phase 2: Dual-Track Identification** — a structured top-down and bottom-up identification process with mandatory reconciliation
- **Phase 3: Assessment and Prioritisation** — four-dimensional scoring, multi-dimensional impact analysis, data quality ratings, and bow-tie analysis
- **Phase 4: Documentation** — the living risk inventory and risk profiles
- **Phase 5: Integration** — linking identification outputs to capital planning, strategic planning, and Board reporting
- **Phase 6: Ongoing Cycle** — quarterly re-identification, event-driven updates, and internal audit assurance

Before diving into the phases, Chapter 2 (The Foundations: Standards and Frameworks) establishes the standards and frameworks that underpin the entire methodology. After that, we move to governance — because without clear ownership and accountability, even the best-designed process will fail.

The methodology is comprehensive. It is detailed. And it works. It has been refined through years of practice, regulatory challenge, and the kind of learning that only comes from doing it wrong before you learn to do it right.

The banks that failed did not lack smart people. They lacked a process. This book gives you the process.

In Chapter 2, we examine the international standards and regulatory frameworks that provide the foundation for the methodology — ISO 31000, ISO 31010, COSO ERM, and the BCBS Corporate Governance Principles — and explain how they relate to each other and to the practical work of identifying risks in a bank.

1. Industry Loss Database, EON Risk Services. See Appendix A: Industry Loss Database Methodology for inclusion criteria, data sources, loss definitions, and the full event list.
2. Lowenstein, R. *When Genius Failed: The Rise and Fall of Long-Term Capital Management*. Random House, 2000. The \$125 billion balance sheet figure and \$35 million VaR estimate are from Lowenstein and from the President's Working Group on Financial Markets, *Hedge Funds, Leverage, and the Lessons of Long-Term Capital Management*, April 1999. The \$553 million single-day loss on 21 August 1998 is from the same source.
3. Board of Banking Supervision, Bank of England, *Report of the Board of Banking Supervision Inquiry into the Circumstances of the Collapse of Barings*, HC 673, July 1995. Internal audit had identified the lack of segregation between Leeson's trading and settlement functions in 1994.
4. US Senate Permanent Subcommittee on Investigations, *JPMorgan Chase Whale Trades: A Case History of Derivatives Risks and Abuses*, 15 March 2013. The CIO was exempted from the bank's standard trading risk limits and VaR monitoring framework.
5. Wells Fargo, independent directors' Sales Practices Investigation Report, April 2017; CFPB Consent Order No. 2016-CFPB-0015, 8 September 2016. The figure of 3.5 million accounts is from Wells Fargo's own expanded remediation analysis.
6. Credit Suisse Group, *Report on Archegos Capital Management*, prepared by Paul, Weiss, Rifkind, Wharton & Garrison LLP, 29 July 2021. The \$20 billion figure refers to Archegos's concentrated equity positions unwound during March 2021.
7. Enron Corp., SEC filings; US Senate Permanent Subcommittee on Investigations, *The Role of the Board of Directors in Enron's Collapse*, 8 July 2002. The \$74 billion figure represents the peak-to-trough decline in Enron's market capitalisation.

8. US Securities and Exchange Commission, Litigation Release No. 17588, 26 June 2002; SEC v. WorldCom Inc. The \$180 billion figure represents the destruction of shareholder value following disclosure of the accounting fraud.
9. EY Special Audit of Wirecard AG, 2020; Wirecard AG insolvency filing, Amtsgericht München, 25 June 2020. The €1.9 billion referred to cash balances that Wirecard claimed were held in trustee accounts in the Philippines but which could not be verified.
10. Financial Crisis Inquiry Commission, *The Financial Crisis Inquiry Report*, January 2011, pp. 238-242. The Bear Stearns High-Grade Structured Credit Fund and Enhanced Leverage Fund collapsed in June-July 2007.
11. Report of the Examiner in the Chapter 11 proceedings of Lehman Brothers Holdings Inc., Anton R. Valukas, Examiner, 11 March 2010. The Repo 105 programme is examined in Sections III.A.4 and III.A.5 of the Examiner's report.

The Foundations: Standards and Frameworks

The First Question Every Regulator Asks

In regulatory examinations — PRA, FINMA, the Fed, the EBA — the conversation about risk identification follows the same pattern. The supervisor opens a copy of the institution's risk identification process document, turns to the first page, and asks a deceptively simple question: *What standards is this process built on?*

The question is not academic. It is a test. The regulator is establishing whether the institution's process is grounded in internationally recognised standards or whether someone has invented something from scratch — and whether the people responsible for the process understand the difference.

When building a risk identification process at any G-SIB, one of the first things to establish is the standards architecture. Not because ISO 31000 or COSO ERM will explain exactly how to identify risks in a global investment bank — they will not — but because every regulator with jurisdiction over the bank will assess the process against those standards. If the methodology cannot trace its key design decisions back to specific provisions of ISO 31000, ISO 31010, and COSO ERM, it will fail the first supervisory test before anyone even looks at the content.

This chapter examines the four standards and frameworks that provide the foundation for the methodology in this book: **ISO 31000** (Risk Management: Principles and Guidelines), **ISO 31010** (Risk Management: Risk Assessment Techniques), the **COSO Enterprise Risk Management Integrated Framework**, and the **BCBS Corporate Governance Principles**. These are not interchangeable. Each serves a different purpose. Together, they form the standards architecture that every bank risk identification process must be built on.

ISO 31000: The Architecture

ISO 31000 is the international standard for risk management. Originally published in 2009 and revised in 2018,¹ it provides principles and guidelines that apply to any organisation, in any sector, managing any type of risk. References in this book follow the current ISO 31000:2018 edition unless otherwise noted. It is not specific to banking — and that is both its strength and its limitation.

The standard is built on a three-layer architecture: **Principles, Framework, and Process**. Understanding this architecture is essential because it determines how a risk identification process should be designed, governed, and operated.

Principles (Clause 3)

ISO 31000 establishes eleven principles that risk management should satisfy.² These are not optional guidelines — they are the criteria against which a regulator will assess whether your process is fit for purpose. The principles most directly relevant to risk identification are:

- **Creates value** (Principle a) — Risk identification is not a compliance exercise. It must contribute to the achievement of objectives and the protection of value. A risk identification process that produces a spreadsheet no one reads creates no value.
- **Is an integral part of organisational processes** (Principle b) — Risk identification cannot exist as a standalone annual exercise disconnected from strategy, capital planning, and business decision-making. It must be embedded.
- **Is part of decision making** (Principle c) — The outputs of risk identification must inform decisions. If identified risks do not change behaviour — if the Board receives the risk register and nothing happens — the process has failed this principle.
- **Explicitly addresses uncertainty** (Principle d) — Risk identification must deal with what is uncertain, not merely catalogue what is already known. This is the principle that demands horizon scanning and emerging risk identification.
- **Is systematic, structured and timely** (Principle e) — Ad hoc identification is not compliant. The process must be repeatable, documented, and executed on a defined schedule.

- **Is based on the best available information** (Principle f) — This principle requires that risk identification draws on multiple information sources — internal data, external intelligence, expert judgement, historical evidence — not just management opinion.
- **Takes human and cultural factors into account** (Principle h) — This is the principle that requires a risk identification process to assess whether the institutional culture enables or suppresses the identification of risks. It is the standards basis for addressing the Cultural Suppression failure mode described in Chapter 1 (Why Banks Fail at Risk Identification).
- **Is dynamic, iterative and responsive to change** (Principle j) — Risk identification is not a one-time activity. The standard requires a process that responds to changes in the internal and external environment — which means event-driven updates, not just scheduled cycles.

These principles are not aspirational. When a regulator examines your risk identification process, they are checking — consciously or not — whether these principles are reflected in the design. A process that runs once a year (violating Principle j), uses only top-down workshops (violating Principle f), and produces outputs that no one acts on (violating Principle c) will fail regulatory scrutiny, even if the content of the risk register looks reasonable.

Framework (Clause 4)

The framework layer describes the organisational arrangements needed to support risk management. ISO 31000, Clause 4 specifies a continuous improvement cycle: **Mandate and Commitment** (4.2) establishes leadership accountability. **Design of Framework** (4.3) requires understanding the organisational context, establishing policy, defining accountability, integrating into organisational processes, allocating resources, and establishing communication mechanisms. **Implementing Risk Management** (4.4) covers both the framework and the process. **Monitoring and Review** (4.5) and **Continual Improvement** (4.6) close the loop.³

For risk identification specifically, the framework layer means that the process must have: - A clear mandate from senior management and the Board - A named process owner with defined accountability - Allocated resources (people, time, budget, technology) - Integration with existing organisational processes — not a parallel universe - Communication mechanisms that ensure stakeholder input and output dissemination - A monitoring and improvement cycle — the process itself must be assessed and refined

This is where most banks fall down. They may have a reasonable risk identification methodology on paper, but the framework is missing. There is no named owner. No allocated budget. No integration with strategic planning. No mechanism for monitoring whether the process is actually working. The methodology exists in a vacuum, and the framework that should sustain it does not exist.

Process (Clause 5)

The process layer is where risk identification lives. ISO 31000, Clause 5 defines the risk management process as a sequence: **Establishing the Context** (5.3), **Risk Assessment** (5.4) — which comprises Risk Identification (5.4.2), Risk Analysis (5.4.3), and Risk Evaluation (5.4.4) — and **Risk Treatment** (5.5), all wrapped in continuous **Communication and Consultation** (5.2) and **Monitoring and Review** (5.6).⁴

Section 5.4.2 — Risk Identification — is the provision most directly relevant to this book. The standard requires that the organisation identify sources of risk, areas of impacts, events and their causes, and their potential consequences. It specifies that the aim is to generate a comprehensive list of risks based on those events that might create, enhance, prevent, degrade, accelerate, or delay the achievement of objectives. It requires that risks be identified whether or not their source is under the control of the organisation, and emphasises the importance of identifying risks associated with not pursuing an opportunity.⁵

The standard also makes a critical point that practitioners often miss: risk identification should include examination of the knock-on effects of particular consequences, including cascade and cumulative effects.⁶ This is the standards basis for the risk interaction analysis covered in Chapter 10 (Risk Interaction: Bow-Ties, Matrices, and Concentration) of this book — the requirement to look beyond individual risks to their interconnections.

Two additional elements of the process layer matter enormously for risk identification:

Establishing the Context (5.3) requires the organisation to define its external context (regulatory, economic, competitive, social, cultural environment), internal context (governance, organisational structure, capabilities, culture, information systems), and the context of the risk management process itself (objectives, scope, risk criteria). This is not preliminary paperwork — it is the foundation that determines what risks are even vis-

ible. A risk identification process that does not first establish context will miss risks that fall outside its implicit assumptions. This is why Phase 1 of the methodology in this book — Foundation Setting — exists before any identification activity begins.

Communication and Consultation (5.2) runs in parallel with every step of the process. The standard requires that stakeholder perceptions be actively sought and that differing views be documented rather than suppressed. This is the standards basis for the dual-track identification approach in this methodology — ensuring that both top-down (senior management perspective) and bottom-up (front-line perspective) views are captured and reconciled.

What ISO 31000 Does Not Do

ISO 31000 is deliberately generic. It provides the architecture — the principles that must be satisfied, the framework that must be in place, the process that must be followed — but it does not tell you *how* to identify risks. It does not specify techniques. It does not prescribe governance structures for banking institutions. It does not address regulatory requirements. It does not provide a risk taxonomy.

This is not a weakness. It is by design. ISO 31000 is intended to be applicable to any organisation in any sector. The specificity required for a bank risk identification process must come from elsewhere — from ISO 31010 (techniques), COSO ERM (enterprise objectives), the BCBS principles (banking governance), and from the practitioner's own experience of what works in a regulated financial institution.

ISO 31010: The Technique Toolkit

If ISO 31000 provides the architecture, ISO 31010 provides the tools. Originally published in 2009 and revised in 2019,⁷ the current ISO 31010:2019 edition catalogues risk assessment techniques and provides guidance on selecting which ones to use for different purposes.

The standard organises these techniques across three stages of risk assessment: identification, analysis, and evaluation. What makes ISO 31010 valuable for practitioners is Table A.1 — the applicability matrix — which maps each of the 31 techniques against these three stages and rates their applicability as “strongly applicable,” “applicable,” or “not applicable.”⁸

For risk identification specifically, the techniques rated as strongly applicable include:

| Technique | What It Does | Where It Appears in This Methodology |
|---|---|--|
| Brainstorming | Stimulates creative thinking in a group to generate a comprehensive list of risks | Phase 2 workshops (Ch 6 (Top-Down Identification: Workshops, SWIFT, and Delphi)), with structured modifications |
| Structured interviews | Systematic questioning of individuals with relevant knowledge | Phase 2 bottom-up templates (Ch 7 (Bottom-Up Identification: Templates, RCSA, and Specialist Processes)) |
| Delphi technique | Iterative anonymous expert consultation to achieve convergence on emerging risks | Phase 2 top-down emerging risk identification (Ch 6 (Top-Down Identification: Workshops, SWIFT, and Delphi)) |
| Checklists | Systematic review against a predefined list of risk categories | Phase 1 starting universe construction (Ch 5 (Setting the Context: External, Internal, and Risk Culture)) |
| SWIFT (Structured What If Technique) | Systematic examination of deviations from normal operations using "what if" prompts | Phase 2 top-down workshops (Ch 6 (Top-Down Identification: Workshops, SWIFT, and Delphi)) |
| HAZOP | Systematic examination of process deviations using guide words | Adapted for operational risk identification (Ch 7 (Bottom-Up Identification: Templates, RCSA, and Specialist Processes)) |
| Scenario analysis | Construction of plausible future scenarios to identify risks | Phase 3 assessment and Phase 2 emerging risks (Ch 6 (Top-Down Identification: Workshops, SWIFT, and Delphi), Ch 9 (Assessment — Scoring, |

| Technique | What It Does | Where It Appears in This Methodology |
|----------------------------|--|---|
| | | Multi-Dimensional Impact, and Data Quality)) |
| Bow-tie analysis | Visual mapping of causes, controls, and consequences for a specific risk event | Phase 3 risk interaction analysis (Ch 10 (Risk Interaction: Bow-Ties, Matrices, and Concentration)) |
| Fault tree analysis | Top-down logical decomposition of how a risk event can occur | Phase 3 detailed analysis of material risks (Ch 10 (Risk Interaction: Bow-Ties, Matrices, and Concentration)) |

The standard also provides guidance on technique selection. Section 6.2 identifies three factors that should drive the choice: the availability of resources and expertise, the nature and degree of uncertainty, and the complexity of the situation.⁹ In practice, this means that a bank risk identification process cannot rely on a single technique. Brainstorming alone will not identify emerging risks — it amplifies groupthink. Checklists alone will not identify novel risks — they only find what is already on the list. Scenario analysis alone will not identify operational risks in individual business units — it operates at too high a level.

This is why the methodology in this book uses a combination of techniques across its six phases. ISO 31010 provides the standards justification for this multi-technique approach: no single technique is sufficient for comprehensive risk identification.

The standard makes another point that practitioners consistently overlook. Section 5.2 on risk identification states that the output should include a structured record not just of the risks identified but also of the information used, the assumptions made, and the limitations of the analysis.¹⁰ This requirement for transparency about methodology limitations is critical — it is the difference between a process that can be audited and improved, and one that produces an opaque list of risks with no traceability.

COSO ERM: The Enterprise Lens

The Committee of Sponsoring Organizations of the Treadway Commission — known as COSO — published its Enterprise Risk Management Integrated Framework in 2004,¹¹ and updated it in 2017 as *Enterprise Risk Management — Integrating with Strategy and Performance*.¹² The 2017 update reorganised the framework around five components and twenty principles, with greater emphasis on strategy and performance integration. This book references the 2004 framework's eight-component structure because it remains the version most widely embedded in banking practice and regulatory expectations, particularly the three-dimensional cube model used as a completeness check in Chapter 4 (The Risk Taxonomy). Where ISO 31000 provides a generic risk management architecture, COSO ERM provides something specific and valuable for banking: a framework that explicitly links risk management to organisational objectives.

Origins and Context

COSO's origins matter for understanding why the framework looks the way it does. The Treadway Commission was established in 1985 in response to a wave of financial fraud and reporting failures in the United States.¹³ Its initial focus was internal controls — resulting in the 1992 COSO Internal Control Integrated Framework,¹⁴ which became the de facto standard for internal controls worldwide and was later enshrined in the Sarbanes-Oxley Act of 2002.¹⁵

The ERM framework evolved from this internal controls heritage, expanding the scope from controls over financial reporting to enterprise-wide risk management. This lineage explains two features of COSO ERM that are particularly relevant to risk identification in banks:

First, COSO ERM retains a strong emphasis on the **internal environment** — the tone at the top, risk culture, integrity, ethical values, and the competence of the organisation's people. For risk identification, this means that the framework explicitly requires assessment of whether the organisational culture supports or undermines the identification of risks. This is not an afterthought in COSO ERM — it is the first of the eight components, the foundation on which everything else depends.

Second, COSO ERM introduced the concept of **four objective categories** against which risks should be identified:

| Objective Category | Description | Risk Identification Implication |
|--------------------|---|--|
| Strategic | High-level goals aligned with the institution's mission | Identifies risks to strategy execution, competitive position, business model viability |
| Operations | Effectiveness and efficiency of operations | Identifies risks to operational performance, process integrity, service delivery |
| Reporting | Reliability of financial and non-financial reporting | Identifies risks to data integrity, regulatory reporting accuracy, disclosure quality |
| Compliance | Adherence to applicable laws and regulations | Identifies risks of regulatory breach, sanctions, licence conditions |

This four-category structure is transformative for risk identification. Most banks, when they conduct risk identification, focus primarily on financial risks — credit, market, liquidity — and operational risks. They frequently miss risks to strategic objectives (business model disruption, competitive erosion) and risks to reporting integrity (data quality degradation, regulatory reporting errors) because their identification process does not systematically ask: *What could prevent us from achieving our strategic objectives? What could compromise the reliability of our reporting?*

The COSO four-category structure forces these questions. It is why the methodology in this book requires that risks be identified against all four objective categories — not just the traditional financial and operational risk types that dominate most bank risk registers.

The Eight Components

COSO ERM defines eight interrelated components that together constitute an effective enterprise risk management system:

- 1. Internal Environment** — The organisation's risk management philosophy, risk appetite, board oversight, integrity, ethical values, and competence. For risk identification, this component asks: *Is the culture one where risks can be identified and escalated without fear?*
- 2. Objective Setting** — The process of defining objectives at strategic, operational, reporting, and compliance levels. Risk identification cannot begin until objectives are defined — because risk is defined as anything that threatens the achievement of those objectives.
- 3. Event Identification** — The component most directly relevant to this book. COSO ERM requires the identification of internal and external events that affect the achievement of objectives. It distinguishes between events that represent risks (negative impact) and opportunities (positive impact), and requires both to be identified. It also requires identification of event interdependencies — how one event can trigger another.
- 4. Risk Assessment** — The evaluation of identified risks in terms of likelihood and impact, on both an inherent (before controls) and residual (after controls) basis.
- 5. Risk Response** — The selection of response strategies: avoid, reduce, share, or accept.
- 6. Control Activities** — The policies and procedures that ensure risk responses are effectively carried out.
- 7. Information and Communication** — The mechanisms for capturing and communicating relevant information in a form and timeframe that enables people to carry out their responsibilities.
- 8. Monitoring** — The ongoing assessment of the ERM system's effectiveness over time.

The Three-Dimensional Model

COSO ERM's distinctive contribution is the three-dimensional model — often represented as a cube — that maps the eight components against the four objective categories against the organisational entity levels (subsidiary, division, business unit, entity). This three-dimensional view forces a critical question: *Have we identified risks across all components, all objective categories, and all organisational levels?*

In practice, this model is one of the most effective tools for identifying gaps in risk identification coverage. Using the COSO cube as a completeness check means mapping every identified risk against the matrix and looking for empty cells — combinations of component, objective, and entity where no risks have been identified. The empty cells are not proof that no risks exist in those spaces. They are proof that nobody has looked. And when someone looks, they invariably find risks that the standard workshop-and-brainstorm approach missed.

BCBS Corporate Governance Principles: The Banking Mandate

The Basel Committee on Banking Supervision published its Corporate Governance Principles for Banks in 2015,¹⁶ establishing expectations for the governance of risk management in banking institutions specifically. Where ISO 31000 and COSO ERM are sector-agnostic, the BCBS principles speak directly to banks and their supervisors.

Principle 7 is the provision most directly relevant to risk identification. It states that risk should be identified, monitored, and controlled on an ongoing bank-wide and individual entity basis, and that the sophistication of the bank's risk management and internal control infrastructure should keep pace with changes to the bank's risk profile, the external risk landscape, and industry practice.¹⁷

Several elements of Principle 7 have direct implications for how a bank designs its risk identification process:

Bank-wide identification. The BCBS principles require risk identification to cover the entire institution — not just individual business lines or risk types. This is the regulatory mandate for the enterprise portfolio view described in Chapter 8 (Reconciliation and the Enterprise Portfolio View) of this book, and the reason the methodology requires reconciliation across top-down and bottom-up identification tracks.

Ongoing identification. The word “ongoing” is significant. The BCBS does not envision risk identification as an annual exercise. It expects a continuous process — which in practice means a combination of scheduled cycles (annual full identification, quarterly re-identification) and event-driven updates. This is the regulatory basis for Phase 6 of the methodology.

Keeping pace with change. The principles require that the risk management infrastructure evolves as the risk profile changes. A risk identification process designed in 2015 that has not been updated to address climate risk, cyber risk, or geopolitical fragmentation is non-compliant — even if it was considered adequate when it was implemented.

Board responsibility. The BCBS principles place ultimate responsibility for risk management with the Board of Directors, including responsibility for approving the risk appetite framework and ensuring that the bank’s risk management function has adequate resources and authority. For risk identification, this means the Board must receive risk identification outputs and must be satisfied that the process is comprehensive.

The BCBS principles also emphasise the importance of the **risk management function’s independence** from the business lines it oversees. For risk identification, this creates a structural requirement: the people conducting the identification must have the authority and independence to identify risks that the business might prefer to ignore. This is the regulatory basis for the governance structures described in Chapter 3 (Governance: Who Owns What).

Beyond Principle 7, the BCBS principles reference several other provisions relevant to risk identification. Principle 1 on the Board’s overall responsibilities includes the duty to approve risk management strategies.¹⁸ Principle 8 on risk communication requires that information about the bank’s risk profile be conveyed to the Board in a timely, accurate, and understandable manner. Principle 9 on compliance requires identification of compliance risks.¹⁹ Together, these principles create a comprehensive regulatory expectation for risk identification that goes well beyond what most banks currently deliver.

How the Four Frameworks Fit Together

A common question from practitioners — and from regulators conducting assessments — is how these four frameworks relate to each other. Are they competing standards? Redundant? Contradictory?

They are none of these things. They are complementary layers that together provide the complete standards architecture for a bank risk identification process:

| Framework | What It Provides | Gap If Missing |
|------------------------|---|--|
| ISO 31000 | The architecture — principles, framework, process structure | No systematic process design; ad hoc approach |
| ISO 31010 | The techniques — which methods to use for identification, analysis, evaluation | Reliance on a single technique (usually brainstorming); systematic blind spots |
| COSO ERM | The enterprise lens — four objective categories, eight components, three-dimensional completeness check | Risks identified only for financial/operational categories; strategic and reporting risks missed |
| BCBS Principles | The banking mandate — bank-wide coverage, ongoing frequency, Board responsibility, regulatory expectation | Process not designed for regulatory examination; gaps in governance and frequency |

The practical implication is that a bank risk identification process must satisfy all four simultaneously. A process that follows ISO 31000's architecture but ignores COSO's four objective categories will miss strategic risks. A process that uses COSO's structure but only employs one technique from ISO 31010 will suffer from methodological blind spots. A process that is standards-compliant but does not meet BCBS expectations for bank-wide coverage and ongoing frequency will fail regulatory examination.

This is how the methodology in this book is designed. Each phase traces back to specific provisions across all four frameworks:

- **Phase 1 (Foundation Setting)** maps to ISO 31000 Section 5.3 (Establishing the Context), COSO ERM Component 1 (Internal Environment) and Component 2 (Objective Setting), and the BCBS requirement for bank-wide scope.
- **Phase 2 (Dual-Track Identification)** uses techniques from ISO 31010 (SWIFT, Delphi, checklists, structured interviews), implements COSO ERM Component 3 (Event Identification), and satisfies the BCBS requirement for comprehensive coverage through both top-down and bottom-up tracks.
- **Phase 3 (Assessment)** follows ISO 31000 Section 5.4.3–5.4.4 and COSO ERM Component 4, using techniques from ISO 31010 (scenario analysis, bow-tie, fault tree).
- **Phase 4 (Documentation)** satisfies ISO 31010's requirement for transparent recording of methods, assumptions, and limitations, and COSO ERM Component 7 (Information and Communication).
- **Phase 5 (Integration)** maps to ISO 31000 Principle b (integral part of organisational processes), COSO ERM Component 7, and the BCBS requirement for Board reporting.
- **Phase 6 (Ongoing Cycle)** satisfies ISO 31000 Principle j (dynamic, iterative, responsive to change), COSO ERM Component 8 (Monitoring), and the BCBS requirement for ongoing identification.

This traceability is not academic. When a regulator examines your risk identification process and asks what standards it is built on, you need to be able to point to each phase and say: *This phase implements these specific provisions of these specific standards.* That traceability is what distinguishes a defensible methodology from an ad hoc collection of risk workshops.

Standards Are Necessary But Not Sufficient

Let me be direct about something the standards will not tell you. They will not tell you how to run a SWIFT workshop with a room of twenty senior bankers who do not want to be there. They will not tell you how to handle the political dynamics when the Chief Risk Officer and the Head of Investment Banking disagree about whether a risk is material.

They will not tell you what to do when the risk taxonomy does not have a category for the risk you have just identified. They will not tell you how to integrate risk identification outputs into a capital planning process that was designed without risk identification in mind.

Standards provide the architecture. They do not provide the craft.

The implementation of Solvency II across European insurers demonstrated this clearly. The institutions that succeeded were not the ones with the most comprehensive standards mapping. They were the ones that understood the standards well enough to know where the standards ended and where practitioner judgement began.

The same is true here. ISO 31000, ISO 31010, COSO ERM, and the BCBS principles give you the architecture, the techniques, the enterprise lens, and the regulatory mandate. They tell you what a good process looks like in structural terms. But turning that structure into a working process that actually identifies the risks that matter — that requires governance, methodology, and the kind of operational detail that no standard can provide.

That is what the remaining chapters of this book deliver: the practitioner's methodology built on these foundations.

The standards tell you what needs to exist. But a process without clear ownership is a process that will fail — regardless of how well it maps to ISO 31000. In Chapter 3, we examine governance: who owns the risk identification process, who participates in it, who has authority to challenge the business, and how findings reach the Board without being filtered by the people who generated the risk.

1. International Organization for Standardization, *ISO 31000:2009 — Risk Management: Principles and Guidelines*, first published November 2009; revised as *ISO 31000:2018 — Risk Management: Guidelines*, February 2018.
2. ISO 31000:2009, Clause 3, enumerated eleven principles labelled (a) through (k). The 2018 revision consolidated these into eight principles (Clause 4). The principle descriptions and letter references in this chapter follow the 2009 edition's enumeration.
3. International Organization for Standardization, *ISO 31000:2009 — Risk Management: Principles and Guidelines*, Clause 4 ("Framework"), Sections 4.2–4.6. Note: The 2018 revision restructured the framework clause as Clause 5 with sub-clauses 5.2 (Leadership and commitment), 5.3 (Integration), 5.4 (Design), 5.5 (Implementation), 5.6 (Evaluation), and 5.7 (Improvement).

4. International Organization for Standardization, *ISO 31000:2009 — Risk Management: Principles and Guidelines*, Clause 5 (“Process”), Sections 5.2–5.6. The 2018 revision retains this process structure in Clause 6 with the same sub-elements.
5. International Organization for Standardization, *ISO 31000:2009 — Risk Management: Principles and Guidelines*, Section 5.4.2 (“Risk Identification”). See also *ISO 31000:2018*, Clause 6.4.2.
6. International Organization for Standardization, *ISO 31000:2009 — Risk Management: Principles and Guidelines*, Section 5.4.2. The requirement to examine “knock-on effects” and “cascade and cumulative effects” appears in the risk identification provisions.
7. International Organization for Standardization, *IEC/ISO 31010:2009 — Risk Management: Risk Assessment Techniques*, first published December 2009; revised as *IEC 31010:2019 — Risk Management: Risk Assessment Techniques*, June 2019.
8. International Organization for Standardization, *IEC 31010:2019 — Risk Management: Risk Assessment Techniques*, Annex A, Table A.1. The table maps techniques against the stages of risk assessment (identification, analysis, evaluation) with applicability ratings.
9. International Organization for Standardization, *IEC 31010:2019 — Risk Management: Risk Assessment Techniques*, Section 6.2 (“Selecting risk assessment techniques”).
10. International Organization for Standardization, *IEC 31010:2019 — Risk Management: Risk Assessment Techniques*, Section 5.2 (“Risk identification”). The requirement for recording information sources, assumptions, and limitations applies to the outputs of the identification stage.
11. Committee of Sponsoring Organizations of the Treadway Commission (COSO), *Enterprise Risk Management — Integrated Framework*, September 2004. Developed by PricewaterhouseCoopers LLP under COSO’s direction.
12. Committee of Sponsoring Organizations of the Treadway Commission (COSO), *Enterprise Risk Management — Integrating with Strategy and Performance*, June 2017. The 2017 update introduced five interrelated components and twenty supporting principles.
13. The National Commission on Fraudulent Financial Reporting (the “Treadway Commission”) was established in 1985 as a joint initiative of five professional accounting and finance organisations: the American Institute of Certified Public Accountants (AICPA), the American Accounting Association (AAA), the Financial Executives Institute (FEI), the Institute of Internal Auditors (IIA), and the National Association of Accountants (now IMA). It issued its final report in October 1987.
14. Committee of Sponsoring Organizations of the Treadway Commission (COSO), *Internal Control — Integrated Framework*, 1992. Updated in May 2013 as the *Internal Control — Integrated Framework (2013 edition)*.
15. The Sarbanes-Oxley Act of 2002 (Pub.L. 107–204, 116 Stat. 745), enacted 30 July 2002. Section 404 requires management to assess the effectiveness of internal controls over financial reporting, with the 1992 COSO Internal Control framework becoming the most widely used evaluation framework for compliance.
16. Basel Committee on Banking Supervision, *Corporate Governance Principles for Banks*, BCBS 328, July 2015.
17. Basel Committee on Banking Supervision, *Corporate Governance Principles for Banks*, BCBS 328, July 2015, Principle 7 (“Risk management”), paragraphs 107–122.
18. Basel Committee on Banking Supervision, *Corporate Governance Principles for Banks*, BCBS 328, July 2015, Principle 1 (“Board’s overall responsibilities”), paragraphs 23–43.
19. Basel Committee on Banking Supervision, *Corporate Governance Principles for Banks*, BCBS 328, July 2015, Principle 8 (“Risk communication”), paragraphs 123–128; Principle 9 (“Compliance”), paragraphs 129–138.

Governance: Who Owns What

The Head of Risk Who Was Fired for Doing His Job

In 2004, Paul Moore was appointed Group Head of Regulatory Risk at HBOS — at the time, the UK's largest mortgage lender and one of its biggest banks.¹ His role included oversight of risk identification across the group. Within months, he identified what he believed was a critical threat: the corporate banking division, under the leadership of Peter Cummings, was pursuing an aggressive commercial real estate and leveraged lending strategy that was creating dangerous concentrations on the balance sheet. Individual loan approvals were bypassing normal credit processes. Growth targets were overriding risk limits. The division was generating enormous short-term revenue, and the institutional culture was rewarding it.

Moore did what a Head of Risk is supposed to do. He raised the concern. He documented it. He escalated it to the Board.

The Board's response was not to investigate the risk. It was to remove the person who had identified it. Moore was dismissed in 2005. The official reason was restructuring. The actual reason, as Moore later testified to a Parliamentary inquiry, was that his risk identification work was inconvenient.² He was telling the Board something it did not want to hear, about a division that was generating the profits the Board wanted to see.

Three years later, HBOS collapsed. The corporate banking division's commercial real estate portfolio imploded. Losses exceeded £10 billion. The bank required an emergency rescue merger with Lloyds TSB. Peter Cummings was banned by the FSA.³ Paul Moore's warnings — documented, escalated, ignored — were vindicated in every detail.

The HBOS case is not primarily a story about credit risk. It is a story about governance. The risk was identified. The methodology, such as it was, worked. What failed was everything around it: the escalation path was blocked, the Board received filtered in-

formation, the person responsible for independent challenge was removed, and the cultural incentive to suppress uncomfortable findings overwhelmed the structural obligation to act on them.

This chapter is about ensuring that cannot happen. Standards provide the architecture. Techniques provide the craft. But without governance — clear ownership, defined roles, structural independence, and unobstructed escalation — the most rigorous identification methodology in the world will fail the moment it produces a finding that someone powerful does not want to hear.

The Three Lines Model — Applied to Risk Identification

The Three Lines Model — updated by the Institute of Internal Auditors in 2020 from the earlier “Three Lines of Defence”⁴ — is well established in banking regulation and enterprise risk management. It provides a simple structural principle: the people who generate risk, the people who oversee risk, and the people who provide independent assurance over both must be organisationally separate. Most banks apply this model to risk management broadly. Far fewer apply it specifically to risk identification.

The distinction matters. Risk management is an ongoing process — monitoring, mitigating, reporting. Risk identification is the discrete act of determining what risks exist in the first place. It requires different skills, different governance arrangements, and different protections.

First Line: Business Units Own Their Risks

The first line of defence in risk identification is the business. Front-line units — trading desks, lending teams, operations centres, branch networks — are closest to the risks. They see what is happening on the ground before anyone else does. A trader knows which positions are stretched. A relationship manager knows which clients are deteriorating. An operations team knows which processes are failing.

The methodology requires that first-line units are not passive recipients of risk identification outputs. They are active participants. Each business unit completes a standardised risk self-assessment for every risk it faces. Business unit heads designate risk as-

sessors within their units and own the risks assigned to their areas. Front-line employees identify and report risks as part of their daily responsibilities, with dedicated reporting channels that do not require permission from their line manager.

This is consistent with the OCC Heightened Standards “Three Units” model,⁵ which makes front-line units accountable for identifying, measuring, and managing the risks associated with their business activities — not merely reporting risks upward to the risk function. Risk ownership resides where risk is generated.

The practical challenge is obvious. Asking a revenue-generating business unit to identify risks in its own activities is asking it to produce information that may constrain its freedom. This is why the first line cannot work alone. It needs the structural counterweight of the second line.

Second Line: The CRO Function and Risk Identification Lead

The second line provides independent oversight, challenge, and coordination of the risk identification process. In the methodology, two roles are critical:

The **Chief Risk Officer** owns the risk identification process end-to-end. The CRO ensures adequate resourcing, monitors process performance, and — critically — maintains the independence of the risk function from the business lines it oversees. This is not a ceremonial role. The CRO must have direct access to the Board Risk Committee, a reporting line that is not subordinate to the CEO (or at minimum, a dual reporting line), and the authority to challenge any business decision on risk grounds.

The **Risk Identification Lead** is the person who executes the process. The Risk Identification Lead designs the communication plan for each cycle, coordinates the top-down workshops, manages the bottom-up template completion, drives the reconciliation between top-down and bottom-up inputs, maintains the risk inventory, and conducts the lessons-learned review. It is a coordination and facilitation role, but it is also a challenge role. The Risk Identification Lead must have the standing and the mandate to push back when business units understate their risks, when workshop participants display group-think, or when the reconciliation reveals gaps that no one wants to acknowledge.

The Risk Identification Lead is not there to rubber-stamp what business units want to report. The role exists to ensure that the risk inventory reflects reality, not comfort. That requires a clear mandate from the CRO, visible support from the Board Risk Committee, and a willingness to have uncomfortable conversations with senior people who do not appreciate being told that their risk self-assessments are incomplete.

Third Line: Internal Audit

Internal Audit provides independent assurance over the risk identification process itself. It does not participate in the identification. It tests whether the process was executed with rigour, whether coverage was comprehensive, whether the reconciliation was genuine, whether data quality ratings were appropriate, and whether the documentation trail is adequate. Internal Audit's annual review of the process is a critical governance safeguard — it is the mechanism by which the Board can be confident that the methodology is being followed, not merely claimed.

The Roles That Matter

The Three Lines model provides the structural architecture. Within that architecture, seven specific roles must be defined with clarity. Ambiguity in role definition is one of the most common governance failures in risk identification. When everyone is responsible, no one is accountable.

Board Risk Committee

The Board Risk Committee sits at the apex of the governance structure. It does not execute the risk identification process — it provides oversight, challenge, and approval. Specifically, the Board Risk Committee:

- **Approves** the risk taxonomy, risk appetite statement, risk criteria, and materiality threshold before each annual cycle begins
- **Receives** the principal risk report at each meeting, including the enterprise portfolio view, emerging risk assessments, and data quality distribution
- **Challenges** the outputs — not by second-guessing technical assessments, but by asking whether the process has looked in the right places, whether uncomfortable

risks have been suppressed, and whether the risk profile is consistent with the strategic direction the Board has approved

- **Ensures independence** — the Board Risk Committee is the ultimate guardian of the risk function's independence from the business

The Board's challenge function is not optional. BCBS Corporate Governance Principle 7 requires that risk identification is conducted on a bank-wide basis and that the Board receives the outputs.⁶ But receiving is not the same as challenging. A Board that receives the risk report, notes it, and moves to the next agenda item is not fulfilling its governance obligation. The challenge must be substantive.

The Cooperative Bank provides a cautionary example of what happens when Board governance fails at this level. When the Co-op Bank attempted to acquire Britannia Building Society in 2009, the merger due diligence failed to identify the true scale of Britannia's commercial property loan impairments. But the deeper failure was structural: the Co-op's mutual governance model had placed individuals without banking expertise in Board oversight roles. The chairman at the time had no financial services background. The Board could not challenge what it could not understand. The result was a £1.5 billion rescue, bondholder bail-in, and the loss of the bank's independence.⁷

Board competence in risk matters is not a nice-to-have. It is a governance prerequisite.

Chief Risk Officer

The CRO owns the process. This means:

- Ensuring adequate resourcing — people, budget, technology, and time
- Appointing the Risk Identification Lead and ensuring they have the mandate to operate independently
- Monitoring process performance through defined KPIs (percentage of business units completing assessments on time, number of new risks identified per cycle, time from identification to inventory inclusion)
- Escalating process deficiencies to the Board Risk Committee
- Ensuring that data management standards are adequate to support the process

The CRO's most important governance function is protection. The CRO must protect the integrity of the process from commercial pressure, from seniority bias, and from the organisational tendency to produce risk assessments that tell management what it wants to hear rather than what it needs to know.

Risk Identification Lead

The operational heart of the process. The Risk Identification Lead:

- Designs and executes the communication plan for each cycle
- Coordinates top-down SWIFT workshops and Delphi exercises
- Manages the bottom-up template completion across all business units
- Drives the reconciliation between top-down and bottom-up inputs — this is where most of the difficult governance conversations happen
- Maintains the central risk inventory
- Conducts the annual lessons-learned review
- Ensures that specialist sub-processes (RCSA, conduct risk, ICT risk, AML/CFT, third-party risk, and others) are integrated into the unified framework

This is a role that requires both technical expertise and political skill. The Risk Identification Lead must understand the methodology in detail, but must also be able to facilitate workshops with reluctant senior bankers, challenge business units diplomatically but firmly, and navigate the institutional dynamics that inevitably surround any process that forces transparency about risk.

Business Unit Heads

Each business unit head participates in the bottom-up risk self-assessment and owns the risks assigned to their business unit. They designate risk assessors within their units and are accountable for ensuring that those assessors have the training, time, and access to complete their assessments properly.

Business unit heads also participate in the top-down workshops, bringing their operational perspective to the strategic risk discussion. Their dual role — as both first-line risk owners and participants in the top-down process — creates a natural bridge between the two tracks of identification.

Risk Assessors and Risk Owners

Risk assessors are designated individuals at all organisational levels responsible for identifying and assessing risks within their area of responsibility. They report to their business unit head and to the Risk Identification Lead. Their quality determines the quality of the bottom-up track.

Risk owners are named individuals — not committees, not functions, not “the business” — accountable for monitoring and managing each identified risk. Every risk in the inventory has one named owner. This is a deliberate design choice. Committee ownership dilutes accountability. When a risk is owned by “the Operational Risk Committee,” no individual is personally accountable for monitoring it, escalating changes, or ensuring that the assessment remains current. Named individual ownership eliminates this ambiguity.

Front-Line Employees

Front-line employees are the earliest sensors in the risk identification system. They observe risks in their daily work that may never surface in a workshop or a template. The governance structure must provide them with:

- Clear channels for reporting observed or potential risks
- Escalation paths that do not require permission from their immediate supervisor
- Protection from retaliation for raising risk concerns
- Training on what constitutes a risk observation and how to report it

This last point connects directly to the Cultural Suppression failure mode identified in Chapter 1 (Why Banks Fail at Risk Identification). At Wells Fargo, front-line employees knew that the aggressive cross-selling targets were driving fraudulent account openings. Many raised complaints. But the governance structure treated these as HR matters — employee dissatisfaction — rather than risk signals. The complaints were absorbed by the system, not escalated through it. The result was an estimated 3.5 million potentially unauthorised accounts,⁸ a \$3 billion DOJ settlement,⁹ an asset-growth restriction imposed by the Federal Reserve in 2018 (not lifted until 2025),¹⁰ and the resignation of the CEO.

The methodology addresses this by requiring dedicated risk reporting channels for front-line employees that feed directly into the risk identification process, not into the HR function. Anonymous reporting must be available. And the Board Risk Committee must receive periodic reporting on the volume and nature of front-line risk observations, so that it has visibility into whether the institution's culture is enabling or suppressing the identification of risks from below.

Independence: The Non-Negotiable

The BCBS Corporate Governance Principles require that the risk management function is independent of the business lines it oversees. This is not a suggestion. It is a regulatory mandate, reflected in Principle 7 and reinforced across every major supervisory framework.

Independence in risk identification means three things:

Structural independence. The Risk Identification Lead and the CRO function must not report to the business. The CRO should report to the CEO with a direct line to the Board Risk Committee, or — preferably — report directly to the Board with administrative reporting to the CEO. The Risk Identification Lead reports to the CRO. At no point should a business unit head have the authority to overrule, modify, or suppress a risk identification finding.

Operational independence. The risk identification process must be able to operate without requiring business unit cooperation to function. This means the Risk Identification Lead must have access to data, systems, and information independently of business unit gatekeepers. If the only way to assess a business unit's risks is through information the business unit chooses to provide, the process is not independent — it is captive.

Intellectual independence. The people conducting risk identification must be free to reach conclusions that the business does not agree with. This sounds obvious. In practice, it is the hardest form of independence to maintain. The social and career pressures to produce risk assessments that align with management's preferred narrative are enormous. A Risk Identification Lead who consistently identifies uncomfortable risks will not always be popular. The governance structure must protect that person's ability to do their job without career consequence.

The Standard Chartered sanctions case illustrates what happens when operational independence exists but is overridden by governance failure. Standard Chartered's compliance function identified the sanctions risk associated with processing approximately \$250 billion in transactions for Iranian clients. The risk was identified at the operational level. But senior management in London overruled the compliance function, treating the Iranian business as commercially important. Risk was identified. Governance failed to act on the identification. The result was a \$667 million fine and a compliance monitorship.¹¹

The HBOS case is the mirror image. Paul Moore identified the risk. He escalated it through the proper governance channels. The Board not only failed to act — it removed the person who had identified the risk. This is governance failure in its purest form: the identification function worked, but the governance structure around it was configured to suppress rather than escalate.

At Banca Monte dei Paschi di Siena, the governance failure went deeper still. Senior management actively used complex derivative transactions to conceal hundreds of millions in losses.¹² The escalation path was compromised because the people who should have been receiving the escalation were the ones generating the risk. When management itself is complicit, the governance structure must provide an alternative path — direct Board access, independent audit, or external whistleblowing — that bypasses the compromised chain entirely.

The methodology addresses this through multiple reinforcing mechanisms: the CRO's direct reporting line to the Board, the Risk Identification Lead's mandate from the Board Risk Committee, anonymous reporting channels, internal audit assurance over the process, and the requirement that the Board receives risk identification outputs directly — not filtered through the management layers that generated the risk.

Escalation: How Findings Reach the Board

The most critical governance question in risk identification is not “who identifies the risks?” It is “how do the findings reach the people who can act on them — without being filtered by the people who generated them?”

Chapter 1 identified Governance Bypass as one of the ten recurring failure modes in bank failures. The pattern is consistent: a risk is identified somewhere in the organisation, but by the time it reaches the Board — if it reaches the Board at all — it has been softened, contextualised, reframed, or buried in an appendix. The Board receives a sanitised version of reality. It makes decisions based on incomplete information. The risk crystallises.

The methodology requires a structured escalation framework with the following elements:

The principal risk report. The Board Risk Committee receives a comprehensive risk report at each meeting. This report contains: all material risks with current residual risk scores; the consequence dimension driving each impact score; trend direction for each risk; risks approaching or breaching appetite; new or emerging risks; risks removed or reclassified; the KRI dashboard; data quality distribution across the portfolio; the enterprise portfolio view showing aggregate position against appetite; and process performance indicators showing whether the identification process itself is functioning.

This is not a summary document that the CRO prepares from memory. It is drawn directly from the risk inventory. Every data point is traceable. Every score has a documented basis. The Board is not reading an interpretation of the risk landscape — it is reading the risk landscape.

Direct access. The CRO and the Risk Identification Lead must have standing access to the Board Risk Committee without requiring permission from the CEO or other executive management. If the CRO believes that management is suppressing or understating a material risk, the CRO must be able to bring that directly to the Board. This is the governance backstop. It is rarely used, but its existence changes behaviour. People are less likely to suppress findings when they know the risk function has an unobstructed path to the Board.

Anonymous reporting. The methodology requires anonymous reporting channels for all employees — not just front-line staff. These channels must feed into the risk identification process, not into HR or compliance alone. Anonymous reports about risk concerns must be reviewed by the Risk Identification Lead and, where material, escalated to the CRO and Board Risk Committee. This addresses the Cultural Suppression failure mode directly: if an employee cannot raise a risk concern through normal channels without career risk, the anonymous channel provides an alternative path.

Challenge sessions. The reconciliation process between top-down and bottom-up tracks includes formal challenge sessions where the Risk Identification Lead facilitates debate between senior management and business unit risk teams. These are not consensus-building exercises. They are structured challenge — designed to surface disagreements, test assumptions, and ensure that the final risk inventory reflects reality rather than the lowest common denominator of institutional comfort.

Process Frequency: Not a Once-a-Year Exercise

One of the most common governance failures in risk identification is treating it as an annual event. The risk register is updated once a year, typically in the weeks before the Board needs to approve the ICAAP or the annual report, and then sits untouched until the next cycle. In a world where risks evolve continuously — where a pandemic, a geopolitical crisis, a cyber attack, or a market dislocation can fundamentally reshape the risk landscape in days — an annual cycle is not governance. It is a filing exercise.

The methodology requires four frequencies of risk identification activity:

Annual full re-identification. The complete process — all six phases, from foundation setting through integration — is re-executed annually. The risk taxonomy is reviewed and updated. Top-down and bottom-up identification is repeated from scratch, not rolled forward from the prior year. The materiality threshold is recalibrated. The risk appetite statement is reviewed in light of new identification outputs. The internal environment assessment is refreshed.

Quarterly re-identification. Consistent with the Fed SR 15-18 (CCAR) framework,¹³ the quarterly cycle is not merely a re-assessment of existing risks. It is an active re-identification: has anything new emerged? Has the nature of an existing risk changed? Has something previously immaterial become material? The quarterly cycle requires multi-stakeholder input from front-line units, independent risk management, and senior management. Its output feeds directly into the Material Risk Inventory used for capital planning.

Monthly KRI monitoring. Key Risk Indicators are monitored monthly against defined thresholds. Breaches trigger investigation and, where warranted, event-driven updates to the risk inventory. KRIs are leading indicators — they signal that a risk may be increasing before it crystallises. Monthly monitoring ensures that signals are not missed between quarterly cycles.

Event-driven updates. The risk inventory is updated immediately — outside any regular cycle — when a material event occurs: a significant loss or near-miss, a major change in the external environment, entry into a new business or market, an acquisition or divestiture, a material control failure identified by internal audit, or a significant change in an outsourcing arrangement. Event-driven identification must employ both backward-looking tools (root cause analysis, loss data review) and forward-looking tools (scenario analysis, horizon scanning) to ensure the response addresses not only what has happened but what may happen next.

The Société Générale case of 2008 illustrates the cost of governance gaps in process frequency and role design. Jérôme Kerviel, a junior trader, built €50 billion in hidden positions on European equity index futures by creating fictitious offsetting trades. He was able to do this because he had previously worked in the back office and understood precisely which controls to circumvent and when. The control framework assumed that a single junior trader could not understand and systematically exploit the full chain of front-to-back verification processes. It was wrong. The resulting loss was €4.9 billion.¹⁴

The governance lesson is not just about segregation of duties — though that was plainly inadequate. It is about the assumption embedded in the risk identification process that certain risk scenarios were implausible. The process did not contemplate a trader who knew both sides of the control chain. There was no mechanism for identifying this as a risk, because the governance framework had not defined it as something to look for. A proper risk identification process — with specialist risk identification sub-processes for traded risk, regular event-driven updates, and challenge sessions that questioned control assumptions — would have forced the question: *what if someone in the front office has detailed knowledge of back-office procedures?*

Communication and Consultation

ISO 31000 Section 5.2 establishes that communication and consultation with stakeholders is not a discrete step in the risk management process — it is a continuous activity that runs in parallel throughout every phase.¹⁵ The methodology applies this principle specifically to risk identification.

Stakeholder identification. At the outset of each cycle, the Risk Identification Lead identifies who will be consulted at each phase: Board, executive committee, business units, the risk function, compliance, internal audit, and external stakeholders where relevant. This is documented in the communication plan.

Perception capture. Different stakeholders perceive risks differently. A business unit head will see risks through the lens of revenue and client relationships. A compliance officer will see regulatory exposure. A front-line employee will see operational vulnerabilities that neither the business head nor the compliance officer are aware of. The methodology requires that these differing perceptions are actively sought and documented, not suppressed. The risk inventory should reflect the full range of stakeholder views, not the consensus of the most senior participants.

Two-way communication. Stakeholders must both contribute to and receive outputs from the process at each phase. This is not a data extraction exercise where the Risk Identification Lead collects inputs from stakeholders and disappears into a back room. Stakeholders should see preliminary outputs, challenge them, and receive final outputs with explanations of how their input was used. This builds ownership and ensures that the risk inventory has institutional credibility.

Documentation. All material stakeholder inputs, challenges, and decisions are recorded as part of the audit trail. This is not bureaucracy — it is governance. If a stakeholder raises a risk concern that is subsequently excluded from the inventory, the reason for exclusion must be documented and defensible. If internal audit later questions why a particular risk was not identified, the audit trail must show whether it was considered, who made the decision, and on what basis.

The Governance Test

Every governance structure looks adequate on paper. The real test is what happens when the process produces a finding that someone does not want to hear.

At HBOS, the governance structure was on paper. Paul Moore had a title, a mandate, and a reporting line. None of it mattered when the Board chose revenue over risk. At Standard Chartered, the compliance function identified the sanctions risk. The governance structure failed when senior management overruled it. At Société Générale, the control framework had defined roles and responsibilities. It failed when the assumptions underlying those definitions proved wrong.

The governance arrangements in this chapter are designed to withstand these pressures. Named individual ownership eliminates ambiguity about accountability. Structural independence protects the risk function from commercial pressure. Direct Board access ensures that findings cannot be filtered. Anonymous reporting provides an alternative path when normal escalation is blocked. Internal audit provides independent validation that the process is being followed. And the frequency structure ensures that governance is continuous, not episodic.

But governance is not self-executing. It requires institutional commitment — from the Board, from the CRO, and from every participant in the process — to maintain these protections even when they are inconvenient. Especially when they are inconvenient.

The standards examined in Chapter 2 (The Foundations: Standards and Frameworks) provide the architecture. The governance arrangements in this chapter provide the structural integrity. In Chapter 4 (The Risk Taxonomy), we turn to the risk taxonomy — the common language that ensures every participant in the process, from the Board to the front line, is identifying and classifying risks in the same way.

-
1. House of Commons Treasury Committee, *Banking Crisis: Reforming Corporate Governance and Pay in the City*, Ninth Report of Session 2008–09 (HC 519), 15 May 2009, paras 53–63. See also Paul Moore's written evidence to the Committee, 10 February 2009.
 2. House of Commons Treasury Committee, oral evidence session, 10 February 2009 (Paul Moore testimony, Q1604–Q1672). Moore testified that he was dismissed after reporting concerns about the corporate division's risk culture to the HBOS Board and to the FSA.

3. FSA, *The Failure of HBOS plc (HBOS): A Report by the Financial Conduct Authority and the Prudential Regulation Authority*, November 2015, Chapter 5 (corporate division losses); FSA Final Notice to Peter Cummings, 12 September 2012 (prohibition and £500,000 fine for mismanagement of the corporate lending book).
4. Institute of Internal Auditors (IIA), *The IIA's Three Lines Model: An Update of the Three Lines of Defense*, July 2020.
5. Office of the Comptroller of the Currency (OCC), *OCC Guidelines Establishing Heightened Standards for Certain Large Insured National Banks, Insured Federal Savings Associations, and Insured Federal Branches* (12 CFR Part 30, Appendix D), September 2014. The "three units" framework (front line, independent risk management, internal audit) is set out in Section III.
6. Basel Committee on Banking Supervision (BCBS), *Corporate Governance Principles for Banks* (BCBS 328), July 2015, Principle 7 ("Risk management"), paras 109–117.
7. The Kelly Review: Sir Christopher Kelly, *Failings in Management and Governance: Report of the Independent Review into the Events Leading to the Co-operative Bank's Capital Shortfall*, 30 April 2014. The report details the Britannia merger, the Board's lack of banking expertise (including the chairman's background), and the £1.5 billion capital shortfall.
8. Wells Fargo Board of Directors, independent directors' *Sales Practices Investigation Report*, 10 April 2017, p. 4 (3.5 million potentially unauthorised accounts figure); CFPB Consent Order No. 2016-CFPB-0015, 8 September 2016 (original 2 million estimate subsequently revised upward).
9. United States Department of Justice, "Wells Fargo Agrees to Pay \$3 Billion to Resolve Criminal and Civil Investigations into Sales Practices Involving the Opening of Millions of Accounts without Customer Authorization," press release 20-146, 21 February 2020.
10. Board of Governors of the Federal Reserve System, enforcement action against Wells Fargo & Company, Cease and Desist Order and \$1.8 billion asset cap, 2 February 2018 (Docket Nos. 18-007-B-HC, 18-007-CMP-HC). The asset cap was lifted on 3 June 2025 per Federal Reserve press release (enforcement20250603a).
11. New York State Department of Financial Services, Consent Order under New York Banking Law §44, *In the Matter of Standard Chartered Bank*, 6 August 2012 (\$340 million settlement); U.S. Department of Justice and the Manhattan District Attorney's Office, Deferred Prosecution Agreement with Standard Chartered, 10 December 2012 (\$327 million, bringing the combined total to \$667 million). The \$250 billion transaction figure is cited in the NYDFS order.
12. Tribunale di Milano, sentencing judgment in the criminal trial of former Banca Monte dei Paschi di Siena executives (including former CEO Giuseppe Mussari and former GM Antonio Vigni), October 2014 (the "Santorini" and "Alexandria" derivative transactions used to conceal approximately €730 million in losses). See also Bank of Italy supervisory findings and ECB Comprehensive Assessment, October 2014.
13. Board of Governors of the Federal Reserve System, Supervision and Regulation Letter SR 15-18, *Federal Reserve Supervisory Assessment of Capital Planning and Positions for LISCC Firms and Large and Complex Firms*, 18 December 2015. The quarterly risk identification and material risk inventory requirements are detailed in the accompanying *Capital Planning at Large Bank Holding Companies* (FR 2014-12) guidance.
14. Société Générale, special committee of the Board of Directors, *Mission Green Report* (the internal investigation report), 20 May 2008; Inspection Générale des Finances, *Report on the Lessons to be Drawn from the Fraud at Société Générale*, February 2008. Kerviel was convicted by the Tribunal correctionnel de Paris on 5 October 2010. The €4.9 billion loss figure is from Société Générale's official announcement of 24 January 2008.
15. International Organization for Standardization, *ISO 31000:2018 Risk Management — Guidelines*, Section 6.2 "Communication and consultation" (renumbered from Section 5.2 in the 2009 edition).

The Risk Taxonomy

The Meeting Where Nobody Agrees

Picture a meeting convened to reconcile risk registers across business units. What should be a straightforward alignment exercise becomes a two-hour argument about what things are called.

The investment bank classifies certain derivative exposures as market risk. The risk function classifies the same exposures as counterparty credit risk. Treasury calls a subset of them liquidity risk. Each is correct within its own framework. None is wrong. But when anyone tries to aggregate these into an enterprise view — to answer the basic question of “what are the top risks facing this institution?” — the numbers do not add up. The same underlying exposure appears in three different categories in three different registers, counted three different ways.

This is not an analytical failure. It is a language failure. The institution does not have a common vocabulary for risk.

The problem extends beyond classification. When business units assess their risks against standardised criteria, the responses are incomparable because each unit has built its own taxonomy. One division may have nearly fifty risk categories. Another just over twenty. A third around thirty. Some categories overlap. Others have no equivalent in other divisions. “Conduct risk” exists in one division’s taxonomy and is entirely absent from another’s — not because the division faces no conduct risk, but because nobody thought to include it.

Without a common taxonomy, reconciliation between top-down and bottom-up identification is impossible. Without reconciliation, the enterprise portfolio view is meaningless. Without an enterprise view, the Board receives a collection of divisional reports rather

than a coherent picture of institutional risk. The entire methodology — from identification through to capital planning — depends on something that does not yet exist: a shared language.

Building that language is the first operational priority. Before an institution can identify risks, it must agree on what it means by them.

Why Taxonomy Matters

A **risk taxonomy** is the hierarchical classification structure that defines how an institution categorises its risks. It is, in the most literal sense, the common language of risk identification.

This is not an administrative convenience. It is a structural prerequisite. Every subsequent step in the methodology depends on the taxonomy being in place and consistently applied:

- **Identification** requires categories to identify against. If the taxonomy does not include a risk type, that risk type will not be systematically looked for.
- **Assessment** requires like-for-like comparison. A risk scored in one business unit must mean the same thing as the same risk scored in another.
- **Reconciliation** — the process of aligning top-down and bottom-up outputs — is impossible without a common classification. You cannot reconcile what you cannot compare.
- **The enterprise portfolio view** depends on aggregation. If business units classify the same underlying exposure differently, aggregation produces noise, not signal.
- **Regulatory reporting** requires mapping internal risk categories to the specific taxonomies used by each regulator. Without a coherent internal taxonomy, every regulatory submission becomes a manual translation exercise.
- **Risk ownership** — the assignment of named individuals to each risk — requires a classification structure that defines what each person owns. As established in the governance framework, risk owners are named individuals, not committees. The taxonomy tells them what they own.

When I describe the taxonomy as a “common language,” I mean it operationally. Every participant in the risk identification process — from the Board Risk Committee approving the principal risk report to the front-line employee completing a risk self-assessment — must be classifying risks in the same way. If they are not, the process produces an illusion of comprehensiveness while the real risks hide in the gaps between definitions.

The Three-Level Structure

The taxonomy is structured in three levels of increasing granularity:

Level 1 (L1) defines the broad risk categories. These are the top-level classifications that correspond to the major risk types a bank faces. A typical banking institution’s L1 taxonomy includes:

| L1 Category | Scope |
|--------------------------------|--|
| Credit Risk | Risk of loss from borrower or counterparty failure to meet obligations |
| Market Risk | Risk of loss from movements in market prices, rates, or volatilities |
| Operational Risk | Risk of loss from inadequate or failed processes, people, systems, or external events (including ICT and cyber risk) |
| Liquidity Risk | Risk that the institution cannot meet obligations as they fall due |
| Strategic Risk | Risk to achieving strategic objectives from adverse business decisions or failure to adapt |
| Compliance and Regulatory Risk | Risk of non-compliance with laws, regulations, or supervisory requirements |
| Reputational Risk | Risk of damage to the institution’s standing with stakeholders |
| Model Risk | Risk from errors in models used for pricing, measurement, or decision-making |

| L1 Category | Scope |
|----------------------------------|---|
| Conduct Risk | Risk of harm to customers, markets, or competition from firm or staff behaviour |
| Third-Party and Outsourcing Risk | Risk from dependence on external providers for critical services |
| Step-in Risk | Risk of financial support to entities beyond contractual obligations |
| Climate and Environmental Risk | Risk from physical and transition climate impacts across all transmission channels |
| Emerging and Systemic Risk | Risks not yet fully understood or categorised, including novel and interconnected threats |

These thirteen categories are not arbitrary. They correspond to the risk types that banking regulators across jurisdictions require institutions to identify and manage. The scope section of the methodology lists all thirteen, and the taxonomy must accommodate every one.

Level 2 (L2) breaks each L1 category into sub-categories. For example:

| L1 | L2 Sub-Categories |
|------------------|--|
| Credit Risk | Counterparty Credit Risk, Credit Concentration Risk, Country and Transfer Risk, Settlement Risk, Sovereign Risk |
| Market Risk | Interest Rate Risk, Foreign Exchange Risk, Equity Risk, Commodity Risk, Credit Spread Risk, Basis Risk |
| Operational Risk | Internal Fraud, External Fraud, Employment Practices, Clients/Products/Business Practices, Physical Asset Damage, Business Disruption/System Failure, Execution/Delivery/Process Management ¹ |

| L1 | L2 Sub-Categories |
|--------------|---|
| Conduct Risk | Product Mis-selling, Market Manipulation, Information Misuse, Unfair Treatment, Conflicts of Interest |

L2 is where the taxonomy begins to have operational specificity. A risk assessor completing a bottom-up template maps each identified risk to an L1/L2 node. A SWIFT workshop facilitator uses L2 categories as prompts. The specialist sub-processes — RCSA, conduct risk assessment, ICT risk assessment, AML/CFT assessment — all feed their outputs into the central inventory using L1/L2 classification.

Level 3 (L3) provides granular risk types within each L2. For example, Counterparty Credit Risk at L2 might include Wrong-Way Risk, Cross-Default Risk, and Margin Call Risk at L3. Credit Concentration Risk at L2 might include Single-Name Concentration, Sector Concentration, and Geographic Concentration at L3.

L3 is where the taxonomy meets the institution's specific business model. Two banks with identical L1 and similar L2 structures will have very different L3 entries depending on their products, markets, and geographic footprint. A bank with a large derivatives book needs detailed L3 entries under Counterparty Credit Risk. A bank with a pure retail mortgage business may need a single L3 entry there but extensive granularity under Credit Concentration Risk.

The three-level structure serves a specific design purpose. L1 provides the language for Board reporting and strategic discussion. L2 provides the working categories for risk identification and assessment. L3 provides the precision for specialist analysis and regulatory mapping. The Board Risk Committee discusses credit risk and market risk. The Risk Identification Lead works at the L2 level when facilitating workshops and reconciling inputs. The credit risk function works at L3 when assessing counterparty exposures. All three levels must be consistent and traceable.

The MECE Principle

The taxonomy must satisfy two properties simultaneously:

Mutually Exclusive — each identified risk maps to one and only one taxonomy node. A risk cannot simultaneously be classified as credit risk and market risk. If an exposure has characteristics of both, the taxonomy must define rules for primary classification. The risk inventory can record secondary classifications and cross-references, but every risk needs a single primary home.

Collectively Exhaustive — no material risk should fall outside the taxonomy. If a risk exists that does not fit any existing category, the taxonomy is incomplete and must be updated.

These two properties — universally known as **MECE** — are the structural test of taxonomy quality. A taxonomy that fails the mutually exclusive test produces double-counting and reconciliation chaos. A taxonomy that fails the collectively exhaustive test produces blind spots.

The collectively exhaustive requirement is the more dangerous to violate, because the failure is invisible. If a risk falls outside the taxonomy, no one is looking for it. No template captures it. No workshop prompt surfaces it. No specialist sub-process owns it. The risk exists in the institution's actual exposure but not in its risk inventory.

This is precisely what happened at JPMorgan Chase in 2012.

When Classification Determines Oversight: JPMorgan and the London Whale

JPMorgan's Chief Investment Office in London built massive synthetic credit positions through credit default swaps. By early 2012, trader Bruno Iksil — nicknamed the "London Whale" for the size of his positions³ — had accumulated exposures large enough to move the market.

The positions were classified as portfolio hedges. This was the critical decision. The CIO was designated as a hedging function, not a proprietary trading desk. That classification determined everything that followed.

Because the CIO was classified as hedging, its positions were exempt from the bank's proprietary trading risk limits. They were exempt from the VaR monitoring applied to trading desks. They were subject to a different, less stringent set of controls. The hedge

designation was never independently challenged. Nobody asked the fundamental question: if these positions are hedges, what specifically are they hedging, and does the hedge relationship actually hold?

The positions were not hedges. They were directional bets on credit indices. When the market moved against them, JPMorgan lost \$6.2 billion.²

The risk identification failure was not analytical. JPMorgan had one of the most sophisticated risk management infrastructures in global banking. The failure was taxonomic. The organisational classification of the activity — hedging versus trading — determined the level of oversight applied to it. The risk existed in the institution's actual exposure. It did not exist in the category structure that governed how that exposure was monitored.

In the language of this methodology: the taxonomy was not collectively exhaustive. It did not account for the possibility that a function designated as hedging could be engaging in proprietary risk-taking. The L2 category structure under Market Risk distinguished between trading book and banking book risk, but the CIO's activities fell into a classification grey zone that neither set of controls adequately covered.

What was missing: A taxonomy that classifies by the nature of the risk, not by the organisational unit that generates it. The methodology requires that every risk maps to an L1/L2/L3 node based on what the risk *is* — its characteristics, drivers, and potential consequences — not on where it sits in the org chart. An independent challenge process — built into the reconciliation phase — would have required someone to ask whether the CIO's positions actually behaved like hedges or like proprietary trades. The enterprise portfolio view would have flagged the sheer size of the positions against the institution's aggregate risk appetite, regardless of what they were called.

When Risk Falls Between Categories: Deutsche Bank and Mirror Trading

The JPMorgan case involved misclassification within the taxonomy. Deutsche Bank's mirror-trading scandal illustrates the opposite problem: risk that falls between categories entirely.

Between 2011 and 2015, Deutsche Bank's Moscow office processed approximately \$10 billion in suspicious transactions through a mirror-trading scheme.⁴ The mechanism was straightforward: a client would buy Russian equities in Moscow using roubles, and a related entity would simultaneously sell the identical equities in London for dollars. The net effect was rouble-to-dollar conversion — money laundering on an industrial scale.

Two separate surveillance functions had partial visibility. Equities trading surveillance saw the matched trades but classified them as routine — matched trades are normal in equities markets. AML monitoring tracked cash flows but did not correlate them with the equities trading patterns. Neither function identified that the combined pattern constituted a massive money-laundering mechanism.

The taxonomy created a structural gap. The equities surveillance function looked at its risk categories and saw nothing unusual within them. The AML function looked at its risk categories and saw nothing unusual within them. The risk existed in the space between the two — in the interaction between a trading pattern and a cash flow pattern that, taken together, constituted something neither function was mandated to find.

The combined UK and US fines totalled \$630 million⁵, and the scandal contributed to years of regulatory scrutiny that reshaped Deutsche Bank's entire strategic direction.

What was missing: A taxonomy designed to force cross-category analysis. The methodology addresses this through two mechanisms. First, the enterprise portfolio view in the reconciliation phase requires explicit identification of common exposures across business units — where multiple functions share exposure to the same underlying driver. Second, risk interaction analysis maps which risks can trigger, amplify, or be triggered by other risks. The mirror-trading pattern was an interaction between market risk (equities trading) and compliance risk (AML) that neither function would catch in isolation. A taxonomy that includes cross-references between L2 categories — combined with a reconciliation process that actively looks for gaps between specialist sub-processes — would have surfaced the question even if no individual function had the answer.

When the Taxonomy Does Not Include the Risk: Lloyds and PPI

Sometimes the failure is not misclassification or inter-category gaps. Sometimes the taxonomy simply does not contain the relevant risk type.

The Payment Protection Insurance scandal — ultimately costing the UK banking industry over £50 billion in redress⁶, with Lloyds Banking Group paying approximately £22 billion⁷ — was a conduct risk event. Banks systematically sold PPI policies alongside loans and credit cards to customers who did not need them, could not claim on them, or were unaware they were paying for them.

The risk identification failure was taxonomic at its core. PPI sales were classified as a profitable cross-selling activity within normal business operations. The risk identification frameworks in use at the time did not include “customer suitability” as a material risk category. Conduct risk, as a distinct taxonomy entry, did not exist in most institutions’ risk classifications. Sales incentive structures that rewarded PPI attachment rates were visible to everyone — they were a feature of the business model, not a bug — but because the taxonomy did not include a category for “harm arising from the institution’s own sales practices,” no one was mandated to assess whether the practice was creating risk.

The risk was there. The information was there. But the taxonomy did not contain a node for it, so no identification process looked for it.

This is the collectively exhaustive requirement in its most consequential form. If a risk type does not exist in the taxonomy, the institution is structurally incapable of identifying it — no matter how sophisticated its workshops, how rigorous its bottom-up templates, or how capable its people. The taxonomy defines the universe of what can be found.

What was missing: A taxonomy maintenance process that evolves ahead of loss events, not after them. The methodology requires annual Board-approved taxonomy review, but more importantly, it requires the Delphi Method and horizon scanning to identify emerging risk types before they crystallise. “Conduct risk” should have entered institutional taxonomies before the PPI scandal, not after it. The taxonomy maintenance process — with CRO amendment authority to propose new categories at any time — exists precisely to close these gaps before the collectively exhaustive test fails in practice.

Building the Taxonomy in Practice

Constructing a risk taxonomy is not an exercise in academic classification. It is a practical design problem with real constraints.

Start With Regulatory Requirements

The first input is what regulators require. Basel frameworks define risk categories for capital purposes. The PRA, ECB, Fed, and OCC each publish their own risk categorisations. The institution’s L1 taxonomy must, at minimum, include every risk category its regulators expect to see assessed and reported.

This sounds straightforward but is immediately complicated by the fact that regulators do not agree with each other. Basel defines operational risk in one way. The EBA’s reporting frameworks use overlapping but non-identical categories. The PRA’s ICAAP requirements emphasise risk types — like step-in risk and pension obligation risk — that other regulators treat differently. A bank operating across multiple jurisdictions must satisfy all of them simultaneously.

The practical solution is to build the internal taxonomy at a level of granularity sufficient to accommodate all regulatory mappings. The L1 and L2 structure must be detailed enough that every regulatory category can be mapped to an internal node without forcing artificial combinations or losing internal coherence. The institution then maintains a regulatory mapping table — a separate document that shows, for each regulator, which internal taxonomy nodes correspond to which regulatory categories.

This mapping is not optional. It is the mechanism by which the institution demonstrates to each supervisor that its risk identification process covers their requirements. When a PRA examiner asks “how do you identify step-in risk?” the answer must trace from the regulatory requirement to a specific L1/L2/L3 node to the identification activities performed against it.

Use COSO Objective Categories as a Completeness Check

As established in Chapter 2 (The Foundations: Standards and Frameworks), the COSO ERM framework defines four objective categories against which risks must be identified: **Strategic, Operations, Reporting, and Compliance.**⁸ Every risk in the taxonomy should be assessable against at least one of these categories.

The COSO three-dimensional model — the cube of eight components, four objectives, and entity levels — serves as a completeness check against the taxonomy. Mapping every L2 risk category against the four objective categories and looking for empty cells reveals gaps. If a risk category has no plausible impact on any objective category, it is probably misclassified. If an objective category has no risks mapped to it for a particular entity, something has probably been missed.

The exercise sounds mechanical but it surfaces real gaps. Most banks’ taxonomies are heavily weighted toward Strategic and Operations risks. Reporting risk — the risk that financial and non-financial reporting is unreliable — is often underdeveloped. Compliance risk, in many institutions, consists of a single L2 entry (“regulatory compliance”) with no further granularity. The COSO mapping forces the taxonomy to be genuinely comprehensive across all four dimensions.

Ground the Taxonomy in Evidence

The third input is empirical. The industry loss database — 179 events across 35 countries and six decades — provides a factual record of what has actually gone wrong in banking. Every L2 risk sub-category in the enrichment data represents a risk type that has caused material loss to at least one institution.

This is not about populating the taxonomy with historical events. It is about ensuring that the taxonomy contains every risk type that has demonstrably caused harm. If the loss database contains entries for “mirror-trading money laundering” and your taxonomy has no node that would capture this, the taxonomy is incomplete.

The starting universe for identification — built in Phase 1 from regulatory categories, industry loss data, and internal incident history — serves the same purpose for the taxonomy itself. If the evidence says a risk type exists, the taxonomy must contain it.

Design for the Process, Not for the Org Chart

The JPMorgan London Whale case illustrates the cardinal error in taxonomy design: classifying risks by where they originate rather than by what they are.

The taxonomy must classify by risk characteristics. A credit exposure arising from the investment bank, the retail bank, and the treasury function are all credit risk. The fact that three different organisational units generate them does not change what the risk is. The business unit dimension is captured elsewhere — in the risk inventory, in the risk owner field, in the bottom-up template — but the taxonomy itself must be independent of organisational structure.

This has practical implications. When a business unit is reorganised, the taxonomy should not change. When a product moves from one division to another, the risks it generates should still map to the same taxonomy nodes. If they do not, the taxonomy is not classifying risk — it is classifying business units.

ISO Guide 73 and Definitional Consistency

Chapter 2 established that ISO 31000 deliberately does not specify risk taxonomies — it provides the architecture, not the vocabulary. That vocabulary comes from **ISO Guide 73: Risk Management — Vocabulary**⁹, which provides standardised definitions for risk management terms.

Definitional consistency matters because taxonomy is fundamentally about shared meaning. If “credit risk” means one thing to the credit risk function and something slightly different to the operational risk function, the taxonomy fails its purpose even if the category names are identical.

The taxonomy should reference ISO Guide 73 definitions as the baseline, adapting them where necessary for the institution’s specific context but maintaining consistency with the standard’s intent. Where the institution’s definition diverges from ISO Guide 73, the divergence should be documented and justified.

This applies particularly to terms that cross category boundaries. “Counterparty credit risk” has a specific meaning in the Basel framework (linked to derivatives and securities financing transactions) that differs from the broader sense in which some institutions use the term. “Operational risk” under Basel includes legal risk but excludes strategic and reputational risk¹⁰ — a narrower definition than many institutions’ internal usage. The taxonomy must be explicit about which definition applies, and the regulatory mapping table must reconcile any differences.

The Regulatory Taxonomy Divergence Problem

No common regulatory risk taxonomy currently exists across jurisdictions. Basel, the EBA, PRA, ECB, Fed, and OCC each use overlapping but non-identical risk categorisations. The EBA has made progress toward standardisation through its reporting frameworks, but full harmonisation does not exist and is unlikely to arrive soon.

This creates a practical challenge for any institution operating across multiple regulatory perimeters. Consider a simple example: a European bank supervised by the ECB, with a UK subsidiary supervised by the PRA, and a US branch supervised by the Fed.

- The ECB expects risk identification against EBA risk categories, including ESG risk with specific transmission channel mapping.
- The PRA expects ICAAP risk identification with emphasis on step-in risk, pension obligation risk, and reverse stress testing.
- The Fed expects a Material Risk Inventory updated quarterly under SR 15-18¹¹, with risk categories aligned to CCAR stress scenario design.

All three regulators expect to see credit risk, market risk, operational risk, and liquidity risk — but the boundaries, sub-categories, and reporting granularity differ. A risk that the ECB classifies under one heading may need to appear under a different heading for the PRA.

The institution’s taxonomy must be the single internal source of truth. Every risk has one L1/L2/L3 classification. The regulatory mapping table — maintained separately — translates the internal taxonomy into each regulator’s language. When the PRA asks about step-in risk, the mapping table points to the specific L1/L2 node where step-in risk lives in the internal taxonomy. When the Fed asks for the Material Risk Inventory, the same underlying data is presented in the categories the Fed expects to see.

This mapping must be active, not aspirational. It must be maintained as regulations change, as the institution enters new jurisdictions, and as regulators revise their own categorisations. The Risk Identification Lead is responsible for ensuring the mapping remains current.

Taxonomy Maintenance

A taxonomy is not a static document. It is a living structure that must evolve as the institution's risk landscape changes.

The governance arrangements for taxonomy maintenance are established in the governance framework from Chapter 3 (Governance: Who Owns What):

- **Annual review and approval** by the Board Risk Committee. The full taxonomy is reviewed at least annually as part of the annual full re-identification cycle. The Board does not need to approve every L3 entry, but it must approve the L1 structure and any material changes to L2.
- **CRO amendment authority.** The CRO can propose amendments to the taxonomy at any time — this is essential for responding to emerging risks that cannot wait for the annual cycle. If a new risk type emerges (as conduct risk did, or as climate risk did more recently), the CRO must be able to add it to the taxonomy promptly so that identification processes can begin capturing it.
- **Event-driven updates.** When a material loss event or near-miss reveals a risk type not currently in the taxonomy, the taxonomy must be updated as part of the event-driven response. The London Whale case, the mirror-trading scheme, the PPI scandal — each would have triggered a taxonomy review if the maintenance process had been functioning.
- **Regulatory-driven updates.** When a regulator introduces a new risk category or revises an existing one, the taxonomy must be updated to ensure the regulatory mapping remains valid.

The maintenance process must be documented and auditable. Internal Audit, as part of its annual assurance over the risk identification process, tests whether the taxonomy is current, whether maintenance actions have been completed, and whether new risk types have been appropriately incorporated.

The Emerging Risk Challenge

The hardest taxonomy maintenance problem is adding risk types that do not yet have a name.

Climate risk did not exist as a distinct taxonomy entry in most banks before 2015. Cyber risk was buried inside “IT risk” or “operational risk — systems” until regulators began requiring dedicated assessment. Conduct risk was invisible until a series of scandals made it impossible to ignore.

In each case, the risk existed long before the taxonomy recognised it. Institutions that added the category early — because their horizon scanning or Delphi processes identified the emerging threat — were better positioned to identify specific exposures before they crystallised. Institutions that waited for regulatory mandate or industry loss events were, by definition, late.

The Delphi Method, described as part of the top-down identification process in Phase 2, serves double duty. It identifies specific emerging risks for the current cycle, and it provides intelligence for taxonomy evolution. When Delphi panellists consistently identify a risk type that does not fit neatly into any existing L2 category, that is a signal that the taxonomy needs a new node.

How the Taxonomy Enables the Methodology

The taxonomy is not a standalone deliverable. It is the infrastructure on which every phase of the methodology operates.

Phase 1 — Foundation Setting: The starting universe is built by mapping regulatory categories and industry loss data to taxonomy nodes. The straw man risk list is organised by L1/L2 classification. The external context assessment (PESTLE) identifies threats that must be mappable to the taxonomy. If a PESTLE finding cannot be mapped, the taxonomy may need extension.

Phase 2 — Dual-Track Identification: Top-down SWIFT workshops use L2 categories as systematic prompts — the facilitator works through each L2 node, applying “what if” questions to ensure comprehensive coverage. Bottom-up templates require every identified risk to be mapped to L1/L2/L3. All specialist sub-processes — RCSA, conduct risk,

ICT risk, AML/CFT, third-party risk, traded risk, treasury risk — must use the common taxonomy. Without this, the reconciliation between top-down and bottom-up outputs is impossible.

Phase 3 — Assessment: Risk scoring is performed within taxonomy categories. The materiality threshold is applied consistently across categories because the taxonomy ensures like-for-like comparison. The risk interaction matrix maps cross-category dependencies using the taxonomy as the structural backbone.

Phase 4 — Documentation: The risk inventory records every risk with its L1/L2/L3 classification. The COSO objective category mapping uses the taxonomy as its row dimension. Risk profiles reference the taxonomy classification as a primary identifier.

Phase 5 — Integration: Capital planning maps material risks to stress scenarios using taxonomy categories. Regulatory reporting translates internal taxonomy to regulatory categories using the mapping table. Board reporting presents the principal risk landscape organised by L1.

Phase 6 — Ongoing Cycle: Quarterly re-identification uses the taxonomy to check coverage systematically. Annual taxonomy review refreshes the structure. Event-driven updates may trigger taxonomy amendments.

The taxonomy is the thread that runs through all six phases. Remove it, and the methodology fragments into disconnected activities. This is why building the common language was the first operational priority at Institution A. Everything else depended on it.

The Taxonomy Test

There is a simple test for whether a taxonomy is fit for purpose. Take any risk that has caused a material loss to any bank in the last thirty years. Map it to the taxonomy. If it maps cleanly to a single L1/L2 node, the taxonomy passes for that risk. If it maps to two categories simultaneously, the taxonomy fails the mutually exclusive test. If it maps to no category, the taxonomy fails the collectively exhaustive test. If it maps to a gap between two categories — where both have partial relevance but neither fully covers it — the taxonomy fails both tests.

Run this test against the industry loss database. Run it against your own institution's incident history. Run it against the emerging risks your Delphi panel has identified. Every failure represents a risk that your identification process cannot systematically find.

The taxonomy is not perfect and never will be. Risk is dynamic, and any static classification will eventually lag behind reality. The question is not whether the taxonomy has gaps — it always will — but whether the institution has a process for finding and closing them before those gaps produce losses.

That process is taxonomy maintenance, supported by evidence from loss data, intelligence from horizon scanning, and governance that allows the CRO to act without waiting for the annual cycle. The Board approves the structure. The CRO maintains it. The Risk Identification Lead uses it. Internal Audit tests it. The entire governance framework described in Chapter 3 exists in part to ensure that this foundation — the common language — remains sound.

The taxonomy is where the methodology begins. Without it, risk identification is a conversation where no one is speaking the same language. With it, every participant — from Board to front line — can identify, classify, and communicate risk in a way that aggregates into an enterprise view.

In Chapter 5 (Setting the Context: External, Internal, and Risk Culture), we turn to what that enterprise view requires before identification can begin: the systematic assessment of the external environment, the internal context, and the risk culture that will either enable or undermine every step that follows.

-
1. The seven Level 2 operational risk event-type categories listed here correspond to the Basel II loss event-type classification defined in Annex 9 of BCBS, *International Convergence of Capital Measurement and Capital Standards: A Revised Framework — Comprehensive Version* (BCBS 128), June 2006.
 2. The \$6.2 billion trading loss was reported in JPMorgan Chase's restated 10-Q filing and detailed in the US Senate Permanent Subcommittee on Investigations, *JPMorgan Chase Whale Trades: A Case History of Derivatives Risks and Abuses*, 15 March 2013.
 3. Bruno Iksil's trading activities and the "London Whale" nickname are documented in the US Senate Permanent Subcommittee on Investigations, *JPMorgan Chase Whale Trades: A Case History of Derivatives Risks and Abuses*, 15 March 2013, and in the SEC Order Instituting Cease-and-Desist Proceedings against JPMorgan Chase & Co., Administrative Proceeding File No. 3-15507, 19 September 2013.
 4. The approximately \$10 billion in mirror trades processed through Deutsche Bank's Moscow office between 2011 and 2015 is documented in the FCA Final Notice to Deutsche Bank AG, Ref. 150018, 31 January 2017, and the NYDFS Consent Order under the New York Banking Law, 30 January 2017.
 5. The combined fines comprised \$425 million from the NYDFS (Consent Order, 30 January 2017) and £163 million (~\$204 million) from the FCA (Final Notice, Ref. 150018, 31 January 2017), totalling approximately \$630 million.
 6. The FCA reported total PPI complaints costs to the industry exceeding £50 billion. See FCA, *PPI Complaints Data*, published periodically through the PPI complaints deadline of 29 August 2019, and FCA Aggregate PPI Complaints Data tables.

7. Lloyds Banking Group's cumulative PPI provisions of approximately £22 billion are reported across its Annual Reports and Accounts (2011–2019) and summarised in Lloyds Banking Group PLC, *Annual Report and Accounts 2019*, p. 58 (provisions note).
8. COSO, *Enterprise Risk Management — Integrated Framework*, September 2004 (the "COSO ERM Cube"), defines the four objective categories as Strategic, Operations, Reporting, and Compliance. These categories are retained in COSO, *Enterprise Risk Management — Integrating with Strategy and Performance*, June 2017, though the updated framework reorganises the component structure.
9. ISO Guide 73:2009, *Risk Management — Vocabulary*, International Organization for Standardization, 2009. This standard provides definitions for generic risk management terms and is designed to be used alongside ISO 31000:2018.
10. BCBS, *International Convergence of Capital Measurement and Capital Standards: A Revised Framework — Comprehensive Version* (BCBS 128), June 2006, paragraph 644: "Operational risk is defined as the risk of loss resulting from inadequate or failed internal processes, people and systems or from external events. This definition includes legal risk, but excludes strategic and reputational risk."
11. Board of Governors of the Federal Reserve System, *Federal Reserve Supervisory Assessment of Capital Planning and Positions for LISCC Firms and Large and Complex Firms*, SR 15-18, 18 December 2015.

Setting the Context: External, Internal, and Risk Culture

The Bank That Didn't See the Crisis Coming

In 2007, Hypo Real Estate completed the acquisition of Depfa Bank, an Irish-domiciled public-sector lender with a balance sheet of approximately €400 billion.¹ Depfa's business model was straightforward: lend long to governments and public-sector entities, fund short through the wholesale interbank market. The credit quality of the assets was excellent. The yields were thin but reliable. And the maturity mismatch between the assets and the funding was catastrophic.

By autumn 2007, the external environment had already changed in ways that made Depfa's business model existentially fragile. BNP Paribas had suspended three funds exposed to US subprime in August.² Northern Rock had experienced the first British bank run since 1866 in September.³ Interbank lending rates were spiking. Central banks were injecting emergency liquidity. The wholesale funding markets on which Depfa depended for its survival were under unprecedented stress.

None of this was secret. It was on the front page of the Financial Times every day. The question is not whether this information was available. The question is whether it was systematically assessed as part of Hypo Real Estate's risk identification process and mapped to the specific vulnerabilities in its own business model.

It was not. Within twelve months, Hypo Real Estate required a €102 billion German government guarantee⁴ — one of the largest financial sector bailouts in European history. The bank was nationalised in 2009⁵ and its assets wound down through a bad bank. The Depfa acquisition, which was supposed to diversify and strengthen the group, destroyed it.

The Hypo Real Estate collapse illustrates a brutal truth: the most technically sophisticated risk measurement in the world is worthless if the foundation — the systematic assessment of the environment in which those risks exist — is absent.

This chapter is about that foundation. Phase 1 of the methodology — **Foundation Setting** — establishes the context in which identification takes place. It requires the institution to formally assess its external environment, its internal context, the quality of its risk culture, the criteria against which risks will be judged, and the evidence base from which identification will begin. Everything that follows in Phases 2 through 6 depends on this work being done properly.

Why Foundation Setting Comes First

ISO 31000 Section 5.3 is explicit: establishing the context precedes risk assessment.⁶ This is not procedural sequencing for its own sake. The context determines what risks are relevant, how they should be evaluated, and what constitutes a material threat. Without that context, identification becomes an exercise in listing whatever risks participants happen to think of on the day — which is exactly how most banks do it, and exactly why they miss what matters.

Phase 1 comprises six activities:

- 1. Establish the external context** — the environment in which the institution operates
- 2. Establish the internal context** — the characteristics of the institution itself
- 3. Assess the internal environment** — the culture, tone, and organisational conditions that enable or undermine risk management
- 4. Define risk criteria** — the scales, thresholds, and boundaries against which risk significance will be evaluated
- 5. Build the starting universe** — the evidence base from which identification will begin
- 6. Prepare the straw man risk list** — the preliminary list that frames the Phase 2 workshops

Each of these is a documented output. Each requires specific inputs and produces specific deliverables. Phase 1 is not a planning meeting. It is a structured analytical process with an audit trail.

External Context: The PESTLE Framework

The external context assessment follows the **PESTLE framework** — Political, Economic, Social, Technological, Legal, Environmental — adapted specifically for banking. This is one of the **checklist** techniques catalogued in ISO 31010⁷ (as mapped in Chapter 2 (The Foundations: Standards and Frameworks)). Used properly, PESTLE ensures that the assessment covers all relevant dimensions of the external environment systematically rather than relying on whatever happens to be in the news.

Each dimension is assessed for its relevance to the institution's specific business model, geographic footprint, and risk profile. The findings must be mappable to the risk taxonomy established in Chapter 4 (The Risk Taxonomy) — a PESTLE finding that cannot be connected to a taxonomy node is either an indication that the taxonomy is incomplete or that the finding is not relevant to this institution.

Political and Geopolitical Environment

Political stability, trade policy, sanctions regimes, armed conflict, regulatory divergence across jurisdictions, and government intervention in financial markets. For an institution operating across multiple jurisdictions, the geopolitical dimension includes the risk that political actions in one jurisdiction directly affect operations in another — sanctions, asset freezes, forced divestment, or the withdrawal of correspondent banking relationships.

Economic and Macroeconomic Environment

GDP growth trajectories, inflation, interest rate cycles, unemployment, commodity prices, currency movements, credit cycles, and sovereign risk. This is the dimension that most banks assess in some form — typically through economics research departments — but the output rarely feeds directly into risk identification. The methodology requires that macroeconomic findings are mapped to specific risk taxonomy categories: an interest rate tightening cycle is not an abstract observation but a trigger for reassessing credit risk in rate-sensitive portfolios, liquidity risk in fixed-rate funding, and strategic risk in lending volumes.

Social and Cultural Environment

Changing customer expectations, workforce demographics, public trust in the banking sector, ESG expectations from investors and the public, and financial inclusion trends. This dimension is frequently overlooked in traditional risk identification because the effects are slow-moving and difficult to quantify. But the PPI scandal across British banks — where shifting public expectations of fair treatment went unrecognised for years — demonstrates that social context changes can produce losses measured in tens of billions.

Technological Environment

Technology dependencies, the cyber threat landscape, digital transformation, artificial intelligence and automation, fintech disruption, and cloud and third-party technology risk. The pace of change in this dimension means that a technology assessment completed twelve months ago may be materially outdated. The Capital One data breach in 2019 — where a misconfigured web application firewall in AWS exposed over 100 million customer records⁸ — illustrates that technology risk identification must extend beyond infrastructure to include configuration, access controls, and the interaction between cloud services and legacy systems.

Legal and Regulatory Environment

Applicable regulatory requirements, including Basel III/IV, ICAAP/ILAAP, CCAR, and jurisdiction-specific mandates. Upcoming regulatory changes and supervisory priorities are assessed not just for compliance purposes but for their impact on business models, product viability, and competitive positioning. Litigation trends, enforcement actions against peers, and regulatory guidance letters provide forward-looking indicators of where supervisory attention is moving.

Environmental and Climate

Climate-related financial risks are assessed through a dedicated framework due to their unique characteristics: long time horizons that exceed normal business planning periods, non-linear transmission mechanisms, and pervasive effects that cut across all traditional risk categories.

Climate and ESG Risk Identification Framework

Climate risk identification follows the EBA ESG Guidelines⁹ and the ECB Guide on Climate-Related and Environmental Risks.¹⁰ The framework maps risks through six **transmission channels**:

| Channel | Description | Examples |
|---------------------------------------|-------------------------------|--|
| Physical risk — acute | Discrete climate events | Extreme weather damaging collateral, supply chain disruption |
| Physical risk — chronic | Long-term climate shifts | Sea-level rise reducing property values, water stress affecting agricultural lending |
| Transition risk — policy | Climate policy changes | Carbon pricing, emissions restrictions, stranded asset regulation |
| Transition risk — technology | Low-carbon technology shifts | Disruption of fossil-fuel-dependent sectors in the lending book |
| Transition risk — market | Changing market preferences | Shifts in consumer and investor sentiment, ESG-linked repricing |
| Transition risk — reputational | Perceived inaction on climate | Loss of customers and investors due to insufficient climate commitments |

Climate risks must be assessed across three **time horizons**: short-term (1-3 years), medium-term (3-10 years), and long-term (10-30 years). The long-term horizon exceeds the normal business planning period but must be considered for strategic and capital planning purposes. A bank with significant exposure to commercial real estate in coastal regions may face minimal short-term physical risk but material long-term risk from sea-level rise and changing flood insurance availability. That long-term risk affects the credit quality of the portfolio today, because it affects the residual value of the collateral over the life of the loan.

Climate risks are not a standalone risk category in the taxonomy. They manifest through existing risk types — credit risk from impaired collateral values, market risk from stranded asset repricing, operational risk from extreme weather events, reputational risk from perceived inaction. The institution must map each climate transmission channel to its internal risk taxonomy and include identified climate-related risks in the central risk inventory, feeding through to the risk appetite statement.

The PESTLE Discipline

The critical discipline is documentation. Each PESTLE dimension produces a written assessment with specific findings, each finding mapped to a taxonomy category, each mapping justified. The assessment is dated, authored, and retained as part of the audit trail. When the institution's environment changes — and it will — the documented baseline allows the change to be identified, not just felt.

When External Context Was Not Assessed: Northern Rock

Northern Rock provides the defining example of what happens when the external context is not systematically assessed and connected to the institution's specific vulnerabilities.

Northern Rock funded 75% of its mortgage lending through securitisation and wholesale markets rather than retail deposits.¹¹ This was not an accident or an oversight. It was the strategy. Northern Rock had explicitly chosen wholesale-funded rapid growth as its competitive model. The board approved it. The regulator was aware of it. Analysts covered it.

The external context was changing throughout 2006 and into 2007. US subprime delinquencies were rising. Structured credit markets were showing stress. Wholesale funding spreads were widening. Each of these developments was individually visible. What was missing was a systematic assessment that connected these external changes to Northern Rock's specific business model — an assessment that would have asked: what happens to an institution that depends on continuous securitisation market access if that market seizes?

The answer, when it came in September 2007, was a bank run. Northern Rock experienced the first run on a British bank since 1866. It was nationalised in February 2008.¹² Taxpayers lost approximately £2 billion.¹³

Northern Rock's business model was itself the unidentified risk. A PESTLE assessment that mapped wholesale funding market conditions to the institution's specific funding structure would have identified the vulnerability. The methodology requires exactly this mapping: external context findings are not abstract observations about the world but specific inputs to the identification of risks that affect this institution.

What was missing: A structured PESTLE assessment that connected macroeconomic and market developments to the institution's business model. The wholesale funding dependency was known but never identified as an existential risk because stress testing used benign historical scenarios that excluded a simultaneous seizure of securitisation markets across all maturities. The methodology's requirement that PESTLE findings be mapped to taxonomy categories and assessed against the institution's specific exposures would have forced this connection to be made explicitly.

Internal Context

Where the external context defines the environment, the internal context defines the institution itself. ISO 31000 Section 5.3.3 requires this assessment.¹⁴ The internal context establishes what the institution is, what it is trying to achieve, and what resources and constraints shape its risk profile.

Six elements are assessed:

Strategic objectives. The bank's stated strategic plan and key priorities for the planning horizon. Strategy defines risk: an institution pursuing aggressive growth in emerging markets faces a fundamentally different risk profile from one focused on cost reduction in mature markets. The strategic objectives are not just context for risk identification — they are themselves a source of risk. A strategy that assumes continuous access to wholesale funding markets, or uninterrupted growth in a single sector, contains embedded risk assumptions that must be identified and challenged.

Organisational structure. The current business unit map, legal entity structure, geographic footprint, and reporting lines. Structure creates risk: complex multi-entity structures across jurisdictions create information asymmetry, regulatory arbitrage opportunit-

ies, and consolidation challenges. The organisational structure also determines where risk identification has gaps — a business unit that does not participate in the enterprise process is a business unit whose risks are not visible at the group level.

Governance. Board composition and engagement, committee structures, and delegation of authority. The governance framework established in Chapter 3 (Governance: Who Owns What) provides the structural architecture. The internal context assessment evaluates whether that architecture is functioning — whether the Board is substantively engaged, whether committees have the information and expertise to challenge, whether delegated authorities are being exercised within their limits.

Capabilities and resources. Skills, expertise, technology, data quality, and risk management capacity. An institution that lacks the analytical capability to model complex structured products should not be holding them. If it does, the capabilities gap is itself a risk that must be identified. The methodology requires honest assessment of whether the institution has the resources to identify and manage the risks it has taken on.

Information systems and data flows. How risk information is captured, aggregated, and reported. If the institution cannot produce a consolidated view of its exposures across business units and legal entities within a reasonable timeframe, that data infrastructure gap is a material risk. BCBS 239 (Principles for effective risk data aggregation and risk reporting)¹⁵ provides the regulatory benchmark against which this assessment should be measured.

Risk appetite statement. The Board-approved qualitative and quantitative appetite metrics. The risk appetite statement is the standard against which identified risks will ultimately be measured. It must exist, it must be specific, and it must be operational — meaning it must contain metrics that can actually be monitored and breached. At many institutions, the risk appetite statement is a single paragraph of aspirational language that provides no practical constraint on risk-taking. That is not a risk appetite. It is a mission statement with the word “risk” in it.

The Internal Environment Assessment: Where Risk Culture Lives

The previous two assessments — external and internal context — are analytical exercises. They involve gathering data, mapping it, and documenting findings. The internal environment assessment is different. It requires the institution to look in the mirror and answer honestly whether its culture, values, and organisational conditions will support or undermine the risk identification process.

This is the assessment that Chapter 1 (Why Banks Fail at Risk Identification) promised when it described the Control Environment Failure mode — institutions where the fundamental organisational conditions for effective risk management were absent. COSO ERM places the **Internal Environment** as the first of its eight components¹⁶ because everything else rests on it. If the internal environment is weak, no amount of process design will produce reliable risk identification.

The assessment evaluates seven elements drawn from the COSO ERM framework:

| Element | Assessment Questions |
|--------------------------------------|---|
| Risk management philosophy | Does the institution have a consistent, articulated attitude toward risk-taking? Is this philosophy understood at all levels? Is there a gap between the stated philosophy and actual behaviour? |
| Board of directors' attitudes | Does the Board actively engage with risk identification outputs? Do independent directors challenge management? Is the Board's risk appetite reflected in operational decisions? |
| Integrity and ethical values | Is there a strong ethical culture? Are codes of conduct enforced, not just published? Does "tone at the top" support transparent risk reporting? Can individuals report risk concerns without fear? |
| Commitment to competence | Does the institution have the skills and expertise to identify and manage its risks? Are competence gaps being addressed through hiring, training, or advisory support? |

| Element | Assessment Questions |
|---------------------------------|--|
| Organisational structure | Do reporting lines support effective risk communication? Can risk information flow freely to decision-makers? Are there structural barriers between risk generators and risk managers? |
| Assignment of authority | Are risk ownership and escalation paths clearly defined? Do individuals understand their risk responsibilities? Is the Risk Identification Lead's mandate understood and respected? |
| Human resource standards | Do hiring, training, compensation, and promotion practices reinforce responsible risk-taking? Or do incentive structures reward risk generation while penalising risk reporting? |

Each element is assessed on a structured scale, with evidence required for each rating. The assessment is not a tick-box exercise. It requires interviews with individuals across the organisation, review of incentive structures and disciplinary records, analysis of historical escalation patterns, and comparison of stated values against observed behaviour.

When the Internal Environment Is the Risk: Anglo Irish Bank

Anglo Irish Bank demonstrates what happens when a weak internal environment corrupts risk identification so completely that the institution cannot see the risks it is creating.

Under chairman Sean FitzPatrick, Anglo Irish pursued aggressive commercial property lending during the Celtic Tiger era. The bank's CRE concentration grew to extraordinary levels relative to its capital base. This was not hidden. The loan book was visible in the financial statements. Analysts commented on the concentration. The risk, in the most basic sense, was identifiable from the outside.

But inside the bank, the internal environment made identification impossible. A culture of deference to the chairman meant that challenging the lending strategy was career-threatening. Growth was equated with success — the bank's share price and profitability reinforced the narrative that concentrated property lending was not a risk but a compet-

itive advantage. FitzPatrick himself concealed €87 million in personal loans by temporarily transferring them to another bank around year-end reporting dates¹⁷, a fraud that persisted because no one was positioned or willing to question the chairman's activities.

The board did not independently challenge the lending concentration. Risk identification, such as it was, treated the property portfolio as core business rather than a systemic exposure. The internal environment — the culture of deference, the alignment of incentives with growth, the absence of independent challenge, the conflation of revenue with safety — rendered the identification process structurally incapable of producing an honest assessment.

Anglo Irish was nationalised in 2009.¹⁸ The taxpayer cost was €29.3 billion.¹⁹ The bank was renamed the Irish Bank Resolution Corporation and liquidated in 2013.²⁰ Executives were prosecuted. FitzPatrick faced multiple trials.

What was missing: A formal internal environment assessment that evaluated whether the organisational culture supported or undermined honest risk identification. An assessment of the seven COSO elements would have flagged: a risk management philosophy subordinated to growth objectives, a board that did not independently challenge the dominant strategy, ethical values compromised by deference to the chairman, compensation structures that rewarded asset growth without adjusting for concentration risk, and human resource standards that penalised dissent. Each of these findings would have been escalated to the Board Risk Committee as a material weakness in the foundation on which risk identification depends.

The Annual Refresh

The internal environment assessment is not a one-time exercise. As established in Chapter 3, the annual full re-identification cycle includes a refresh of the internal environment assessment. Culture changes — sometimes deliberately, sometimes through drift. An institution that had a strong risk culture three years ago may have weakened through leadership changes, acquisition integration, strategic pivots, or simple attrition. The annual refresh ensures that the foundation on which the identification process rests is re-evaluated, not assumed.

Material weaknesses identified in the internal environment assessment must be documented as findings and escalated to the CRO and Board Risk Committee. A weak internal environment can compromise the entire risk identification process regardless of how well-designed the methodology is. If the assessment reveals that individuals cannot report risk concerns safely, that incentive structures penalise risk identification, or that the Board is not substantively engaged — these are not background observations. They are material risks to the process itself.

Defining Risk Criteria

Before identification and assessment can begin, the criteria against which risk significance will be evaluated must be formally defined and approved. This is not optional. Without agreed criteria, every participant in the process brings their own implicit definitions of “significant”, “material”, and “acceptable” — and the result is an identification exercise where different people are applying different standards to different risks, producing outputs that cannot be compared, aggregated, or prioritised.

Risk criteria comprise five components:

Impact scales. Definitions for each level of consequence across multiple dimensions. The methodology uses a multi-dimensional approach — a risk’s impact is assessed not only in financial terms but across financial, reputational, regulatory, customer, and operational dimensions. The COSO ERM practice guidance recommends five-point scales²¹ as providing meaningful differentiation without implying false precision. The highest applicable dimension determines the impact rating. A risk that has moderate financial impact but extreme reputational impact receives the extreme rating. Impact scales must be calibrated to the institution’s size and complexity — the financial thresholds for a global systemically important bank differ from those for a regional institution.

Likelihood scales. Definitions for each level of probability, expressed as both annual frequency and probability over a defined horizon. Likelihood assessment in risk identification is inherently qualitative — we are not running Monte Carlo simulations at this stage. The scales provide a common language for expressing professional judgement about how often a risk event might occur.

Materiality threshold. The point at which a risk is classified as material and receives full treatment in the risk inventory, including assignment of a risk owner, detailed assessment, and inclusion in Board reporting. The materiality threshold is a governance de-

cision — it determines the boundary between the risks that the Board sees and the risks that are managed at lower levels. Setting it too high creates blind spots. Setting it too low overwhelms the Board with noise.

Risk appetite boundaries. The levels of risk the institution is willing to accept, defined by risk category and mapped to the taxonomy. Appetite boundaries must be specific and measurable. A statement that the institution has “low appetite for operational risk” provides no practical constraint. A statement that the institution will not accept more than a specified value of operational risk losses in any quarter, measured against a defined threshold, provides a boundary that can be monitored and breached.

Aggregation rules. How individual risk scores combine to form enterprise-level views. Aggregation is not simple addition — a bank with ten moderate risks may face a different aggregate exposure than a bank with one extreme risk, depending on correlations and concentrations. The aggregation rules must be defined before identification begins so that the enterprise portfolio view produced in Phase 2 (Chapter 8 (Reconciliation and the Enterprise Portfolio View)) rests on a consistent methodology.

Risk criteria must be approved by the CRO and endorsed by the Board Risk Committee before the identification cycle begins. They are reviewed annually and recalibrated as needed. The criteria are not static — when the institution’s risk profile changes, when market conditions shift, or when regulatory expectations evolve, the criteria may need updating. But they must exist as a defined, documented baseline before the first workshop is convened or the first assessment template is distributed.

The detailed scoring methodology — including the four-dimensional assessment framework and the data quality overlay — is covered in Chapter 9 (Assessment — Scoring, Multi-Dimensional Impact, and Data Quality). Here, the requirement is that the criteria are defined and agreed as part of Foundation Setting, before identification begins.

Building the Starting Universe

The most common failure in risk identification workshops is the blank-page problem. A group of senior bankers sits in a room and the facilitator asks: “What are the risks facing this institution?” The result is a list shaped by whatever is top of mind — the last crisis, the latest regulatory letter, the risk that caused a loss last quarter. Risks that are not recent, not visible, or not fashionable are missed. The identification process becomes an exercise in recall bias.

The methodology eliminates this problem by building the **starting universe** — the evidence-based foundation from which identification begins. The starting universe is seeded from three sources:

Regulatory Categories

Every risk type required by applicable regulatory frameworks must be included. The regulatory mapping table introduced in Chapter 4 provides the source. If a regulator expects the institution to identify and assess a particular risk type — whether credit concentration, model risk, step-in risk, or climate risk — it must appear in the starting universe regardless of whether current management believes it is material.

This is not compliance for its own sake. Regulatory categories represent the collective supervisory judgement of what has gone wrong across the industry. When the Basel Committee includes conduct risk in its supervisory expectations, or the ECB mandates climate risk identification, these additions reflect actual losses at actual institutions. Ignoring regulatory categories is ignoring evidence.

Industry Loss Data

The 179-event industry loss database described in Chapter 1 provides empirical evidence of what has gone wrong and at what scale. The database covers failures across 35 countries over six decades, with aggregate losses exceeding \$2.3 trillion. Each event is classified by risk type, root cause, and risk identification failure mode.

The starting universe includes a mapping of industry loss events to the institution's risk taxonomy. For each L1 and L2 taxonomy category, the question is: have there been material industry losses in this category? If yes, this category receives enhanced attention in the identification cycle. The industry loss data provides an evidence-based antidote to the institutional tendency to assume that risks which have not materialised internally are not relevant.

Internal Incident History

The bank's own loss events, near-misses, and audit findings from the prior period. Internal incident history is the most institution-specific input — it reflects the risks that have actually crystallised or come close to crystallising within this organisation. Near-misses are particularly valuable because they reveal risks that the existing control environment caught but that the identification process may not have formally recognised.

Additional Data Sources

Beyond the three core sources, the starting universe should draw on additional inputs that are monitored on an ongoing basis:

- **Social media and public sentiment** — monitoring of social platforms, news feeds, and public forums for early signals of reputational, conduct, or emerging risks
- **Vendor and third-party assessments** — reviews of outsourcing partner performance, concentration in critical suppliers, and fourth-party dependencies
- **Technology and system logs** — cybersecurity event logs, system availability data, and access control records as inputs to operational and cyber risk identification
- **Customer data and feedback** — complaints data, NPS scores, service quality metrics, and customer conduct indicators
- **Peer institution benchmarking** — comparison of the institution's risk profile, loss experience, and process maturity against industry peers and published benchmarks such as ORX data and industry loss studies

The starting universe is not the risk inventory. It is the evidence base from which identification proceeds. It ensures that the process begins with what is known — what regulators require, what has gone wrong in the industry, what has gone wrong internally — rather than with the unaided judgement of whoever happens to be in the room.

The Straw Man Risk List

The starting universe provides the evidence. The straw man risk list translates that evidence into a working document that frames the Phase 2 workshops.

Before top-down workshops are convened, the **Risk Identification Lead** prepares the straw man — a preliminary, deliberately incomplete list of potential principal risks drawn from the starting universe, the prior-year risk inventory, and the current PESTLE and internal context assessments. The list is organised by **L1 and L2 taxonomy classification**, ensuring it uses the common language established in Chapter 4.

Purpose

The straw man serves three functions:

First, it prevents blank-page syndrome. When workshop participants receive a structured starting point, the discussion begins at a higher level than “what risks can you think of?” Participants can focus on challenging, refining, and adding to an evidence-based list rather than generating one from scratch under time pressure.

Second, it counters recall bias. The starting universe ensures that the straw man includes risks evidenced by regulatory expectations, industry losses, and internal incidents — not just risks that happen to be top of mind. This is particularly important for risks that are slow-moving, historically infrequent, or outside the direct experience of the participants.

Third, it provides a baseline against which the workshop’s output can be compared. If the workshop removes a risk from the straw man, that removal must be justified and documented. If it adds a risk not on the straw man, that addition demonstrates the value of the dual-track approach. In both cases, the straw man creates an audit trail showing how the identification process moved from evidence to judgement.

Rules of Use

The straw man is powerful precisely because it is used correctly. Used incorrectly, it becomes a draft answer that participants rubber-stamp:

- The straw man is circulated to all workshop participants in advance, alongside the pre-workshop independent assessment
- It is clearly labelled as a discussion prompt, not a draft answer
- The facilitator emphasises at the start of every workshop that the straw man is meant to be torn apart — participants are expected to add, remove, modify, and challenge every item
- The final output must reflect the workshop’s collective judgement, not the straw man

Framing matters as much as content. When the straw man is presented as “the proposed risk list,” participants defer to it. When it is presented as “here is what the evidence shows — now tell me what is wrong with it,” workshops produce genuine challenge, debate, and identification of risks the straw man missed.

How Hypo Real Estate Could Have Been Different

Return to October 2007 and the acquisition of Depfa Bank. Apply the Foundation Setting methodology to what Hypo Real Estate faced.

A **PESTLE assessment** would have documented: subprime contagion spreading through structured credit markets (Economic), central bank emergency interventions (Political/Economic), BNP Paribas fund suspensions and Northern Rock bank run (Economic/Social), wholesale funding market stress visible in LIBOR-OIS spreads (Economic/Legal), and growing regulatory concern about bank liquidity risk management (Legal/Regulatory). Each finding would have been mapped to the taxonomy. The funding market stress would have mapped directly to Liquidity Risk — Wholesale Funding Dependence.

An **internal context assessment** would have documented: the Depfa acquisition had fundamentally changed the institution's strategic objectives, organisational structure, and risk profile. A €400 billion balance sheet funded almost entirely through short-term wholesale markets had been added to the group. The capabilities assessment would have asked whether Hypo Real Estate had the systems and expertise to manage this combined exposure in a stressed environment.

An **internal environment assessment** would have examined whether the acquisition due diligence culture prioritised deal completion over risk identification — a common pattern when strategic transactions are driven by senior management ambition. It would have assessed whether the integration team had the authority and independence to challenge the assumptions underlying the transaction.

Risk criteria would have defined the materiality threshold and appetite boundaries against which the combined entity's exposures would be measured. A maturity mismatch of the scale embedded in Depfa's business model would have exceeded any reasonable appetite boundary for liquidity risk.

The **starting universe** would have included industry loss events involving wholesale-funding-dependent institutions, regulatory expectations for liquidity risk management (including the Basel Committee's work on liquidity coverage ratios), and the internal incident history from the integration process.

The **straw man risk list** would have included, as a matter of evidence-based discipline, the risk that wholesale funding markets could seize across all maturities simultaneously — because this risk was visible in the industry loss data and was, by October 2007, beginning to materialise in real time.

Would this have prevented the loss? I cannot know. But I can say with certainty that the methodology would have produced a documented assessment that identified the existential vulnerability in the funding model, mapped it to specific external environment changes, assessed it against defined risk criteria, and escalated it to the Board before the crisis became irreversible. The information was there. The methodology to process it was not.

The Foundation Is Set

Phase 1 produces six documented outputs: the PESTLE assessment, the internal context assessment, the internal environment assessment, the approved risk criteria, the evidence-based starting universe, and the straw man risk list. Together, these outputs create the foundation on which the rest of the methodology rests.

The foundation is not neutral. It shapes everything that follows. An external context assessment that misses a critical macroeconomic shift will leave the institution blind to the risks that shift creates. An internal environment assessment that gives the culture a clean bill of health when it should not will allow cultural weaknesses to corrupt the identification process. Risk criteria that are poorly calibrated will misclassify material risks as immaterial. A starting universe that relies on management judgement instead of evidence will reproduce the biases it was designed to eliminate.

This is why Foundation Setting is a formal phase, not a planning step. It has defined inputs, documented outputs, quality standards, and governance requirements. The CRO approves the risk criteria. The Board Risk Committee endorses them. The internal environment assessment findings are escalated if they reveal material weaknesses. The PESTLE assessment is dated, authored, and retained for audit.

With the foundation set and the straw man prepared, the actual identification begins. In Chapter 6 (Top-Down Identification: Workshops, SWIFT, and Delphi), we turn to the dual-track approach — the top-down workshops using SWIFT and Delphi techniques that bring senior management's strategic perspective to bear on the risk landscape, combined with the bottom-up analysis that captures what only the people closest to the

business can see. Neither track alone is sufficient. Together, they produce the comprehensive identification that the standards require and that the evidence shows banks consistently fail to achieve on their own.

1. Depfa Bank plc's balance sheet size at the time of the Hypo Real Estate acquisition is reported in Hypo Real Estate Holding AG, *Annual Report 2007*, and in BaFin's subsequent supervisory assessments. See also German Bundestag, *Parlamentarischer Untersuchungsausschuss zur Hypo Real Estate*, documentation of the Depfa acquisition.
2. BNP Paribas announced the suspension of three funds (Parvest Dynamic ABS, BNP Paribas ABS EURIBOR, and BNP Paribas ABS EONIA) on 9 August 2007, citing an inability to value structured assets. See BNP Paribas press release, 9 August 2007, and Financial Crisis Inquiry Commission (FCIC), *The Financial Crisis Inquiry Report*, January 2011, p. 248.
3. House of Commons Treasury Committee, *The Run on the Rock*, Fifth Report of Session 2007–08, HC 56-I, 26 January 2008, paragraph 1. The report confirmed this was the first bank run in the United Kingdom since the failure of Overend, Gurney & Company in 1866.
4. The €102 billion federal guarantee to Hypo Real Estate Holding AG was authorised by the German government through the Finanzmarktstabilisierungsgesetz (Financial Market Stabilisation Act) in October 2008. See German Federal Ministry of Finance press releases, October 2008, and BaFin supervisory records.
5. Hypo Real Estate Holding AG was nationalised through the Rettungsübernahmegesetz (Rescue Acquisition Act) enacted by the German Bundestag in April 2009, with the Federal Republic acquiring full ownership via the Financial Market Stabilisation Fund (SoFFin).
6. ISO 31000:2018, *Risk Management — Guidelines*, International Organization for Standardization, February 2018, Section 6.3 (Scope, context and criteria). Note: the 2018 edition restructured the numbering from the 2009 edition's Section 5.3 but retains the requirement that context establishment precedes risk assessment.
7. ISO 31010:2019, *Risk Management — Risk Assessment Techniques*, International Organization for Standardization, June 2019. Checklists are described as a risk identification technique in Section 6.3.2 and Table A.1.
8. US Department of Justice, *Seattle Woman Arrested for Data Theft Involving Large Capital One Financial Data Breach*, press release, 29 July 2019. Capital One Financial Corporation disclosed that the breach affected approximately 106 million individuals. See also OCC Consent Order AA-EC-2020-44, 6 August 2020, which assessed an \$80 million civil money penalty against Capital One for the security failures.
9. European Banking Authority (EBA), *Report on Management and Supervision of ESG Risks for Credit Institutions and Investment Firms*, EBA/REP/2021/18, 23 June 2021, and EBA, *Guidelines on the Management of ESG Risks*, issued under the Capital Requirements Directive.
10. European Central Bank, *Guide on Climate-Related and Environmental Risks: Supervisory Expectations Relating to Risk Management and Disclosure*, November 2020.
11. House of Commons Treasury Committee, *The Run on the Rock*, Fifth Report of Session 2007–08, HC 56-I, 26 January 2008, paragraph 18: Northern Rock's reliance on securitisation and wholesale funding rather than retail deposits is documented as approximately 75% of total funding.
12. Northern Rock plc was taken into temporary public ownership on 22 February 2008 under the Banking (Special Provisions) Act 2008. See HM Treasury press release, 22 February 2008.
13. The National Audit Office reported the net cost to the taxpayer of the Northern Rock intervention. See National Audit Office, *HM Treasury: The Nationalisation of Northern Rock*, HC 298, Session 2008–09, 20 March 2009, and subsequent NAO updates on the disposal of Northern Rock assets.
14. ISO 31000:2018, *Risk Management — Guidelines*, International Organization for Standardization, February 2018, Section 6.3 (Scope, context and criteria), which covers internal and external context establishment.

15. BCBS, *Principles for Effective Risk Data Aggregation and Risk Reporting* (BCBS 239), January 2013. The 14 principles cover governance, data architecture, accuracy, completeness, timeliness, and adaptability of risk data aggregation and reporting.
16. COSO, *Enterprise Risk Management — Integrated Framework*, September 2004. The Internal Environment is the first of eight components in the COSO ERM framework (the others being Objective Setting, Event Identification, Risk Assessment, Risk Response, Control Activities, Information & Communication, and Monitoring).
17. The concealment of Sean FitzPatrick's personal loans through temporary transfers to Irish Nationwide Building Society around year-end dates was established in the investigation by the Office of the Director of Corporate Enforcement (ODCE), Ireland, and documented in the subsequent criminal proceedings. See also Nyberg Commission, *Misjudging Risk: Causes of the Systemic Banking Crisis in Ireland*, Report of the Commission of Investigation into the Banking Sector in Ireland, March 2011, Chapter 4.
18. Anglo Irish Bank Corporation was nationalised under the Anglo Irish Bank Corporation Act 2009, enacted by the Oireachtas on 21 January 2009.
19. The €29.3 billion net cost to the Irish taxpayer from Anglo Irish Bank is documented in the National Asset Management Agency (NAMA) records and the Central Bank of Ireland's financial stability reports. See also Nyberg Commission, *Misjudging Risk: Causes of the Systemic Banking Crisis in Ireland*, March 2011, and the Comptroller and Auditor General, *Special Report on Anglo Irish Bank*, various years.
20. Anglo Irish Bank was merged with Irish Nationwide Building Society to form the Irish Bank Resolution Corporation (IBRC) in July 2011. IBRC was placed into special liquidation on 7 February 2013 under the Irish Bank Resolution Corporation Act 2013.
21. COSO, *Enterprise Risk Management — Integrating with Strategy and Performance: Compendium of Examples*, June 2018, provides illustrative impact and likelihood scales. The five-point scale recommendation is also reflected in COSO, *Enterprise Risk Management — Integrated Framework: Application Techniques*, September 2004.

Top-Down Identification: Workshops, SWIFT, and Delphi

The Workshop Where Everyone Agrees

The scene is universal. Twelve senior people sit around a conference table. The facilitator — a well-meaning colleague from the risk function — opens with a question: “What are the key risks facing our business?”

The most senior person in the room speaks first. He identifies three risks. They are the same three risks that appeared at the top of the previous year’s risk register: regulatory change, market volatility, and credit deterioration. The next person agrees with all three and adds a fourth. The person after that agrees with all four. Within thirty minutes, the room has produced a list of eight risks, each essentially a restatement of the existing register, and the facilitator is already moving to the assessment phase.

No one mentions counterparty concentration in prime brokerage. No one raises the question of whether the bank’s exposure to leveraged family offices is adequately visible. No one asks what would happen if a single client defaulted across multiple desks simultaneously. At Credit Suisse, Archegos Capital Management would eventually demonstrate exactly why those questions mattered — at a cost of \$5.5 billion.¹

This kind of workshop is not unusual. It is typical. And it is worse than useless — it is actively dangerous, because it creates a documented record showing that senior management has “identified” the institution’s risks, when in reality the process has done nothing more than confirm existing assumptions. The risk register looks comprehensive. The governance box is ticked. And the bank remains blind to the exposures that could destroy it.

Traditional brainstorming, applied to risk identification in a hierarchical institution, does not identify risks. It identifies consensus. And consensus, in risk management, is the enemy.

Why Brainstorming Fails

The risk identification literature, and much of industry practice, defaults to brainstorming as the primary workshop technique. ISO 31010 includes brainstorming in its catalogue of thirty-one risk assessment techniques, describing it as suitable for the identification stage.² On paper, it should work: assemble knowledgeable people, ask them to generate ideas, capture the output.

In practice, brainstorming produces predictable pathologies when applied to risk identification in a banking context. The academic literature on group decision-making has documented these for decades, but the risk management profession has been slow to absorb the implications.

Authority bias is the most damaging. In a hierarchical organisation — and every bank is hierarchical — the views expressed by the most senior person in the room anchor the entire discussion. Junior participants self-censor. Middle managers calibrate their contributions to what they believe their superiors want to hear. The risks that get identified are the risks that the most powerful person in the room already knows about. This is not identification. It is ratification.

Groupthink compounds the problem. Once an initial set of risks has been articulated, the group converges rapidly. Disagreement carries social cost. Proposing a risk that contradicts the emerging consensus requires the proposer to implicitly suggest that their colleagues have missed something — and in a room with the CRO, the head of the investment bank, and the treasurer, few people are willing to take that position.

Recall bias adds a third distortion. Under time pressure, participants default to risks they can most easily retrieve from memory — typically the risks that have materialised recently, the risks that are currently in the news, or the risks that appear in the existing register. Novel risks, emerging risks, and risks that sit in the gaps between established categories are structurally disadvantaged.

Production blocking — the constraint that only one person can speak at a time — means that in a two-hour workshop with twelve participants, each person has on average ten minutes of speaking time. Subtract the facilitator's contributions, the introduc-

tions, the administrative overhead, and you are left with perhaps five minutes per participant to identify, articulate, and defend the risks they believe matter most. This is not a methodology. It is a lottery.

The result is a workshop that produces an output indistinguishable from what a competent risk analyst could have generated alone in an afternoon by reading last year's register and the morning's Financial Times. The collective intelligence of twelve senior leaders, each with deep institutional knowledge and strategic perspective, has been wasted.

This is the problem that the **Structured What-If Technique** was designed to solve.

The Dual-Track Rationale

Before describing the top-down techniques in detail, it is important to understand why the methodology requires two parallel identification tracks — top-down and bottom-up — and why neither alone is sufficient.

ISO 31000 Section 5.2 establishes communication and consultation as a continuous parallel activity across all phases of risk management.³ The standard requires that multiple perspectives are sought, that differing perceptions are captured rather than averaged away, and that the risk identification process draws on the knowledge of people at every level of the organisation.

Top-down identification — the subject of this chapter — captures risks that are visible from the strategic vantage point of senior management. These are the risks that arise from the institution's overall business model, its competitive position, the macroeconomic environment, its interconnectedness with counterparties and markets, and its exposure to emerging trends. Senior leaders see these risks because they sit at the intersection of strategy, regulation, and market dynamics.

Bottom-up identification — the subject of Chapter 7 (Bottom-Up Identification: Templates, RCSA, and Specialist Processes) — captures risks that are visible only to the people closest to the business. The trader who knows that the risk system does not capture a particular exposure. The operations manager who knows that the reconciliation process has a gap. The compliance officer who has seen a pattern in customer complaints. These risks are granular, operational, and often invisible to senior management.

As Oliver Wyman observed in their analysis of risk identification practices, the use of parallel top-down and bottom-up processes provides a higher likelihood of identifying an organisation's full suite of risks than either process in isolation. The top-down process captures strategic and emerging risks; the bottom-up process achieves comprehensive, granular coverage. The reconciliation between them — which Chapter 8 (Reconciliation and the Enterprise Portfolio View) will address — is where the methodology creates its most distinctive value.

The straw man risk list, prepared by the Risk Identification Lead during Phase 1, bridges both tracks. It is circulated to workshop participants in advance, providing an evidence-based starting point that prevents blank-page syndrome while being explicitly designed to be challenged, modified, and torn apart. The framing of this document matters enormously — present it as a “proposed risk list” and participants defer to it; present it as a document to be dismantled and they engage critically.

Designing the Top-Down Workshop

The workshop is the primary vehicle for top-down identification. Its design determines whether it produces genuine insight or merely ratifies existing assumptions. Every element — from participant selection to the questioning methodology to how outputs are recorded — must be deliberately structured to overcome the pathologies described above.

Participant Selection

The workshop requires participants who collectively hold the strategic perspective of the institution. At a minimum, this means:

- **CRO or deputy** — owns the process, provides independent challenge
- **Business unit heads** (or their designated senior representatives) — hold the strategic knowledge of each major business
- **Treasury / ALM** — understands the institution's funding and liquidity position
- **Compliance / Legal** — brings the regulatory and legal perspective
- **Finance / CFO office** — understands the financial reporting and capital implications
- **Risk Identification Lead** — facilitates, documents, and challenges

The participant list should be reviewed annually. Permanent membership creates its own form of groupthink — the same people, asking the same questions, producing the same answers. Rotating one or two participants each cycle introduces fresh perspective.

External participants should be considered where the institution faces genuinely novel risks. The Delphi Method, described later in this chapter, provides the structured mechanism for incorporating external expertise.

Pre-Workshop Independent Assessment

This is the single most important design element. Before the workshop convenes, every participant independently identifies their top ten risks, in writing, submitted to the Risk Identification Lead.

The independent assessment serves three purposes. First, it captures each participant's uncontaminated view before group dynamics take effect — eliminating the authority bias that destroys unstructured brainstorming. Second, it provides the facilitator with a complete map of where participants agree and, more importantly, where they disagree. The facilitator uses this map to design the workshop's questioning strategy, focusing time on the areas of disagreement rather than the areas of consensus. Third, it creates a documented baseline that can be compared to the workshop's final output, providing the audit trail to demonstrate that the process genuinely challenged and refined initial views rather than simply rubber-stamping them.

Each participant receives, alongside the straw man risk list, a **briefing pack** containing:

- The current PESTLE assessment (from Phase 1)
- The internal context and internal environment summaries
- The prior-year principal risk list with trend indicators
- Any relevant regulatory communications or supervisory findings
- Industry loss events from the period since the last identification cycle

The briefing pack ensures that participants arrive with a common factual base. The independent assessment ensures they apply their own judgement to it before the group discussion begins.

The Workshop Structure

A properly structured top-down workshop follows a defined sequence:

| Stage | Duration | Purpose |
|--------------------------------------|------------|--|
| Opening | 15 min | Facilitator states objectives, ground rules, and the principle that the straw man is to be challenged |
| Independent assessment review | 30 min | Facilitator presents the aggregated results of pre-workshop assessments (anonymised), highlighting areas of disagreement |
| SWIFT identification rounds | 90–120 min | Systematic what-if questioning across all strategic objectives and risk categories |
| Emerging risk session | 30 min | Forward-looking identification of risks not yet in the taxonomy |
| Multivoting prioritisation | 20 min | Each participant allocates votes to the risks they consider most significant |
| Gap review | 15 min | Facilitator checks identified risks against taxonomy and PESTLE for completeness |

Total duration: approximately four hours. This is not a meeting that can be compressed into a lunch slot. Institutions that attempt to run risk identification workshops in ninety minutes are signalling that they do not take the process seriously.

SWIFT: The Structured What-If Technique

SWIFT — the Structured What-If Technique — is described in ISO 31010 (Section B.7) as a systematic method for identifying risks using a set of prepared guide words applied to each element of the system under examination.⁴ Originally developed for the process safety industry, SWIFT has proven highly effective for banking risk identification because it imposes structure on the questioning process without constraining the scope of what can be identified.

The key difference between SWIFT and unstructured brainstorming is that SWIFT uses **prepared prompts** applied systematically. The facilitator does not ask “what are the risks?” — a question that invites recall bias and authority anchoring. Instead, the facilitator works through a structured set of guide words, applying each one to specific elements of the institution’s operations, strategy, and environment.

Guide Words and Prompt Structure

The SWIFT facilitator prepares a prompt matrix before the workshop, combining guide words with the institution’s specific context. The standard guide words for banking risk identification include:

| Guide Word | Purpose | Example Prompt |
|---|--|--|
| What if... | Explores specific scenarios | “What if wholesale funding markets close for 30 days?” |
| What would happen if... | Traces consequences | “What would happen if our largest counterparty defaulted?” |
| Could someone / something... | Probes internal vulnerabilities | “Could someone in the CIO function take positions that bypass trading risk limits?” |
| What has changed since last year that could... | Identifies emerging and evolving risks | “What has changed in the regulatory environment that could require us to hold additional capital?” |
| What are we assuming that might not be true? | Challenges embedded assumptions | “What are we assuming about the correlation between our credit portfolios?” |
| Where are the gaps between... | Identifies interface and silo risks | “Where are the gaps between what equities surveillance monitors and what AML monitors?” |

The facilitator applies these guide words systematically to each of the following domains:

1. **Strategic objectives** — risks to each of the institution's stated strategic goals
2. **Business areas** — risks specific to each major business line
3. **Risk categories** — risks within each L1 taxonomy category, with particular attention to cross-category risks
4. **External environment** — risks arising from the PESTLE factors identified in Phase 1
5. **Interconnections** — risks arising from dependencies between business areas, counterparties, or systems

Why SWIFT Works Where Brainstorming Fails

SWIFT overcomes each of the brainstorming pathologies:

- **Authority bias:** The facilitator controls the questioning. No single participant dominates the direction of discussion. The prompts ensure that every domain is covered regardless of what the most senior person wants to discuss.
- **Groupthink:** The what-if structure invites disagreement by design. When the facilitator asks "what are we assuming that might not be true?", the question itself legitimises challenge. Participants are responding to a prompt, not contradicting a colleague.
- **Recall bias:** The systematic coverage of all strategic objectives, business areas, and risk categories means that the process does not depend on what participants happen to remember. The prompt matrix ensures comprehensive coverage.
- **Production blocking:** While only one person speaks at a time in the main discussion, the pre-workshop independent assessment has already captured every participant's views. The SWIFT session builds on that foundation rather than starting from scratch.

Facilitation: The Craft That Standards Cannot Teach

ISO 31010 describes SWIFT as a technique. It does not describe how to facilitate it in a room full of senior bankers who would rather be doing something else, who are accustomed to being the most important person in every meeting they attend, and who may have active incentives to suppress certain risks.

The Risk Identification Lead, who facilitates these workshops, requires a specific set of skills that go beyond technical risk expertise:

Managing dominant voices. In every workshop, one or two participants will attempt to control the discussion. The facilitator must redirect without confrontation: “Thank you — I want to hear from treasury on this point before we move on.” The pre-workshop independent assessment provides the tool: “I note that several participants identified a risk in this area that we haven’t discussed yet — let me bring that forward.”

Drawing out dissent. The most valuable contributions often come from the quietest participants. Structured turn-taking — going around the table systematically rather than relying on voluntary contributions — ensures that every voice is heard. The facilitator should watch for non-verbal disagreement: the head of compliance who frowns but does not speak, the business unit head who shifts uncomfortably when a particular risk is discussed. These signals must be pursued.

Maintaining challenge culture. The facilitator must model the behaviour they want to see. When a risk is identified, the facilitator should ask: “Who disagrees? What am I missing? What’s the alternative view?” Silence should not be interpreted as agreement. The facilitator should explicitly state: “I’m interpreting silence as agreement — if anyone holds a different view, now is the time.”

Protecting the process from political interference. The most difficult facilitation challenge occurs when a senior participant attempts to remove a risk from the list or downgrade its significance for reasons that are political rather than analytical. The facilitator must ensure that the risk remains documented, the challenge is recorded, and the final decision reflects genuine analytical judgement rather than political pressure. This requires the structural independence described in Chapter 3 (Governance: Who Owns What) — the Risk Identification Lead must have the mandate and the protection to maintain the integrity of the output.

These skills are learned through practice, not training. No course prepares a facilitator for the moment when a business unit head declares, in front of twelve colleagues, that a risk does not exist. The risk may well exist — and may appear in the final inventory. But the political skill required to navigate that moment — to document the disagreement, to seek supporting evidence from other participants, to escalate through the governance structure when necessary — is the craft that distinguishes a process that works from a process that merely exists on paper.

Scenario Analysis as an Identification Technique

ISO 31010 includes scenario analysis as a technique applicable to both the identification and analysis stages. Within the methodology, scenario analysis serves a specific identification purpose: it reveals risks that are invisible under normal conditions but become apparent under stress.

The Phase 1 PESTLE assessment produces a set of external factors that could affect the institution. Scenario analysis takes these factors and constructs plausible combinations — not to quantify their impact (that is the assessment function of Phase 3, covered in Chapter 9 (Assessment — Scoring, Multi-Dimensional Impact, and Data Quality)) but to ask: “Under this scenario, what risks would crystallise that we have not yet identified?”

The CCAR framework, as implemented through Fed SR 15-18, provides the most rigorous regulatory application of this approach.⁵ CCAR requires institutions to develop firm-specific stress scenarios that target their particular vulnerabilities. As the Fed has increasingly emphasised, this requires risk identification that is much more granular than generic categories — institutions must identify specific risk events, their drivers, and the transmission mechanisms through which they affect the balance sheet.

The workshop facilitator uses scenario prompts to drive this identification:

- “Under a scenario of prolonged interest rate inversion, what risks would emerge that are not on our current list?”
- “If our three largest counterparties defaulted simultaneously, what would the second-order effects be?”
- “If a major operational failure disabled our payments system for 48 hours, what risks beyond direct operational loss would materialise?”

Scenario analysis in the identification phase is deliberately forward-looking. It is not about predicting the future. It is about ensuring that the risk inventory is robust to a range of plausible futures rather than calibrated only to the current environment.

Multivoting: Transparent Prioritisation

After the identification phase of the workshop is complete, the methodology uses **multivoting** to produce an initial, transparent prioritisation. Each participant receives a set number of votes — typically five to ten — and allocates them to the risks they consider most significant. Multiple votes may be placed on a single risk, allowing participants to signal strength of conviction.

Multivoting serves two purposes. First, it produces a prioritisation that is transparent and traceable — the distribution of votes is documented, providing the audit trail for why certain risks were prioritised over others. Second, it eliminates single-voice dominance. In an unstructured prioritisation, the most senior person's view prevails by default. With multivoting, a risk that every participant votes for once outranks a risk that one participant votes for heavily. Collective judgement prevails over individual authority.

The multivoting output is preliminary. The detailed four-dimensional assessment framework described in Chapter 9 provides the rigorous scoring methodology. But the workshop prioritisation shapes the allocation of analytical effort in the phases that follow.

The Archegos Failure: What a SWIFT Workshop Would Have Asked

On 26 March 2021, Archegos Capital Management — a family office run by Bill Hwang — collapsed, triggering approximately \$10 billion in net losses across its prime brokers as an estimated \$20 billion in concentrated positions were unwound.⁶ Credit Suisse absorbed \$5.5 billion of those losses, a blow that contributed materially to the institution's eventual demise.⁷

The mechanism was straightforward. Hwang used total return swaps to build massive concentrated positions in a handful of stocks, leveraged at five to eight times, across multiple prime brokers.⁸ No single prime broker had visibility into his total exposure. Each bank assessed its bilateral exposure to Archegos as manageable. None asked the question that mattered: what is this client's aggregate position across all brokers?

The risk identification failure was not a failure of information — it was a failure of questioning methodology. The information existed within each institution. Margin calls were rising. Position sizes were growing. Concentration in single names was visible on each desk's books. But the workshop process — to the extent one existed — did not include prompts designed to surface this type of cross-counterparty, cross-institution concentration risk.

Consider what a SWIFT workshop would have asked:

- **“What if our largest family office client has equivalent or larger positions with every other prime broker?”** — this prompt, applied to the prime brokerage business, would have immediately raised the question of aggregate leverage visibility.
- **“Where are the gaps between what our prime brokerage desk sees and what other banks' prime brokerage desks see?”** — the gap analysis guide word would have surfaced the information asymmetry.
- **“What are we assuming about the risk of concentrated leverage in family offices that might not be true?”** — the assumption-challenge prompt would have exposed the belief that bilateral exposure management was sufficient.
- **“What has changed in the family office space since last year that could create new risks?”** — the evolution prompt would have flagged the growth in total return swap usage to avoid 13F disclosure requirements.⁹

The independent pre-workshop assessment would have added another layer. If the head of prime brokerage, the head of market risk, and the CRO each independently identified their top risks before the workshop, the facilitator would have had visibility into whether anyone had flagged concentrated family office leverage — and, critically, whether the absence of that risk from any participant's list was itself a signal.

None of these questions are difficult. None require information that was unavailable. What they require is a methodology that forces them to be asked — systematically, regardless of what the most senior person in the room wants to discuss.

What was missing: A structured workshop methodology that applies systematic prompts across all counterparty types, that challenges assumptions about bilateral versus aggregate exposure, and that uses pre-workshop independent assessments to surface divergent views before group dynamics suppress them. SWIFT, applied to the prime brokerage business with concentration-focused guide words, would have made the Archegos exposure a discussion topic rather than an unasked question.

The Delphi Method: When Workshops Are Not Enough

SWIFT workshops are effective for identifying risks within the current strategic and operational landscape. They are less effective for a specific category of risk: emerging risks where no historical data exists, where the risk may not yet fit within the established taxonomy, and where the conventional wisdom of the institution may be actively wrong.

For these risks, the methodology uses the **Delphi Method** — a structured, multi-round, anonymous expert consultation process described in ISO 31010 (Section B.4).¹⁰

How Delphi Works

The Delphi Method operates through a defined sequence:

- 1. Panel assembly.** A panel of ten to twenty experts is assembled. The panel should include internal experts from across the institution and, where appropriate, external experts — academics, regulators, industry consultants, or specialists in relevant fields. The diversity of the panel is critical: Delphi's value lies in aggregating perspectives that no single individual holds.
- 2. Round 1: Independent identification.** Each panellist independently identifies emerging risks they believe could affect the institution within a three-to-five-year horizon. Responses are submitted anonymously to the Risk Identification Lead, who acts as the **oracle** — the only person who knows which responses came from which panellist.

- 3. Aggregation and circulation.** The oracle aggregates and summarises all responses, identifying common themes, areas of agreement, and areas of divergence. This summary is circulated back to the panel without attribution.
- 4. Round 2: Revision.** Each panellist reviews the aggregated results, considers perspectives they had not previously encountered, and submits a revised assessment. Panellists may change their views, strengthen their convictions, or add new risks prompted by others' contributions.
- 5. Further rounds.** The process repeats for two to three rounds until reasonable convergence emerges. Convergence does not mean unanimity — persistent disagreement is itself a valuable signal, indicating genuine uncertainty about whether a particular risk will materialise.
- 6. Output.** A prioritised list of emerging risks with consensus descriptions, dissenting views documented, integrated into the top-down risk list and flagged for taxonomy review.

Why Anonymity Matters

The anonymity of the Delphi process is not a convenience — it is the mechanism that makes the technique work for precisely the risks where workshops fail.

In a workshop, the social dynamics of the room determine what gets said. A junior analyst who believes that the institution's business model is fundamentally exposed to a particular emerging risk will not say so in front of the CEO. An external consultant who suspects that the bank's exposure to a new asset class is inadequately understood will not challenge the head of trading in a face-to-face setting. A compliance officer who believes the bank's regulatory posture is more fragile than leadership acknowledges will keep that view private.

Delphi eliminates these constraints. The junior analyst's view carries equal weight to the CEO's. The external consultant's challenge is heard without attribution. The compliance officer's concern enters the record without career risk. The oracle ensures that every perspective is represented in the aggregated output, and the iterative rounds allow participants to revise their views in light of perspectives they would never have heard in a workshop setting.

Delphi's Double Duty

As established in Chapter 4 (The Risk Taxonomy), the Delphi Method serves two purposes within the methodology. Its primary purpose is to identify specific emerging risks that the institution should be monitoring and, where appropriate, including in its risk inventory. Its secondary purpose is to provide intelligence for **taxonomy evolution** — when multiple Delphi panellists independently identify a risk that does not fit within the current taxonomy structure, that is a signal that the taxonomy needs to be updated.

Climate risk provides a historical example. Before climate and environmental risk became an established L1 taxonomy category, institutions that used Delphi-style processes were identifying physical risk to asset portfolios, transition risk from regulatory changes, and stranded asset risk from energy policy shifts. These identifications — emerging from anonymous expert consultation rather than workshop consensus — provided the evidence base for adding climate risk to the taxonomy years before regulatory mandates required it.

The Delphi output feeds into two downstream processes: the risk inventory (for risks that are sufficiently developed to warrant inclusion) and the taxonomy maintenance process (for signals that the classification structure itself needs to evolve).

Wirecard: When the Identification Ecosystem Inverts

The Wirecard fraud, which finally collapsed in June 2020 with €1.9 billion in fabricated cash balances,¹¹ provides the most extreme illustration of why the Delphi Method's inclusion of external experts is not optional but essential.

Wirecard's internal risk identification processes were captured by the fraud itself. The company fabricated revenues from third-party acquiring partners in Asia, inflated cash balances at Philippine trustee banks, and sustained the deception for years through deliberate intimidation of anyone who questioned the accounts.

The extraordinary feature of the Wirecard case was not merely that internal identification failed — it was that the external identification ecosystem was actively inverted. Short-sellers and investigative journalists at the Financial Times identified the fraud through forensic analysis of publicly available data. They raised specific, documented concerns about fabricated revenues and phantom bank balances. The German financial

regulator, BaFin, responded not by investigating Wirecard but by filing criminal complaints against the journalists.¹² The regulator defended the fraudster and prosecuted those who identified the risk.

This is the scenario that the Delphi Method is designed to address — not in its extreme form, but in its underlying dynamic. In every institution, there are risks that the internal consensus is actively wrong about. Risks where the prevailing view is not merely incomplete but inverted — where the institution believes it is safe precisely in the area where it is most exposed. Traditional workshops, conducted entirely with internal participants who share the institutional consensus, cannot surface these risks. The consensus itself is the obstacle.

A Delphi panel that included external participants — analysts with no institutional loyalty, academics with no career risk, consultants with cross-industry perspective — would have introduced the dissenting view into the process. The short-sellers' analysis of Wirecard's third-party revenue was available from 2015.¹³ The FT's investigative reporting began in 2019.¹⁴ An external Delphi panellist with access to this analysis would have submitted it anonymously, and the aggregation process would have ensured it reached decision-makers without the political dynamics that allowed Wirecard's management to suppress it through legal threats and regulatory capture.

What was missing: A structured mechanism for incorporating external expert perspectives into the risk identification process, protected by anonymity and institutional independence. The Delphi Method, with external panellists and anonymous submission, would have surfaced the dissenting analysis that internal processes were structurally incapable of producing because the fraud had captured the identification function itself.

Connecting Top-Down Identification to the CCAR Framework

For US banking institutions subject to the Federal Reserve's Comprehensive Capital Analysis and Review, top-down identification has a specific regulatory dimension. Fed SR 15-18 requires that institutions maintain a **Material Risk Inventory** that is updated quarterly — not annually — through a process of active re-identification.¹⁵

The connection between SWIFT workshops and CCAR is direct. The Fed expects institutions to identify risks that may appear only under stress, to understand how scenarios that break historical patterns would affect the balance sheet, and to demonstrate that the risk identification process is forward-looking rather than backward-looking. SWIFT's scenario-based prompts — “what if wholesale funding markets close?”, “what would happen if credit spreads widened beyond historical precedent?” — are precisely the questioning methodology that produces the identification outputs CCAR requires.

The quarterly re-identification cycle established in Chapter 3 means that SWIFT workshops are not annual events. The full annual workshop covers all domains comprehensively. Quarterly workshops are more focused — targeting areas where the risk landscape has changed, where new risks have emerged, or where prior identifications need to be revisited in light of new information. The straw man for quarterly workshops draws on the prior quarter's inventory, updated PESTLE factors, and any event-driven triggers that have occurred since the last cycle.

This integration between the identification methodology and the CCAR framework is not incidental. At institutions subject to CCAR, the risk identification process must feed directly into the submission. The Material Risk Inventory produced by the methodology — through SWIFT workshops, Delphi consultations, and the bottom-up processes described in Chapter 7 — becomes the foundation for scenario design and stress testing. Risk identification is not a standalone exercise. It is the front end of the capital planning process.

Top-Down Outputs

A properly conducted top-down identification — combining pre-workshop independent assessments, SWIFT workshops, scenario analysis, Delphi consultation, and multivoting — produces a defined set of outputs:

| Output | Description |
|-------------------------------|--|
| Principal risk list | 15–30 risks identified as the institution's most significant, with preliminary descriptions and proposed risk owners |
| Emerging risk register | |

| Output | Description |
|----------------------------|--|
| | Risks from the Delphi process that are not yet material but require monitoring; flagged for taxonomy review where they do not map to existing categories |
| Assumption register | Key assumptions that the workshop identified and challenged — documented for the audit trail and for revisiting in subsequent cycles |
| Disagreement log | Areas where participants held materially different views — not averaged away but documented for further analysis |
| Taxonomy gap list | Risk types identified in the workshop or Delphi process that do not map cleanly to the current taxonomy — fed into the taxonomy maintenance process |

The disagreement log deserves particular emphasis. In most institutions, disagreement is treated as a problem to be resolved — the workshop must produce a consensus output. The methodology treats disagreement differently. Where two senior participants hold genuinely different views about whether a risk is material, that divergence is itself information. It signals uncertainty, complexity, or genuine analytical disagreement that the assessment phase (Chapter 9) must resolve through evidence rather than through the social dynamics of a meeting room.

The Limits of Top-Down Identification

Top-down identification, however well designed, has structural limitations. Senior leaders see the landscape from altitude. They understand strategic risks, macroeconomic forces, and competitive dynamics. They do not see the trader who has found a way to circumvent a position limit. They do not see the operations process that has a reconciliation gap. They do not see the customer complaint pattern that signals an emerging conduct risk.

These risks — granular, operational, and often invisible to anyone who does not work directly in the affected area — can be just as destructive as the strategic risks identified in workshops. The industry loss database is populated with institutions that failed not because of strategic miscalculation visible to the Board, but because of operational failures visible only to the people closest to the business.

The methodology addresses this through the second identification track: bottom-up identification using standardised templates, Risk and Control Self-Assessments, and specialist sub-processes for specific risk types. Neither track alone captures the full picture. The top-down track identifies the twenty risks that could destroy the institution. The bottom-up track identifies the two hundred risks that collectively define the institution's operational risk profile. The reconciliation between them — the subject of Chapter 8 — is where the methodology produces its enterprise portfolio view.

Chapter 7 describes the bottom-up track: how standardised templates capture granular risks across every business unit, how RCSA and specialist sub-processes feed into the central inventory, and how the Risk Identification Lead ensures that nothing falls through the gaps between specialist functions and the enterprise process.

-
1. Credit Suisse Group, *Report of the Special Committee of the Board of Directors on the Supply Chain Finance Funds Matter and Related Matters* (prepared by Paul, Weiss, Rifkind, Wharton & Garrison LLP), 29 July 2021. The report documented the \$5.5 billion loss from the Archegos default.
 2. ISO 31010:2019, *Risk management — Risk assessment techniques*, International Organization for Standardization. Table A.1 lists 31 techniques and their applicability to the identification, analysis, and evaluation stages.
 3. ISO 31000:2018, *Risk management — Guidelines*, International Organization for Standardization, Section 5.2 ("Communication and consultation").
 4. ISO 31010:2019, *Risk management — Risk assessment techniques*, International Organization for Standardization, Annex B.7 ("Structured 'What if?' Technique (SWIFT)").
 5. Board of Governors of the Federal Reserve System, *SR 15-18: Federal Reserve Supervisory Assessment of Capital Planning and Positions for LISCC Firms and Large and Complex Firms*, 18 December 2015.
 6. Multiple prime brokers disclosed Archegos-related losses in Q1 2021 earnings: Nomura (~\$2.9 billion), Morgan Stanley (~\$911 million), and others. Aggregate net losses across counterparties were widely reported as approximately \$10 billion. See also: U.S. Securities and Exchange Commission, *SEC Charges Archegos Capital Management and its Founder Bill Hwang*, 27 April 2022.
 7. Credit Suisse Group, *Report of the Special Committee of the Board of Directors* (Paul, Weiss), 29 July 2021, p. 4. The report confirmed the \$5.5 billion charge from the Archegos default.
 8. Credit Suisse Group, *Report of the Special Committee of the Board of Directors* (Paul, Weiss), 29 July 2021, pp. 10-15. The report detailed Hwang's use of total return swaps and the leverage ratios across prime brokerage counterparties.

9. Total return swaps were not reportable under SEC Form 13F at the time of the Archegos collapse. The SEC subsequently adopted amendments to Form PF (effective 2023) and proposed amendments to Regulation 13D-G and Form 13F to address this disclosure gap. See SEC Release No. 34-93784 (proposed rule), 15 February 2022.
10. ISO 31010:2019, *Risk management — Risk assessment techniques*, International Organization for Standardization, Annex B.4 ("Delphi technique").
11. Wirecard AG filed for insolvency on 25 June 2020 after disclosing that EUR 1.9 billion in cash balances reported in its accounts "probably did not exist." See: *Wirecard AG — Insolvency filing*, Amtsgericht Munich, 25 June 2020; EY Wirecard audit disclaimer, 18 June 2020.
12. BaFin filed a criminal complaint against *Financial Times* journalists with the Stuttgart public prosecutor's office in February 2019, alleging market manipulation. BaFin also imposed a two-month short-selling ban on Wirecard shares (18 February to 18 April 2019). See: German Parliamentary Inquiry Committee on Wirecard (*Wirecard-Untersuchungsausschuss*), Bundestag, 2021.
13. Zatarra Research published a short-seller report on Wirecard in February 2016 alleging fraud in the company's Asian operations. Earlier concerns were raised by short sellers as early as 2008 and 2015.
14. Dan McCrum, *Financial Times*, "House of Wirecard" series, beginning January 2019. McCrum's investigative reporting into Wirecard's third-party acquiring business in Asia ran from early 2019 through the company's collapse in June 2020.
15. Board of Governors of the Federal Reserve System, *SR 15-18: Federal Reserve Supervisory Assessment of Capital Planning and Positions for LISC Firms and Large and Complex Firms*, 18 December 2015. The letter requires firms to maintain a comprehensive firm-wide risk identification process that is forward-looking and updated at least quarterly.

Bottom-Up Identification: Templates, RCSA, and Specialist Processes

The Risk Assessment That Has Not Changed in Three Years

The first round of bottom-up risk assessments arrives from the business units. A standardised template was distributed with detailed guidance. Eight weeks were allowed for completion. The submissions arrive on schedule. The first one — from a major trading division — is compared against the prior year's submission retrieved from the archives.

It is identical. Not similar. Identical. The same risks, in the same order, with the same scores, the same control descriptions, the same owner names — one of whom left the firm six months earlier. The dates have been updated. Nothing else has.

The second submission follows the same pattern. The third is no different. The business units have treated the bottom-up risk identification exercise as an administrative burden — a form to be completed and returned, not an analytical process to be performed. They took last year's spreadsheet, changed the date in the header, and sent it back.

This is **compliance theatre**: the appearance of risk identification without the substance. And it is the single most common failure mode in bottom-up processes. The template exists. The governance requires its completion. The business unit complies. But no one actually identifies anything.

The problem is not that bottom-up identification is unnecessary. It is essential. The top-down workshops described in Chapter 6 (Top-Down Identification: Workshops, SWIFT, and Delphi) identify the twenty risks that could destroy the institution — the strategic, emerging, and cross-cutting risks visible from senior management's vantage point. But those workshops cannot see the two hundred operational risks that collectively define the institution's risk profile: the process failures, the control weaknesses, the technology

vulnerabilities, the third-party dependencies, the conduct risks embedded in incentive structures. Only the people closest to the business can see those. The bottom-up track exists to capture what the top-down track structurally cannot.

The challenge is making it work. A bottom-up process that produces compliance theatre is worse than no process at all, because it creates a documented record suggesting that identification has occurred when it has not. The institution believes it has a comprehensive risk inventory. It does not. It has last year's inventory with this year's date.

This chapter describes how to design a bottom-up identification process that produces genuine, granular risk identification — through standardised templates, structured techniques, specialist sub-processes, and an integration discipline that ensures nothing falls through the gaps.

The Standardised Risk Assessment Template

The foundation of bottom-up identification is a **Standardised Risk Assessment Template** that every business unit completes for every risk it faces. Standardisation is not bureaucracy. It is the mechanism that makes aggregation, comparison, and reconciliation possible. If each business unit submits risk information in its own format, with its own definitions, using its own scales, the Risk Identification Lead cannot produce an enterprise view. The template enforces a common language.

The template captures eleven fields:

| Field | Description |
|--------------------------------|---|
| Risk ID | Unique identifier assigned by the central function, ensuring traceability across cycles |
| Taxonomy Classification | L1 / L2 / L3 mapping from the institution's risk taxonomy (Chapter 4 (The Risk Taxonomy)) |
| Risk Definition | Plain-language description of the risk event — what could happen, not what category it belongs to |

| Field | Description |
|--------------------------------------|--|
| Underlying Drivers — Direct | Factors that directly cause or trigger the risk |
| Underlying Drivers — Indirect | Factors that amplify or enable the risk without directly causing it |
| Quantitative Metrics | Measurable indicators with current values and limits, where available |
| Qualitative Information | Narrative assessment where quantification is not possible or meaningful |
| Current Controls | Existing mitigants, their type (prevent / detect / correct), and assessed effectiveness |
| Risk Owner | Named individual — not a committee, not a function (Chapter 3 (Governance: Who Owns What)) |
| Emerging Risk Indicators | Early warning signals that this risk may be increasing |
| Data Quality Rating | Assessment of confidence in the underlying data (High / Medium / Low / Very Low) |

Three of these fields deserve particular attention because they are where most bottom-up submissions fail.

Underlying Drivers — both direct and indirect — force the business unit to think about causation, not just classification. A risk definition that says “operational loss from technology failure” is a taxonomy label, not an identification. The driver fields require the assessor to specify *what* could fail, *why* it might fail, and *what conditions* would make failure more likely. This is where identification actually happens. In practice, risk definitions are often adequate while driver fields are empty or contain single-word entries. The business units have classified their risks but have not identified them.

Current Controls must specify the control type — preventive, detective, or corrective — and an honest assessment of effectiveness. The natural tendency is to list controls as “effective” because admitting a control weakness feels like admitting a management failure. The template must explicitly require an effectiveness rating with supporting evidence, not just a binary effective/ineffective judgement. A control that exists on paper but is routinely overridden, bypassed, or under-resourced is not an effective control. The assessment must reflect operational reality, not design intent.

Data Quality Rating is the field most likely to be overlooked and most important for downstream assessment. A risk scored as “high impact, low likelihood” based on expert judgement alone carries fundamentally different information from the same score based on ten years of loss data. The four-level rating (High, Medium, Low, Very Low) creates transparency about the evidential basis for each assessment. Chapter 9 (Assessment — Scoring, Multi-Dimensional Impact, and Data Quality) will describe how this rating feeds into the four-dimensional assessment framework. For now, the critical point is that the rating must be assigned honestly. In practice, business units default to “Medium” for everything. The Risk Identification Lead must challenge this — a risk with no quantitative data and no historical precedent is not “Medium” quality. It is “Low” or “Very Low,” and that information matters.

Techniques for Bottom-Up Identification

A template is a capture mechanism, not an identification technique. Handing a business unit a blank template and asking them to fill it in is an invitation to produce compliance theatre. The template must be supported by structured identification techniques that guide the assessor through a systematic process.

Chapter 2 (The Foundations: Standards and Frameworks) mapped the ISO 31010 techniques to their primary chapters in this methodology. For bottom-up identification, four techniques are recommended.

Structured Interviews

Structured interviews use a predetermined set of questions applied consistently across interviewees to elicit risk information that might not surface in workshops or self-assessment. Unlike the top-down SWIFT workshops described in Chapter 6, which bring senior management together in a group setting, structured interviews are conducted one-on-one with individuals across different levels and functions within a business unit.

The value of structured interviews in bottom-up identification is threefold. First, they capture perspectives from individuals who would never speak in a workshop — the mid-level operations manager who knows the reconciliation process breaks every month-end, the compliance officer who has raised concerns three times without response, the technology specialist who knows the disaster recovery plan has never been tested. Second, the one-on-one setting removes the authority bias and groupthink that Chapter 6 identified as brainstorming pathologies. Third, the structured format ensures consistency — the same questions asked of the head of settlements and the junior operations analyst produce comparable outputs.

The Risk Identification Lead should develop interview protocols that cover: what risks the interviewee sees in their area, what has changed since the last assessment, what controls they rely on and whether those controls actually work, what they would worry about if they were responsible for the whole business unit, and what risks they believe are not being captured by existing processes. That last question — what is being missed — is often the most productive.

Checklists

Checklists drawn from the risk taxonomy, regulatory requirements, and the industry loss database provide a systematic prompt for identification. They prevent the assessor from relying solely on recall — which, as Chapter 6 established, defaults to recent and familiar risks at the expense of systemic or slow-moving ones.

For bottom-up identification, the checklist should be organised by L2 taxonomy category and supplemented with specific prompts drawn from the starting universe built in Phase 1 (Chapter 5 (Setting the Context: External, Internal, and Risk Culture)). The assessor works through each category, asking: does this risk exist in our business unit? If yes, what form does it take? What drives it? What controls it?

The checklist is a prompt, not a constraint. If the assessor identifies a risk not on the checklist, that risk must be captured — and the gap reported to the Risk Identification Lead for potential taxonomy update. The checklist ensures coverage. The template captures the output.

Cause-and-Effect Diagrams

Cause-and-effect diagrams (also called **Ishikawa** or fishbone diagrams) are particularly valuable for operational risk identification. The technique starts with a known or potential risk event and works backwards to identify all contributing causes, organised into categories such as people, process, technology, data, and external factors.

Where the standard template asks for direct and indirect drivers, Ishikawa analysis provides a structured method for discovering them. A business unit that writes “technology failure” as a driver has not completed the analysis. An Ishikawa diagram forces the question: *which* technology? *What kind* of failure? Caused by *what*? Enabled by *what*? The output is a visual map of causal chains that populates the driver fields with genuine analytical content.

FMEA and HAZOP for Process Risks

Failure Mode and Effects Analysis (FMEA) examines each step in a business process, identifies the ways each step can fail, assesses the consequences of each failure mode, and evaluates the controls in place. For process-intensive operations — payments, settlements, trade booking, client onboarding, regulatory reporting — FMEA is the most thorough bottom-up identification technique available.

HAZOP (Hazard and Operability Study), like SWIFT, originated in process safety engineering. Applied to banking operations, HAZOP uses guide words (no, more, less, reverse, other than) applied to process parameters (flow, timing, sequence, composition) to identify deviations from intended operation. Where SWIFT asks “what if?” about strategic risks, HAZOP asks “what if?” about process steps. A HAZOP analysis of a payments process might ask: what if the payment is sent to the wrong counterparty (*other than*)? What if the payment is duplicated (*more*)? What if the authorisation step is skipped (*no*)? What if the payment is processed in the wrong sequence (*reverse*)?

FMEA and HAZOP are resource-intensive. They are not appropriate for every risk in every business unit. But for critical processes — those where failure would produce material financial loss, regulatory breach, or customer harm — they produce a level of granularity that no other technique matches.

The choice of technique depends on the risk type, the available resources, and the complexity of the business unit's operations. The Risk Identification Lead should provide guidance on which techniques are appropriate for which contexts, drawing on the Technique Selection Guide. The critical principle is that some structured technique must be used. A business unit that completes the template without applying any identification technique has not performed identification. It has performed data entry.

The Ten Specialist Sub-Processes

Certain risk types cannot be adequately identified through a general-purpose template process. They require dedicated expertise, specialised data sources, regulatory-specific methodologies, and practitioners who understand the technical domain. These are the **specialist risk identification sub-processes** — ten in total — that operate alongside the standardised template process and feed their outputs into the bottom-up track.

Each specialist sub-process is owned by the relevant specialist function. Each must use the institution's common risk taxonomy, apply the same four-dimensional scoring methodology, submit identified risks to the central inventory via the standardised template, and participate in the reconciliation process described in Chapter 8 (Reconciliation and the Enterprise Portfolio View). These integration requirements are non-negotiable. A specialist sub-process that operates in isolation — with its own taxonomy, its own scales, its own reporting format — creates precisely the silos that the enterprise methodology is designed to eliminate.

1. RCSA (Risk and Control Self-Assessment)

RCSA is the most widely deployed operational risk identification tool in banking, and the one most frequently reduced to compliance theatre. A well-designed RCSA is a collaborative process where business units and risk functions jointly identify operational risks and assess the effectiveness of controls. A poorly designed RCSA is a spreadsheet that gets rolled forward every quarter.

The methodology requires RCSA to incorporate the full **BIS Principles for the Sound Management of Operational Risk (PSMOR)** toolbox¹: the self-assessment itself, key risk indicators (KRIs), external loss data from industry databases, business process mapping, and event management. Internal and external loss data must be used as inputs to identification — not just as validation after the fact. Event management provides the critical feedback loop: when an operational risk materialises, that event must feed back into the RCSA to update the identification and control assessment.

The difference between effective and ineffective RCSA is whether the business unit treats it as an analytical exercise or a filing exercise. Effective RCSA involves structured workshops within the business unit, uses Ishikawa diagrams to map control weaknesses, challenges existing control effectiveness ratings against actual incident data, and produces genuinely new risk identifications each cycle. Ineffective RCSA involves one person updating a spreadsheet.

2. TSRA (Threat Scenario-led Risk Assessment)

TSRA constructs hypothetical adverse event scenarios to evaluate the institution's preparedness and resilience. Unlike the scenario analysis described in Chapter 6 — which operates at a strategic level in the top-down workshops — TSRA develops detailed scenarios at the operational level: a major counterparty default coinciding with a systems outage, a cyber attack during a regulatory reporting deadline, a pandemic affecting a critical outsourcing location.

TSRA is owned by the Risk Identification Lead or CRO function. Its value lies in identifying risks that exist only in combination — risks that no individual business unit would identify because each component sits in a different function. TSRA scenarios are designed to stress the interactions between risk types, making it a natural complement to the risk interaction analysis described in Chapter 10 (Risk Interaction: Bow-Ties, Matrices, and Concentration).

3. RSR (Reputational and Sustainability Risk Assessment)

RSR identifies risks affecting public perception and long-term viability, including ESG factors. Reputational risk is unusual in the taxonomy because it rarely materialises independently — it is almost always a consequence of another risk crystallising. A conduct failure becomes a reputational crisis. An environmental incident becomes a reputational crisis. A cyber breach becomes a reputational crisis.

RSR must therefore look across all other risk categories and ask: if this risk materialises, what is the reputational consequence? This cross-cutting nature means RSR cannot be conducted in isolation. It requires inputs from conduct risk, compliance, operations, and the external context assessment (Chapter 5). The ESG dimension — environmental, social, and governance factors — maps to the climate risk transmission channels established in Chapter 5 and the sustainability taxonomy categories from Chapter 4.

4. Conduct Risk Assessment

Conduct risk identifies risks arising from behaviours by the firm or its staff that result in poor outcomes for customers, market integrity, or competition. The FCA's **five Conduct Questions** framework provides the regulatory basis:²

1. What could cause harm to customers or markets?
2. What could undermine market integrity?
3. What could affect competition?
4. How could the firm's culture contribute to poor outcomes?
5. What are the inherent conduct risks in the firm's business model?

Conduct risk identification must be **business-led at desk-by-desk granularity** — not a central function exercise applied generically. It should be **reverse-engineered from potential customer harm and market integrity impact**, not limited to regulatory breach. The PPI scandal described in Chapter 4 demonstrated what happens when conduct risk is absent from the taxonomy entirely. Even where the taxonomy includes conduct risk, identification frequently fails because it looks backward at past regulatory actions rather than forward at embedded business model risks.

Incentive structures are the single most important input to conduct risk identification. Where compensation rewards volume without reference to customer outcomes, conduct risk is structural, not behavioural. The identification must examine: what are salespeople incentivised to do? What happens when those incentives conflict with customer interest? What controls exist to detect the conflict?

5. Model Risk Assessment

Model risk arises from inaccuracy, misuse, or inappropriate application of financial models. In an industry that runs on models — for pricing, risk measurement, capital calculation, stress testing, and regulatory reporting — model risk identification must cover the full lifecycle: development assumptions, data inputs, calibration methodology, validation process, and the governance around model use and override.

The model risk assessment must identify not just the risk that a model is wrong, but the risk that a model is *used for the wrong purpose*. The JPMorgan London Whale case described in Chapter 4 was partly a model risk failure: the CIO's VaR model was changed to reduce reported risk, and the change was not independently challenged. Model risk identification must examine: which models are used for decisions that could produce material loss? How are model limitations documented and communicated? Where are models being used outside their validated scope?

6. ICT and Cyber Security Assessment

ICT risk identification must be consistent with **EBA GL/2019/04** (Guidelines on ICT and Security Risk Management).³ The institution must classify and maintain an inventory of all information assets and ICT systems, conduct annual ICT risk assessments, and identify risks arising from ICT change management, data integrity, and business continuity of ICT services.

Cyber risk identification has a characteristic that distinguishes it from most other risk types: the threat landscape changes continuously and adversarially. Unlike credit risk, where the drivers are economic conditions and borrower behaviour, cyber risk is driven by intelligent adversaries who actively seek and exploit vulnerabilities. This means the identification process must incorporate external threat intelligence — not just internal vulnerability assessment — and must be refreshed more frequently than the annual cycle. ICT risk identification should connect to the event-driven update triggers defined in Chapter 3: a material cyber incident anywhere in the industry should prompt reassessment of the institution's own exposure.

7. AML/CFT and Financial Crimes Assessment

Anti-money laundering and counter-terrorist financing risk identification follows **EU AM-LD6** and local AML/CFT regulations.⁴ The institution must conduct entity-level ML/TF risk identification and assessment aligned with national risk assessments. Critically, AML/CFT risks must be identified for all new products, services, and delivery channels **before launch** — not retrospectively.

Financial crimes assessment covers money laundering, terrorist financing, sanctions evasion, fraud, bribery, and corruption. The identification challenge is that these risks are deliberately concealed by the actors creating them. Unlike operational risk, where failures are usually visible after the fact, financial crime risks may persist for years before detection — as Deutsche Bank’s mirror-trading scheme, described in Chapter 4, demonstrated. The specialist function must therefore look for structural indicators: product features that facilitate anonymity, delivery channels that reduce oversight, customer segments with inherently higher ML/TF risk, and geographic exposures to high-risk jurisdictions.

8. Third-Party and Outsourcing Risk Assessment

Per the **EBA Outsourcing Guidelines**,⁵ the institution must maintain a comprehensive outsourcing register, identify risks associated with each material outsourcing arrangement — including sub-outsourcing and fourth-party dependencies — and assess concentration risk where multiple critical services depend on the same provider or geographic location.

Third-party risk identification has become materially more important as banks outsource critical functions to a decreasing number of cloud providers, technology vendors, and service companies. The concentration dimension is critical: if five critical functions depend on a single cloud provider, the failure of that provider is not five independent operational risks — it is one systemic risk. The identification must look through the bilateral relationship to the aggregate dependency.

9. Traded Risk Assessment

Traded risk covers market, credit, and operational risks arising from trading activities. This sub-process is owned by the Market Risk or Trading Risk function and must identify risks specific to the trading book: position concentration, liquidity risk in held instru-

ments, counterparty credit risk in OTC derivatives, basis risk between hedges and underlying positions, and the operational risks embedded in trade execution, booking, and settlement.

Traded risk identification interfaces directly with the front office and must incorporate trader input — not just risk function analysis. The traders themselves are often the earliest identifiers of market dislocations, liquidity deterioration, or counterparty stress. The identification process must create channels for that information to flow into the formal inventory.

10. Treasury Risk Assessment

Treasury risk covers risks to liquidity, funding, and capital structure. The Treasury or Asset-Liability Management (ALM) function owns this sub-process and must identify: funding concentration risk (the Northern Rock failure described in Chapter 5), interest rate risk in the banking book, intraday liquidity risk, contingent liquidity demands from off-balance-sheet commitments, and currency mismatch risk.

Treasury risk identification has a structural challenge: many of the risks it must identify are created by the commercial activities of other business units. Loan growth creates funding requirements. Derivative positions create contingent collateral demands. International expansion creates currency mismatches. The treasury sub-process must therefore receive inputs from all business units — not just identify risks within its own operations — and must assess aggregate institutional exposure, not just treasury-specific risk.

When Bottom-Up Identification Fails: AIB and John Rusnak

In February 2002, Allied Irish Banks disclosed that a foreign exchange trader named John Rusnak at its US subsidiary, Allfirst Financial, had concealed \$691 million in losses through fictitious option trades.⁶ The losses had accumulated over several years. Rusnak was sentenced to seven and a half years in prison.⁷ AIB absorbed the full loss. Allfirst was subsequently sold to M&T Bank.⁸

The Rusnak case is a textbook failure of bottom-up identification across geographic boundaries. Allfirst operated with significant autonomy from AIB's Dublin headquarters. The subsidiary had its own risk management function, its own reporting lines, and its own operational processes. AIB's central risk function relied on local management assurances that risks were being properly identified and controlled.

Rusnak exploited weak back-office confirmation processes for FX options. He created fictitious option trades that offset his real positions, making his portfolio appear hedged when it was massively exposed. The back-office processes that should have independently confirmed these trades with counterparties were inadequate — confirmations were not obtained, or were fabricated by Rusnak himself.

Every element of this failure maps to a bottom-up identification deficiency:

- **No standardised template:** Allfirst's risk reporting to AIB used locally developed formats that did not capture the granular driver and control information a standardised template would require. The parent bank received summary data, not analytical detail.
- **No structured technique applied:** No one had performed an FMEA on the FX options booking and confirmation process. A systematic examination of each process step would have identified that the confirmation process depended on a single point of failure — the same trader whose positions were being confirmed.
- **No independent control effectiveness assessment:** The current controls field — had it existed in a standardised template — would have required someone to assess whether the back-office confirmation process actually worked. In Allfirst, the answer was that it did not, but no one was required to ask the question.
- **Geographic separation:** The specialist traded risk assessment for the FX desk was conducted locally, with no integration into AIB's central risk identification process. The Risk Identification Lead — had one existed with the mandate described in Chapter 3 — would have been responsible for ensuring that the subsidiary's risk assessment met the same standards as every other business unit.

What was missing: A standardised bottom-up process requiring Allfirst to complete the same risk assessment template as every other AIB business unit, using the same taxonomy, the same scoring methodology, and the same control effectiveness assessment — with the outputs submitted to a central function that could identify gaps, challenge local assurances, and reconcile subsidiary risks against the enterprise view.

The Same Failure, Repeated: UBS and Kweku Adoboli

Nine years after Rusnak, in September 2011, UBS disclosed that a trader named Kweku Adoboli on its London Delta One desk had concealed \$2.3 billion in losses from unauthorised ETF trading by creating fictitious hedging positions.⁹ Adoboli exploited gaps between the bank's trading and settlement systems — he booked fictitious trades that offset his real positions, and the reconciliation processes that should have detected the discrepancies were insufficiently automated. He was sentenced to seven years.¹⁰

The UBS case is remarkable because it occurred *after* the global financial crisis, *after* UBS had suffered \$37 billion in losses from structured credit exposures,¹¹ *after* the Swiss National Bank had created a special purpose vehicle to absorb \$39 billion in toxic assets¹², and *after* UBS had undertaken a comprehensive risk management overhaul. Despite all of this, the bottom-up identification process still contained the same fundamental gap: fictitious trades booked against genuine positions were not independently verified.

The RCSA for the Delta One desk — which UBS undoubtedly conducted, given the regulatory requirements — had not identified that the control framework assumed trade authenticity rather than testing it. The control assessment said reconciliation processes existed. It did not ask whether those processes could detect deliberately fabricated trades. The FSA fined UBS £29.7 million for systems and controls failures.¹³ The incident accelerated UBS's strategic exit from large parts of investment banking.

The Adoboli case, following Kerviel at Societe Generale in 2008¹⁴ and Rusnak at AIB in 2002, demonstrates that rogue trading is not an unpredictable individual event. It is a recurring operational risk with identifiable structural preconditions: a trader who under-

stands the control chain, inadequate independent verification of trade authenticity, and an RCSA process that assesses controls against design specification rather than against determined circumvention.

What was missing: An RCSA process that incorporated external loss data — Kerviel and Rusnak were both available as inputs — to challenge the assumption that existing controls were adequate. A traded risk assessment that examined not just whether controls existed, but whether they could withstand deliberate exploitation. And a bottom-up template that required the business unit to assess control effectiveness against specific failure scenarios drawn from the industry loss database.

The Integration Challenge

Ten specialist sub-processes, each owned by a different function, each with its own expertise and regulatory requirements, each producing risk identification outputs that must feed into a single unified inventory. This is the integration challenge, and it is where the Risk Identification Lead earns their title.

The failure mode is predictable: each specialist function develops its own methodology, its own templates, its own scoring scales, and its own reporting cadence. The conduct risk assessment uses a five-point likelihood scale while the ICT assessment uses a three-point scale. The RCSA classifies risks using Basel operational risk categories while the AML/CFT assessment uses FATF risk categories. The third-party risk assessment tracks 200 outsourcing arrangements while the treasury risk assessment has no visibility into which of those arrangements affect its liquidity contingency plans.

The result is not a risk inventory. It is ten separate inventories with no common language and no means of aggregation.

The four integration requirements from the methodology are designed to prevent this:

- 1. Common taxonomy.** Every specialist sub-process must classify its risks using the institution's L1/L2/L3 taxonomy. The RCSA does not get to use Basel categories as a substitute. The conduct risk assessment does not get to use FCA categories as a

substitute. The institution's taxonomy — mapped to regulatory categories through the regulatory mapping table described in Chapter 4 — is the single classification standard.

- 2. Same four-dimensional scoring methodology.** Every risk, regardless of which sub-process identified it, must be scored using the same impact dimensions (financial, reputational, regulatory, customer, operational) and the same likelihood scales. Detailed scoring methodology is described in Chapter 9. The principle here is uniformity: a "high impact" conduct risk and a "high impact" ICT risk must mean the same thing.
- 3. Standardised template submission.** Every specialist sub-process must submit its identified risks via the eleven-field template. This forces each function to express its specialist analysis in common terms — taxonomy classification, drivers, controls, data quality rating — that the Risk Identification Lead can aggregate and compare.
- 4. Participation in reconciliation.** Every specialist function must participate in the reconciliation process described in Chapter 8, where top-down and bottom-up outputs are compared, gaps identified, and coverage confirmed. This is where the Risk Identification Lead challenges: has the RCSA captured the same operational risks that the top-down workshop identified? Has the traded risk assessment captured the risks that the treasury assessment's liquidity scenarios depend on? Are there gaps between specialist functions where risks could fall?

The Risk Identification Lead is responsible for ensuring that these four requirements are met. This is not a coordination role. It is a quality assurance role with challenge authority. When a specialist function submits risks classified using its own taxonomy rather than the institution's, the Risk Identification Lead sends it back. When an RCSA produces the same risk list as last year with no reference to recent loss events, the Risk Identification Lead challenges it. When the third-party risk assessment does not include an assessment of concentration across providers, the Risk Identification Lead escalates the gap.

From Compliance Theatre to Genuine Identification

The difference between a bottom-up process that works and one that produces compliance theatre comes down to three things.

First, the techniques must be applied. A template without a technique is a data entry exercise. Every business unit must use at least one structured identification technique — checklists for comprehensive coverage, Ishikawa for causal analysis, FMEA for critical processes, structured interviews for capturing distributed knowledge. The Risk Identification Lead should specify minimum technique requirements based on the risk profile and complexity of each business unit.

Second, the inputs must be current. The starting universe built in Phase 1 (Chapter 5), the PESTLE assessment, the internal context analysis, and the industry loss database must all feed into the bottom-up process — not just the top-down workshops. Business units completing their self-assessments must receive the same briefing materials that workshop participants receive. If a peer institution has suffered a material loss in the past quarter, every relevant business unit should be prompted to assess whether the same risk exists in their operations.

Third, the output must be challenged. The Risk Identification Lead reviews every submission against three tests: Has the risk list changed since the last cycle? If not, why not — has genuinely nothing changed, or has the assessor simply rolled forward? Are the driver fields populated with specific, analytical content, or with generic descriptions? Are the control effectiveness ratings supported by evidence, or by assumption?

The methodology requires that every business unit identify at least three *new* risks or *materially changed* risks each annual cycle. Not because three is a magic number, but because the external environment, the business model, the regulatory landscape, and the technology infrastructure all change continuously. An honest assessment will always find something new. If it does not, the assessor has not looked.

The Output

A well-functioning bottom-up process produces a **comprehensive risk register of 100 to 500 risks** across all business units, incorporating outputs from all ten specialist sub-processes. The register is organised by the institution's risk taxonomy, scored using the common methodology, documented in standardised templates, and accompanied by data quality ratings that create transparency about the evidential basis for each assessment.

This register is not the final risk inventory. It is one of two inputs — alongside the top-down outputs described in Chapter 6 — into the reconciliation process that produces the enterprise portfolio view. The bottom-up register provides granular coverage. The top-down outputs provide strategic coverage. Neither alone is sufficient.

The register must be designed for reconciliation. Every risk must carry its taxonomy classification, enabling comparison with the top-down principal risk list. Every risk must identify its business unit of origin, enabling the Risk Identification Lead to check coverage across all units. Every specialist sub-process output must be traceable to its source function, enabling the reconciliation to verify that all ten sub-processes have contributed.

The Bridge to Reconciliation

The bottom-up track has now produced its output: hundreds of granular risks, identified through structured techniques, captured in standardised templates, submitted by business units and specialist functions across the institution. The top-down track, described in Chapter 6, has produced its output: a principal risk list of fifteen to thirty strategic risks, an emerging risk register, an assumption register, a disagreement log, and a taxonomy gap list.

These two outputs must now be brought together. Chapter 8 describes the reconciliation process — the iterative comparison between top-down and bottom-up that identifies gaps, resolves conflicts, escalates bottom-up risks of strategic significance, assigns top-down risks to business unit owners, and ultimately produces the enterprise portfolio view that gives the Board and CRO a single, integrated picture of the institution's risk landscape. It is in reconciliation that most banks fail. They perform either top-down or bottom-up, but not both — and they almost never iterate between them. The methodology requires both tracks precisely because the reconciliation between them is where the most valuable risk identification occurs.

1. Basel Committee on Banking Supervision, *Principles for the Sound Management of Operational Risk (PSMOR)*, BCBS 195, June 2011 (revised). The principles establish expectations for identification tools including RCSA, KRIs, external loss data, business process mapping, and event management.

2. Financial Conduct Authority, *Five Conduct Questions*, FCA Approach to Supervision, 2015. The FCA introduced the five conduct questions as part of its supervisory framework to assess firms' conduct risk.
3. European Banking Authority, *Guidelines on ICT and Security Risk Management*, EBA/GL/2019/04, 28 November 2019. The guidelines set expectations for institutions' ICT risk management frameworks, including asset classification, risk assessment, and business continuity.
4. Directive (EU) 2018/1673 of the European Parliament and of the Council of 23 October 2018 on combating money laundering by criminal law (6th Anti-Money Laundering Directive, AMLD6). Transposition deadline: 3 December 2020.
5. European Banking Authority, *Guidelines on Outsourcing Arrangements*, EBA/GL/2019/02, 25 February 2019. The guidelines require institutions to maintain a register of all outsourcing arrangements and to identify and manage associated risks, including sub-outsourcing and concentration risk.
6. Allied Irish Banks, public disclosure and SEC filings, February 2002. See also: Promontory Financial Group and Wachtell, Lipton, Rosen & Katz, *Report to the Boards of Directors of Allied Irish Banks, p.l.c., Allfirst Financial Inc., and Allfirst Bank Concerning Currency Trading Losses*, 12 March 2002 (the Ludwig Report).
7. United States District Court, District of Maryland, *United States v. John M. Rusnak*, Case No. 1:02-cr-00049, sentencing January 2003. Rusnak pleaded guilty to one count of bank fraud and was sentenced to 7.5 years in federal prison.
8. AIB completed the sale of Allfirst to M&T Bank Corporation in April 2003 in exchange for a 22.5% stake in M&T. See: AIB Group Annual Report, 2003.
9. UBS AG press release, 15 September 2011, disclosing the unauthorised trading loss. The initial estimate of \$2 billion was subsequently revised to approximately \$2.3 billion. See also: FSA Final Notice to UBS AG, 26 November 2012.
10. Southwark Crown Court, *R v Kweku Adoboli*, sentencing 20 November 2012. Adoboli was convicted of two counts of fraud by abuse of position and sentenced to seven years' imprisonment.
11. UBS AG, *Shareholder Report on UBS's Write-Downs*, 18 April 2008. UBS disclosed cumulative write-downs of approximately \$37 billion related to US residential mortgage-backed securities and CDO positions. See also: Swiss Federal Banking Commission (SFBC), *Subprime Crisis: SFBC Investigation into the Causes of the Write-downs of UBS AG*, 30 September 2008.
12. Swiss National Bank established the StabFund SPV in October 2008 to acquire up to \$39.1 billion in illiquid assets from UBS. See: Swiss National Bank press release, 16 October 2008; UBS AG press release, 16 October 2008.
13. Financial Services Authority, *Final Notice to UBS AG*, 26 November 2012. The FSA imposed a financial penalty of £29.7 million on UBS AG for systems and controls failures relating to the unauthorised trading loss caused by Kweku Adoboli.
14. Societe Generale disclosed a EUR 4.9 billion loss in January 2008 from unauthorised trading by Jerome Kerviel on European equity index futures. See: Societe Generale, *General Inspection Department Mission Green: Summary Report*, 20 May 2008; Cour d'appel de Versailles, *Ministere Public c. Kerviel*, 2012.

Reconciliation and the Enterprise Portfolio View

The Risk That Nobody Owns

Consider a G-SIB completing its first dual-track identification cycle. On one side, the output of the top-down SWIFT workshop — just over twenty risks identified by the CRO, business unit heads, and senior risk officers through the structured methodology described in Chapter 6 (Top-Down Identification: Workshops, SWIFT, and Delphi). On the other, the consolidated bottom-up submissions from every business unit — nearly two hundred risks documented using the standardised template and specialist sub-processes described in Chapter 7 (Bottom-Up Identification: Templates, RCSA, and Specialist Processes).

The task is to reconcile them.

The top-down list includes a risk labelled “cross-counterparty concentration in prime brokerage services.” The SWIFT workshop surfaced it through the guide word prompt “Where are the gaps between what we see at the business unit level and what we would see at the enterprise level?” A participant observed that several business lines appeared to hold exposure to the same cluster of counterparties through different product types — equity financing, derivatives, and securities lending — but that no single view of the aggregate existed.

Cross-referencing this against the bottom-up submissions reveals: not a single business unit has included it.

This is not because the bottom-up process has failed. Each business unit has performed genuine identification. The equity financing desk reported counterparty credit risk against its positions. The derivatives desk reported market risk and counterparty risk within its own book. The securities lending desk reported operational risk in its collateral

management. Each submission is complete, well-documented, and accurate within its scope. But the risk that the workshop identified — the aggregate, cross-product exposure to a common set of counterparties — exists between business units, not within them.

The question reconciliation is designed to answer: who owns this risk?

No one. No business unit's risk assessment captures it because no business unit's remit encompasses it. It is not a credit risk, a market risk, or an operational risk as any individual desk would define those terms. It is an enterprise risk — visible only when the two lists are placed side by side and someone looks for what neither track captured alone.

That moment — when the gap between top-down and bottom-up reveals something that neither track could have found independently — is the purpose of reconciliation. Without it, the methodology produces two lists. With it, it produces an enterprise view. And the enterprise view is where the most dangerous risks become visible.

In March 2021, the collapse of Archegos Capital Management cost Credit Suisse \$5.5 billion.¹ The risk that materialised was precisely the kind of cross-counterparty, cross-product concentration that reconciliation is designed to identify — total return swap exposure aggregated across multiple prime brokers, invisible to any single business unit's assessment.²

This chapter describes the two final steps of Phase 2 — the reconciliation process and the enterprise portfolio view — that transform two separate identification tracks into a single, consolidated picture of institutional risk.

Where Most Banks Fail

Chapter 6 established that the top-down track identifies approximately twenty risks of strategic significance. Chapter 7 established that the bottom-up track identifies one hundred to five hundred risks of operational and specialist detail. Both tracks are necessary. Neither alone is sufficient.

Most banks fail at what comes next.

What passes for reconciliation at the majority of institutions is compilation. Someone in the risk function takes both lists, removes obvious duplicates based on risk name, combines them into a single spreadsheet, and sends the result to the CRO. The top-down risks sit at the top of the page. The bottom-up risks fill the remainder. The document is called a “reconciled risk inventory.” It is nothing of the sort.

This pattern is pervasive. The two tracks exist. The outputs are produced. But the analytical work between receiving those outputs and producing an enterprise view — the work that transforms two lists into genuine institutional intelligence — is either absent or superficial. At a typical institution, the reconciliation documentation from a prior cycle consists of a merged spreadsheet with colour-coding to indicate source: blue for top-down, green for bottom-up, yellow for both. No gap analysis. No escalation decisions documented. No evidence of challenge sessions. No enterprise portfolio view. The institution has two identification tracks. It does not have reconciliation.

Compilation tells you what two tracks have produced. Reconciliation tells you what neither track has found. The difference is the analytical work between receiving the two outputs and producing the enterprise view — the gap analysis, the escalation decisions, the ownership assignments, the challenge sessions, and the iterative refinement that converts two lists into a single, coherent picture of institutional risk.

The reconciliation process asks five categories of question that compilation never poses. First, what did the top-down workshop identify that no business unit captured — and why? Second, what did the bottom-up process surface that senior management did not see from altitude — and does it have enterprise significance? Third, which risks have no owner, and who should own them? Fourth, where do the two tracks disagree, and what does the evidence show? Fifth, when you aggregate all identified risks across business units, what picture emerges that no individual unit could have seen?

The methodology requires both tracks precisely because the reconciliation between them is where the most valuable risk identification occurs. This is not a philosophical point. It is an operational one. The risks that live between business units — the aggregate concentrations, the correlated exposures, the systemic dependencies — are the risks that have caused the largest losses in banking history. And they are structurally invisible to any process that does not reconcile top-down and bottom-up identification into an enterprise portfolio view.

The Five Steps of Reconciliation

Step 1: Gap Analysis

The Risk Identification Lead performs a systematic comparison between the top-down and bottom-up outputs using the common taxonomy as the mapping key.

Top-down gaps are risks identified in the SWIFT workshop that no business unit has captured in its bottom-up submission. These fall into two categories. The first is risks that exist within a business unit's scope but have not been identified — a genuine bottom-up identification failure. The second is risks that are genuinely enterprise-level and do not reside within any single business unit's remit. Cross-counterparty concentration, systemic infrastructure dependency, and macroeconomic regime change are examples. Both categories require investigation, not assumption. The Risk Identification Lead cannot simply assume that a top-down gap represents a bottom-up failure without determining whether the risk has a natural business unit home.

Bottom-up gaps are risks identified in business unit submissions that the top-down workshop did not surface. Again, two categories. Senior management may be unaware of operational risks visible only at process level — a genuine top-down blind spot that the dual-track methodology is specifically designed to catch. Alternatively, the risk may be genuinely local with no strategic significance. The Risk Identification Lead assesses which bottom-up risks are candidates for escalation.

The mapping uses the common taxonomy as its primary key. Every risk in both outputs carries an L1/L2 classification — this was established as a non-negotiable integration requirement in Chapter 7. The Risk Identification Lead maps each top-down risk to its taxonomy node and searches the bottom-up submissions for risks at the same node. Where a match exists, the two entries are linked. Where no match exists, the gap is recorded and categorised.

At Institution A, the first reconciliation cycle produced a dozen top-down risks with no bottom-up match and nearly fifty bottom-up risks with no top-down equivalent. Each required investigation. Some top-down gaps reflected genuinely enterprise-level risks with no natural BU owner — cross-border regulatory divergence, for instance, which no single business unit experienced as a risk but which the SWIFT workshop had identified as a strategic concern. Some reflected bottom-up identification failures that required re-

mediation — risks that clearly fell within a business unit's scope but had not been captured, often because the prior year's submission had been rolled forward without genuine re-identification. Some bottom-up gaps contained risks of strategic significance that the workshop had not considered — process-level vulnerabilities that, when aggregated, constituted material operational exposure.

This is the normal output of a functioning reconciliation process. A clean reconciliation with no gaps should be treated with suspicion, not satisfaction. The gaps are the signal, not the noise.

Step 2: Escalation

Bottom-up risks of strategic significance must be elevated to the principal risk list.

A business unit submits a risk that, when examined in the context of the full institution, represents a material enterprise exposure. The operational risk of a single cloud provider outage, for example, might appear manageable when reported by one business unit. But when the Risk Identification Lead discovers that five critical functions across three business units depend on the same provider, the risk transforms from a local operational concern into a systemic enterprise dependency. Chapter 7 established this principle — the failure of that provider is not five independent operational risks but one systemic risk.

The escalation threshold is straightforward: does this risk, individually or when aggregated with related risks across other business units, potentially exceed the institution's materiality threshold or breach the risk appetite boundaries defined in Phase 1? The Risk Identification Lead makes the initial assessment. The CRO approves the escalation. The risk is added to the principal risk list with full documentation of the escalation rationale.

Step 3: Assignment

Top-down risks with no business unit owner must be assigned.

Senior management identified the risk in the workshop. The gap analysis confirmed that no business unit captured it. Someone must own it. Risk owners are named individuals — not committees, not functions. This was established in Chapter 3 (Governance: Who Owns What) and applies with full force here. A risk without an owner is a risk without accountability, and a risk without accountability is a risk that will not be managed.

Assignment follows a simple hierarchy. If the risk has a natural business unit home, it is assigned to the head of that unit. If it spans multiple units, it is assigned to the unit with the largest share of the exposure, with documented obligations for other affected units to provide input. If it is genuinely enterprise-level — systemic, macroeconomic, reputational contagion — it is owned at CRO or Group Risk level, but with a requirement for business-unit-level input to inform assessment.

The assignment step sounds administrative. It is not. It is one of the most politically sensitive moments in the entire methodology. No business unit head wants to own a risk they did not identify. The Risk Identification Lead must have the mandate — established in Chapter 3's governance framework — to assign risks based on analytical judgement, not organisational willingness.

Step 4: Challenge

The Risk Identification Lead facilitates structured challenge sessions between senior management and business unit risk teams. These are not consensus-building exercises. They are structured confrontations designed to test the completeness and accuracy of both tracks.

Challenge works both ways. Business unit teams challenge whether top-down risks are real — does the evidence support the workshop's assessment, or is the risk theoretical? Senior management challenges whether bottom-up submissions are complete — are the driver fields populated, are the control ratings evidenced, has the risk list genuinely changed since last year?

The disagreement log produced in the top-down workshop informs this step directly. Where workshop participants held materially different views on a risk, the challenge session must resolve the disagreement through evidence — bottom-up data, loss history, control effectiveness assessment, external benchmarking — not through social dynamics. The assumption register from the workshop is similarly tested: assumptions challenged in the SWIFT process are now examined against the granular evidence the bottom-up track has produced. If a risk was disputed in the workshop, the bottom-up evidence either supports or refutes it.

Effective challenge sessions are the most productive — and the most uncomfortable — part of the entire process. A business unit head who dismissed a top-down risk as theoretical may discover that his own unit's bottom-up submission contains the evidence for it — the driver fields in the template, when examined, point directly to the mechanism

the workshop identified. A senior risk officer who championed a risk in the workshop may learn that the bottom-up data shows effective controls already in place, with documented evidence of preventive and detective measures that the top-down discussion had not considered.

Consider an investment banking division challenging the workshop's identification of model risk in a specific product class. The bottom-up submission rates the relevant controls as effective. The challenge session requires the division to present the evidence for that rating. When examined, the "effective" rating turns out to be based on a model validation performed two years earlier, against data from a different market regime. The challenge session produces a reassessment, a revised data quality rating, and a commitment to revalidate. Neither the top-down workshop nor the bottom-up template would have generated that outcome alone. The structured confrontation between them does.

Challenge sessions are where analysis displaces opinion, and where the reconciled inventory earns its credibility.

Step 5: Iteration

The reconciliation cycle repeats until both tracks are aligned and the CRO is satisfied that coverage is comprehensive. A typical annual cycle requires two to three iterations. Quarterly re-identification cycles, being more focused on changes and emerging risks, typically require one to two.

Each iteration narrows the gap between tracks. The first iteration identifies the major gaps and disagreements. The second resolves the majority through evidence. The third, where needed, addresses residual issues that require CRO judgement.

The process is documented at every stage. Each gap identified, each escalation decision, each ownership assignment, each challenge session outcome, each iteration's movement — all recorded as part of the audit trail. When a regulator asks how the institution arrived at its risk inventory, the reconciliation documentation provides the answer. It shows not just what was identified but how the identification process converged on its conclusions.

The CRO's sign-off on the reconciled inventory represents the governance conclusion of Phase 2. The institution now has a single, integrated risk inventory that reflects both the strategic perspective of senior management and the operational detail of the people closest to the business.

Citigroup and the Forty-Five Billion Dollar Gap

Before the 2008 financial crisis, Citigroup was one of the largest and most diversified financial institutions in the world. Its business divisions — investment banking, consumer banking, transaction services, and wealth management — each maintained their own risk reporting, their own risk assessments, and their own view of the exposures they held.³ Each operated within its own risk limits. Each produced its own bottom-up risk submission. And each was, by its own internal standards, within acceptable parameters.

The investment banking division had created a series of **Structured Investment Vehicles** — off-balance-sheet entities that held portfolios of subprime-linked collateralised debt obligations. The SIVs funded themselves by issuing short-term commercial paper to investors. Citigroup had written liquidity puts to the SIVs — contractual commitments to provide funding if the commercial paper market became unavailable.⁴ These puts were classified as contingent liabilities.

Here is where the enterprise portfolio view was absent, and where its absence proved catastrophic.

The investment banking division reported its on-balance-sheet CDO warehousing positions as part of its structured credit risk assessment. Treasury reported funding risk and contingent liquidity commitments. The consumer banking division reported mortgage origination quality and delinquency trends. Each risk assessment was internally coherent. None was wrong in isolation.

But no enterprise view existed that connected these exposures into a single picture. The CDO warehouse positions on the balance sheet. The liquidity puts to the SIVs off the balance sheet. The mortgage origination pipeline that was feeding the same CDO structures. The commercial paper funding model that would collapse if investor confidence faltered. These were not four independent risks. They were four manifestations of a

single underlying exposure: Citigroup's total position in US residential mortgage credit quality, viewed across on-balance-sheet, off-balance-sheet, direct, and contingent channels.

The off-balance-sheet SIV exposures were excluded from the bank's risk aggregation framework because the liquidity puts were considered remote contingencies. The probability of the entire commercial paper market closing simultaneously was assessed as negligible. This was an assumption that a reconciliation process would have tested — and that an enterprise portfolio view would have challenged by asking: what happens to our aggregate exposure if this assumption proves wrong?

When the subprime crisis struck, all the SIV liquidity puts triggered simultaneously. Citigroup had to bring more than fifty billion dollars in SIV assets back onto its balance sheet.⁵ Combined with its existing CDO warehouse losses and mortgage-related writedowns, the institution required a forty-five billion dollar TARP bailout and a US government equity stake to survive.⁶

The risk was not hidden. It was distributed — spread across business units, product types, and accounting classifications in a way that no single business unit's risk assessment captured the aggregate. Each division had identified its portion of the exposure. None had identified the whole. The enterprise portfolio view that would have aggregated all subprime-linked exposures — on-balance-sheet CDO positions, off-balance-sheet SIV liquidity commitments, mortgage origination pipeline, commercial paper funding dependency — into a single consolidated picture did not exist.

What was missing was not identification within business units. It was reconciliation between them. A structured reconciliation process would have forced the question: what is our total exposure — direct, indirect, contingent, on-balance-sheet and off — to a decline in US residential mortgage credit quality? That question, applied through the five reconciliation steps and the enterprise portfolio view, would have produced a number that no individual business unit's assessment could have generated. And that number would have exceeded Citigroup's risk appetite by a margin that demanded immediate Board attention.

The Enterprise Portfolio View

After reconciliation produces the unified risk inventory, the Risk Identification Lead constructs the **enterprise portfolio view** — the aggregated picture of institutional risk that prevents precisely the failure Citigroup experienced.

The enterprise portfolio view is not a summary of individual business unit assessments. It is an analytical product that treats the institution's entire risk inventory as an interconnected portfolio and examines it for properties that are invisible at the business unit level. It rests on the aggregation rules defined in Phase 1 — Chapter 5 (Setting the Context: External, Internal, and Risk Culture) established that these rules must be set before identification begins, so that the enterprise portfolio view is built on a consistent methodology rather than improvised after the fact.

The portfolio view assesses four dimensions.

Common Exposures

Where multiple business units share exposure to the same underlying driver, the enterprise portfolio view identifies and quantifies the aggregate.

Common drivers include interest rate movements affecting multiple business lines simultaneously, counterparty exposures spanning credit, market, and operational risk categories across different desks, shared technology platforms creating correlated operational risk across functions, geographic concentrations where multiple business lines depend on the same regional economy, and regulatory changes that affect several business activities at once.

The analysis operates at two levels. First, taxonomy-based aggregation: all risks sharing the same L2 or L3 classification are aggregated across business units to reveal total institutional exposure to that risk type. Second, cross-taxonomy analysis: risks classified under different taxonomy nodes but sharing the same underlying driver are identified and connected. The example from this chapter's opening illustrates the second type — counterparty credit risk in one business unit, market risk from the same counterparty's instruments in another, and operational risk from the same entity as a technology vendor in a third. Three different L1 categories, one underlying driver.

Simultaneous Crystallisation

The enterprise portfolio view must assess how a single event could trigger multiple risks across the institution at the same time. This is where **concentration blindness** — the failure mode identified in Chapter 1 (Why Banks Fail at Risk Identification) — becomes visible.

The analysis is scenario-based. The Risk Identification Lead constructs a set of plausible stress events — drawn from the PESTLE assessment (Chapter 5), the scenario analysis performed in the SWIFT workshop (Chapter 6), and the industry loss database — and traces which risks across the reconciled inventory would activate under each scenario. A sudden interest rate shock, for example, might simultaneously trigger market risk losses in the trading book, credit risk deterioration in the lending portfolio, liquidity risk as funding costs increase, and operational risk as systems designed for a low-rate environment encounter processing exceptions.

This assessment is identification, not quantification. The purpose is not to calculate the combined loss — that is the work of the assessment phase in Chapter 9 (Assessment — Scoring, Multi-Dimensional Impact, and Data Quality) and the stress testing programmes it feeds. The purpose is to identify which risks are linked, so that the institution knows which apparently independent entries in its risk inventory would crystallise together. A risk inventory that contains two hundred individually assessed risks without an analysis of which ones would activate simultaneously is not a portfolio view. It is a catalogue.

The simultaneous crystallisation assessment directly addresses the **silos thinking** failure mode identified in Chapter 1. Silo thinking is not a failure of individual business units. It is the structural absence of an enterprise-level analysis that connects what each business unit sees into what the institution faces. The enterprise portfolio view is the methodological response to that structural absence.

Aggregate Position Against Appetite

Each business unit may be within its own risk limits while the institution as a whole exceeds its risk appetite. This is the defining insight of the enterprise portfolio view — the one that Citigroup's absence of such a view failed to deliver.

The risk appetite boundaries established in Phase 1 are applied to the aggregated exposure. Where the portfolio view reveals that total institutional exposure to a risk driver exceeds the appetite boundary, the breach is escalated immediately to the CRO and Board Risk Committee. This escalation is not discretionary. An appetite breach identified through the enterprise portfolio view is a finding of the same significance as any other appetite breach — it simply could not have been detected at the business unit level.

Diversification and Correlation

The portfolio view must assess where genuine diversification exists and where risks are more correlated than individual business unit assessments suggest.

This is the assessment that distinguishes a genuine enterprise view from a simple aggregation. Individual business units naturally assume their risks are independent — they operate in different markets, serve different clients, use different products. The enterprise portfolio view challenges that assumption by examining correlation under stress conditions.

HSH Nordbank provides a precise illustration. Before the 2008 crisis, the German regional bank was the world's largest shipping lender.⁷ It had also built a substantial portfolio of US subprime structured products. The two exposures were treated as diversified — shipping finance and structured credit occupied different taxonomy categories, different desks, different geographies. The risk assessments for each were internally sound.

But both exposures were correlated to global economic activity. The same crisis that impaired CDO valuations also collapsed global trade volumes, which collapsed shipping rates, which impaired the shipping loan book. Losses materialised simultaneously across what the institution had believed were diversified portfolios. HSH Nordbank required a ten billion euro state guarantee from Hamburg and Schleswig-Holstein,⁸ was restructured as Hamburg Commercial Bank, and was eventually privatised in 2018.⁹

What was missing was an enterprise portfolio view that assessed correlation between apparently unrelated asset classes under stress conditions. The diversification assumption itself was a risk — one that only an enterprise-level analysis, looking across the full portfolio rather than within individual business units, could have identified and challenged. A reconciliation process that included the fourth dimension of portfolio assess-

ment — diversification and correlation — would have asked the question: under what conditions do these exposures move together? The answer would have changed the institution's risk appetite for the combined position.

This is not an exotic analytical requirement. It is the application of basic portfolio theory to risk identification. Any quantitative analyst understands that correlations increase under stress — the diversification benefits that hold in normal conditions evaporate precisely when they are needed most. The enterprise portfolio view applies this principle to the risk inventory as a whole: where does the institution believe it is diversified, and what evidence supports that belief under adverse conditions? If the answer is "historical correlations in normal markets," the diversification assumption has not been stress-tested, and the portfolio view must flag it as a concentration risk in disguise.

The Output

The enterprise portfolio view produces an **enterprise risk map** — a documented analytical product showing the aggregated risk landscape, identifying clusters of correlated exposure, and flagging areas where aggregate risk exceeds institutional appetite.

This is the deliverable that justifies the entire dual-track methodology. It feeds directly into the principal risk report that the Board Risk Committee receives (established in Chapter 3), the CCAR Material Risk Inventory that the quarterly re-identification cycle maintains (established in Chapter 6), and the integration processes that connect risk identification to capital planning, strategy, and regulatory reporting (Chapter 12 (Integration: Capital Planning, Strategy, and the Board)).

The enterprise risk map is not static. As the ongoing cycle described in Chapter 13 (The Ongoing Cycle: Refresh, Events, and Audit) refreshes the risk inventory through quarterly re-identification, event-driven updates, and annual full re-identification, the enterprise portfolio view is updated to reflect the current aggregated position. Correlation assumptions are tested against emerging data. Concentration patterns are monitored. New common exposures are identified as the institution's business mix evolves.

Phase 2 of the methodology — dual-track identification, reconciliation, and the enterprise portfolio view — is now complete. The institution has moved from two separate identification processes to a single, reconciled, aggregated picture of its risk landscape.

Every risk in the inventory has a taxonomy classification, a named owner, a documented source (top-down, bottom-up, or both), and a position within the enterprise portfolio view.

What Comes Next

The reconciled inventory and enterprise portfolio view provide the foundation for Phase 3: Assessment and Prioritisation. The institution now knows what its risks are. It must now determine how significant each risk is — and which ones demand the most urgent attention.

Chapter 9 describes how each identified risk is scored across four dimensions — financial impact, regulatory and legal impact, reputational impact, and customer and operational impact — using the scales and criteria established in Phase 1. It introduces the data quality overlay that adjusts confidence in those scores based on the quality of underlying information. And it establishes the materiality determination process that separates the risks requiring Board attention from those managed at business unit level. The four-dimensional assessment framework that Chapter 5 introduced and this chapter's enterprise portfolio view requires is the subject of what follows.

-
1. Credit Suisse Group, *Report on Archegos Capital Management* (Paul, Weiss, Rifkind, Wharton & Garrison LLP, 29 July 2021), p. 1. The report states total losses to Credit Suisse of approximately \$5.5 billion.
 2. Credit Suisse Group, *Report on Archegos Capital Management* (Paul, Weiss, 29 July 2021), pp. 4–12. The report details how Archegos's total return swap positions were spread across multiple prime brokers, with no single broker having visibility of the aggregate exposure.
 3. Financial Crisis Inquiry Commission, *The Financial Crisis Inquiry Report* (Washington, DC: U.S. Government Printing Office, January 2011), pp. 298–303. The FCIC documents Citigroup's divisional structure and the fragmented risk reporting across its business units.
 4. FCIC, *The Financial Crisis Inquiry Report* (January 2011), pp. 299–301. The report describes the SIV structure, the commercial paper funding mechanism, and the liquidity put commitments that Citigroup had written to its off-balance-sheet vehicles.
 5. FCIC, *The Financial Crisis Inquiry Report* (January 2011), p. 301. Citigroup consolidated approximately \$49 billion in SIV assets back onto its balance sheet when the commercial paper market froze. See also Citigroup, Inc., Form 10-K for fiscal year ended 31 December 2007, filed with the SEC.
 6. U.S. Department of the Treasury, *Troubled Asset Relief Program (TARP): Monthly Report to Congress*, December 2008 and subsequent reports. Citigroup received \$45 billion in TARP capital injections (\$25 billion in October 2008 and \$20 billion in December 2008), plus a government loss-sharing arrangement on approximately \$301 billion of assets.

7. HSH Nordbank AG, Annual Report 2007. HSH Nordbank was widely reported as the world's largest provider of shipping finance, with a shipping loan portfolio of approximately EUR 30 billion at its peak.
8. European Commission, State Aid Decision SA.29338, *HSH Nordbank — Restructuring Aid* (20 September 2011). The guarantee was provided by the Free and Hanseatic City of Hamburg and the State of Schleswig-Holstein, totalling EUR 10 billion.
9. HSH Nordbank was renamed Hamburg Commercial Bank (HCOB) following its privatisation and sale to a consortium led by Cerberus Capital Management and J.C. Flowers & Co., completed in November 2018. See Hamburg Commercial Bank, press release, 28 November 2018.

Assessment — Scoring, Multi-Dimensional Impact, and Data Quality

The Risk That Was Assessed at Near-Zero

In 2005, AIG Financial Products was the most profitable division of the largest insurance company in the world. Its primary revenue engine was the sale of credit default swaps — insurance contracts that paid out if mortgage-linked CDOs defaulted. By the peak, AIG had sold protection on a CDS portfolio with \$527 billion in notional exposure¹ — of which the lethal concentration was approximately \$78 billion in multi-sector CDO protection,² the tranche most directly tied to subprime mortgage performance.

The risk assessment was straightforward. The CDOs AIG insured were rated AAA by every major agency. The historical default rate on AAA-rated securities was near zero. AIG's internal models, calibrated to that historical data, assigned negligible probability to a scenario in which the insured tranches would suffer losses. The financial impact, assessed on a single dimension, appeared minimal. Likelihood, assessed against historical defaults, was vanishingly small.

What the assessment did not capture was everything else.

It did not assess the regulatory consequence of a ratings downgrade triggering tens of billions in collateral calls that AIG could not meet. It did not assess the reputational destruction of the world's largest insurer requiring emergency government intervention. It did not assess the customer and operational impact of a counterparty failure that would cascade through every major bank that had purchased AIG's protection. And it did not assess the quality of the data underlying the entire analysis — models calibrated to a period in which nationwide housing prices had never declined, using correlation assumptions that had never been tested under stress.

The \$85 billion Federal Reserve bailout³ that followed was not a failure of risk identification. AIG knew it was selling CDS. The counterparties were documented. The notional exposure was calculable. The failure was in how that exposure was assessed — a single-dimensional scoring methodology, calibrated to benign historical data of questionable relevance, that produced a near-zero risk rating for what turned out to be the single largest risk event in the history of insurance.

Phase 3 of the methodology exists to prevent that outcome. Every risk that has survived the identification and reconciliation process of Phases 1 and 2 must now be scored, assessed for data quality, evaluated on both an inherent and residual basis, and tested against a materiality threshold. The four-dimensional assessment framework that Chapter 5 (Setting the Context: External, Internal, and Risk Culture) introduced and Chapter 8 (Reconciliation and the Enterprise Portfolio View)'s enterprise portfolio view requires is the subject of this chapter.

Why Assessment Is Not Measurement

Before examining the scoring methodology, a distinction must be drawn. Assessment, as used in this methodology, is the structured evaluation of identified risks using defined scales and criteria. It is not quantitative risk measurement. It does not replace the bank's credit risk models, market risk VaR engines, operational risk capital calculations, or stress testing programmes. Those are measurement tools that serve different purposes — capital adequacy, regulatory reporting, limit-setting.

Assessment serves the risk identification process. Its purpose is threefold: to prioritise the risk inventory so that material risks receive appropriate attention, to enable comparison across fundamentally different risk types using a common language, and to provide the Board Risk Committee with an enterprise view of the institution's risk profile that is neither so granular as to be unusable nor so aggregated as to be meaningless.

The multivoting exercise in the top-down workshop (Chapter 6 (Top-Down Identification: Workshops, SWIFT, and Delphi)) produced a preliminary prioritisation based on collective senior judgement. That prioritisation was useful for structuring the workshop's outputs, but it was not rigorous. It did not use defined scales, it did not assess non-financial consequences, it did not account for data quality, and it did not distinguish inherent from residual risk. Phase 3 replaces that preliminary view with a structured assessment that can withstand regulatory scrutiny and Board challenge.

ISO 31000 Section 5.4.3 requires that risk analysis consider consequences and their likelihood, the nature and magnitude of those consequences, complexity and connectivity, time-related factors, and the effectiveness of existing controls.¹⁵ ISO 31010 provides the technique framework. The four-dimensional scoring methodology that follows is the practitioner implementation of those requirements.

The Four-Dimensional Impact Framework

The single most important design decision in the assessment framework is that impact is not assessed on one dimension. A risk that produces a modest financial loss but triggers licence revocation is not a moderate risk. A risk that costs nothing in direct financial terms but destroys client confidence and generates sustained international press coverage is not an incidental risk. A risk that causes no immediate financial impact but results in total service failure affecting millions of customers is not a minor risk.

The methodology assesses impact across four consequence dimensions. The highest-scoring dimension determines the overall impact rating. This design principle — **the dominant dimension rule** — ensures that risks with severe non-financial consequences are never understated by a benign financial assessment.

Financial Impact

Financial impact is anchored to CET1 capital, providing a consistent denominator that scales with institutional size.

| Score | Rating | Description |
|-------|------------|---|
| 1 | Incidental | Less than 0.1% of CET1 capital |
| 2 | Minor | 0.1–1% of CET1 capital |
| 3 | Moderate | 1–5% of CET1 capital |
| 4 | Major | 5–15% of CET1 capital |
| 5 | Extreme | Greater than 15% of CET1 capital; threatens viability |

The CET1 anchor achieves two things. First, it makes the scale objective and auditable — a dispute about whether a risk is “moderate” or “major” can be resolved by reference to the institution’s published capital position. Second, it prevents the common failure where a fixed monetary threshold (say, \$100 million) is applied identically to institutions of vastly different sizes, making the scale meaningless at either end.

Regulatory and Legal Impact

| Score | Rating | Description |
|-------|------------|--|
| 1 | Incidental | No regulatory impact |
| 2 | Minor | Informal supervisory feedback, minor findings |
| 3 | Moderate | Formal regulatory investigation or enforcement warning |
| 4 | Major | Formal enforcement action, significant fines, restrictions on activities |
| 5 | Extreme | Licence revocation, criminal prosecution, forced restructuring |

This dimension captures what financial impact alone cannot. Standard Chartered’s \$667 million sanctions fine⁴ (Chapter 3 (Governance: Who Owns What)) was a financial event, but the regulatory dimension — formal enforcement, activity restrictions, sustained supervisory scrutiny — constrained the institution’s strategic options for years beyond the fine itself. Deutsche Bank’s \$630 million combined UK/US fines for mirror-trading⁵ (Chapter 4 (The Risk Taxonomy)) carried a regulatory impact that extended well beyond the payment: enhanced monitoring requirements, consent orders, and reputational damage with supervisors that affected every subsequent regulatory interaction.

Reputational Impact

| Score | Rating | Description |
|-------|------------|----------------------|
| 1 | Incidental | No external coverage |

| Score | Rating | Description |
|-------|----------|---|
| 2 | Minor | Limited trade or local press, quickly contained |
| 3 | Moderate | National press coverage, sustained for days |
| 4 | Major | International press, sustained coverage, loss of key client relationships |
| 5 | Extreme | Total loss of market confidence, Board or CEO departure, sustained international coverage |

Reputational impact is the dimension most frequently omitted from assessment frameworks and most frequently decisive in determining whether a risk event threatens institutional survival. Wells Fargo's unauthorised accounts scandal (Chapter 3) cost \$3 billion in direct financial terms⁶ — large, but manageable for an institution of that size. The reputational destruction — CEO departure, Fed asset cap, sustained Congressional scrutiny, customer attrition — was the existential dimension. A single-dimensional financial assessment would have rated it as serious but manageable. The reputational score places it at the extreme end of the scale.

Customer and Operational Impact

| Score | Rating | Description |
|-------|------------|---|
| 1 | Incidental | Negligible operational disruption, no customer impact |
| 2 | Minor | Limited operational disruption, minimal customer impact |
| 3 | Moderate | Significant operational disruption, noticeable customer service degradation |
| 4 | Major | Major service disruption, large-scale customer impact, material complaints |
| 5 | Extreme | |

| Score | Rating | Description |
|-------|--------|-------------|
|-------|--------|-------------|

Total service failure, mass customer harm, systemic operational breakdown

This dimension captures the operational and customer consequences that may precede, accompany, or exist independently of financial loss. The PPI scandal (Chapter 4) was a customer impact event before it was a financial event — millions of customers were sold unsuitable insurance products over years before the redress programme began. An assessment framework that waited for financial crystallisation would have missed the customer harm dimension entirely.

Applying the Dominant Dimension Rule

Consider a cyber attack that exfiltrates customer data but causes no direct financial loss, attracts moderate press attention, and triggers a formal regulatory investigation. Scored individually: Financial = 1 (Incidental), Regulatory = 3 (Moderate), Reputational = 3 (Moderate), Customer/Operational = 4 (Major — large-scale customer impact). The overall impact score is 4. Without the four-dimensional framework, a single financial impact assessment would rate this risk at 1.

The dominant dimension rule is the mechanism that prevents the AIG failure from recurring. AIG's CDS exposure, assessed on financial impact alone using benign historical assumptions, appeared incidental. Assessed across four dimensions — with the regulatory consequence of collateral calls, the reputational destruction of government bailout, and the customer/operational impact of counterparty failure cascading through the financial system — the risk would have scored at the extreme end of the scale on at least three of the four dimensions.

Likelihood

| Score | Rating | Description |
|-------|--------|-------------|
|-------|--------|-------------|

| | | |
|---|------|--|
| 1 | Rare | Less than 1% probability within the assessment horizon |
|---|------|--|

| | | |
|---|----------|-------------------|
| 2 | Unlikely | 1–10% probability |
|---|----------|-------------------|

| Score | Rating | Description |
|-------|----------|---|
| 3 | Possible | 10–50% probability |
| 4 | Likely | 50–90% probability |
| 5 | Frequent | Greater than 90% probability, or has already occurred |

The probability ranges provide guidance, not false precision. Expert judgement is required to place each risk on the scale, and the data quality rating (below) discloses how much confidence should be placed in that judgement.

Three common failures in likelihood assessment deserve explicit attention.

Anchoring to recent experience. If a risk has not materialised in the assessor’s tenure, it defaults to “Rare.” This is the complacency failure mode from Chapter 1 (Why Banks Fail at Risk Identification). LTCM’s models were calibrated to a period in which correlations between global fixed income markets had remained moderate. The 1998 Russian crisis caused correlations to spike to one across all markets simultaneously⁷ — an event the models assigned near-zero probability because it had not occurred in the calibration window. The methodology’s response to this failure is twofold: the Data Quality Rating (below) forces disclosure of the evidence basis, and the scenario analysis outputs from the top-down workshop (Chapter 6) provide structured challenge to likelihood assumptions by constructing plausible scenarios that differ from recent experience.

Conflating likelihood with perceived controllability. Assessors rate a risk as unlikely because controls exist. But likelihood in this framework is assessed at the inherent level first — before controls are applied. A risk with high inherent likelihood and strong controls has high inherent risk and (potentially) lower residual risk. Collapsing both into a single “unlikely” rating conceals the degree to which the institution depends on controls that may themselves fail.

Treating independence as default. Multiple risks are each rated as individually unlikely, but the scenario in which they crystallise simultaneously is not assessed. Simultaneous crystallisation belongs to the enterprise portfolio view (Chapter 8) and risk interaction analysis (Chapter 10 (Risk Interaction: Bow-Ties, Matrices, and Concentration)), but likelihood assessors must be aware that their individual ratings will be used in aggregate analyses where independence assumptions matter.

Vulnerability

Impact and likelihood are the traditional two dimensions of risk assessment. The methodology adds two more. The first is vulnerability.

| Score | Rating | Description |
|-------|-----------|---|
| 1 | Very Low | Robust controls, strong resilience, proven under stress |
| 2 | Low | Adequate controls, minor gaps identified |
| 3 | Medium | Controls exist but with known weaknesses or untested elements |
| 4 | High | Material control gaps, limited resilience |
| 5 | Very High | No effective controls, fully exposed |

Vulnerability captures the institution's preparedness independently of whether the risk event occurs. Two institutions may face the same risk with the same likelihood and the same potential impact, but one has tested controls, proven resilience, and documented recovery procedures while the other has untested controls and no recovery capability. They are not in the same position, and their assessment scores should reflect that difference.

Vulnerability scoring draws on the control effectiveness information captured in the standardised risk assessment template (Chapter 7 (Bottom-Up Identification: Templates, RCSA, and Specialist Processes)), which requires assessors to specify control type (preventive, detective, corrective) and provide evidence-based effectiveness ratings. The vulnerability score is not a repetition of the control assessment — it is an institutional-level judgement about the degree to which the institution is exposed if the risk materialises.

The critical distinction is between controls that have been tested and controls that are assumed to work. UBS's internal shareholder report following its \$37.4 billion in structured credit write-downs⁸ identified that the risk function had relied on the same models and assumptions as the front office, providing no independent challenge. VaR

limits were repeatedly increased at traders' requests. The controls existed on paper. The vulnerability was extreme because no independent challenge mechanism had ever tested whether those controls would function under stress. A vulnerability rating of 1 (Very Low) requires controls that are not merely documented but **proven under stress** — a standard that most institutions cannot honestly claim for most of their risk exposures.

Speed of Onset

The fourth dimension captures how quickly a risk can move from latent to crystallised.

| Score | Rating | Description |
|-------|-----------|-----------------------------------|
| 1 | Very Slow | Years — ample time to respond |
| 2 | Slow | Quarters — time to plan and act |
| 3 | Moderate | Months — urgent but manageable |
| 4 | Fast | Weeks — limited response window |
| 5 | Immediate | Days or less — no time to respond |

Speed of onset determines whether the institution will have time to implement risk responses when a risk begins to crystallise. A risk rated "Major" on impact and "Possible" on likelihood presents a fundamentally different management challenge depending on whether onset is measured in years or days.

Bear Stearns illustrates the operational significance of speed of onset. Its concentrated MBS exposure was a known risk that materialised not through gradual credit deterioration but through a sudden loss of counterparty confidence in March 2008. The institution went from operational to acquired in a matter of days.¹⁴ No risk response plan designed for a "Slow" onset would have been adequate. A speed-of-onset rating of 5 (Immediate) for a counterparty confidence scenario would have flagged the need for pre-positioned liquidity and contingency arrangements — not as a response to the risk crystallising, but as a standing requirement for a risk that, when it moves, moves too fast for real-time intervention.

Climate risk provides the contrasting example. Physical risk from chronic sea-level rise affecting property portfolios may score “Major” on impact but “Very Slow” on speed of onset — the institution has years to adjust exposures. Transition risk from sudden regulatory policy change (carbon tax, stranded asset rules) may score the same impact but “Fast” on speed of onset — the policy announcement creates immediate mark-to-market consequences and strategic constraints.

The four dimensions together — Impact, Likelihood, Vulnerability, Speed of Onset — constitute the **four-dimensional risk score**. Every risk in the inventory carries all four scores for both its inherent and residual assessment.

Data Quality and Confidence Rating

The four-dimensional score is only as reliable as the evidence underlying it. A risk rated “Major impact, Possible likelihood” based on five years of validated internal loss data is a fundamentally different proposition from the same rating based on a single expert’s opinion with no supporting data. Both may be the best assessment available, but the degree of confidence the Board should place in each is entirely different.

Every risk assessment must include a **Data Quality Rating** that discloses the reliability of the evidence underlying the scores. This is not optional. It is a methodological requirement grounded in ISO 31010, Section 4.3, which requires transparent recording of the basis for risk assessments.

| Rating | Description | Typical Basis |
|---------------|---|--|
| High | Based on robust quantitative data, validated models, or extensive historical evidence | Internal loss data, market data, regulatory capital models |
| Medium | Based on limited quantitative data supplemented by structured expert judgement | Partial loss history, industry benchmarks, scenario analysis |
| Low | Based primarily on expert judgement with limited or no historical data | Workshops, interviews, peer institution experience |

| Rating | Description | Typical Basis |
|----------|--|--|
| Very Low | Speculative — no relevant data, precedent, or experience | Emerging risks, novel exposures, unprecedented scenarios |

Chapter 7 identified the Data Quality Rating as one of the three fields where bottom-up submissions most commonly fail — business units default to “Medium” for everything, avoiding both the uncomfortable admission that evidence is weak and the analytical effort required to justify a “High” rating. The Risk Identification Lead must challenge these defaults. A submission that rates data quality as “Medium” must demonstrate that it has the limited quantitative data and structured expert judgement the definition requires. If the only basis is a single person’s opinion, the rating is “Low.”

How Data Quality Affects the Assessment

The Data Quality Rating is not a standalone curiosity reported alongside the score. It has three operational consequences.

Conservatism adjustment. Where data quality is Low or Very Low, the methodology applies a conservatism principle: uncertainty about a risk should not reduce the attention it receives. In practice, this means that risks with low data quality are not permitted to carry low impact or likelihood ratings without explicit CRO approval. If the evidence basis is too weak to support a firm assessment, the default assumption is that the risk is more severe than the limited evidence suggests — not less. This directly addresses the AIG pattern, where a risk scored as near-zero was based on historical data that had never included the relevant stress scenario.

Sensitivity testing. Risks with Low or Very Low data quality ratings are flagged for sensitivity analysis: how would the overall assessment change if the impact were one level higher? If the likelihood were one level higher? If the sensitivity analysis moves the risk across the materiality threshold, it is treated as potentially material regardless of the base assessment. This is not a quantitative exercise requiring Monte Carlo simulation. It is a structured what-if: the same SWIFT questioning discipline applied in the top-down workshop (Chapter 6), now applied to individual risk scores.

Board transparency. The Board Risk Report includes a summary of data quality distribution across the material risk portfolio. If thirty per cent of material risks carry Low or Very Low data quality ratings, the Board needs to know that. It changes how the Board should interpret the risk profile, it identifies where the institution needs to invest in better data, and it prevents the false precision that arises when a risk heatmap presents all risks as if they were assessed with equal confidence.

The Data Quality Problem at Wachovia

Wachovia's 2008 collapse illustrates the data quality failure in its purest form. The bank's acquisition of Golden West Financial in 2006 brought \$122 billion in option ARMs⁹ — adjustable-rate mortgages where borrowers could choose to pay less than the interest accruing, with the difference added to the loan principal. The acquisition due diligence accepted Golden West's internal assessment of credit quality: the loans had performed well historically in California, with low default rates and strong recovery values.

That assessment deserved a Data Quality Rating of Very Low. The evidence basis was a historical track record that did not include a single period of nationwide housing price decline. The option ARM product was relatively new, meaning even the California-specific track record was short. And the portfolio was about to be deployed into a fundamentally different economic environment — from a rising housing market with loose lending standards to whatever came next.

Had the assessment carried a Very Low data quality rating, the conservatism principle would have required higher assumed losses. The sensitivity analysis would have tested what happened if default rates exceeded anything in the historical record. The Board would have seen that the single largest acquisition in the institution's history was supported by the weakest possible evidence base. Instead, the historical performance data was taken at face value, the assessment methodology did not require disclosure of data limitations, and Wachovia absorbed approximately \$25 billion in losses — reflecting the convergence of quarterly operating losses, goodwill impairment on the Golden West acquisition, and option ARM portfolio writedowns — before its emergency sale to Wells Fargo.¹⁰

Inherent and Residual Risk

Each risk is scored twice across all four dimensions.

Inherent risk is the risk level assuming no controls are in place. It answers the question: how exposed is the institution to this risk in the absence of any mitigating action? This is not a hypothetical exercise. It establishes the magnitude of the risk that the control environment must contain.

Residual risk is the risk level after accounting for current controls. This is the risk the institution is actually running. The difference between inherent and residual risk reflects the institution's dependence on its control environment — and by implication, the consequence if those controls fail or are circumvented.

Between the two sits **control effectiveness** — an assessment rated 1 to 5 (highly effective to ineffective) that draws on the control information captured in the bottom-up template (Chapter 7). The control assessment must specify the type of each control (preventive, detective, corrective) and provide evidence for the effectiveness rating.

The dual assessment matters for three reasons.

It reveals control dependency. A risk with inherent impact of 5 and residual impact of 2 tells the Board that the institution depends heavily on the controls reducing that risk by three levels. If those controls are untested, recently changed, or dependent on a single system or individual, the Board may conclude that the residual assessment is optimistic. The vulnerability dimension captures this concern, but the inherent-residual gap makes it visible.

It identifies where controls mask exposure. One of the most productive exercises in assessment is mapping inherent-residual gaps against control effectiveness evidence. Risks where the gap is large but the control evidence is weak — effectiveness ratings based on design documentation rather than operational testing — should be systematically flagged for deeper review. These frequently turn out to be risks where the institution is running significantly more exposure than the residual score suggests.

It prevents assessment gaming. Without the inherent-residual separation, assessors can produce moderate residual scores by implicitly assuming that controls work, without ever being required to assess the underlying risk magnitude. The dual assessment forces transparency: state the inherent risk, state the controls and their evidence, derive the residual risk. Each step is auditable.

The Independent Challenge Requirement

Chapter 1 identified Model Overreliance as one of the ten failure modes — the pattern where institutions substitute model outputs for independent analytical judgement, treating quantitative models as facts rather than as one input to assessment. The methodology's response is explicit: **models are an input to assessment, not a substitute for it.**

Where quantitative models inform the risk score — credit risk models, VaR, LGD models, operational risk capital models — the assessment must include an independent challenge of the model's assumptions and limitations. This challenge is documented in the risk inventory alongside the score.

Independent challenge means three things in practice. First, the model's key assumptions are stated explicitly — not buried in technical documentation but visible in the assessment record. Second, the scenarios in which the model's assumptions break down are identified. Third, management judgement about the model's applicability to the current environment is recorded. UBS's post-crisis shareholder report found that its risk function relied on the same models and assumptions as the front office, providing no independent challenge — VaR limits were repeatedly increased at traders' requests. The methodology requires that risk assessment exists independently of the models the front office uses to manage its positions.

This is not a prohibition on using models. It is a prohibition on using models uncritically. A credit risk model's output is valuable evidence for the financial impact dimension. But the model's output must be accompanied by: What does the model assume? Under what conditions do those assumptions fail? What is the Data Quality Rating applicable to the model's inputs? A model calibrated to benign historical conditions, using data rated Very Low for relevance to current conditions, should not produce an assessment score that the Board treats with the same confidence as a model validated against the specific stress scenario the institution faces.

Resolving the Disagreement Log

Chapter 6 introduced the disagreement log — a record of materially different views expressed during the top-down workshop that could not be reconciled through discussion. Chapter 6 stated that the assessment phase must resolve these disagreements through evidence, not social dynamics. This is where that resolution occurs.

The four-dimensional scoring framework provides the structured basis for resolution. When two senior participants disagree about the severity of a risk, the assessment forces specificity: on which dimension do you disagree? Financial impact? Regulatory consequence? Likelihood? What evidence supports each position? What is the Data Quality Rating applicable to that evidence?

In many cases, the disagreement dissolves once the dimensions are separated. A CRO and a business unit head may both be correct — the CRO assessing severe regulatory impact while the business head assesses manageable financial impact. Under a single-dimension framework, one must “win.” Under the four-dimensional framework, both assessments are captured and the dominant dimension rule determines the overall score.

Where genuine analytical disagreement persists after dimensional separation, the methodology preserves it. The minority view is recorded in the risk inventory alongside the consensus score, with the evidence basis for each position documented. This is not a compromise — it is transparency. The Board Risk Committee receives the consensus assessment and the dissenting view, with the data quality ratings applicable to each. The Board decides how much weight to place on the minority position. Risks where senior participants genuinely disagree after structured analysis are disproportionately likely to be the risks the institution most needs to understand.

Materiality Determination

Not all risks require the same level of attention. A bank with two hundred risks in its inventory cannot — and should not — subject each to full risk profiling, dedicated KRIs, and integration into capital and strategic planning. Materiality determination is the governance mechanism that focuses resources on what matters.

The process is straightforward:

1. Plot all risks on an **Impact x Likelihood matrix** — the risk heatmap — using the inherent scores and the residual scores separately
2. Apply the **materiality threshold** defined in Phase 1 (Chapter 5) — calibrated to the institution's risk appetite and approved by the CRO
3. Risks above the threshold are classified as **material**
4. Material risks receive full risk profiles (Chapter 11 (Documentation: The Living Risk Inventory)), dedicated KRIs (Chapter 13 (The Ongoing Cycle: Refresh, Events, and Audit)), and integration into capital and strategic planning (Chapter 12 (Integration: Capital Planning, Strategy, and the Board))
5. Non-material risks remain in the inventory and are monitored for changes through the ongoing cycle (Chapter 13)

Typically, 20–60 risks out of the full inventory are classified as material. This range reflects institutional size and complexity — a large universal bank with global operations will have more material risks than a regional commercial bank, because it faces a broader range of risk types at a scale that exceeds materiality thresholds.

The materiality threshold is not a line on a heatmap and nothing more. It incorporates three factors from the four-dimensional framework:

The dominant dimension. A risk that scores below threshold on financial impact but above threshold on regulatory impact is material. The four dimensions are assessed independently and the highest determines materiality.

Aggregation. Individual risks below threshold may be material when aggregated with related risks. The enterprise portfolio view (Chapter 8) identifies common exposures and correlated clusters. If three individually sub-material risks share a common driver and would crystallise simultaneously, their aggregate position is assessed against the materiality threshold. This is the mechanism that catches the Citigroup pattern — SIV exposures distributed across divisions, each individually within limits, collectively catastrophic.

Data quality adjustment. A risk assessed as sub-material based on a Very Low data quality rating — meaning the evidence basis is speculative — is flagged for CRO review rather than automatically excluded from the material risk portfolio. The sensitivity testing described above is specifically designed to prevent risks from falling below the materiality threshold simply because the institution lacks the data to assess them properly.

What the Heatmap Does Not Show

The risk heatmap — the Impact x Likelihood matrix — is the most widely used output of risk assessment in banking. It is also the most frequently misused.

A heatmap plots individual risks as points on a two-dimensional grid. It does not show vulnerability or speed of onset. It does not show data quality. It does not show inherent-residual gaps. It does not show correlations or common exposures. A heatmap with forty risks evenly distributed across the grid looks reassuring. The same forty risks, annotated with their vulnerability ratings, speed-of-onset scores, data quality ratings, and correlation clusters, may present a fundamentally different picture.

The heatmap is a screening tool. It is the starting point for materiality determination, not the conclusion. Material risks identified through the heatmap proceed to full risk profiling (Chapter 11), risk interaction analysis (Chapter 10), and integration into capital planning and strategy (Chapter 12). The heatmap alone does not tell the Board what it needs to know. It tells the Risk Identification Lead where to direct the deeper analysis that will.

Merrill Lynch and the AAA Assumption

Merrill Lynch accumulated approximately \$40 billion in gross subprime exposure, of which some \$32 billion consisted of super-senior CDO tranches retained on its balance sheet¹¹ — the most senior positions in the CDO structure, designed to suffer losses only after every junior tranche had been wiped out. These tranches carried AAA ratings from every major agency. Merrill's risk assessment treated the AAA rating as equivalent to independent analysis: if the rating was AAA, the financial impact of loss was negligible, and the position did not appear as material in internal risk reports.

The assessment failure was multi-dimensional. Financially, the \$51.8 billion in eventual write-downs¹² made the super-senior retention one of the largest single-risk losses in banking history. But the regulatory dimension was equally severe — the forced emer-

agency sale to Bank of America at half the peak share price, negotiated over a weekend under intense regulatory pressure.¹³ The reputational dimension produced the departure of CEO Stan O’Neal and years of franchise damage. A four-dimensional assessment that independently scored each consequence dimension — rather than relying on the AAA rating as a proxy for all of them — would have produced a very different risk profile.

The Data Quality Rating is equally revealing. What was the evidence basis for Merrill’s assessment? External ratings produced by agencies whose revenue depended on the CDO issuance business. Internal models that had not been validated against a nationwide housing price decline because no such decline existed in the data. The appropriate Data Quality Rating was Low at best — expert judgement supplemented by models of unproven relevance. Under the methodology’s conservatism principle, that data quality rating would have prevented the near-zero risk assessment that allowed \$40 billion in concentrated exposure to accumulate without triggering materiality thresholds.

Putting It Together: The Assessment Record

When Phase 3 is complete, every risk in the inventory carries a structured assessment record:

| Field | Content |
|-----------------------------------|---|
| Inherent Risk Score | Impact (with dominant dimension noted), Likelihood, Vulnerability, Speed of Onset |
| Control Effectiveness | Rating (1–5), control types (preventive/detective/corrective), evidence basis |
| Residual Risk Score | Impact, Likelihood, Vulnerability, Speed of Onset |
| Data Quality Rating | High / Medium / Low / Very Low, with evidence basis stated |
| Materiality Classification | Material / Non-material, with basis for determination |
| Model Dependency | Where models inform the score: key assumptions, limitations, independent challenge documented |

| Field | Content |
|----------------------------|--|
| Disagreement Record | Where applicable: minority view, evidence basis, dimensional breakdown |

This record integrates with the standardised risk assessment template fields established in Chapter 7. The taxonomy classification, risk definition, underlying drivers, current controls, risk owner, and emerging risk indicators from the bottom-up process are now supplemented with the structured assessment from Phase 3. Together, they form the basis for the risk inventory that Chapter 11 will describe.

How This Could Have Changed AIG

Return to the opening case. Apply the methodology to AIG’s CDS portfolio as of 2006.

Four-dimensional impact. Financial: the notional exposure was enormous, but the perceived loss was near-zero based on default assumptions. Under the four-dimensional framework, the assessor must also score the regulatory, reputational, and customer/operational dimensions. The regulatory dimension alone — collateral call triggers upon ratings downgrade, potential for government intervention — would have scored at least Major (4). The reputational dimension of the world’s largest insurer requiring emergency bailout would have scored Extreme (5). The dominant dimension rule produces an overall impact score of 5, regardless of the financial impact assumption.

Likelihood. Under AIG’s models, the likelihood of AAA tranches defaulting was assessed at Rare (1). The Data Quality Rating for that assessment was, at best, Low — models calibrated to a period without the relevant stress scenario. Under the conservatism principle, a Rare likelihood rating based on Low data quality requires CRO approval and sensitivity testing. The sensitivity test — “what if the likelihood is one level higher?” — would have moved the risk to Unlikely (2) with Extreme (5) impact. That combination exceeds any reasonable materiality threshold.

Vulnerability. AIG’s controls against a systemic CDS crystallisation scenario were effectively non-existent. There was no hedging programme for the written CDS, no collateral reserve adequate for a mass downgrade, and no pre-positioned liquidity for collateral calls at scale. Vulnerability: Very High (5).

Speed of onset. The collateral call mechanism meant that onset would be Immediate (5) — triggered by ratings action, not by actual defaults.

The four-dimensional assessment produces: Impact 5, Likelihood 1 (disputed — Low data quality), Vulnerability 5, Speed of Onset 5. With Data Quality at Low or Very Low, the conservatism adjustment applies. The sensitivity test moves likelihood to 2, placing the risk firmly in the material category. The risk profile that would follow (Chapter 11) would document the collateral call mechanism, the concentration in mortgage-linked CDOs, and the absence of hedging. The Board would have seen it. The CRO would have been required to explain why the institution was writing \$527 billion in notional unhedged protection — including \$78 billion in the most toxic multi-sector CDO tranche — on a single risk factor with Extreme impact, Very High vulnerability, Immediate speed of onset, and evidence quality that deserved no confidence.

The \$85 billion bailout was not inevitable. It was the consequence of a single-dimensional assessment methodology applied to a risk that required four.

The Bridge to Risk Interaction

Phase 3 has now scored every risk in the inventory on four dimensions, overlaid each score with a data quality rating, separated inherent from residual assessment, and applied a materiality threshold to focus resources. But each risk has been assessed individually. The enterprise portfolio view in Chapter 8 identified correlated clusters and common exposures. The next step is to analyse those interactions systematically — using bow-tie analysis to map causal chains, risk interaction matrices to identify which risks amplify each other, and concentration analysis to determine where the institution's aggregate exposure exceeds what any individual risk score would suggest. That analysis — risk interaction — is the subject of Chapter 10.

-
1. AIG, Form 10-K for fiscal year ended 31 December 2007, filed with the U.S. Securities and Exchange Commission, pp. 122–127. The filing reports AIGFP's credit default swap portfolio at a notional amount of approximately \$527 billion as of year-end 2007.
 2. AIG, Form 10-K, fiscal year 2007, pp. 122–127. Of the total CDS portfolio, approximately \$78 billion was referenced to multi-sector CDOs, including subprime residential mortgage-backed securities. See also FCIC, *The Financial Crisis Inquiry Report* (January 2011), pp. 265–270.

3. Board of Governors of the Federal Reserve System, press release, 16 September 2008. The Federal Reserve Bank of New York authorised a secured credit facility of up to \$85 billion to AIG to prevent the company's disorderly failure. The facility was subsequently restructured multiple times; total government support ultimately exceeded \$180 billion across all programmes.
4. Standard Chartered Bank, Deferred Prosecution Agreement with the New York County District Attorney's Office and consent orders with the New York State Department of Financial Services (NYDFS), 2012–2014. The combined penalties across multiple settlements totalled approximately \$667 million for violations of U.S. sanctions laws.
5. U.K. Financial Conduct Authority, Final Notice to Deutsche Bank AG, 31 January 2017 (£163 million penalty); New York State Department of Financial Services, Consent Order, 30 January 2017 (\$425 million penalty). The combined UK/US fines related to mirror-trading schemes that moved approximately \$10 billion out of Russia.
6. U.S. Department of Justice, press release, 21 February 2020. Wells Fargo agreed to pay \$3 billion to resolve criminal and civil investigations into the bank's practice of opening millions of unauthorised accounts. The settlement covered conduct from 2002 to 2016.
7. President's Working Group on Financial Markets, *Hedge Funds, Leverage, and the Lessons of Long-Term Capital Management* (Washington, DC: U.S. Department of the Treasury, April 1999). The report documents how the August 1998 Russian default triggered simultaneous liquidity crises across previously uncorrelated markets, causing LTCM's portfolio losses to exceed what its models predicted as possible.
8. UBS AG, *Shareholder Report on UBS's Write-Downs* (Zurich: UBS, 18 April 2008), pp. 13–14, 30–32. The report documents \$37.4 billion in write-downs on subprime-related positions and identifies the failure of independent risk challenge as a root cause, noting that the risk function relied on the same models and assumptions as the front office.
9. FCIC, *The Financial Crisis Inquiry Report* (January 2011), pp. 137–138. The report documents Wachovia's 2006 acquisition of Golden West Financial Corporation for approximately \$25.5 billion, acquiring its \$122 billion portfolio of option adjustable-rate mortgages. See also Golden West Financial Corporation, Form 10-K for fiscal year ended 31 December 2005.
10. FDIC, press release, 29 September 2008, and subsequent regulatory filings. Wachovia's losses reflect the combination of quarterly operating losses from the option ARM portfolio, a \$18.7 billion goodwill impairment charge related to the Golden West acquisition (Q4 2007 through Q3 2008), and continued mortgage portfolio write-downs. The FDIC facilitated Wachovia's emergency acquisition by Wells Fargo in October 2008.
11. FCIC, *The Financial Crisis Inquiry Report* (January 2011), pp. 255–261. The report documents Merrill Lynch's accumulation of CDO exposure, including approximately \$40 billion in gross subprime exposure and approximately \$32 billion in super-senior CDO tranches retained on balance sheet. See also Merrill Lynch & Co., Inc., Form 10-K for fiscal year ended 28 December 2007.
12. Merrill Lynch & Co., Inc., SEC filings, 2007–2008. Cumulative write-downs on CDO and subprime-related positions totalled approximately \$51.8 billion across multiple quarters. See also FCIC, *The Financial Crisis Inquiry Report* (January 2011), pp. 255–261.
13. FCIC, *The Financial Crisis Inquiry Report* (January 2011), pp. 380–386. The report details the emergency negotiations over the weekend of 13–14 September 2008 that resulted in Bank of America's acquisition of Merrill Lynch at \$29 per share, approximately half Merrill's peak share price.
14. FCIC, *The Financial Crisis Inquiry Report* (January 2011), pp. 280–291. Bear Stearns experienced a rapid loss of counterparty confidence in March 2008; by 14 March its liquidity pool had effectively been depleted. The Federal Reserve facilitated JPMorgan Chase's acquisition, announced on 16 March 2008 at \$2 per share (later revised to \$10 per share). See also SEC, Office of Inspector General, *SEC's Oversight of Bear Stearns and Related Entities*, Report No. 446-A, 25 September 2008.
15. International Organization for Standardization, *ISO 31000:2018 — Risk Management: Guidelines*, Section 6.4.3 (renumbered from 5.4.3 in the 2009 edition). See also ISO, *IEC 31010:2019 — Risk Management: Risk Assessment Techniques*, Section 4.3, which requires transparent recording of the basis for risk assessments.

Risk Interaction: Bow-Ties, Matrices, and Concentration

The Bank That Was Destroyed by Five Risks at Once

Lehman Brothers did not fail because of credit risk. It did not fail because of liquidity risk, or market risk, or operational risk, or reputational risk. It failed because all five crystallised simultaneously, each amplifying the others in a chain reaction that no individual risk assessment had captured.

By early 2008, Lehman held approximately \$85 billion in residential mortgage-backed securities — the largest component of a \$111 billion total real estate portfolio¹ — on a balance sheet leveraged at roughly 30:1.² That was a credit concentration problem. But it was also a market risk problem — those positions were marked to market, and falling valuations eroded the capital buffer. The capital erosion triggered a liquidity problem, as counterparties demanded additional collateral and prime brokerage clients began withdrawing funds. The liquidity pressure created an operational problem, as the firm struggled to unwind illiquid positions in a market with no buyers. And the visible distress created a reputational problem, as confidence evaporated among the counterparties, clients, and creditors whose continued engagement Lehman needed to survive. Each risk fed the next. The credit losses drove the market losses. The market losses drove the liquidity crisis. The liquidity crisis drove the operational paralysis. The operational paralysis drove the reputational collapse. And the reputational collapse drove further counterparty withdrawals, accelerating the liquidity crisis back into the credit losses. The entire chain took less than six months from the Bear Stearns rescue in March to Lehman's bankruptcy filing on 15 September 2008 — \$639 billion in assets, the largest bankruptcy in American history.³

Lehman's risk function had assessed each of these risks. Mortgage-backed securities were in the credit risk inventory. Leverage was monitored. Liquidity was reported. Counterparty exposure was tracked. But each was assessed individually, in its own silo, using

its own models, reported to its own committee. No one had mapped the causal chains between them. No one had asked: if credit losses reach a threshold, what does that trigger in market risk? If market risk drives margin calls, what does that trigger in liquidity? If liquidity forces asset sales, what does that do to credit valuations? The interaction between the risks — the amplification, the feedback loops, the cascade — was invisible because the methodology did not look for it.

Chapter 9 (Assessment — Scoring, Multi-Dimensional Impact, and Data Quality) established the four-dimensional assessment framework for scoring individual risks. But a risk inventory where every risk is assessed in isolation — however sophisticated the scoring methodology — remains fundamentally incomplete. The heatmap produced by individual assessment does not show correlations. It does not show common exposures. It does not show which risks amplify each other under stress. ISO 31000 requires examination of “knock-on effects of particular consequences, including cascade and cumulative effects.”⁴ That requirement is the standards basis for everything in this chapter.

Risk interaction analysis is the systematic examination of how risks relate to each other — which risks trigger other risks, which risks amplify the impact of other risks, and where the institution’s aggregate exposure exceeds what any individual risk score would suggest. It operates through three complementary tools: the **risk interaction matrix**, which maps relationships across the entire material risk portfolio; **bow-tie analysis**, which traces causal chains for the most critical individual risks; and **concentration analysis**, which identifies where multiple risks share common drivers, exposures, or dependencies. Together, they transform the individually assessed risk inventory from Chapter 9 into the interconnected risk landscape that the Board and regulators actually need to see.

The Risk Interaction Matrix

The risk interaction matrix is a square matrix with the institution’s material risks listed on both axes. For each pair of risks, the matrix records whether a directional relationship exists: does Risk A trigger, amplify, or have no effect on Risk B? The matrix is not symmetric. Credit risk may trigger liquidity risk, but liquidity risk may not trigger credit risk — or it may, through a different mechanism. Both directions must be assessed independently.

The matrix serves three purposes. First, it identifies **risk clusters** — groups of risks that are connected through trigger or amplification relationships and could therefore crystallise together. Second, it reveals **cascade pathways** — chains where one risk triggers a second, which triggers a third, potentially reaching risks that appear unrelated to the original event. Third, it highlights **amplification mechanisms** — relationships where the crystallisation of one risk does not just trigger but actively worsens another.

In practice, for an institution with 30 material risks, the matrix contains 870 directional pairs (30 x 29). This is not a theoretical exercise run from a spreadsheet. It requires structured facilitation — the Risk Identification Lead works through the matrix with risk owners, CRO, and relevant specialists in dedicated sessions. The question for each cell is specific: “If Risk A materialises at its assessed impact level, does it trigger or amplify Risk B? Through what mechanism? With what time lag?”

Three categories of relationship are recorded:

| Relationship | Definition | Example |
|-------------------|--|---|
| Triggers | Crystallisation of Risk A directly causes Risk B to materialise | Credit losses on concentrated portfolio trigger counterparty confidence withdrawal (liquidity risk) |
| Amplifies | Crystallisation of Risk A increases the impact or likelihood of Risk B | Market volatility amplifies funding cost risk through collateral calls |
| Correlated | Risks A and B share a common driver and would crystallise simultaneously under the same conditions | Interest rate shock affecting both trading book positions and pension fund obligations |

The completed matrix produces a network map of the institution’s risk landscape. Risks with many outgoing trigger relationships are **propagation nodes** — their crystallisation cascades across the portfolio. Risks with many incoming relationships are **vulnerable nodes** — they are exposed to crystallisation from multiple sources. Both warrant additional analysis and more conservative assessment.

At a typical G-SIB, building the first interaction matrix will reveal that counterparty credit risk connects — through trigger or amplification relationships — to ten or more other material risks. Market risk typically connects to a similar number. These connections are not surprises to experienced practitioners. What is new is making them explicit, documented, and traceable. The interaction matrix turns intuition into auditable methodology.

The interaction matrix also connects directly to the **thematic stress and risk assessment** (TSRA) described in Chapter 7 (Bottom-Up Identification: Templates, RCSA, and Specialist Processes) as a complement to this chapter's analysis. Where the TSRA constructs hypothetical adverse scenarios at the operational level and traces which risks activate together, the interaction matrix provides the structural map of those activation pathways. The TSRA tests the matrix under stress; the matrix explains why the TSRA scenarios produce the cascades they do.

The matrix also directly informs the enterprise portfolio view established in Chapter 8 (Reconciliation and the Enterprise Portfolio View). The simultaneous crystallisation assessment described there — which risks would activate together under plausible stress scenarios — now has a structured analytical basis rather than relying solely on workshop judgement. The interaction matrix provides the mechanism mapping; the SWIFT scenarios from Chapter 6 (Top-Down Identification: Workshops, SWIFT, and Delphi) provide the triggering events; together they produce a rigorous assessment of correlated risk clusters.

Bow-Tie Analysis: Mapping the Causal Chain

For the institution's five to ten most critical risks — as determined by the materiality assessment in Chapter 9 — the risk interaction matrix provides the landscape, but a more granular tool is needed to map the full causal architecture of each individual risk. That tool is **bow-tie analysis**, catalogued in ISO 31010 Section B.21.⁵

A bow-tie diagram is a visual representation of the pathways from causes to consequences for a single risk event, with the controls that interrupt those pathways explicitly mapped at each stage. The diagram takes its name from its shape: causes fan out on the left, consequences fan out on the right, and the risk event sits at the centre — the knot of the bow tie.

The Left Side: Causes and Prevention

On the left side, every identified cause of the risk event is listed. For each cause, the mechanism by which it leads to the risk event is described — not just “market downturn” but the specific transmission: “market downturn reduces asset valuations below margin thresholds, triggering collateral calls that exceed available liquid assets.”

Preventive barriers are shown as vertical bars across each causal pathway. These are the controls that prevent the cause from leading to the risk event. For a counterparty credit risk bow-tie, preventive barriers might include: credit limits, collateral requirements, netting agreements, portfolio concentration limits, independent credit assessment.

Critically, the bow-tie then identifies **escalation factors** — conditions that could cause a preventive barrier to fail. A credit limit is a preventive barrier, but if the limit can be overridden by a single individual without independent approval, that override authority is an escalation factor. A collateral requirement is a preventive barrier, but if the collateral is correlated with the underlying exposure (as it was in many GFC-era structured products), that correlation is an escalation factor.

Escalation controls are then mapped against each escalation factor — the additional controls that prevent escalation factors from degrading barriers. Four-eyes approval for limit overrides. Independent collateral valuation. Correlation analysis for collateral-exposure relationships.

This layered structure — cause, preventive barrier, escalation factor, escalation control — is what distinguishes bow-tie analysis from a simple risk-and-control register. It forces the institution to examine not just whether controls exist, but whether the conditions under which those controls would fail have been identified and addressed.

The Centre: The Risk Event

The risk event itself sits at the centre of the diagram. It should be defined precisely — not “credit risk” but “counterparty default on OTC derivative portfolio exceeding \$500 million net exposure.” Precision matters because the causes and consequences change depending on how the event is specified.

The Right Side: Consequences and Mitigation

On the right side, all potential consequences of the risk event are listed — and this is where the four-dimensional framework from Chapter 9 directly applies. A major counterparty default has financial consequences (direct loss), regulatory consequences (reporting obligations, supervisory attention, potential enforcement action), reputational consequences (market confidence, counterparty willingness to trade), and customer/operational consequences (service disruption if the counterparty provides critical services).

Mitigating barriers are shown as vertical bars across each consequence pathway. These are the controls that reduce the severity of consequences after the risk event has occurred. For counterparty default: close-out netting, collateral liquidation, credit insurance, business continuity plans for operational dependencies.

Recovery controls are separately identified — controls that support recovery after the event. Insurance claims processes, legal recovery proceedings, replacement counterparty sourcing, regulatory communication protocols.

The Supporting Layer

Beneath the bow-tie sits the **management function layer** — the training, inspection, maintenance, and testing activities that support the effectiveness of each barrier. A preventive barrier is only as reliable as the management system that maintains it. If credit limits are reviewed annually but market conditions change quarterly, the management function (review frequency) is inadequate for the barrier (credit limit) it supports.

Why Bow-Ties Work for Board Communication

Bow-tie diagrams serve a dual purpose. For risk practitioners, they provide the granular causal analysis that informs control design and testing. For the Board and senior management, they provide a visual summary that communicates the full risk architecture in a single page. A Board member looking at a bow-tie for the institution's most critical risk can immediately see: what could cause this? What stops it? What could make those stops fail? If it happens, what are the consequences? What limits the damage?

This is why the methodology requires bow-ties for the top five to ten risks specifically — not for the full inventory. Bow-tie analysis is resource-intensive. Each diagram requires facilitated sessions with risk owners, control owners, and relevant specialists. But for the risks that could threaten the institution's viability, the investment is essential.

The diagrams are drawn directly from facilitated sessions — not from desk-based analysis. The Risk Identification Lead facilitates, but the content comes from the people who understand the risk: traders, operations managers, technology specialists, compliance officers, and risk analysts. The facilitation approach mirrors the SWIFT methodology from Chapter 6 — structured prompts, systematic coverage, challenge and dissent encouraged.

Lehman Brothers Through the Bow-Tie Lens

Consider what a bow-tie analysis of Lehman's credit concentration risk would have revealed. On the left side, the causes would include: mortgage market deterioration, mark-to-market accounting requirements on illiquid positions, leverage ratio creating thin capital buffer. The preventive barriers would include: position limits, VaR monitoring, stress testing, capital adequacy reporting. The escalation factors would include: Repo 105 transactions temporarily removing \$50 billion in assets from the balance sheet at each quarter-end⁶ — a mechanism that actively undermined the capital adequacy barrier by concealing the true leverage ratio. The escalation control — independent verification of reported leverage — was absent.

On the right side, the consequences would cascade across every dimension: financial (direct losses), regulatory (capital adequacy breach, supervisory intervention), reputational (counterparty confidence loss), and operational (inability to roll short-term funding). The mitigating barriers — access to Federal Reserve facilities, potential acquisition by a stronger institution — would have been assessed as uncertain at best. The bow-tie would have made visible what individual risk assessments concealed: the causal chain from credit concentration through leverage to liquidity to institutional collapse was not a remote scenario. It was a mapped pathway with identified control gaps.

Fault Tree Analysis: Working Backwards from Failure

Where bow-tie analysis maps from causes through the risk event to consequences, **fault tree analysis** works backwards from a defined failure event to identify all the combinations of conditions that could produce it. ISO 31010 catalogues it alongside bow-tie analysis, and for complex risks with multiple potential failure modes, it provides a complementary perspective.

A fault tree uses Boolean logic — AND gates and OR gates — to decompose a top-level failure into its constituent conditions. An OR gate means any one of several conditions can cause the failure. An AND gate means multiple conditions must be present simultaneously.

The practical application in risk identification is for risks where the failure requires a combination of factors. Rogue trading, for example, is an AND-gate failure: it requires a trader with knowledge of the control chain AND inadequate independent verification AND RCSA that assesses controls against design specification rather than determined circumvention (the three structural preconditions established in Chapter 7). A fault tree makes explicit that the risk materialises only when all three conditions coexist — which means that addressing any one condition breaks the chain.

Fault trees are more analytical than bow-ties and less intuitive for Board communication. They are most valuable for operational risks with complex failure modes — technology failures, process breakdowns, fraud scenarios — where understanding the logical structure of the failure helps prioritise control investment. The methodology does not require fault trees for all material risks, but the Risk Identification Lead should deploy them where the logical structure of a risk warrants it.

Concentration Analysis: The Risks Between the Risks

Chapter 8 established the enterprise portfolio view with its four assessments: common exposures, simultaneous crystallisation, aggregate position against appetite, and diversification and correlation. Concentration analysis in Chapter 10 (Risk Interaction: Bow-Ties, Matrices, and Concentration) provides the systematic analytical method for the first of those four — identifying where multiple risks share common drivers, exposures, or dependencies.

Concentration takes three forms, each requiring a different analytical approach:

Single-name concentration is the most visible: excessive exposure to a single counterparty, sector, geography, or asset class. The interaction matrix often reveals single-name concentration that individual risk assessments missed, because the exposure appears in different taxonomy categories. A bank may have credit exposure to a major technology firm, market risk from equity positions in the same firm, operational dependency on the firm's cloud infrastructure, and conduct risk from the firm's role as a distribution partner. Four risks, four taxonomy categories, one underlying concentration. The interaction matrix, cross-referenced with the enterprise portfolio view, makes this visible.

Structural concentration is less visible but often more dangerous. This is concentration not in a single name but in a structural assumption — a funding model, a correlation assumption, a regulatory interpretation, a technology platform. Fortis provides a defining example. In 2007, Fortis led the consortium that acquired ABN AMRO for €71 billion — the largest banking acquisition in history at that point.⁷ The acquisition was funded by leverage, reducing Fortis's capital buffers. Simultaneously, Fortis held subprime exposure across both its legacy portfolio and the acquired ABN AMRO assets. Risk identification had assessed the acquisition risk and the subprime risk separately. What it had not identified was the **interaction**: the acquisition reduced the capital buffer at precisely the moment subprime losses required that capital. Each risk was individually manageable. Together, they created a compounding spiral. Fortis required an €11.2 billion government rescue and was broken up — the largest corporate failure in Benelux history.⁸ The two risks were not merely correlated. They amplified each other through a direct mechanism: the acquisition consumed the capital that the subprime losses subsequently demanded.

Systemic concentration extends beyond the institution to the system in which it operates. The Icelandic banking collapse of October 2008 is the most extreme illustration. Kaupthing, Landsbanki, and Glitnir collectively grew to ten times Iceland's GDP through aggressive international expansion funded by wholesale markets.⁹ Each bank's risk assessment focused on its individual balance sheet. Each assessed its own liquidity, its own credit quality, its own capital adequacy. What no individual bank's risk function identified was the systemic concentration: the entire banking system had grown beyond the fiscal capacity of the sovereign to backstop. There was no credible lender of last

resort at the scale required. When wholesale funding markets froze globally, all three banks failed within a single week. Iceland imposed capital controls, the currency collapsed, and the country required an IMF bailout.¹⁰

The Icelandic failure is particularly instructive because the concentration was not in any individual institution's risk inventory. It existed at the system level — between the banking system's size and the sovereign's fiscal capacity. Neither the banks nor the regulator had identified it. The Central Bank of Iceland's financial stability reports assessed individual bank metrics without aggregating to the system-level question: can this country backstop its banking system? Concentration analysis within the methodology requires asking not just "where are our internal concentrations?" but "what system-level concentrations is our institution embedded within?" This connects directly to the PESTLE assessment from Chapter 5 (Setting the Context: External, Internal, and Risk Culture) — the external context assessment should identify system-level vulnerabilities that individual risk assessments cannot see.

Identifying Hidden Concentration

The most dangerous concentrations are those disguised by apparent diversification. HSH Nordbank, referenced briefly in Chapter 8, believed it was diversified because its portfolio combined shipping finance — where it was the world's largest lender — with US subprime structured products. Two different asset classes, two different geographies, two different risk categories. But both were correlated to global economic activity. When global trade collapsed, shipping values fell. When the US housing market collapsed, structured product values fell. Both happened simultaneously because both were driven by the same underlying factor: the global economic contraction of 2008-2009. The €10 billion state guarantee that followed demonstrated the cost of false diversification.¹¹

Concentration analysis must therefore go beyond taxonomy-level aggregation. It requires identifying the **underlying drivers** — the economic, market, political, and operational factors that connect risks across taxonomy categories. The driver fields in the standardised template from Chapter 7 are the raw material. When the same driver appears across multiple risks in different taxonomy categories, that is a concentration signal regardless of the apparent diversification of the portfolio.

Cost-Benefit Assessment and the ALARP Principle

Risk interaction analysis identifies how risks relate to each other. It does not, by itself, determine what to do about them. For material risks where new controls or mitigants are being considered — particularly where the interaction analysis has revealed cascade pathways or amplification mechanisms requiring additional barriers — the methodology requires a **cost-benefit assessment** of proposed risk responses.

ISO 31010 Section B.30 provides the framework.¹² The approach is straightforward in principle:

1. **Quantify the cost** of the proposed control or mitigant — both implementation and ongoing operation
2. **Estimate the risk reduction** achieved — expressed as change in residual risk score and, where possible, reduction in expected loss
3. **Where quantification is possible**, express the comparison as a net present value or benefit-cost ratio
4. **Where quantification is not feasible** — and for many operational, reputational, and conduct risks it is not — produce a qualitative trade-off table comparing costs, benefits, and residual uncertainties

The cost-benefit assessment is not a decision rule. It is an input to decision-making. A negative NPV does not automatically mean the control should not be implemented — regulatory requirements, reputational considerations, and Board risk appetite may all override a purely financial calculation. But without the assessment, the institution is making control investment decisions without evidence.

The assessment is governed by the **ALARP principle** — As Low As Reasonably Practicable — borrowed from safety engineering and well-established in operational risk management:

- **Above the upper threshold:** the risk is intolerable. It must be reduced regardless of cost, unless reduction is genuinely impracticable. For a bank, risks above the upper threshold are those that threaten viability, breach regulatory requirements, or exceed Board-approved appetite boundaries.
- **Within the ALARP region:** the risk should be reduced if the cost of reduction is not grossly disproportionate to the benefit. This is where cost-benefit analysis does its

work. “Grossly disproportionate” sets a deliberately high bar — the cost must substantially exceed the benefit, not merely exceed it marginally.

- **Below the lower threshold:** the risk is broadly acceptable and need only be monitored to ensure it remains at this level. The risk criteria established in Phase 1 (Chapter 5) define this threshold.

ALARP prevents two opposite errors. It prevents the institution from tolerating risks that should be reduced because the cost of reduction appears high. And it prevents the institution from investing disproportionately in reducing risks that are already at acceptable levels — resources that could be deployed more effectively against the risks that the interaction analysis has identified as cascade nodes or amplification points.

The cost-benefit assessment is documented in the risk profile for each material risk and reviewed as part of the quarterly refresh cycle described in Chapter 13 (The Ongoing Cycle: Refresh, Events, and Audit). This ensures that the assessment evolves as conditions change — a control that was cost-effective when market conditions were stable may become essential when the interaction matrix reveals new cascade pathways under changed circumstances.

Bringing Interaction Analysis Together

The three tools — risk interaction matrix, bow-tie analysis, and concentration analysis — are complementary, not alternative. The interaction matrix maps the landscape of relationships across the full material portfolio. Bow-tie analysis provides granular causal architecture for the most critical individual risks. Concentration analysis identifies where multiple risks share common drivers or dependencies that taxonomy-level assessment would miss.

Together with the cost-benefit assessment, they complete the second half of Phase 3. Every material risk now carries not just its individual four-dimensional score from Chapter 9, but a documented set of interaction relationships, causal chain mapping for the most critical risks, identified concentrations, and evidence-based assessment of control proportionality.

The outputs feed directly into the documentation and inventory that Chapter 11 (Documentation: The Living Risk Inventory) will describe. The risk interaction matrix becomes a standing document, updated with each quarterly refresh. Bow-tie diagrams for the top risks are maintained as living documents, with barriers and escalation factors reas-

sessed as controls change. Concentration analysis informs both the enterprise portfolio view (Chapter 8) and the Board's principal risk report (Chapter 3 (Governance: Who Owns What)). And the cost-benefit assessments provide the evidence base for risk response decisions that the Board must approve.

Consider, finally, what these tools would have revealed at Lehman Brothers. The risk interaction matrix would have shown credit risk triggering market risk through mark-to-market, market risk triggering liquidity risk through collateral calls, liquidity risk triggering operational risk through forced asset sales, and reputational risk amplifying liquidity risk through counterparty withdrawal — a documented cascade chain. The bow-tie for credit concentration would have identified Repo 105 as an escalation factor undermining the capital adequacy barrier. Concentration analysis would have flagged the single-sector, single-asset-class exposure across mortgage-backed securities. And the ALARP assessment would have identified the leverage ratio as above the upper threshold — intolerable regardless of cost of reduction.

None of this required information that was unavailable. The mortgage positions were known. The leverage was reported. The Repo 105 transactions were documented — they were even reviewed by external auditors. What was missing was the systematic methodology to connect the individually known facts into a coherent picture of interacting risks. That is what risk interaction analysis provides.

What Comes Next

Phase 3 is now complete. Every risk in the material inventory has been scored across four dimensions, overlaid with data quality ratings, separated into inherent and residual assessments, tested for materiality, and analysed for interactions, concentrations, and causal chains. The analytical work of risk identification — from the foundation setting of Phase 1 through the dual-track identification of Phase 2 to the assessment and interaction analysis of Phase 3 — has produced a comprehensive, interconnected risk landscape.

But analysis without documentation is ephemeral. The risk inventory must be captured in a living document that preserves the full analytical record — not just the scores, but the evidence, the judgements, the disagreements, the interaction relationships, and the

control assessments that produced them. That documentation is the subject of Chapter 11: the fourteen-field risk inventory, the risk profiles for material risks, and the standards that ensure the inventory remains a decision tool rather than a compliance artefact.

1. Report of Anton R. Valukas, Examiner, *In re Lehman Brothers Holdings Inc.*, United States Bankruptcy Court, Southern District of New York, 11 March 2010, Volume 1, Section III.A.1, pp. 2–3. The Examiner documented Lehman's residential mortgage-backed securities holdings at approximately \$85 billion within a total real estate portfolio of \$111 billion as of late 2007.
2. Valukas Examiner's Report, Volume 1, Section III.A.4, pp. 46–59. Lehman's net leverage ratio was approximately 16:1 by its own reported metrics, but the gross leverage ratio — before netting — exceeded 30:1. See also Financial Crisis Inquiry Commission, *The Financial Crisis Inquiry Report*, U.S. Government Printing Office, January 2011, pp. 177–178.
3. Lehman Brothers Holdings Inc. filed for Chapter 11 bankruptcy protection on 15 September 2008 in the United States Bankruptcy Court, Southern District of New York (Case No. 08-13555). Total assets of \$639 billion were listed in the filing, making it the largest bankruptcy in U.S. history. See FCIC Report, January 2011, p. 339.
4. ISO 31000:2018, *Risk management — Guidelines*, Section 6.4.3 (Risk analysis), which requires consideration of "knock-on effects of particular consequences, including cascade and cumulative effects."
5. ISO 31010:2019, *Risk management — Risk assessment techniques*, Section B.21 (Bow tie analysis), pp. 81–84.
6. Valukas Examiner's Report, Volume 3, Section III.A.4 (Repo 105), pp. 732–1027. The Examiner found that Lehman used Repo 105 transactions to temporarily remove approximately \$50 billion in assets from its balance sheet at the end of the first and second quarters of 2008, reducing reported net leverage. Ernst & Young, Lehman's external auditor, was aware of the Repo 105 programme.
7. The Royal Bank of Scotland-led consortium (RBS, Fortis, and Santander) completed the acquisition of ABN AMRO on 10 October 2007 for approximately €71 billion (\$98 billion), the largest banking takeover in history at that time. Fortis's share was approximately €24 billion. See FCIC Report, January 2011, p. 157; see also Dutch Parliamentary Commission on the Financial System (De Wit Commission), *Final Report*, April 2010.
8. On 28 September 2008, the governments of Belgium, the Netherlands, and Luxembourg injected €11.2 billion into Fortis to prevent collapse. Fortis was subsequently broken up, with the Dutch banking and insurance operations nationalised and the Belgian banking operations sold to BNP Paribas. See De Wit Commission Final Report, April 2010.
9. Special Investigation Commission (SIC) of the Icelandic Parliament (*Althingi*), *Report of the Special Investigation Commission*, 12 April 2010, Chapter 2 (The Growth of the Banks). Combined assets of Kaupthing, Landsbanki, and Glitnir reached approximately ten times Iceland's GDP by end-2007.
10. All three Icelandic banks were placed into receivership in the first week of October 2008: Glitnir (29 September), Landsbanki (7 October), and Kaupthing (9 October). Iceland imposed emergency capital controls on 28 November 2008 and entered into a Stand-By Arrangement with the International Monetary Fund on 19 November 2008 (\$2.1 billion programme). See SIC Report, Chapter 17; IMF Country Report No. 08/362, November 2008.
11. HSH Nordbank received a €10 billion guarantee from the German states of Hamburg and Schleswig-Holstein in 2009, plus a €3 billion capital injection from the Financial Market Stabilization Fund (*SoFFin*). HSH Nordbank was the world's largest shipping finance lender. See European Commission State Aid Decision SA.29338, 20 September 2011.
12. ISO 31010:2019, *Risk management — Risk assessment techniques*, Section B.30 (Cost/benefit analysis), pp. 107–109.

Documentation: The Living Risk Inventory

The Risk Register That Nobody Read

In 2018, when Denmark's largest financial scandal finally became undeniable, investigators at Danske Bank confronted a familiar paradox. The bank had a risk inventory. It had risk policies. It had documented governance frameworks and named risk owners across its operations. Yet approximately €200 billion in suspicious transactions had flowed through its Estonian branch over nearly a decade¹, processed by non-resident customers — primarily from Russia and former Soviet states — through a small peripheral operation that represented a fraction of the group's balance sheet.

The Estonian branch had its own risk documentation. Group-level compliance had its own risk documentation. The problem was not that risks were undocumented. The problem was that the documentation existed in fragments that no one had assembled into a single, coherent picture. The branch's non-resident portfolio was documented locally as a profitable niche business. At group level, the Estonian operation appeared as a small but high-performing subsidiary. Nowhere in Danske Bank's risk inventory did a single entry read: *concentrated high-risk AML exposure in a peripheral branch operating under disconnected compliance oversight, generating returns inconsistent with legitimate transaction volumes*. The risk was identified in pieces. It was documented in pieces. And because the documentation never forced those pieces into a single inventory with enterprise-level fields, the institution processed €200 billion in suspicious flows before the outside world forced a reckoning. The CEO resigned. Fines exceeded €2 billion. The Estonian licence was withdrawn. The stock price halved.²

Danske Bank did not lack documentation. It lacked a **living risk inventory** — a single, structured, enterprise-wide register that imposes enough discipline on documentation that risks identified in one part of the institution cannot be invisible to the people responsible for the whole.

This chapter describes what that inventory looks like, how it is built, and — most importantly — how it is maintained so that it remains a decision tool rather than a compliance artefact.

Phase 4 in Context

The methodology's first three phases have now completed their work. Phase 1 (Chapter 5 (Setting the Context: External, Internal, and Risk Culture)) established the context — external environment, internal context, risk culture, risk criteria, starting universe, and straw man risk list. Phase 2 (Chapters 6 (Top-Down Identification: Workshops, SWIFT, and Delphi) – 8 (Reconciliation and the Enterprise Portfolio View)) identified risks through the dual-track process — top-down workshops using SWIFT and Delphi, bottom-up templates and specialist sub-processes, reconciliation, and the enterprise portfolio view. Phase 3 (Chapters 9 (Assessment — Scoring, Multi-Dimensional Impact, and Data Quality) – 10 (Risk Interaction: Bow-Ties, Matrices, and Concentration)) assessed those risks — four-dimensional scoring, data quality ratings, materiality determination, risk interaction analysis, bow-ties, and concentration analysis.

Phase 4 is where all of that work is captured in a structured, auditable, and — critically — maintainable form. Documentation is not an administrative afterthought. It is the mechanism through which identification becomes institutional knowledge. Without it, the work of Phases 1 through 3 exists only in the minds of the people who did it, and institutional memory in banking has a half-life measured in bonus cycles.

ISO 31010 requires a “structured record not just of risks identified but also of information used, assumptions made, and limitations.”³ That transparency requirement is the standards basis for everything in this chapter. The risk inventory is how you deliver it.

The Risk Inventory: Fourteen Fields

The risk inventory is the central register of all identified risks. It is a living document, not a point-in-time snapshot. Every risk that survives the identification and assessment process — whether surfaced through the top-down workshops, the bottom-up templates, the specialist sub-processes, or the reconciliation — enters the inventory with a complete set of structured fields.

I specify fourteen fields. This is not arbitrary. Each field exists because its absence has, in documented cases, contributed to a failure of risk identification or risk management. Some of these fields were established in Chapter 7 (Bottom-Up Identification: Templates, RCSA, and Specialist Processes)'s standardised risk assessment template (eleven fields for bottom-up submissions) and Chapter 9's assessment record (seven fields). The inventory integrates and extends both, creating a single authoritative record for each risk.

The Fourteen Fields

| # | Field | Description | Source Phase |
|---|--------------------------------|---|--------------|
| 1 | Risk ID | Unique identifier, persistent across cycles | Phase 2 |
| 2 | Taxonomy Classification | L1 / L2 / L3 classification per institutional taxonomy | Phase 2 |
| 3 | Risk Definition | Plain-language description of the risk, its nature, and scope | Phase 2 |
| 4 | COSO Objective Category | Strategic / Operations / Reporting / Compliance | Phase 1 |
| 5 | Risk Owner | Named individual — not a committee, not a function | Phase 2 |
| 6 | | | Phase 3 |

| # | Field | Description | Source Phase |
|----|---------------------------------|---|--------------|
| | Inherent Risk Score | Four-dimensional: Impact (with dominant dimension noted), Likelihood, Vulnerability, Speed of Onset | |
| 7 | Control Summary | Key controls with type (preventive/detective/corrective) and assessed effectiveness (1-5 scale) | Phase 2-3 |
| 8 | Residual Risk Score | Four-dimensional score after control effectiveness applied | Phase 3 |
| 9 | Material (Y/N) | Whether the risk exceeds the materiality threshold defined in Phase 1 | Phase 3 |
| 10 | Data Quality Rating | High / Medium / Low / Very Low — with implications for conservatism adjustments | Phase 3 |
| 11 | Risk Interaction Summary | Cross-references to correlated, triggering, or amplifying risks from the interaction matrix | Phase 3 |
| 12 | Key Risk Indicators | KRIs with current values and thresholds (green / amber / red) | Phase 2-3 |
| 13 | Trend Indicator | Increasing / Stable / Decreasing — direction of travel since last assessment | Phase 3 |
| 14 | Date of Last Review | When this risk was most recently assessed, with next scheduled review date | Ongoing |

Every field serves a purpose. Remove any one and you create a gap that, under pressure, becomes the path of least resistance for the institution to ignore what it has found.

Why Each Field Matters

Risk ID provides persistence. When a risk is re-assessed in the quarterly cycle (Chapter 13 (The Ongoing Cycle: Refresh, Events, and Audit)), the Risk ID connects the current assessment to its history. Without it, the same risk can appear as “new” each cycle, losing trend data and making it impossible to track whether the institution’s exposure is growing or shrinking.

Taxonomy Classification ensures the risk is findable. If an analyst searches the inventory for all L1 Credit Risks, every credit risk must appear. If a risk is misclassified — or worse, unclassified — it falls outside the enterprise portfolio view. The JPMorgan London Whale failure, where the CIO’s positions were classified as hedging rather than proprietary trading, demonstrates what happens when classification determines oversight.⁴

Risk Definition forces clarity. A risk defined as “market conditions” tells the reader nothing. A risk defined as “a sustained increase in credit spreads on the institution’s senior unsecured debt funding, triggered by rating agency action or counterparty confidence loss, increasing wholesale funding costs beyond the level assumed in the treasury plan” tells the reader exactly what is at stake, what drives it, and why it matters. The definition must be specific enough that two independent analysts would agree on whether a given event constitutes a materialisation of this risk.

COSO Objective Category (Strategic, Operations, Reporting, Compliance) prevents the institution from identifying only financial risks while ignoring strategic, reporting, and compliance exposures. Chapter 4 (The Risk Taxonomy) established the COSO cube as a completeness check. The inventory encodes that check at the individual risk level.

Risk Owner means a named individual. Chapter 3 (Governance: Who Owns What) established this principle: “committee ownership dilutes accountability.” The inventory enforces it. When a risk materialises, the Board must be able to identify who was accountable for monitoring it. When an owner changes role, the inventory must be updated within the same cycle — orphaned risks are invisible risks.

Inherent Risk Score captures the four-dimensional assessment from Chapter 9: Impact (with the dominant dimension — financial, regulatory, reputational, or customer/operational — noted), Likelihood, Vulnerability, and Speed of Onset. Recording all four dimen-

sions, not just a single composite score, preserves the analytical richness that the assessment methodology was designed to produce. A risk scored Moderate on financial impact but Extreme on regulatory impact tells a different story from a composite “High.”

Control Summary records not just what controls exist but what type they are (preventive, detective, corrective) and how effective they are (the 1–5 scale from Chapter 9). Chapter 7 identified that one of the three most common bottom-up submission failures is controls “defaulting to effective without evidence.” The inventory field demands specificity.

Residual Risk Score is the four-dimensional score after controls. The gap between inherent and residual scores reveals control dependency — how much the institution is relying on controls to manage this risk. A large gap with a control effectiveness rating of 3 or below is a signal that warrants attention.

Material (Y/N) is the gatekeeper. Material risks — typically twenty to sixty in a large institution — receive the full treatment: dedicated risk profiles, bow-tie diagrams, assigned KRIs, capital planning integration, and Board reporting. The materiality threshold was defined in Phase 1 (Chapter 5) and applied in Phase 3 (Chapter 9). The inventory records the outcome.

Data Quality Rating is one of the most important fields in the inventory, precisely because it is the one institutions most want to ignore. A rating of “Very Low” attached to a risk scored as non-material creates discomfort — it means the institution has classified something as unimportant based on evidence it does not trust. Chapter 9 established three operational consequences for low data quality: conservatism adjustments, sensitivity testing, and Board transparency. The inventory is where those consequences become visible.

Risk Interaction Summary captures the output of Chapter 10’s interaction matrix — which other risks this risk can trigger, amplify, or be triggered by. Without this field, the inventory is a list of independent items. With it, the inventory becomes a map of the institution’s interconnected risk landscape. The Lehman Brothers cascade — credit triggering market triggering liquidity triggering operational triggering reputational⁹ — was invisible precisely because no documentation connected individually assessed risks into causal chains.

Key Risk Indicators with current values and traffic-light thresholds transform the inventory from a periodic assessment document into a continuous monitoring tool. A KRI in amber is an early warning. A KRI in red demands immediate investigation. Without KRIs attached to specific risks in the inventory, monitoring becomes a separate activity disconnected from identification — and the gap between identification and action widens.

Trend Indicator answers the question the Board asks most often: “Is this getting better or worse?” An Increasing trend on a material risk with amber KRIs and a Data Quality Rating of Low is a combination that should trigger immediate escalation. No single field in isolation produces that signal. The inventory structure produces it by placing all relevant information in one record.

Date of Last Review creates accountability for currency. A risk last reviewed fourteen months ago in an institution with quarterly cycles has been missed. The field makes that failure visible to anyone who opens the inventory — including Internal Audit, which Chapter 3 established as providing independent assurance over the process itself.

The Compliance Artefact Problem

If the fourteen fields sound like a bureaucratic exercise, consider what happens without them.

Chapter 7 described **compliance theatre** — the appearance of risk identification without the substance. Bottom-up submissions that roll forward prior year without genuine analysis. Risk definitions copied verbatim from one cycle to the next. Control effectiveness ratings that never change. The inventory equivalent of compliance theatre is a risk register with the right number of columns but the wrong kind of content: generic definitions, stale assessments, no interaction data, no KRI values, and review dates that tell you the register was updated on the same day by the same person for every risk in the book.

This pattern is common. At many institutions, the reconciliation documentation from a prior cycle is a merged spreadsheet with colour-coding but no gap analysis, no challenge sessions, and no enterprise portfolio view. The institution has two identification tracks. It does not have reconciliation. And the documentation — which should make the absence of reconciliation visible — instead conceals it, because the spreadsheet looks complete to anyone who does not know what to look for.

The fourteen-field structure is designed to make the absence of substance visible. An empty Risk Interaction Summary field is a gap you can see. A Data Quality Rating of “Medium” applied to every risk in the inventory is a pattern that Internal Audit can challenge. A Trend Indicator that has been “Stable” for twelve consecutive quarters on a risk whose external environment has materially changed is a flag. The structure creates the conditions for accountability.

Risk Profiles: The Material Risk Deep Dive

The fourteen-field inventory captures every identified risk. For the twenty to sixty risks classified as material, the inventory entry alone is not sufficient. Each material risk requires a **risk profile** — a structured narrative that provides the depth of analysis necessary for Board-level decision-making, capital planning integration, and regulatory engagement.

The risk profile contains fourteen elements:

- 1. Risk definition** — the same plain-language description as the inventory, expanded where necessary to capture nuance that a single field cannot convey
- 2. Taxonomy classification** — L1 / L2 / L3
- 3. COSO objective category** — Strategic / Operations / Reporting / Compliance
- 4. Underlying drivers** — direct and indirect causes, using Ishikawa analysis where appropriate; this is where the driver fields from the bottom-up template (Chapter 7) are synthesised with the top-down analysis from workshops
- 5. Current exposure** — quantitative metrics where available; dollar exposure, notional amounts, portfolio concentrations, customer counts
- 6. Risk appetite** — the Board-approved appetite for this risk and the current position against it; breaches highlighted
- 7. Key controls** — expanded from the inventory’s Control Summary; includes control type, effectiveness rating, evidence basis for the rating, and any control gaps identified in bow-tie analysis (Chapter 10)
- 8. Cost-benefit assessment** — for risks where new controls have been proposed or existing controls are being reconsidered, the proportionality assessment from Chapter 10

- 9. Key Risk Indicators** — leading indicators with thresholds, current values, and trend; expanded from the inventory's KRI field to include data source, refresh frequency, and escalation protocol
- 10. Scenario linkage** — which stress scenarios (ICAAP, ILAAP, CCAR) map to this risk; how the risk's assessment scores inform scenario severity; this is the bridge between risk identification and capital planning that Chapter 12 (Integration: Capital Planning, Strategy, and the Board) will describe
- 11. Data quality rating** — the confidence level of the assessment, what evidence supports it, and specifically what additional data or analysis would improve it
- 12. Enterprise interactions** — the full interaction analysis from Chapter 10: which other risks this risk can trigger, amplify, or be triggered by; bow-tie diagram reference; concentration analysis findings
- 13. Trend and outlook** — direction of travel with forward-looking assessment; not just "Increasing" but why it is increasing, what would reverse the trend, and what the risk owner is doing about it
- 14. Risk owner** — the named individual accountable, with escalation path to CRO and Board Risk Committee

The risk profile is not a form to be filled in. It is a structured analysis. Introducing risk profiles for the material risk inventory typically follows a three-cycle maturation. In the first cycle, initial submissions from some business units will be precisely the compliance theatre described above — boilerplate definitions, controls listed as "effective" without evidence, and outlook sections that read like reassurance rather than analysis. The first cycle requires significant push-back from the Risk Identification Lead. By the second cycle, business units understand that risk profiles with thin content will be challenged and returned. By the third cycle, quality improves materially — not because people have become better risk analysts overnight, but because the structure and the challenge process create accountability for substance.

Risk Profiles as Board Communication

The risk profile serves dual purposes. It is a practitioner tool — detailed enough for the risk owner and the Risk Identification Lead to manage the risk actively. It is also a Board communication tool. The principal risk report, which Chapter 3 defined with ten specific contents, draws directly from these profiles. When the Board Risk Committee asks

“What is our current position on cyber risk?” — the risk profile provides the answer: the definition, the exposure, the appetite position, the controls, the KRIs, the trend, and the interactions with other risks.

A Board member reading a well-constructed risk profile should understand within five minutes: what the risk is, how bad it could get, what is being done about it, how confident the institution is in its assessment, and what has changed since the last report. If the profile cannot deliver that in a structured one-to-two-page document, it needs to be rewritten.

Bow-Tie Integration

For the most critical material risks — typically the five to ten risks that the interaction matrix identified as propagation nodes or that received Extreme ratings on any dimension — the risk profile includes or references a full bow-tie diagram. Chapter 10 described bow-tie analysis in detail: causes on the left, preventive barriers, escalation factors, the precisely defined risk event at the centre, consequences mapped to the four-dimensional framework on the right, mitigating barriers, and recovery controls.

The bow-tie diagram is referenced in the risk profile rather than reproduced in full, because bow-ties are living documents that are updated as controls change. The profile records the date of the most recent bow-tie review, the key findings, and any escalation factors or barrier weaknesses that have been identified. This ensures the profile remains current without duplicating the bow-tie’s content.

The Audit Trail: Every Change Documented

The risk inventory is not a static document that is produced once and filed. It is updated through every quarterly re-identification cycle (Chapter 13), every event-driven update, and every time a KRI breaches a threshold. Each of those updates creates a change to the inventory — a risk score revised upward, a new risk added, an existing risk reclassified, a risk owner changed, a control assessment modified.

Every change must be documented with three elements:

- 1. Date** — when the change was made

2. Author — who made the change (and under what authority — Risk ID Lead, risk owner, CRO approval)

3. Reason — why the change was made, with reference to the evidence that prompted it

This is not bureaucracy. It is the mechanism through which the inventory becomes auditable. When Internal Audit reviews the risk identification process — as Chapter 3 established it should — the audit trail is the evidence base. Can the auditor trace a risk from its first identification through every subsequent assessment, seeing how the score changed, who changed it, and why? If yes, the process has integrity. If no, the inventory is a document with no provenance.

The audit trail also creates institutional memory. When a new CRO arrives and asks “Why was this risk downgraded two quarters ago?” — the trail provides the answer. When a regulator asks “When did you first identify this risk, and how has your assessment evolved?” — the trail provides the answer. Without it, the institution is relying on the memory of individuals who may no longer be there.

The Disagreement Log and Assumption Register

Two outputs from earlier phases require specific documentation treatment in the inventory.

The **disagreement log**, introduced in Chapter 6 and referenced in the challenge sessions of Chapter 8, records instances where senior participants held materially different views about a risk’s nature, severity, or even existence. Chapter 9 described how the four-dimensional framework often resolves disagreement — participants may agree on different dimensions — but where disagreement persists, it must be preserved in the inventory. The minority view, the evidence basis for both positions, and the data quality ratings are all documented. The Board receives both views.

Disagreement is information. A risk where the CRO and the head of the investment bank genuinely disagree on severity after structured analysis is disproportionately likely to be one of the risks the institution most needs to understand. Averaging their views destroys that signal. The inventory preserves it.

The **assumption register**, also introduced in Chapter 6, records key assumptions that were identified and challenged during workshops. Assumptions about correlations, about the effectiveness of hedging strategies, about the stability of funding markets,

about the behaviour of counterparties under stress. Each assumption is linked to the risks it underpins. When the assumption is revisited in a subsequent cycle — as Chapter 13's ongoing process requires — the inventory provides the reference point.

The Inventory as Enterprise Infrastructure

The fourteen-field structure and the risk profile format are not just documentation standards. They are the infrastructure through which the risk identification process connects to the institution's management processes. Every downstream use of risk identification outputs depends on the inventory being complete, current, and structured.

Feeding the Principal Risk Report

Chapter 3 defined the principal risk report with ten specific items: material risks with scores, consequence dimensions, trends, appetite breaches, new or emerging risks, re-classified risks, KRI dashboard, data quality distribution, enterprise portfolio view, and process performance indicators. Every one of those items is drawn from the inventory. The principal risk report is not a separate analytical exercise. It is a presentation layer over the inventory. If the inventory is incomplete, the Board report is incomplete. If the inventory is stale, the Board is making decisions on outdated information.

Feeding Capital Planning

Chapter 12 will describe how the inventory integrates with ICAAP, ILAAP, and CCAR capital planning. The connection is direct: each material risk in the inventory maps to one or more stress scenarios. The inventory's four-dimensional scores inform scenario severity. The Data Quality Rating informs the level of conservatism applied to management judgement overlays. The scenario linkage field in the risk profile is where identification meets capital adequacy.

Feeding Regulatory Reporting

The regulatory mapping table, introduced in Chapter 4, translates the institution's internal taxonomy into each regulator's categories. The inventory is what gets translated. When the PRA examines the ICAAP, it assesses the quality and comprehensiveness of

the risk inventory. When the Fed reviews the CCAR Material Risk Inventory, it evaluates whether the quarterly re-identification has been substantive. The inventory is the artefact that regulators examine.

Feeding the Ongoing Cycle

Chapter 13 will describe the quarterly re-identification cycle, the annual full re-identification, and event-driven updates. Each of these processes takes the current inventory as its starting point — the prior quarter’s inventory is the baseline against which new identification is compared. Risks that were in the inventory but no longer appear must be explicitly retired, with documented reasons. New risks must be explicitly added, with documented sources. Changed assessments must show what changed and why. The inventory is the continuity mechanism that prevents each cycle from starting with a blank page.

What Good Looks Like — and What Bad Looks Like

Signs of a Living Inventory

A living inventory has observable characteristics:

- **Risk definitions vary in specificity and language** — they were written by different people at different times about genuinely different risks, not copied from a template
- **Scores change between quarters** — the external environment changes, controls are strengthened or weakened, new information emerges; a stable inventory in a changing world is a sign that nobody is looking
- **Data Quality Ratings are distributed** — some risks have High ratings (good data, validated models, extensive history) and some have Very Low (emerging risk, limited data, high uncertainty); an inventory where everything is rated Medium is an inventory where nobody has made a judgement
- **Interaction summaries are populated** — risks reference other risks; the inventory reads as a network, not a list
- **Review dates are staggered** — different risks were reviewed at different times, reflecting the quarterly cycle, event-driven updates, and the natural rhythm of an active process

- **The disagreement log has entries** — if no one has ever disagreed about any risk in the inventory, the institution is not challenging hard enough

Signs of a Compliance Artefact

A compliance artefact also has observable characteristics:

- **All definitions use the same sentence structure** — “the risk that [X] could [Y]” repeated two hundred times suggests a template-filling exercise, not genuine analysis
- **All scores are clustered in the middle of the scale** — Moderate impact, Possible likelihood, Medium data quality — the path of least resistance for a risk assessor who does not want to draw attention
- **No risk has an Extreme rating on any dimension** — the institution has identified nothing that could threaten its viability, which is either a remarkable achievement or a remarkable failure of candour
- **Control effectiveness is uniformly high** — every control is effective, which contradicts the operational experience of every institution that has ever existed
- **The Risk Interaction Summary field is blank** — risks exist in isolation, which they do not
- **Every risk was last reviewed on the same date** — the entire inventory was updated in a single batch, probably in the week before the Board Risk Committee meeting, probably by one person

Both exist in practice. The difference is not the structure — both may have fourteen fields. The difference is whether the structure is enforced through the challenge process (the Risk Identification Lead sending back non-compliant submissions, as Chapter 7 described), the governance framework (the Board Risk Committee asking substantive questions, as Chapter 3 described), and the audit process (Internal Audit testing quality, not just existence).

Documentation Failures in Practice

The Danske Bank case that opened this chapter illustrates documentation failure at the enterprise level — a fragmented inventory that never forced the institution to see its aggregate AML exposure. Two further cases illustrate different dimensions of the same problem.

AMP: When Complaints Data Is Not a Risk Indicator

AMP, one of Australia's largest financial services companies, systematically charged financial advice fees to customers who received no service — including deceased customers — over a prolonged period. The failures were eventually exposed through the Australian Royal Commission into Misconduct in the Banking, Superannuation and Financial Services Industry in 2018.⁵

AMP had risk identification frameworks. It had a risk register. What it did not have was an inventory that aggregated customer complaints as a risk indicator. Complaints data existed — customers had complained about being charged for services they did not receive. But the complaints were categorised as individual customer service issues, handled through operational channels, and never aggregated or escalated into the risk identification process. The risk register contained entries for conduct risk in generic terms. It did not contain an entry that said: *systemic pattern of fee charging without service delivery, evidenced by complaint volumes across multiple advice channels, representing regulatory, reputational, and customer detriment exposure.*

The inventory structure matters here. If AMP's risk inventory had required KRIs with thresholds attached to its conduct risk entries — and if one of those KRIs had been complaint volumes by category — the systemic pattern would have become visible as a red indicator long before the Royal Commission made it front-page news. If the Data Quality Rating had forced the institution to declare the evidence basis for its conduct risk assessment, the absence of aggregated complaints data would have been flagged as a gap. The CEO resigned. Remediation exceeded AUD 600 million. The share price halved.⁶

What was missing: An inventory structure that required granular KRIs attached to conduct risk entries, that aggregated complaints data as a risk indicator rather than treating individual complaints as operational noise, and that forced a Data Quality Rating disclosure on the evidence basis for conduct risk assessment.

SNS Reaal: When the Subsidiary Reports Separately

SNS Reaal, a Dutch banking and insurance group, was nationalised in February 2013 at a cost of €3.7 billion to the Dutch taxpayer.⁷ The cause was its property finance subsidiary, SNS Property Finance, which had accumulated €7.4 billion in commercial real estate loans that became severely impaired in the Dutch property downturn.⁸

The documentation failure was structural. SNS Property Finance operated with significant autonomy from the parent bank's risk framework. It managed its own risk identification. It reported its own risk inventory. The concentration risk it created relative to group capital was not identified at the enterprise level because the subsidiary's growth was managed and reported separately. The parent bank's inventory did not contain an entry that aggregated the subsidiary's property exposure against the group's capital base. The subsidiary's inventory did not contextualise its own exposure against the group's capacity to absorb losses.

This is a documentation architecture problem. If the group had maintained a single enterprise-wide inventory — or, at minimum, an inventory with explicit consolidation rules that required subsidiary exposures to be surfaced at group level — the concentration would have been visible. The fourteen-field structure addresses this directly: the Risk Interaction Summary field requires cross-references to correlated risks, and the enterprise portfolio view (Chapter 8) requires aggregation across all entity levels. The documentation was the gap between identification and action.

What was missing: An enterprise-level inventory that consolidated subsidiary risk exposures against group capital, with interaction summaries linking property concentration in the subsidiary to credit risk and capital adequacy at the parent level.

Inventory Governance

The inventory requires its own governance. Who can create a new entry? Who can modify an existing one? Who can retire a risk?

Creation

New risks enter the inventory through one of four routes:

1. **Top-down workshops** — risks identified through SWIFT and scenario analysis in the annual or quarterly cycle
2. **Bottom-up submissions** — risks identified through business unit templates, RCSA, and specialist sub-processes
3. **Reconciliation** — risks identified through the gap analysis between top-down and bottom-up tracks
4. **Event-driven updates** — risks surfaced by material events between scheduled cycles

In every case, the Risk Identification Lead is responsible for ensuring the new entry meets the fourteen-field standard before it enters the inventory. Incomplete entries are not accepted. A risk without a named owner, without a taxonomy classification, or without an initial assessment does not enter the inventory — it enters the Risk ID Lead's queue for completion.

Modification

Risk owners can propose modifications to their entries — updated scores, revised control summaries, new KRI values. All modifications require the Risk ID Lead's review before they become effective. This is not a bottleneck; it is quality assurance. The Risk ID Lead confirms that the modification is supported by evidence, that cross-references remain consistent, and that any change to a material risk is flagged for CRO attention.

Material changes — a risk moving from non-material to material, a score increasing by more than one level on any dimension, a Data Quality Rating being downgraded — require CRO approval and are reported to the Board Risk Committee in the next principal risk report.

Retirement

Risks do not simply disappear from the inventory. Retirement requires a documented justification: the exposure no longer exists (a business line has been exited, a counterparty relationship has been terminated), or the risk has been absorbed into another inventory entry through reclassification. The retirement reason is recorded, the retirement is approved by the Risk ID Lead (or CRO for material risks), and the retired entry remains in the inventory archive — visible for audit purposes and for trend analysis across cycles.

An inventory that shrinks without documented retirements should be treated with the same suspicion as a clean reconciliation with no gaps. Risks do not resolve themselves. If the inventory is getting smaller, either the institution is genuinely reducing its risk profile, or the process is losing track of risks it previously identified.

Technology and the Inventory

The fourteen-field structure can be implemented in a spreadsheet. In practice, many institutions build their initial inventory in a structured spreadsheet before investing in dedicated GRC (governance, risk, and compliance) technology. A well-structured spreadsheet with defined fields, data validation, change tracking, and version control can serve as a functional inventory for the early cycles of a risk identification programme.

But a spreadsheet has limitations that become material as the inventory matures. It does not enforce referential integrity — if a risk owner leaves and the name is updated in the inventory but not in the risk profile, there is no automatic check. It does not support the kind of network queries that the Risk Interaction Summary field demands — “show me all risks that are triggered by a liquidity event” requires either manual searching or custom formulae. It does not provide role-based access control — anyone with file access can modify any entry.

Chapter 14 (Technology: AI, ML, and Data Analytics) will address technology in detail. For the purposes of this chapter, the principle is: **the inventory structure must be designed independently of the technology that will host it**. Define the fourteen fields, the risk profile elements, the audit trail requirements, and the governance rules first. Then select or build the technology to deliver them. Institutions that start with the technology and let the tool define the structure end up with inventories shaped by software capabilities rather than risk management requirements.

The Living Document Standard

The title of this chapter uses the word “living” deliberately. A living document changes. It responds to new information. It reflects the institution’s current understanding, not its understanding six months ago. The mechanisms that keep the inventory alive are the subject of Chapter 13 — the quarterly re-identification cycle, the annual full re-identification, event-driven updates, and the internal audit assurance process. But the inventory’s design must anticipate those mechanisms.

Every field in the fourteen-field structure is designed to change. The Trend Indicator changes quarterly. The KRI values change as data is refreshed. The Data Quality Rating changes as evidence improves or deteriorates. The Risk Interaction Summary changes as the interaction matrix is updated. The inherent and residual scores change as the external environment evolves and controls are strengthened or weakened. A field that never changes is a field that is not being maintained.

The difference between a living inventory and a compliance artefact is not structural. It is cultural. It is whether the institution treats the inventory as its best current understanding of its risk landscape — updated, challenged, used — or as a document that satisfies a governance requirement with minimum effort. The methodology provides the structure. The governance framework (Chapter 3) provides the accountability. The ongoing cycle (Chapter 13) provides the rhythm. But the commitment must come from the institution itself.

Danske Bank had documentation. AMP had a risk register. SNS Reaal had subsidiary risk reporting. What none of them had was a living inventory — a structured, enterprise-wide, continuously maintained register that forced the institution to confront what it had found, connect what it had found to what it already knew, and act on the result.

Phase 4 is complete. The risk inventory now contains every identified risk with fourteen structured fields, every material risk with a detailed risk profile, and an audit trail that records every change. Chapter 12 turns to the question that gives all of this work its purpose: how the inventory integrates with the institution's capital planning, strategic planning, regulatory reporting, and Board governance — the mechanisms through which risk identification becomes institutional action.

1. Bruun & Hjejle (external law firm), *Report on the Non-Resident Portfolio at Danske Bank's Estonian Branch*, 19 September 2018, p. 11. The report estimated that approximately €200 billion in total transaction flows passed through the non-resident portfolio of the Estonian branch between 2007 and 2015, a significant portion of which was deemed suspicious.
2. Danske Bank CEO Thomas Borgen resigned on 19 September 2018 upon publication of the Bruun & Hjejle report. The Danish FSA withdrew the Estonian branch's authorisation effective February 2019. Danske Bank's share price fell from approximately DKK 230 in early 2018 to approximately DKK 115 by late 2018. Cumulative regulatory penalties exceeded €2 billion, including a \$2 billion settlement with the U.S. Department of Justice and SEC announced in December 2022.
3. ISO 31010:2019, *Risk management — Risk assessment techniques*, Section 4.7 (Documentation), which requires that risk assessment produce "a structured record not just of risks identified but also of information used, assumptions made, and limitations."
4. United States Senate Permanent Subcommittee on Investigations, *JPMorgan Chase Whale Trades: A Case History of Derivatives Risks and Abuses*, 15 March 2013. The report detailed how the Chief Investment Office's Synthetic Credit Portfolio, which generated losses exceeding \$6.2 billion, was classified and managed as a hedging portfolio rather than a proprietary trading position, placing it outside the risk oversight applied to the Investment Bank's trading activities.
5. Royal Commission into Misconduct in the Banking, Superannuation and Financial Services Industry (Commissioner Kenneth Hayne), *Final Report*, 4 February 2019, Volume 2, Section 2.3 (Fees for no service). The Royal Commission documented that AMP had charged ongoing advice fees to approximately 4,600 customers, including deceased customers, who received no financial advice services.
6. AMP CEO Craig Meller resigned on 20 April 2018 during the Royal Commission hearings. AMP disclosed total customer remediation provisions exceeding AUD 600 million across multiple regulatory findings. AMP's share price fell from approximately AUD 5.40 in early 2018 to approximately AUD 2.50 by end-2018. See also ASIC enforcement proceedings against AMP Financial Planning Pty Ltd.
7. Dutch Minister of Finance Jeroen Dijsselbloem announced the nationalisation of SNS Reaal on 1 February 2013 under the Intervention Act (*Intervetiewet*), at a direct cost to the Dutch state of €3.7 billion (comprising a €2.2 billion equity injection and €1.5 billion write-down of subordinated debt). See Dutch Ministry of Finance, *Letter to Parliament on the Nationalisation of SNS Reaal*, 1 February 2013.
8. SNS Property Finance held a commercial real estate loan portfolio of approximately €7.4 billion, of which a significant portion was concentrated in Dutch and international property developments that became severely impaired. See De Nederlandsche Bank, *Assessment of SNS Reaal*, January 2013; Evaluation Committee Nationalisation SNS Reaal (*Commissie Scheltema*), *Report*, February 2014.
9. The interconnected cascade of risk categories at Lehman Brothers is documented extensively in the Valukas Examiner's Report (March 2010), Volume 1, Sections III.A.1 through III.A.5, and in the FCIC Report (January 2011), Chapter 18 (The Bankruptcy of Lehman Brothers), pp. 324–343.

Integration: Capital Planning, Strategy, and the Board

The Bank That Was Rescued Twice

In September 2008, Dexia — the Franco-Belgian municipal lending giant — received its first government bailout. Belgium, France, and Luxembourg injected €6.4 billion in capital and provided €150 billion in state guarantees.¹ The institution's problem was familiar: large positions in structured credit products funded through short-term wholesale markets, a maturity mismatch that the credit crisis had made fatal.

What happened next is what makes Dexia's story relevant to this chapter.

After the first rescue, Dexia reconstituted its risk framework. It identified the risks that had caused its near-failure. It rebuilt its risk inventory. It documented the maturity mismatch vulnerability, the wholesale funding dependency, the concentration in structured products. By any reasonable measure, the institution had learned from its crisis. The risks were identified.

Three years later, in October 2011, Dexia required a second bailout.² This time the institution was broken up entirely — banking operations separated across Belgium, France, and Luxembourg, with the residual portfolio placed into orderly wind-down. The cause was sovereign debt concentration, particularly Greek and peripheral European bonds, funded through the same short-term wholesale markets that had nearly destroyed the institution three years earlier.

The risk identification outputs existed. The sovereign debt positions were visible. The wholesale funding dependency was documented. What was missing was the integration that would have connected those identification outputs to the institution's capital planning, its strategic decisions, and its Board governance. Dexia's reconstituted risk framework identified risks. It did not ensure that identified risks changed institutional behaviour.

This is the problem Phase 5 solves. The first four phases of the methodology produce a comprehensive, assessed, documented risk inventory. Phase 5 ensures that inventory becomes institutional action — through capital planning, strategic planning, regulatory reporting, and Board governance. Without Phase 5, the methodology produces a very sophisticated filing cabinet.

Why Integration Is the Test of the Methodology

At many institutions, the risk identification process is rigorous and the capital planning process is rigorous, and the two operate as parallel universes. A risk inventory may contain thirty-seven material risks while the ICAAP addresses twelve. The gap is not deliberate — it is structural. Different teams, different timelines, different reporting lines, different systems.

The inventory documented in Chapter 11 (Documentation: The Living Risk Inventory) is not an end in itself. Its fourteen fields, its risk profiles, its interaction summaries, its data quality ratings — all of this analytical work has value only to the extent that it reaches the processes where institutional decisions are made. Capital allocation. Strategic direction. Regulatory submissions. Board oversight.

ISO 31000 Principle b states that risk management must be “an integral part of all organisational processes.”³ COSO ERM Component 7 — Information and Communication — requires that relevant risk information is “identified, captured, and communicated in a form and timeframe that enables people to carry out their responsibilities.”⁴ BCBS Corporate Governance Principle 7 requires that risk identification outputs reach the Board and inform capital adequacy.⁵ These are not aspirational statements. They are the standards against which regulators assess whether the risk identification process is functioning.

Phase 5 has five integration points: COSO ERM alignment, capital planning, strategic planning, regulatory reporting, and Board reporting. Each connects the inventory to a specific institutional process. Together, they ensure that the methodology produces outcomes, not documents.

COSO ERM Framework Alignment

The methodology maps to all eight components of the COSO ERM framework. This mapping was introduced in Chapter 2 (The Foundations: Standards and Frameworks) as part of the standards architecture. Here, in the context of integration, the mapping serves a practical purpose: it demonstrates to regulators and stakeholders that the risk identification process satisfies the requirements of an internationally recognised ERM framework, and that no component of enterprise risk management has been inadvertently omitted.

| COSO ERM Component | Process Mapping |
|--|---|
| Internal Environment | Phase 1: Internal Environment Assessment — 7 COSO elements (Ch 5 (Setting the Context: External, Internal, and Risk Culture)) |
| Objective Setting | Phase 1: Strategic objectives, risk criteria, risk appetite (Ch 5 (Setting the Context: External, Internal, and Risk Culture)) |
| Event Identification | Phase 2: Dual-track identification, SWIFT, Delphi (Ch 6 (Top-Down Identification: Workshops, SWIFT, and Delphi)–7) |
| Risk Assessment | Phase 3: Four-dimensional scoring, data quality (Ch 9 (Assessment — Scoring, Multi-Dimensional Impact, and Data Quality)) |
| Risk Response | Phase 3: Cost-benefit assessment, control effectiveness (Ch 9 (Assessment — Scoring, Multi-Dimensional Impact, and Data Quality)–10) |
| Control Activities | Phase 2: Bottom-up control assessment; Phase 3: Bow-tie control mapping (Ch 7 (Bottom-Up Identification: Templates, RCSA, and Specialist Processes), 10) |
| Information & Communication | Cross-cutting: Communication and consultation (Ch 3 (Governance: Who Owns What)); Phase 4: Documentation (Ch 11 (Documentation: The Living Risk Inventory)) |

| COSO ERM Component | Process Mapping |
|--------------------|--|
| Monitoring | Phase 6: Quarterly refresh, KRI monitoring, event-driven updates (Ch 13 (The Ongoing Cycle: Refresh, Events, and Audit)) |

The table is not decorative. When a regulator conducts a supervisory assessment, the first question is whether the institution’s risk management framework maps to a recognised standard. The second question is whether that mapping is substantive or cosmetic. A risk identification process that can trace every design decision to a specific COSO component, ISO 31000 provision, or BCBS principle — as the traceability table in Chapter 2 demonstrated — passes both tests. A process that claims COSO alignment but cannot show how its internal environment assessment maps to Component 1, or how its event identification maps to Component 3, does not.

The practical implication is that the Risk Identification Lead must maintain this mapping as a living document. When the process changes — a new specialist sub-process is added, or the assessment methodology is refined — the COSO mapping must be updated to reflect the current state of the process, not the state at which it was originally designed.

Capital Planning: ICAAP, ILAAP, and CCAR

This is where the inventory earns its cost. The connection between risk identification and capital planning is the single most consequential integration point in the methodology, and the one where failure carries the most severe regulatory and institutional consequences.

The Regulatory Imperative

BCBS Corporate Governance Principle 7 requires that risks are identified and monitored on a firm-wide basis. The quality of the risk inventory directly determines the quality of the capital assessment. If a material risk is absent from the inventory, it will be absent from the capital assessment. If a risk is present but assessed with inadequate rigour, the capital set aside against it will be inadequate.

The supervisory consequences for deficient risk identification in the capital context are material and specific. Under the PRA SREP framework (SS31/15), identified weaknesses in risk identification can result in **Pillar 2A capital add-ons** — additional capital requirements imposed by the supervisor above the minimum regulatory requirement.⁶ PRA buffer requirements can include scalars of up to 40% of Pillar 2A for institutions where the supervisor judges that risk identification is deficient. This is not a theoretical penalty. It is a direct, quantifiable cost of inadequate risk identification, paid in capital that cannot be deployed commercially.

Under the Fed CCAR framework (SR 15-18), the **Material Risk Inventory** must be comprehensive and updated quarterly, with direct linkage to stress scenario design and capital estimation.⁷ The quarterly re-identification cycle described in Chapter 6 (Top-Down Identification: Workshops, SWIFT, and Delphi) — SWIFT workshops producing updated risk lists that feed into the Material Risk Inventory — exists precisely to satisfy this requirement. At any institution subject to CCAR, the risk identification process must feed directly into the submission. The Material Risk Inventory becomes the foundation for scenario design and stress testing. This connection is not optional; it is the primary use case that justifies the investment in the process.

How the Inventory Feeds Capital Planning

The integration operates through six specific mechanisms:

Material risk to stress scenario mapping. Each material risk in the inventory is mapped to one or more stress scenarios. The scenario linkage field in the risk profile — described in Chapter 11 as “the bridge between risk identification and capital adequacy” — documents which ICAAP, ILAAP, or CCAR scenario tests each risk. A material risk with no scenario linkage is a risk that the capital framework has not addressed.

Severity informed by four-dimensional scores. The assessment scores from Chapter 9 (Assessment — Scoring, Multi-Dimensional Impact, and Data Quality) directly inform scenario severity. A risk scored 5 on Impact and 4 on Likelihood demands a severe scenario. A risk scored 3 on Impact but 5 on Speed of Onset requires a scenario that tests the institution’s ability to respond rapidly, not just its ability to absorb loss. The four dimensions provide the structured basis for calibrating scenarios beyond the generic “severe but plausible” formulation that regulators require but institutions often interpret loosely.

Loss estimates tied to the risk inventory. Under both ICAAP and CCAR, loss estimates under each scenario must be traceable to the risk inventory. This means the institution must be able to demonstrate that the risks generating losses in the stress test are the same risks documented in the inventory, assessed using the same methodology, and scored using the same four-dimensional framework. A stress test that produces loss estimates from risks not in the inventory — or that omits risks in the inventory from the loss estimation — fails the integration test.

Management judgement overlays with data-quality-informed conservatism. Not all material risks can be quantified through models. Risks that cannot be modelled — emerging risks, conduct risks, reputational risks, some operational risks — are addressed through management judgement overlays. These overlays must be documented and justified. The **Data Quality Rating** from the inventory directly informs the level of conservatism applied: a risk rated Low or Very Low on data quality requires more conservative assumptions in the management overlay than a risk rated High. This is the operational consequence of the data quality framework described in Chapter 9 — it prevents institutions from assigning low data quality ratings without accepting the capital implications.

Reverse stress testing. The ICAAP must include reverse stress testing — identifying the scenarios that would render the institution non-viable. These scenarios are constructed by working backwards from the inventory: which risks, if they materialised simultaneously at their worst credible severity, would breach the institution's viability threshold? The risk interaction analysis from Chapter 10 (Risk Interaction: Bow-Ties, Matrices, and Concentration) — cascade pathways, propagation nodes, simultaneous crystallisation assessment — provides the analytical basis for constructing these scenarios. Reverse stress testing is identification in reverse: instead of asking “what risks exist?”, it asks “which combination of identified risks could destroy us?”

Specific risk categories requiring ICAAP/ILAAP attention. Certain risk categories demand dedicated treatment in capital planning: credit concentration risk, interest rate risk in the banking book (IRRBB), pension obligation risk, step-in risk, securitisation risk, excessive leverage risk, and climate-related financial risk. The taxonomy from Chapter 4 (The Risk Taxonomy) ensures these categories are captured in the inventory. The ILAAP additionally requires that liquidity risks — intraday liquidity, contingent liquidity demands, currency mismatch — receive dedicated scenario treatment informed by the treasury risk sub-process described in Chapter 7 (Bottom-Up Identification: Templates, RCSA, and Specialist Processes).

Washington Mutual: When Risk Identification Exists but Integration Does Not

Washington Mutual illustrates the catastrophic consequence of identification without integration. By 2007, WaMu was the largest savings institution in the United States with \$307 billion in assets.⁸ Its business model was built on aggressive origination of subprime and option adjustable-rate mortgages, driven by a corporate culture that called itself the “power of yes.”

The risk identification outputs existed. Internal reports documented deteriorating loan quality. Credit risk metrics showed rising delinquencies in the option ARM portfolio. The data was available. But WaMu’s sales culture systematically prevented that data from reaching the capital planning process. The Chief Enterprise Risk Officer role was marginalised. Risk appetite limits were repeatedly raised to accommodate growth targets. The credit risk function lacked authority to decline loans that met minimum origination criteria, even when portfolio-level indicators signalled deterioration.

The integration failure was not that risk identification did not occur. It was that the identification outputs were severed from the processes where they should have had consequences — capital allocation, strategic planning, Board governance. A material risk identified and documented but excluded from the capital assessment is functionally equivalent to a risk that was never identified at all.

WaMu was seized by the Office of Thrift Supervision on 25 September 2008 — the largest bank failure in United States history.⁹ Its banking operations were sold to JPMorgan Chase for \$1.9 billion, a fraction of the institution’s former value.¹⁰

What was missing: An integration framework requiring that every material risk in the inventory maps to a capital scenario, with governance mechanisms preventing management from overriding the connection between identification and capital. The methodology’s integration of the inventory with capital planning through mandatory scenario linkage, data-quality-informed conservatism, and Board-level reporting of unlinked material risks would have made WaMu’s selective blindness structurally impossible.

The Regulatory vs Economic Risk Gap

Chapter 1 (Why Banks Fail at Risk Identification) identified **Regulatory Arbitrage Masking** as one of the ten recurring failure modes — the pattern where institutions structure activities to minimise regulatory capital requirements in ways that obscure the economic risk. This chapter delivers the methodology's response.

The regulatory mapping table introduced in Chapter 4 translates the institution's internal taxonomy into each regulator's categories. That translation serves a reporting function. But it also serves an analytical function: it reveals the gaps between what the regulator requires the institution to hold capital against and what the institution's own risk identification process has identified as its actual risk exposure.

The **regulatory vs economic risk gap analysis** is a structured comparison conducted as part of Phase 5 integration. For each material risk in the inventory:

- 1. Identify the regulatory capital treatment.** What capital charge does the regulator impose for this risk? Under which framework (Pillar 1 standardised, Pillar 1 internal models, Pillar 2A)? Is the risk captured by regulatory capital at all?
- 2. Identify the economic risk assessment.** What does the institution's own risk identification process say about this risk? What is the four-dimensional score? What is the data quality rating? What does the risk interaction analysis show about correlations and concentration?
- 3. Map the gap.** Where is the regulatory treatment less conservative than the economic assessment? Where does the regulatory framework not capture the risk at all? Where do regulatory categories fail to aggregate risks that the enterprise portfolio view has identified as connected?
- 4. Document and escalate.** Gaps are documented in the risk profile and reported to the Board Risk Committee through the principal risk report. Gaps that exceed defined thresholds — a risk assessed as material by the institution but carrying zero regulatory capital, or a concentration identified by the enterprise portfolio view but invisible in regulatory returns — require explicit Board acknowledgement and, typically, additional capital through Pillar 2A or management buffers.

The gap analysis directly addresses the regulatory arbitrage failure mode. If an institution structures a portfolio to minimise regulatory capital — for example, classifying activities as hedging rather than proprietary trading, or booking exposures in jurisdictions with lower capital requirements — the gap analysis will show the divergence between regulatory treatment and economic risk. The Board then has the information it needs to decide whether the institution is comfortable with that divergence or whether additional capital should be held.

The ICAAP process requires exactly this comparison. Risks identified through the enterprise-wide process are mapped against the Pillar 1 regulatory capital framework, and any risk assessed as material that is not captured — or is inadequately captured — by Pillar 1 must be explicitly addressed in the Pillar 2A assessment. The regulatory mapping table is the translation mechanism; the gap analysis is the analytical output.

This is not a theoretical exercise. Step-in risk — the risk that an institution provides financial support to an unconsolidated entity beyond contractual obligations — was introduced as a taxonomy category in Chapter 4 precisely because it falls outside the standard regulatory perimeter. An institution that identifies step-in risk in its inventory but does not assess its capital implications in the ICAAP has identified the risk without integrating it. The gap analysis closes that loop.

Strategic Planning: Risk Identification as Strategic Input

In most institutions, strategic planning and risk identification operate on separate tracks. The strategy team develops growth plans, market entry proposals, and M&A opportunities. The risk function is consulted — usually late in the process — to confirm that the proposed strategy does not obviously breach risk appetite. This sequence is backwards.

Risk identification outputs should inform strategic planning, not merely validate it. The principal risk report, the emerging risk register from the Delphi process, and the enterprise portfolio view all contain information that is directly relevant to strategic direction. An emerging risk identified through the Delphi method may constrain or redirect growth plans. A concentration identified in the enterprise portfolio view may argue against further expansion in a particular market or product line. A cost-benefit assessment from Chapter 10 may demonstrate that a proposed control investment generates insufficient risk reduction to justify the capital expenditure.

Three specific integration mechanisms connect risk identification to strategic planning:

New product and market approval. Every new product, market, or business line must be assessed through the risk identification framework before launch. This is not a new requirement — EBA guidelines, PRA expectations, and the AML/CFT sub-process in Chapter 7 all require pre-launch risk assessment. The integration point is that the assessment uses the same taxonomy, the same four-dimensional scoring, the same data quality framework, and the same enterprise portfolio view as the existing inventory. The new product is not assessed in isolation; it is assessed in the context of the institution's existing risk landscape. Does it create a new concentration? Does it introduce a risk type not currently in the taxonomy? Does it interact with existing risks in ways the interaction matrix would flag?

Mergers and acquisitions. M&A due diligence must include a risk identification assessment using the methodology. The acquiring institution maps the target's risk profile against its own inventory, identifies new risks that the acquisition would introduce, and assesses the aggregate risk position after completion. The enterprise portfolio view is recalculated with the combined exposure.

Risk appetite as strategic constraint. The risk appetite statement — which Chapter 5 required to be “specific and operational” rather than aspirational — functions as a strategic constraint when properly integrated. A strategy that would breach appetite boundaries identified in Phase 1 requires either a change in strategy or a change in appetite, both requiring Board approval. The appetite framework described in Chapter 5 (Setting the Context: External, Internal, and Risk Culture) becomes the mechanism through which risk identification constrains strategic risk-taking.

Bankia: When Strategy Ignores Risk Identification

Bankia illustrates what happens when strategic decisions are made without risk identification integration. In 2010, seven struggling Spanish savings banks — *cajas* — were merged to create Bankia.¹¹ The strategic rationale was that combining weak institutions would create a stronger one. The risk identification reality was the opposite: each *caja* held massive property developer loan portfolios that were inadequately provisioned, and combining them aggregated concentrated property exposures without recapitalisation.

No assessment identified that merging seven institutions with concentrated real estate exposure simply created a larger concentration. The enterprise portfolio view that Phase 2 would have produced — aggregating all property-related exposures across the seven

entities into a single consolidated picture — was not conducted. The strategic decision was made on the basis of operational synergies and cost reduction, not on the basis of risk identification outputs.

Bankia was listed through an IPO in 2011 with a prospectus that materially misrepresented the institution's financial health. Within months, the true extent of the property losses became apparent. The Spanish government injected €22.4 billion — Spain's largest bank bailout.¹² Former chairman Rodrigo Rato was convicted of fraud.¹³ Bankia was eventually merged with CaixaBank in 2021.¹⁴

What was missing: A risk identification framework integrated with the strategic planning process. The methodology requires that M&A due diligence includes a full risk identification assessment — mapping the target's risk profile against the existing inventory, calculating the combined enterprise portfolio view, and identifying concentrations that the combination would create. The seven cajas' property exposures, assessed individually, were each known. Assessed in aggregate through an enterprise portfolio view, they would have been identified as a concentration that the combined institution's capital could not support.

Regulatory Reporting: The Inventory as Submission Foundation

The risk inventory provides the foundation for regulatory submissions. The integration is not cosmetic — regulators examine whether the institution's risk identification process produces outputs that are consistent with, and directly inform, its regulatory returns.

Four regulatory reporting channels draw from the inventory:

Principal risk disclosures. The Annual Report and Pillar 3 disclosures require identification and description of the institution's principal risks. These disclosures draw directly from the risk inventory — specifically from the material risks with their risk definitions, four-dimensional scores, and trend indicators. An institution whose Pillar 3 principal risk disclosures do not match its internal risk inventory has a consistency problem that supervisors will identify.

ICAAP and ILAAP submissions. The capital and liquidity adequacy assessments reference the risk identification process and its outputs. The risk inventory provides the foundation — the comprehensive list of material risks, their scores, their data quality ratings, their interaction relationships. The scenario linkage field connects each risk to its capital treatment. The Data Quality Rating informs the conservatism of management overlays. A supervisor reviewing the ICAAP will trace backward from the capital number to the risk inventory to determine whether the institution has identified all material risks and assessed them with appropriate rigour.

Recovery plans. Recovery planning requires identification of scenarios that could threaten the institution's viability. These scenarios are drawn from the reverse stress testing described earlier in this chapter, which in turn is drawn from the risk inventory, the interaction matrix, and the enterprise portfolio view. The recovery plan's risk indicators should map to the Key Risk Indicators in the inventory.

Supervisory engagement. Beyond formal submissions, the risk inventory grounds day-to-day supervisory engagement. When a supervisor asks about a specific risk, the institution should be able to produce the inventory entry: its taxonomy classification, its four-dimensional score, its interaction relationships, its trend, its data quality rating, its scenario linkage. The regulatory mapping table from Chapter 4 provides the translation between the institution's internal taxonomy and the supervisor's categories.

The COSO ERM alignment mapping from earlier in this chapter also serves a regulatory reporting function: it demonstrates to supervisors that the institution's risk identification process satisfies the requirements of an internationally recognised framework, providing confidence that no dimension of enterprise risk management has been overlooked.

Board Reporting: The Principal Risk Report

The final integration point is the one that matters most for institutional governance: the Board Risk Committee receives the risk identification outputs through the **principal risk report**.

Chapter 3 (Governance: Who Owns What) defined the principal risk report as containing ten items:

1. Summary of all material risks with current residual risk scores
2. The consequence dimension driving the Impact score for each risk

3. Trend direction for each risk (increasing / stable / decreasing)
4. Risks approaching or breaching appetite
5. New or emerging risks added since the last report
6. Risks removed or reclassified since the last report
7. Key Risk Indicator dashboard (RAG status)
8. Data quality distribution across the material risk portfolio
9. Enterprise portfolio view — aggregate risk position against institutional appetite
10. Process performance indicators

Every item on this list is drawn directly from the risk inventory and risk profiles described in Chapter 11. The principal risk report is a **presentation layer over the inventory**, not a separate analytical exercise. This distinction matters because it means the Board receives the same analytical output that informs capital planning, strategic decisions, and regulatory submissions. There is one risk picture, not multiple versions tailored to different audiences.

Board Substantive Challenge

The principal risk report is designed not merely to inform the Board but to enable the Board to challenge. Three features of the report structure support substantive challenge:

The dominant dimension. Item 2 — the consequence dimension driving each risk's Impact score — tells the Board whether a risk is primarily a financial threat, a regulatory threat, a reputational threat, or a customer/operational threat. This matters because the Board's risk response may differ depending on the dimension. A financial risk may be addressed through capital allocation; a reputational risk may require strategic communication; a regulatory risk may demand immediate remediation. The dominant dimension framework from Chapter 9 gives the Board the information it needs to challenge whether management's proposed response addresses the right dimension.

The data quality distribution. Item 8 — the distribution of data quality ratings across the material risk portfolio — tells the Board how much confidence to place in the overall risk picture. A portfolio where 60% of material risks carry High or Medium data quality ratings is more reliable than one where 40% are Low or Very Low. The Board should challenge management on Low-rated risks: what is being done to improve the evidence

base? What conservatism adjustments have been made in the capital assessment? The data quality framework from Chapter 9 becomes, through this reporting mechanism, a Board governance tool.

The enterprise portfolio view. Item 9 — the aggregate risk position against institutional appetite — is the single most important item in the report for Board governance. It answers the question: “considering all identified risks and their interactions, is this institution within the boundaries the Board has set?” If the answer is no, the Board must act — by directing risk reduction, by adjusting appetite (with documented justification), or by allocating additional capital. The enterprise portfolio view from Chapter 8 (Reconciliation and the Enterprise Portfolio View), updated through the ongoing cycle described in Chapter 13 (The Ongoing Cycle: Refresh, Events, and Audit), provides this answer at every Board meeting.

Risk Appetite Governance

The Board’s governance of risk appetite is the mechanism through which risk identification constrains institutional behaviour at the highest level. The risk appetite statement approved in Phase 1 — which Chapter 5 required to be specific and operational — is tested against the enterprise portfolio view at every reporting cycle.

When the enterprise portfolio view shows the institution approaching or breaching appetite boundaries, the principal risk report escalates this to the Board immediately. The Board then faces a governance decision: either the institution reduces its risk exposure to return within appetite, or the Board adjusts the appetite statement with full documentation of the rationale and the new boundaries. The second option is legitimate — risk appetite should evolve with strategy and context — but it requires conscious governance action, not passive drift.

The methodology’s integration of risk identification with Board reporting ensures that appetite is not aspirational. It is operational. The inventory measures risk. The enterprise portfolio view aggregates it. The principal risk report presents it to the Board in a form that demands action when boundaries are approached.

Dexia Revisited: How Integration Would Have Changed the Outcome

Return to Dexia. After the 2008 rescue, the institution's risk identification improved. Sovereign debt positions were documented. Wholesale funding dependency was known. But the reconstituted risk framework did not ensure that these identification outputs reached the capital planning process with sufficient force to change behaviour. Sovereign debt was again treated as effectively risk-free for capital purposes — a regulatory convention that the institution's own risk identification, properly integrated, should have challenged through the regulatory vs economic risk gap analysis.

The methodology's Phase 5 integration would have required three things that Dexia's reconstituted framework did not provide:

First, **mandatory scenario linkage**. Every material risk — including sovereign concentration and wholesale funding dependency — mapped to ICAAP stress scenarios with severity informed by four-dimensional scores. The risk interaction analysis would have shown that sovereign credit risk and wholesale funding risk were connected through the same transmission mechanism that had nearly destroyed the institution three years earlier.

Second, **the regulatory vs economic risk gap analysis**. Sovereign debt carried zero risk weight under standard regulatory treatment.¹⁵ Dexia's own risk identification — if it had conducted the gap analysis this chapter describes — would have shown that the economic risk of concentrated peripheral European sovereign exposure was materially higher than the regulatory capital treatment suggested. That gap would have been reported to the Board, requiring either additional capital or exposure reduction.

Third, **Board reporting with enterprise portfolio view**. The principal risk report would have shown the Board that the institution's aggregate sovereign concentration, combined with its wholesale funding dependency, recreated the same structural vulnerability that had required the 2008 rescue. The Board would have been confronted with a documented, analytically rigorous picture of an institution walking back into the same trap.

Dexia's second failure was not a failure of identification. It was a failure of integration — the institutional machinery that should have connected what was known to what was decided.

The Integration Standard

Phase 5 connects the inventory to the institution. The connections are not optional or discretionary — each one is mandated by either regulatory requirement or standards compliance:

| Integration Point | Regulatory/Standards Basis | Mechanism |
|----------------------------|--|---|
| COSO ERM Alignment | COSO ERM 2004, ISO 31000 Principle b | Eight-component mapping table |
| Capital Planning | PRA SS31/15, Fed SR 15-18, BCBS Principle 7 | Scenario linkage, 4D severity calibration, DQ-informed conservatism |
| Regulatory vs Economic Gap | ISO 31000 Principle b, BCBS Principle 7 | Gap analysis per material risk, Board escalation |
| Strategic Planning | ISO 31000 Principle b, c; COSO Objective Setting | New product assessment, M&A due diligence, appetite constraints |
| Regulatory Reporting | PRA SS31/15, Fed SR 15-18, EBA Guidelines | Inventory-sourced submissions, regulatory mapping table |
| Board Reporting | BCBS Principle 7, COSO Component 7 | Principal risk report (10 items), appetite governance |

The Risk Identification Lead is responsible for maintaining these integration points. The CRO is responsible for ensuring they are functioning — that identified risks are reaching capital planning, that the Board is receiving the enterprise portfolio view, that regulatory submissions are consistent with the inventory. Internal Audit, as part of the third line described in Chapter 3, tests whether the integration is substantive or cosmetic.

An inventory that is comprehensive but disconnected from capital planning is Dexia. An inventory that exists but is overridden by sales culture is Washington Mutual. A strategic merger that proceeds without risk identification integration is Bankia. In each case, the identification work was done — or could have been done — but the institutional machinery to connect identification to action was absent.

The methodology's Phase 5 provides that machinery. It is not the most analytically demanding phase — that distinction belongs to Phase 3. It is not the most politically challenging — that is Phase 2's reconciliation. But it is the phase that determines whether all the preceding work changes institutional outcomes or merely populates a register.

The inventory is built. The risks are identified, assessed, documented, and integrated with the institution's core management processes. But risk identification is not a one-time exercise. The risk landscape changes — new risks emerge, existing risks evolve, controls degrade, the external environment shifts. Chapter 13 describes the ongoing cycle that keeps the inventory current: the quarterly re-identification, the annual full refresh, event-driven updates, and the internal audit assurance that ensures the process maintains its rigour over time.

-
1. European Commission, State Aid Decision SA.33760 (2008/C), *Dexia — Restructuring Aid*, 26 February 2010. The initial rescue comprised a €6.4 billion recapitalisation by Belgium (€3B), France (€3B), and Luxembourg (€376M), plus state guarantees on borrowings of up to €150 billion. See also Belgian Federal Government press release, 30 September 2008.
 2. European Commission, State Aid Decision SA.33760 (2012/NN), *Orderly Resolution of Dexia*, 28 December 2012. The 2011 restructuring involved Belgium acquiring Dexia Bank Belgium (renamed Belfius), France assuming Dexia Municipal Agency (Dexia Crédit Local), and the residual Dexia SA placed into managed wind-down with additional state guarantees of €85 billion.
 3. ISO 31000:2018, *Risk Management — Guidelines*, Section 4, Principle (b): "Risk management is an integral part of all organisational processes."
 4. Committee of Sponsoring Organizations of the Treadway Commission (COSO), *Enterprise Risk Management — Integrated Framework*, 2004, Component 7: Information and Communication.
 5. Basel Committee on Banking Supervision, *Corporate Governance Principles for Banks* (BCBS 328), July 2015, Principle 7: "Risks should be identified and monitored on an ongoing firm-wide and individual entity basis."
 6. Prudential Regulation Authority, Supervisory Statement SS31/15, *The Internal Capital Adequacy Assessment Process (ICAAP) and the Supervisory Review and Evaluation Process (SREP)*, updated December 2020. Pillar 2A capital add-ons are set through the SREP where the PRA identifies weaknesses in risk management or capital adequacy.
 7. Board of Governors of the Federal Reserve System, SR 15-18 / CA 15-14, *Federal Reserve Supervisory Assessment of Capital Planning and Positions for LISC Firms and Large and Complex Firms*, 18 December 2015. Requires firms to maintain a comprehensive Material Risk Inventory updated at least quarterly.

8. FDIC, *Washington Mutual Bank — Institution Profile*, and Office of Thrift Supervision records. Total assets of approximately \$307 billion as of 30 June 2008. WaMu was the largest savings and loan association in the United States.
9. FDIC Press Release PR-79-2008, “JPMorgan Chase Acquires Banking Operations of Washington Mutual,” 25 September 2008. The OTS closed Washington Mutual Bank and appointed the FDIC as receiver. It was the largest bank failure in U.S. history by total assets.
10. FDIC, *Failed Bank Information — Washington Mutual Bank*, 25 September 2008. JPMorgan Chase acquired the banking operations in a transaction facilitated by the FDIC for \$1.888 billion.
11. Bank of Spain, *Financial Stability Report*, May 2011. Bankia was formed through the merger of Caja Madrid, Bancaja, and five smaller cajas (Caja de Canarias, Caja de Ávila, Caixa Laietana, Caja Segovia, and Caja Rioja) via a Systemic Institutional Protection Scheme (SIP) in 2010, with the listed entity Bankia S.A. created in 2011.
12. European Commission, State Aid Decision SA.35253 (2012/NN), *Restructuring of BFA/Bankia*, 28 November 2012. Total public capital injections into BFA-Bankia amounted to approximately €22.4 billion, funded through the European Stability Mechanism (ESM) financial assistance programme for Spain.
13. Audiencia Nacional (Spanish High Court), Criminal Chamber, Section 4, Judgment in the Bankia IPO fraud trial, 24 February 2017. Rodrigo Rato and other former directors were convicted of offences related to the Bankia IPO prospectus. Rato received a prison sentence of four years and six months.
14. European Central Bank, press release, 24 September 2020, approving the merger of CaixaBank S.A. and Bankia S.A. The merger was completed on 26 March 2021, creating Spain’s largest domestic bank by assets.
15. Under the Basel II/III Standardised Approach for credit risk (BCBS 128, para. 53–54, and CRR Article 114), exposures to EU member state sovereigns denominated and funded in the domestic currency of that sovereign may be assigned a 0% risk weight. This treatment was applied to Greek and other peripheral European sovereign bonds held by Dexia despite escalating default risk.

The Ongoing Cycle: Refresh, Events, and Audit

The Bank That Stopped Looking

Silicon Valley Bank held \$209 billion in assets on the day it failed.¹ It was the sixteenth-largest bank in the United States. Its collapse on 10 March 2023 was the largest American bank failure since Washington Mutual in 2008, and it happened in roughly thirty-six hours.²

The proximate cause was a bank run — depositors withdrew \$42 billion in a single day after SVB disclosed a \$1.8 billion loss on the sale of its available-for-sale securities portfolio.³ But the underlying cause was interest rate risk in the banking book that had been accumulating for two years in plain sight.

SVB's business model concentrated deposits from venture capital-backed technology firms. During 2020 and 2021, as technology funding surged, deposits nearly tripled to \$189 billion.⁴ Management invested the surplus in long-duration mortgage-backed securities and US Treasuries, locking in yields that seemed attractive at the time. The held-to-maturity portfolio reached \$91 billion.⁵ As the Federal Reserve raised rates by 500 basis points across 2022 and 2023,⁶ the market value of those holdings fell dramatically. By the end of 2022, SVB's unrealised losses on held-to-maturity securities exceeded \$15 billion — more than the bank's total equity.⁷

Every one of these facts was publicly available. The unrealised losses appeared in the financial statements. The deposit concentration was visible to anyone who examined the customer base. The duration mismatch between long-dated assets and potentially volatile deposits was a textbook **interest rate risk in the banking book** scenario that any competent risk identification process would flag.

Yet SVB's risk identification process did not flag it — or, more precisely, the process that should have been updating the risk landscape in response to a fundamentally changed interest rate environment had stopped functioning. The Chief Risk Officer departed in April 2022 and was not replaced until January 2023⁸ — eight months during which the most consequential monetary policy shift in a generation was underway. The risk committee met infrequently. Internal risk reports continued to circulate, but the active re-identification that the changing environment demanded did not occur.

SVB did not fail because it lacked a risk identification process. It failed because the ongoing cycle — the quarterly re-identification, the event-driven updates, the continuous recalibration that keeps the process current — had degraded to the point of irrelevance. The methodology described in the preceding twelve chapters of this book produces a comprehensive risk landscape at a point in time. This chapter describes what keeps that landscape current.

Phase 6 in Context

The five phases of the risk identification methodology produce, at their conclusion, a fully documented risk inventory: every material risk scored across four dimensions, data quality rated, interaction relationships mapped, concentration analysis completed, risk profiles maintained for the most critical exposures, and the entire landscape integrated with capital planning, strategic planning, regulatory reporting, and Board governance. Phases 1 through 5, executed properly, represent the most thorough risk identification any institution is likely to undertake.

And yet, the moment the annual cycle concludes, the inventory begins to age. External conditions change. Internal strategies evolve. New products launch. Acquisitions complete. Regulations shift. Competitors fail. Markets move. The risk landscape that was accurate on the day the Board signed off the principal risk report may be materially different three months later.

This is why ISO 31000 Principle j requires the risk management process to be **dynamic, iterative, and responsive to change**.⁹ It is why COSO's eighth component — **Monitoring** — exists as the mechanism that ensures the preceding seven components continue to function.¹⁰ And it is why BCBS Corporate Governance Principles require risk identification on an **ongoing** basis, not merely an annual one.

Phase 6 is the methodology's response to these requirements. It comprises four distinct mechanisms operating at different frequencies: quarterly re-identification, annual full re-identification, event-driven updates, and internal audit assurance. Supporting these mechanisms are two cross-cutting enablers: a training and awareness programme that maintains the capability of process participants, and a framework monitoring function that measures and improves the process itself.

Without Phase 6, the methodology produces a sophisticated point-in-time exercise. With it, the methodology produces a living risk identification capability. The difference is the difference between a photograph and a surveillance system.

Quarterly Re-Identification

The quarterly cycle is the heartbeat of the ongoing process. It is also the area where most institutions fall short, because the temptation to reduce it to a re-assessment of existing risks — simply updating scores on a static list — is almost irresistible.

The methodology requires something fundamentally different. Consistent with Fed SR 15-18,¹¹ the quarterly cycle is an **active re-identification** — not a passive review. At any institution subject to CCAR, the Material Risk Inventory must be comprehensive and current, and that means every quarter the institution must ask not just “have our existing risks changed?” but “are there new risks we did not identify last quarter?”

The quarterly cycle encompasses six activities, each building on the infrastructure established in earlier phases.

First, the external context is updated. The PESTLE assessment from Phase 1 (Chapter 5 (Setting the Context: External, Internal, and Risk Culture)) is refreshed to reflect material developments since the prior quarter. This is not a full PESTLE exercise — it is a targeted update focusing on what has changed: new regulatory announcements, macroeconomic shifts, geopolitical developments, technology disruptions, legal precedents, environmental events. Each development is mapped to the taxonomy and assessed for its implications across the risk landscape.

Second, the straw man for the quarterly workshop is prepared. As described in Chapter 6 (Top-Down Identification: Workshops, SWIFT, and Delphi), the quarterly straw man draws on the prior quarter's inventory, the updated PESTLE assessment, event-

driven triggers since the last cycle, and any emerging risk intelligence from the Delphi process. The straw man frames the workshop around change — what is different, what is new, what has escalated.

Third, a focused SWIFT workshop is conducted. The quarterly workshop is more targeted than the annual workshop described in Chapter 6. Rather than systematically covering all five SWIFT domains, the quarterly workshop concentrates on changes since the prior quarter. The guide words are adapted: “What has changed since last quarter that could...?”, “What are we now assuming that we were not assuming three months ago?”, “Where have our controls been tested since last quarter, and what did we learn?” The pre-workshop independent assessment remains — each participant independently identifies changes and new risks before the workshop convenes. The workshop typically runs two to three hours rather than the four hours of the annual exercise.

Fourth, all material risks are re-assessed. The four-dimensional scoring methodology from Chapter 9 (Assessment — Scoring, Multi-Dimensional Impact, and Data Quality) is applied to every material risk. Scores that have changed require documented justification. The inherent-residual gap is re-examined. The interaction matrix from Chapter 10 (Risk Interaction: Bow-Ties, Matrices, and Concentration) — which is a standing document, updated with each quarterly refresh — is reviewed for new relationships triggered by changed conditions. Bow-tie diagrams for the five to ten most critical risks are revisited, with particular attention to whether escalation factors have changed or barrier effectiveness has been tested.

Fifth, Data Quality Ratings are reviewed. Where ratings are Low or Very Low, the quarterly cycle tracks progress on improving the evidence base. A risk that carried a Low data quality rating for three consecutive quarters without an improvement plan is itself a risk identification failure — it signals that the institution is making decisions about a risk it does not adequately understand, and has accepted that state of ignorance. The CRO must approve the continued acceptance of any Very Low rating beyond two consecutive quarters.

Sixth, the assumption register is revisited. The assumption register established during the annual identification cycle (Chapter 6) and documented in the risk inventory (Chapter 11 (Documentation: The Living Risk Inventory)) records the assumptions underpinning key risk assessments — assumptions about correlations, hedging effectiveness, counterparty behaviour, funding availability, market access. The quarterly cycle tests each assumption against current evidence. An assumption that was reasonable twelve

months ago may have become untenable. The assumption register is not an administrative record — it is an early warning mechanism, and it functions only if it is actively reviewed.

The output of the quarterly cycle feeds directly into the Material Risk Inventory used for capital planning (Chapter 12 (Integration: Capital Planning, Strategy, and the Board)). New risks identified are assessed and added to the inventory. Risks that have reduced below the materiality threshold are noted but retained — the methodology does not permit silent disappearance. KRI thresholds from Chapter 11 are reviewed and recalibrated where the evidence warrants.

The quarterly cycle typically requires two to three weeks from PESTLE update to CRO sign-off. Institutions that attempt to compress it into a single afternoon meeting are performing re-assessment, not re-identification — and the distinction matters. Re-assessment asks whether existing risks have changed. Re-identification asks whether the risk landscape has changed. The first updates a list. The second updates an understanding.

Annual Full Re-Identification

Once per year, the full methodology is re-executed from the beginning. All six phases. From scratch.

This is the most demanding exercise in the ongoing cycle, and the point where compliance theatre is most likely to take hold. The temptation to roll forward last year's inventory with minor updates — adjusting a few scores, adding one or two risks that have become obvious, retiring one or two that have resolved — is considerable. It saves time, avoids difficult conversations, and produces a document that looks substantially similar to last year's, which many participants interpret as evidence of stability rather than evidence of stagnation.

It is precisely the compliance theatre pattern described in Chapter 7 — bottom-up submissions that have not changed in three years. Compliance theatre is not merely ineffective. It is dangerous, because it creates a documented record that identification has occurred, providing false assurance to the Board, regulators, and capital planning functions that rely on the inventory's currency.

The annual re-identification is designed to prevent this. It requires:

- **Taxonomy review and update.** The L1 structure and material L2 changes are presented to the Board for approval. Emerging risk intelligence from the Delphi process and event-driven updates during the year inform proposed amendments. The taxonomy maintenance process described in Chapter 4 (The Risk Taxonomy) is the governance mechanism.
- **Fresh top-down and bottom-up identification.** Not rolled forward. The SWIFT workshops and bottom-up template exercises from Chapters 6 (Top-Down Identification: Workshops, SWIFT, and Delphi) and 7 (Bottom-Up Identification: Templates, RCSA, and Specialist Processes) are conducted as though the institution were identifying risks for the first time. The prior year's inventory serves as one input — through the straw man — but the process must be capable of producing a materially different output if conditions warrant.
- **Recalibrated materiality threshold and risk appetite.** The Phase 1 risk criteria from Chapter 5 are reviewed. Materiality thresholds that were appropriate for last year's balance sheet may not be appropriate for this year's. Risk appetite boundaries are tested against the enterprise portfolio view and adjusted where the Board determines that the institution's risk tolerance has changed.
- **Refreshed internal environment assessment.** The seven COSO elements assessed in Phase 1 (Chapter 5) are reassessed. Risk culture changes through leadership transitions, acquisitions, strategic shifts, and attrition. An internal environment assessment that has not changed in three years is not evidence of a stable culture — it is evidence that the assessment is not being conducted.
- **Full reconciliation.** The five-step reconciliation process from Chapter 8 (Reconciliation and the Enterprise Portfolio View) is executed in full, with challenge sessions resolving gaps between the top-down and bottom-up tracks. The reconciliation is where the most valuable annual identification occurs — the gap analysis reveals what neither track found.
- **Updated enterprise portfolio view.** Common exposures, simultaneous crystallisation scenarios, aggregate position against appetite, and diversification assumptions are all reassessed from the current inventory. The cost-benefit assessments and ALARP determinations from Chapter 10 are reviewed and updated where the risk landscape or control environment has changed.

The annual cycle is the methodology's mechanism for preventing institutional complacency — the tenth failure mode from Chapter 1 (Why Banks Fail at Risk Identification). Complacency sets in when the absence of recent loss events is interpreted as evidence that the risk landscape is benign. The annual re-identification forces the institution to confront the question afresh: what risks are we exposed to, and are we identifying them?

Monthly KRI Monitoring

Between quarterly cycles, the risk landscape is monitored through the Key Risk Indicators established in the risk inventory (Chapter 11). Each material risk carries defined KRIs with green, amber, and red thresholds. Monthly monitoring serves as the early warning system that bridges the gap between quarterly identification exercises.

KRI monitoring is not identification — it is surveillance. It detects movement in known risk indicators. An amber breach triggers investigation. A red breach triggers escalation to the CRO and, depending on the risk, to the Board Risk Committee. Multiple amber breaches across related risks may signal a pattern that warrants an event-driven update before the next quarterly cycle.

The discipline of monthly KRI monitoring also generates trend data that feeds the quarterly and annual cycles. A KRI that has been drifting steadily from green toward amber over six months is a different signal than one that jumped from green to red overnight. The trend indicator in the risk inventory (Increasing, Stable, Decreasing) is derived from this monitoring data.

Event-Driven Updates

The scheduled cycles — quarterly and annual — operate on a fixed calendar. But risk does not follow calendars. The event-driven update mechanism ensures that the inventory is updated immediately when circumstances demand, without waiting for the next scheduled cycle.

Six categories of event trigger an immediate risk identification update:

A material loss event or near-miss occurs. This is the most obvious trigger. When a risk materialises — whether as a financial loss, a regulatory sanction, a reputational incident, or an operational failure — the event-driven process activates. But the response is not

merely backward-looking. The methodology requires both **root cause analysis** of what happened and **forward-looking assessment** of what may happen next as a consequence.

A significant change in the external environment occurs. Regulatory announcements, macroeconomic shifts, geopolitical events, technology disruptions. Each Federal Reserve rate increase during 2022 and 2023 was an event-driven trigger that should have activated SVB's risk identification update mechanism. The ECB's introduction of new supervisory expectations, a major cyber attack on a peer institution, a sovereign debt downgrade in a country where the institution has material exposure — each of these changes the risk landscape and warrants an update that cannot wait for the next quarterly cycle.

A new business, product, or market is entered. Per EU AMLD6,¹² new products and delivery channels must undergo AML/CFT risk assessment **before launch**. The methodology extends this requirement to all risk types: a new product should be assessed against the taxonomy, scored using the four-dimensional framework, and positioned within the enterprise portfolio view before it enters the market. The risk identification update is a precondition for launch, not a retrospective exercise.

An acquisition, divestiture, or restructuring takes place. M&A activity is a particularly high-risk trigger. The EBA Internal Governance Guidelines require comprehensive risk identification of the target entity, integration risks, and the combined risk profile. As Chapter 12 described, Bankia's merger of seven struggling cajas aggregated concentrated property exposures without assessing the combined enterprise portfolio view. Every acquisition changes the risk landscape fundamentally. The event-driven update must recalculate the enterprise portfolio view for the combined entity.

Internal audit identifies a material control failure. A control that was rated effective in the last assessment and is now found to be ineffective changes the residual risk score for every risk that depends on it. The inherent-residual gap analysis from Chapter 9 must be updated, and the vulnerability dimension reassessed.

A material outsourcing arrangement is entered into, significantly changed, or terminated. Per the EBA Outsourcing Guidelines, changes to critical outsourcing arrangements require risk re-assessment. If the institution's cloud provider changes its terms of service, or a critical vendor is acquired by a competitor, or a fourth-party dependency is discovered in the outsourcing chain, the risk landscape has changed.

The event-driven update process employs the same tools as the scheduled cycles — SWIFT prompts, structured templates, four-dimensional scoring, interaction analysis — but applies them to the specific event and its consequences rather than the full risk landscape. The scope is targeted, but the rigour is identical.

The case of National Australia Bank illustrates a subtler form of event-driven failure. In 2018, Commonwealth Bank of Australia received a record AUD 700 million AUSTRAC fine for failing to report over 53,000 suspicious cash transactions through its intelligent deposit machines.¹³ This was not merely CBA's problem. It was an event-driven trigger for every institution operating similar technology. Yet NAB failed to identify that its own transaction monitoring systems had analogous gaps. When AUSTRAC brought civil penalty proceedings against NAB in 2020¹⁴, the risk identification failure was not that NAB's own systems had weaknesses — it was that a peer institution's public, catastrophic failure had not triggered a re-identification of the same risk within NAB's own operations.

The ongoing cycle requires that material events at peer institutions are treated as event-driven triggers for internal re-identification. The starting universe described in Chapter 5 includes industry loss data for precisely this reason. When a competitor fails, the question is not “could that happen to us?” — it is “what specific aspects of our risk landscape share the structural characteristics that produced that failure?”

Internal Audit Assurance

Internal Audit is the third line of defence, and its role in Phase 6 is distinct from all other participants. Internal Audit does not identify risks, does not participate in workshops, does not complete templates, and does not score assessments. Internal Audit provides **independent assurance** over the process itself.

The annual audit covers seven areas:

1. Completeness of risk coverage against the taxonomy and regulatory requirements.

Are there taxonomy nodes with no identified risks? Are regulatory risk categories from the Chapter 4 mapping table all represented in the inventory?

- 2. Quality and consistency of risk assessments** across business units. Are the four-dimensional scores applied consistently? Are there systematic differences between business units that suggest calibration drift? Are the three most commonly failed fields from Chapter 7 — underlying drivers, control effectiveness, data quality rating — populated with genuine analysis rather than default responses?
- 3. Effectiveness of the reconciliation** between top-down and bottom-up. Did the reconciliation produce genuine gap analysis, or was it compilation masquerading as reconciliation? Were challenge sessions conducted? Were disagreements resolved through evidence?
- 4. Appropriateness of Data Quality Ratings** and their use in decision-making. Are the three operational consequences — conservatism adjustment, sensitivity testing, Board transparency — being applied? Or are Low and Very Low ratings being carried without the conservatism that the methodology requires?
- 5. Adequacy of documentation and audit trail.** Does every inventory entry carry a complete history of changes with date, author, and reason? Are retired risks documented and archived? Are the assumption register and disagreement log maintained as living records?
- 6. Integration with capital planning and reporting.** Are all material risks linked to ICAAP/ILAAP/CCAR stress scenarios? Is the principal risk report a genuine presentation layer over the inventory, or has it become a separate document with different content? Is the regulatory vs economic risk gap analysis from Chapter 12 current?
- 7. Effectiveness of the continual improvement loop.** Are the process performance indicators tracked? Are lessons-learned reviews conducted? Are improvement actions implemented before the next cycle?

Internal Audit's findings are reported to the Board Audit Committee and the Board Risk Committee. A finding that the risk identification process has degraded is itself a material risk — because every downstream function that depends on the inventory's accuracy (capital planning, regulatory reporting, strategic planning, Board governance) is compromised.

Training and Awareness

The risk identification process is only as effective as the people executing it. A methodology built on SWIFT workshops, structured templates, four-dimensional scoring, and challenge sessions requires that every participant understands their role and possesses the skills to fulfil it.

The training programme operates at five levels:

| Audience | Content | Frequency |
|--|---|-------------------------------|
| Board / Risk Committee | Risk landscape overview, emerging risks, risk appetite interpretation, challenge techniques for the principal risk report | Annual, at minimum |
| CRO function and Risk Identification Lead | Methodology updates, technique mastery (SWIFT facilitation, Delphi management, bow-tie construction), regulatory developments | Annual plus ad hoc |
| Business Unit Heads and Risk Assessors | Standardised template completion, scoring calibration, data quality requirements, driver analysis techniques | Annual plus before each cycle |
| Front-Line Employees | Risk awareness, how to identify and report risks, escalation channels, examples relevant to their function | Annual plus at onboarding |
| Internal Audit | Process methodology, assessment criteria, common weaknesses to test for | Annual |

Scoring calibration deserves particular attention. To ensure consistency across business units, annual calibration exercises present a set of sample risks to all risk assessors, who independently score them using the four-dimensional framework. Results

are compared, and systematic differences in interpretation are identified and resolved. Without calibration, “Major” impact to one business unit may equate to “Moderate” in another, and the enterprise portfolio view becomes unreliable.

Beyond formal training, awareness activities sustain engagement between cycles. Case studies from the industry loss database shared as learning materials at the start of each annual cycle. Post-cycle communications summarising key findings and new risks identified. Recognition of business units and individuals who make high-quality contributions to the identification process. The objective is to prevent risk identification from being perceived as a compliance burden imposed by the second line — it must be understood as a discipline that protects the institution.

Framework Monitoring and Continual Improvement

The risk identification process is itself subject to monitoring and improvement. This is distinct from Internal Audit assurance — it is management’s own self-improvement loop, consistent with ISO 31000 Sections 4.5 and 4.6.

The CRO function tracks process performance indicators — KPIs for the process itself, not for the risks it identifies:

| Indicator | Target |
|---|----------------------------|
| Percentage of BUs completing bottom-up assessments on time | 100% |
| Percentage of material risks with named, active risk owners | 100% |
| Number of risks reclassified (upgraded or downgraded) per cycle | Tracked for trend |
| Number of new risks identified per cycle | Tracked for trend |
| Time from risk identification to inclusion in inventory | Less than 10 business days |
| Stakeholder satisfaction scores (post-workshop survey) | Tracked for trend |

| Indicator | Target |
|---|-------------------|
| Number of process improvement actions implemented per cycle | Tracked for trend |

These indicators serve as diagnostic tools. An institution where no new risks are identified across three consecutive quarterly cycles is either operating in a static environment — unlikely — or running a process that has succumbed to compliance theatre. An institution where every business unit submits on time but data quality ratings cluster uniformly at “Medium” has achieved administrative compliance without analytical rigour.

At the conclusion of each annual cycle, the Risk Identification Lead conducts a **lessons-learned review**: what worked well, what gaps or delays were identified, what changes to methodology, templates, governance, or tooling should be implemented. Actions are documented, assigned to named owners, and tracked to completion before the next cycle. The continual improvement loop closes when the next cycle incorporates the changes and the review evaluates their effectiveness.

Silicon Valley Bank: When the Ongoing Cycle Fails

SVB’s failure is the defining illustration of what happens when the ongoing cycle stops functioning. The risk was not exotic. IRRBB — interest rate risk in the banking book — is one of the specific risk categories that Chapter 12 identified as requiring dedicated ICAAP attention. It is a standard taxonomy entry at L2 under Market Risk. Every prudential regulator in the world expects banks to identify, measure, and manage it. SVB had a risk function, risk reports, and a risk committee. What it did not have, during the critical period, was an operational ongoing cycle.

Apply the methodology retrospectively.

Quarterly re-identification would have flagged IRRBB as an escalating risk from Q1 2022 onward. Each Federal Reserve rate increase — March, May, June, July, September, and November 2022 — would have appeared in the PESTLE update as a significant external environment change. The SWIFT prompt “What has changed since last quarter that could alter our risk profile?” would have directed attention to the growing unrealised losses in the held-to-maturity portfolio. The four-dimensional re-assessment would have shown Impact increasing (financial losses growing as a percentage of CET1), Speed of

Onset shifting from Slow to Fast (deposit concentration meant outflows could be sudden), and Vulnerability rising (no hedging programme in place for the duration exposure).

Event-driven updates should have been triggered multiple times. The Fed's rate hikes were not surprises — they were publicly announced, widely anticipated, and extensively analysed. Each constituted a significant change in the external environment. A new business model analysis was warranted when SVB's deposit base shifted dramatically during 2020-2021, concentrating further in technology-sector clients whose funding was itself sensitive to interest rate conditions.

Monthly KRI monitoring should have shown the unrealised loss position drifting steadily from green toward red throughout 2022. The held-to-maturity portfolio's market value decline was a measurable, quantifiable indicator that required no sophisticated analytics — only attention.

The internal environment had deteriorated. The CRO's departure in April 2022 left the risk function without its most senior leader during the period of greatest change. The seven COSO internal environment elements from Chapter 5 — risk management philosophy, Board attitudes, commitment to competence, assignment of authority — were all compromised by the leadership vacancy.

Internal Audit assurance should have identified that the ongoing cycle had degraded. The seven audit areas described earlier in this chapter — completeness, assessment consistency, reconciliation quality, data quality, documentation, integration, continual improvement — would have revealed that the quarterly re-identification was not functioning.

SVB's failure was a \$209 billion illustration of the Complacency failure mode from Chapter 1. The bank had experienced years of growth, profitability, and recognition. The absence of recent loss events bred confidence that the risk landscape was benign. The ongoing cycle — the mechanism designed to prevent precisely this complacency — had stopped functioning.

A second illustration from the same month reinforces the point. Signature Bank, with \$110 billion in assets, failed on 12 March 2023, two days after SVB.¹⁵ Signature had built significant deposit concentration in the cryptocurrency sector. As the crypto market deteriorated through 2022 — the collapse of TerraUSD in May, Celsius Network in June, FTX in November¹⁶ — each event was a trigger that should have activated event-driven

re-identification of Signature's concentration risk. The quarterly cycle should have shown deposit concentration shifting from a strategic advantage to an existential vulnerability. It did not. Signature became the third-largest bank failure in American history.¹⁷

Both SVB and Signature Bank illustrate the same principle: the ongoing cycle is not an administrative afterthought. It is the mechanism that determines whether the methodology continues to protect the institution after the initial identification is complete. Without it, the risk inventory becomes a historical document — accurate as of the date it was produced, and progressively less relevant with each passing day.

The Bridge to Technology

The ongoing cycle described in this chapter — quarterly re-identification, annual refresh, event-driven updates, KRI monitoring, internal audit, training, framework monitoring — is demanding. It requires sustained institutional commitment, skilled practitioners, and governance discipline.

It also generates an enormous volume of data: PESTLE assessments, workshop outputs, template submissions, four-dimensional scores, Data Quality Ratings, KRI readings, interaction matrices, bow-tie updates, assumption register entries, audit findings, process performance indicators. The previous thirteen chapters have described a methodology that is comprehensive and rigorous. Chapter 14 (Technology: AI, ML, and Data Analytics) examines how technology — artificial intelligence, machine learning, natural language processing, and data analytics — can enhance the process: automating the surveillance that supports event-driven triggers, identifying patterns in the data that human analysts might miss, and scaling the methodology across complex, multi-entity institutions without sacrificing the analytical rigour that the process demands.

-
1. Board of Governors of the Federal Reserve System, *Review of the Federal Reserve's Supervision and Regulation of Silicon Valley Bank*, April 2023, p. 6. SVB reported \$209 billion in total assets as of 31 December 2022.
 2. FDIC Press Release PR-16-2023, "FDIC Creates a Deposit Insurance National Bank of Santa Clara to Protect Insured Depositors of Silicon Valley Bank," 10 March 2023. The California Department of Financial Protection and Innovation closed SVB and appointed the FDIC as receiver. It was the second-largest bank failure in U.S. history by total assets, after Washington Mutual (2008).
 3. Board of Governors of the Federal Reserve System, *Review of the Federal Reserve's Supervision and Regulation of Silicon Valley Bank*, April 2023, pp. 3, 39. On 9 March 2023, customers initiated withdrawals of approximately \$42 billion. The \$1.8 billion after-tax loss on the AFS portfolio sale was disclosed on 8 March 2023.

4. Board of Governors of the Federal Reserve System, *Review of the Federal Reserve's Supervision and Regulation of Silicon Valley Bank*, April 2023, p. 11. SVB's total deposits grew from approximately \$62 billion at year-end 2019 to \$189 billion at year-end 2021.
5. Board of Governors of the Federal Reserve System, *Review of the Federal Reserve's Supervision and Regulation of Silicon Valley Bank*, April 2023, p. 14. SVB's held-to-maturity securities portfolio reached approximately \$91 billion by year-end 2021.
6. Board of Governors of the Federal Reserve System, Federal Open Market Committee (FOMC) meeting statements, March 2022 through May 2023. The federal funds rate target range was raised from 0–0.25% to 5.00–5.25%, a cumulative increase of 500 basis points across eleven rate increases.
7. SVB Financial Group, 2022 Annual Report (Form 10-K), filed with the SEC, 24 February 2023. Unrealised losses on held-to-maturity securities were approximately \$15.2 billion as of 31 December 2022. SVB's total shareholders' equity was approximately \$16.0 billion at the same date.
8. Board of Governors of the Federal Reserve System, *Review of the Federal Reserve's Supervision and Regulation of Silicon Valley Bank*, April 2023, p. 10. The CRO departed in April 2022. SVB did not fill the role on a permanent basis until January 2023, leaving the position vacant for approximately eight months.
9. ISO 31000:2018, *Risk Management — Guidelines*, Section 4, Principle (j): "The process of managing risk shall be dynamic, iterative, and responsive to change."
10. Committee of Sponsoring Organizations of the Treadway Commission (COSO), *Enterprise Risk Management — Integrated Framework*, 2004, Component 8: Monitoring. Requires ongoing monitoring activities and separate evaluations to ensure ERM components continue to function.
11. Board of Governors of the Federal Reserve System, SR 15-18 / CA 15-14, *Federal Reserve Supervisory Assessment of Capital Planning and Positions for LISC Firms and Large and Complex Firms*, 18 December 2015. Requires comprehensive and current Material Risk Inventories with at least quarterly updates.
12. Directive (EU) 2018/1673 of the European Parliament and of the Council of 23 October 2018 on combating money laundering by criminal law (6th Anti-Money Laundering Directive / AMLD6), and Directive (EU) 2015/849 (AMLD4), Article 8, requiring obliged entities to conduct risk assessments of new products and delivery channels prior to launch.
13. AUSTRAC and Commonwealth Bank of Australia, *Statement of Agreed Facts and Penalty*, Federal Court of Australia, 20 June 2018 (NSD 864/2017). CBA agreed to a civil penalty of AUD 700 million for serious and systemic non-compliance with the AML/CTF Act, including failure to report over 53,000 threshold transactions through its intelligent deposit machines.
14. AUSTRAC, civil penalty proceedings against National Australia Bank Limited, filed in the Federal Court of Australia, November 2020 (NSD 1215/2020). AUSTRAC alleged systemic failures by NAB in customer due diligence and ongoing monitoring obligations under the AML/CTF Act.
15. FDIC Press Release PR-17-2023, "Joint Statement by the Department of the Treasury, Federal Reserve, and FDIC," 12 March 2023. The New York Department of Financial Services closed Signature Bank and appointed the FDIC as receiver on 12 March 2023. Total assets of approximately \$110 billion as of 31 December 2022.
16. TerraUSD (UST) lost its dollar peg on 9 May 2022, triggering a collapse of the Terra/Luna ecosystem. Celsius Network filed for Chapter 11 bankruptcy on 13 July 2022 (U.S. Bankruptcy Court, SDNY, Case No. 22-10964). FTX Trading Ltd. filed for Chapter 11 bankruptcy on 11 November 2022 (U.S. Bankruptcy Court, District of Delaware, Case No. 22-11068).
17. FDIC, *Failed Bank List*, updated 2023. Signature Bank was the third-largest bank failure in U.S. history by total assets, after Washington Mutual (\$307 billion, 2008) and Silicon Valley Bank (\$209 billion, 2023).

Technology: AI, ML, and Data Analytics

The Question That Should Trouble Every Risk Function

Consider a thought experiment. In 2015, the Financial Times began publishing detailed investigative journalism questioning Wirecard's reported revenues from third-party acquiring partners in Asia.¹ Short-sellers had been flagging discrepancies in the company's accounts for years before that. By 2019, the FT's reporting had produced a forensic body of evidence that €1.9 billion in cash balances at Philippine trustee banks did not exist.² The information was public, specific, and damning.

Now imagine that every bank with exposure to Wirecard — as lender, counterparty, or investor — had deployed a natural language processing system scanning financial media for risk signals. Not a sophisticated system. A basic one, monitoring a curated list of financial news sources for mentions of institutions in its counterparty universe, flagging articles that contained words like "fraud", "investigation", "short-seller", "accounting irregularity", or "fabricated". Such a system, operating in 2015, would have surfaced the FT's reporting to the risk identification function within hours of publication. The risk would have appeared on the emerging risk register. The Delphi panel described in Chapter 6 (Top-Down Identification: Workshops, SWIFT, and Delphi) would have had documented external evidence to consider. The event-driven trigger framework from Chapter 13 (The Ongoing Cycle: Refresh, Events, and Audit) would have been activated.

None of this happened. Not because the technology did not exist — NLP systems capable of this kind of media scanning have been commercially available for over a decade. It did not happen because the institutions had not integrated technology into their risk identification methodology.

This chapter examines how technology — artificial intelligence, machine learning, natural language processing, and data analytics — enhances the six-phase methodology described in Chapter 13. The emphasis is on *enhances*. Technology does not replace the methodology. It does not replace the SWIFT workshops, the Delphi panels, the re-

conciliation process, or the judgement of experienced practitioners. What it does is extend the reach and speed of human analysis, automate surveillance that no team of analysts could sustain manually, detect patterns in data volumes that exceed human cognitive capacity, and scale the methodology across complex multi-entity institutions without sacrificing analytical rigour.

Technology is a cross-cutting enabler. It is not a seventh phase.

Why Technology Comes After Methodology, Not Before

Chapter 11 (Documentation: The Living Risk Inventory) established a principle that bears repeating here: the inventory structure must be designed independently of the technology that will host it. Define the fourteen fields, the risk profiles, the audit trail requirements, the governance protocols — then select the technology. This sequence is not arbitrary. It prevents a common failure: institutions purchase a GRC platform, configure it according to the vendor's default taxonomy and workflow, and then declare that they have a risk identification process.

They do not. They have a technology platform. The process — the workshops, the templates, the reconciliation, the challenge sessions, the enterprise portfolio view — must exist as a methodology before technology can enhance it. The right approach is to build the initial risk inventory in a spreadsheet. It will not be elegant. It will have no referential integrity, no network queries, no role-based access control. But it captures the methodology accurately and allows the institution to refine the process through two or three complete cycles before specifying technology requirements. By the time the institution engages with GRC platform vendors, it knows exactly what it needs the technology to do, because it has done it manually first.

The institutions that get this wrong — and there are many — buy the platform first. They configure it for risk assessment, not risk identification. They populate it with whatever risk categories the vendor's template provides. And they discover, two years and several million dollars later, that they have an expensive repository with no process to fill it.

Technology Across the Six Phases

Technology serves every phase of the methodology, but the nature of its contribution differs at each stage. The mapping below is not exhaustive — it identifies the highest-value applications where technology extends human capability most significantly.

Phase 1: Foundation Setting

NLP-powered PESTLE scanning. The external context assessment described in Chapter 5 (Setting the Context: External, Internal, and Risk Culture) requires systematic monitoring of political, economic, social, technological, legal, and environmental developments relevant to the institution. Performed manually, this depends on the Risk Identification Lead and a small team reading regulatory announcements, financial media, central bank publications, and industry reports. The coverage is necessarily limited by human reading capacity.

NLP systems can automate this surveillance across thousands of sources simultaneously — regulatory websites, central bank publications, financial news services, social media, academic journals, industry bodies. The technology does not replace the PESTLE assessment. It feeds it. The Risk Identification Lead still structures the analysis, maps findings to the taxonomy, and determines relevance to the institution's specific context. But the raw material is vastly more comprehensive.

Data analytics for the starting universe. The starting universe described in Chapter 5 draws on regulatory categories, industry loss data, and internal incident history. Analytical tools can process the institution's own loss and near-miss data, correlate it with external databases such as ORX, and identify patterns — risk concentrations that have increased, incident categories with rising frequency, emerging loss types not currently in the taxonomy. This is pattern detection applied to historical evidence, producing a richer starting universe than manual compilation can achieve.

Phase 2: Dual-Track Identification

Automated briefing pack generation. The pre-workshop briefing pack — PESTLE assessment, internal context summaries, prior-year risk list with trends, regulatory communications, industry loss events — can be partially automated. Systems that track reg-

ulatory announcements, extract key provisions, and flag changes relevant to the institution's risk profile reduce preparation time and increase coverage. The Risk Identification Lead curates and contextualises the output, but the data gathering is automated.

NLP analysis of bottom-up submissions. When 100-500 bottom-up risk assessments arrive from business units (Chapter 7 (Bottom-Up Identification: Templates, RCSA, and Specialist Processes)), NLP tools can analyse the text for quality indicators: Are driver fields populated with genuine causal analysis, or do they contain single-word entries? Have risk definitions changed from the prior cycle, or are they copied verbatim? Are control effectiveness ratings supported by specific evidence, or do they default to "effective"? This analysis does not replace the Risk Identification Lead's quality challenge — it accelerates it, flagging submissions that require deeper review and identifying patterns of compliance theatre across the organisation.

Phase 3: Assessment

Anomaly detection in risk scoring. When assessors across multiple business units score risks using the four-dimensional framework (Chapter 9 (Assessment — Scoring, Multi-Dimensional Impact, and Data Quality)), ML algorithms can identify scoring anomalies: a risk scored as "Minor" financial impact by one unit that is scored "Major" by another for the same underlying exposure; a Data Quality Rating of "High" where the underlying data has known gaps; a Vulnerability rating of "Very Low" where the institution has no proven controls. These anomalies are flags for the calibration process described in Chapter 13, not automated corrections.

Predictive analytics for emerging risk signals. ML models trained on historical loss data and market indicators can identify conditions that have historically preceded risk crystallisation. This is not prediction in the precise sense — it is pattern recognition. When a combination of market conditions, portfolio positions, and external factors resembles patterns that preceded historical losses in the industry database, the system flags it for human review. The practitioner determines whether the pattern is relevant; the technology ensures the pattern is noticed.

Phase 4: Documentation

Automated inventory maintenance. The fourteen-field risk inventory (Chapter 11) requires continuous maintenance: audit trails for every change, cross-references between risks and their interaction analysis, KRI threshold monitoring, trend calculations. Techno-

logy automates the clerical burden of inventory management — date-stamping changes, enforcing mandatory fields, calculating trends, generating alerts when review dates pass without updates. The inventory remains a living document; the technology ensures it stays alive.

Phase 5: Integration

Dashboard reporting for Board and capital planning. The principal risk report, the enterprise portfolio view, the capital planning integration described in Chapter 12 (Integration: Capital Planning, Strategy, and the Board) — all require synthesis of inventory data into executive-level presentations. Dashboard technology transforms the inventory into visual representations: risk heatmaps updated in real time, concentration maps, trend charts, appetite breach alerts. The Board receives the same underlying data that the Risk Identification Lead works with, presented at the appropriate level of aggregation.

Phase 6: Ongoing Cycle

Real-time event-driven trigger monitoring. Chapter 13 defined six event-driven triggers — material losses, external environment changes, new business entry, M&A, control failures, outsourcing changes. Several of these can be monitored continuously through technology rather than relying on human observation:

- **External environment scanning:** NLP systems monitoring news, regulatory announcements, and social media for events relevant to the institution's risk profile — the Wirecard scenario described in this chapter's opening
- **KRI threshold monitoring:** automated comparison of key risk indicators against green/amber/red thresholds, with immediate escalation when breaches occur
- **Peer institution event detection:** automated scanning for material loss events at comparable institutions, triggering the peer learning mechanism from Chapter 13
- **Transaction anomaly detection:** algorithms monitoring transaction flows for patterns that deviate from established baselines — unusual volumes, atypical counterparties, unexpected geographic patterns

This is where technology delivers its most distinctive value to the ongoing cycle. No team of analysts, however skilled, can monitor thousands of data points across hundreds of sources twenty-four hours a day. Technology can. The Wirecard reporting was available for five years before the fraud was exposed. The FT's journalism on the company was continuous and increasingly specific. NLP scanning of financial media would

have detected it. The question is whether the institution had the technology deployed and, critically, whether the technology's output was connected to the risk identification process.

The Four Core Technology Applications

The methodology identifies four technology applications for risk identification. Each has specific implementation requirements and limitations.

Natural Language Processing

What it does: Automated scanning of regulatory publications, news feeds, social media, and internal documents to identify emerging risk signals and sentiment shifts.

Where it adds most value: External context monitoring (Phase 1 PESTLE), event-driven trigger surveillance (Phase 6), bottom-up submission quality analysis (Phase 2), regulatory change tracking.

What it does not do: NLP identifies *signals*. It does not determine whether those signals represent genuine risks to the institution. A news article mentioning "credit risk" at a peer institution may be relevant or irrelevant — that determination requires human judgement informed by institutional context. NLP produces a filtered, prioritised information feed. The Risk Identification Lead and the methodology's analytical processes determine what to do with it.

Machine Learning and Predictive Algorithms

What they do: ML models trained on historical loss data and market indicators to forecast emerging risk conditions — loan default patterns, market dislocation precursors, operational failure indicators.

Where they add most value: Pattern detection in the starting universe (Phase 1), anomaly identification in KRI monitoring (Phase 6), scoring calibration support (Phase 3).

What they do not do: ML models are, by definition, trained on historical data. They detect patterns that have occurred before. The most dangerous risks — the ones this methodology is specifically designed to identify — are those that have no historical precedent in the institution's own data. The Delphi Method (Chapter 6) exists precisely be-

cause conventional analysis, including ML-based analysis, cannot identify risks that fall outside its training data. ML augments the methodology's evidence-based components. It does not substitute for the methodology's forward-looking components.

Anomaly Detection

What it does: Algorithms identifying unusual patterns in transaction data, trading activity, system logs, or financial statements that may signify previously unidentified risks.

Where it adds most value: Bottom-up identification support (Phase 2), KRI monitoring (Phase 6), ongoing surveillance between formal identification cycles.

Limitations: Anomaly detection identifies deviations from established patterns. An institution that has never experienced a particular type of fraud will have no baseline against which to detect it. This is where the industry loss database becomes essential — anomaly detection parameters informed by how losses have materialised at *other* institutions extend detection beyond the institution's own experience. The Punjab National Bank case described below illustrates what happens when systems operate in isolation without cross-system anomaly detection.

Robotic Process Automation

What it does: Automation of routine risk assessment tasks — data gathering from multiple source systems, template population with current metrics, threshold monitoring against defined limits, report generation.

Where it adds most value: Reducing the clerical burden of Phase 4 documentation, Phase 6 KRI monitoring, and Phase 5 Board reporting. Every hour that a risk analyst spends compiling data from disparate systems is an hour not spent on the analytical work that the methodology demands — driver analysis, control effectiveness assessment, interaction mapping.

What it does not do: RPA automates tasks. It does not perform analysis. An automated system that populates a risk template with data from source systems has not performed risk identification. It has performed data entry — precisely the compliance theatre that Chapter 7 warned against. The analytical content — the risk definition, the driver analysis, the control assessment — must come from human practitioners applying the techniques described in Chapters 6 (Top-Down Identification: Workshops, SWIFT, and Delphi) and 7 (Bottom-Up Identification: Templates, RCSA, and Specialist Processes).

Data Management Infrastructure: The Foundation Beneath the Technology

Every technology application described above depends on a foundation that most institutions underinvest in: data management infrastructure. Without it, AI and ML applications produce unreliable outputs, anomaly detection generates false positives that overwhelm human reviewers, and NLP systems surface noise rather than signal.

Chapter 5 introduced **BCBS 239** — Principles for effective risk data aggregation and risk reporting — as the regulatory benchmark for data infrastructure assessment. BCBS 239 was published in 2013, with a compliance deadline of January 2016 for global systemically important banks.³ More than a decade later, supervisory assessments consistently find material gaps in compliance. The principles that matter most for risk identification are data accuracy, completeness, timeliness, and adaptability — the ability to aggregate data across the institution and produce risk reports that reflect the enterprise position.

The methodology requires seven data management components:

| Component | Requirement |
|-------------------------------|--|
| Data governance | Formal policies defining how risk data is collected, stored, accessed, used, and retired |
| Data quality controls | Automated and manual checks for accuracy, completeness, timeliness, and consistency |
| Data security | Protection of sensitive risk data from unauthorised access, compliant with information security policies |
| Data integration | Processes to combine risk data from multiple source systems into a unified repository |
| Master data management | Single authoritative source for key entities: counterparties, products, legal entities, risk events |

| Component | Requirement |
|----------------------------------|---|
| Data lifecycle management | Retention, archival, and disposal policies compliant with regulatory requirements |
| Data auditing | Regular audits of governance compliance, quality metrics, and security controls |

The CRO function is responsible for ensuring that data management standards are adequate to support the risk identification process. Data management deficiencies should be recorded as findings in the risk inventory and tracked to remediation. This is not an IT responsibility — it is a risk management responsibility that requires IT capabilities.

The Equifax breach of 2017 illustrates what happens when data management is treated as a routine IT matter. Equifax failed to patch a known Apache Struts vulnerability for over two months.⁴ The vulnerability was in an internet-facing system holding personal data on 147 million consumers.⁵ The estimated \$1.38 billion in total breach costs — comprising the \$575–700 million regulatory settlement plus over \$1 billion in technology remediation and legal expenses⁶ — the departure of the CEO, CIO, and CISO, and the comprehensive cybersecurity overhaul that followed were consequences of a risk identification failure: software patching was treated as an operational IT task rather than a critical risk requiring executive-level urgency and a defined escalation path. The institution’s risk identification process did not flag that an unpatched vulnerability in a system holding this volume of sensitive data was itself a material risk. The data management infrastructure assessment — had the methodology’s framework been in place — would have identified this gap.

GRC Platforms: From Spreadsheets to Systems

At some point in an institution’s risk identification maturity journey, spreadsheets become inadequate. That point arrives earlier than most institutions recognise.

As Chapter 11 described, the initial risk inventory is typically built in a spreadsheet. This is the right decision at the start — the process is new, the methodology is being refined, and the cost of configuring a technology platform before the process is stable would be prohibitive. But spreadsheets have structural limitations that become acute as the methodology scales:

- **No referential integrity.** A risk owner's name in the inventory is a text string, not a link to an organisational structure. When that person changes role, the spreadsheet does not flag orphaned risks.
- **No network queries.** The risk interaction analysis from Chapter 10 (Risk Interaction: Bow-Ties, Matrices, and Concentration) requires mapping relationships between risks — triggers, amplifiers, correlations. A spreadsheet cannot represent or query a network.
- **No role-based access control.** Every participant with access to the spreadsheet has access to everything. The methodology requires that business units see their own risks and the enterprise portfolio view, but not the detail of other units' confidential submissions.
- **No concurrent access control.** Multiple users editing simultaneously creates version conflicts and data loss.
- **No automated audit trail.** The three-element audit trail (date, author, reason) must be manually maintained. In practice, it decays within cycles.

A **Governance, Risk, and Compliance (GRC) platform** addresses these limitations by providing a structured database designed for risk management workflows. The best platforms support the methodology's requirements: configurable taxonomy structures, workflow-driven risk assessment cycles, automated KRI monitoring, interaction mapping, dashboard reporting, role-based access, and full audit trails.

The selection criteria should be derived from the methodology, not from the vendor's feature list:

1. **Taxonomy flexibility.** Can the platform support the institution's three-level taxonomy (L1/L2/L3) with the ability to modify structure through the governance process described in Chapter 4 (The Risk Taxonomy)?
2. **Assessment methodology support.** Does it accommodate four-dimensional scoring (impact, likelihood, vulnerability, speed of onset) with the dominant dimension rule?

- 3. Reconciliation workflow.** Can it support the five-step reconciliation process (Chapter 8 (Reconciliation and the Enterprise Portfolio View)) — gap analysis, escalation, assignment, challenge, iteration — with documented outcomes?
- 4. Risk interaction mapping.** Can it represent the directional interaction matrix (Chapter 10) and support network analysis queries?
- 5. Integration capability.** Can it receive data feeds from source systems (trading, core banking, compliance, HR) and provide data to downstream consumers (capital planning, regulatory reporting, Board dashboards)?
- 6. Audit trail.** Does every change record date, author, and reason automatically?
- 7. Regulatory reporting.** Can it generate outputs mapped to multiple regulatory frameworks using the regulatory mapping table described in Chapter 4?

The migration from spreadsheet to platform should follow the same structured pathway as any technology implementation: feasibility assessment, data preparation, configuration aligned to the established methodology, pilot testing with a subset of the inventory, full deployment, and ongoing monitoring. The critical requirement is that the platform is configured to match the methodology — not the other way around.

When Technology Creates the Risk: Knight Capital and the 45-Minute Catastrophe

On 1 August 2012, Knight Capital Group deployed a software update to its market-making system. The deployment activated dormant legacy code that had not been removed from the production environment. Within forty-five minutes, the system executed millions of erroneous trades across 148 stocks on the New York Stock Exchange, accumulating \$440 million in pre-tax trading losses before the system was shut down — rising to \$460 million in total corporate cost including the subsequent SEC penalty and legal fees.⁷ Knight Capital, a firm with \$365 million in equity, had lost more than its entire capital base in less time than it takes to conduct a risk identification workshop.⁸

The root cause was a software deployment error — a change management failure in the technology function. But the risk identification failure was structural: the institution's pre-deployment testing and change management processes did not identify that the code release contained a configuration that would activate legacy trading logic. More critically, no automated kill-switch existed to halt trading when position accumulation exceeded predefined thresholds at the speed that algorithmic systems can generate.

This case illustrates a principle that runs through this entire chapter: technology is simultaneously an enabler of risk identification and a source of the risks being identified. The methodology must account for both.

What the methodology would have required:

The **ICT/Cyber specialist sub-process** (Chapter 7) would have included algorithmic trading systems in the information asset inventory, with change management classified as a critical operational control. The **FMEA analysis** applied to the deployment process would have examined failure modes including activation of legacy code. The **Data Quality Rating** for the control effectiveness assessment of pre-deployment testing would have been questioned if the testing protocol did not include regression testing against dormant code paths. The **KRI monitoring** framework (Chapter 13) would have included real-time position accumulation thresholds with automated circuit-breakers — not as a technology feature but as a risk control identified through the bottom-up process and documented in the risk inventory.

Most importantly, the **four-dimensional assessment** would have scored the speed of onset dimension as “Immediate” — losses accumulating at a rate that exceeded any human intervention capability. When speed of onset is Immediate and the control environment has not been tested against that speed, the Vulnerability rating cannot credibly be “Low”. The methodology forces these dimensions to be assessed explicitly, rather than allowing them to be subsumed into a single “operational risk” score that obscures the catastrophic speed at which algorithmic systems can generate losses.

Knight Capital required an emergency capital injection from a consortium of investors. It was acquired by Getco (now Virtu Financial) the following year.⁹ The firm ceased to exist — destroyed in forty-five minutes by a technology risk that its risk identification process had not identified.

When Systems Cannot Talk to Each Other: Punjab National Bank

If Knight Capital illustrates what happens when technology moves faster than human oversight, Punjab National Bank illustrates what happens when critical systems are not integrated.

Between 2011 and 2018, a rogue employee at PNB's Brady House branch in Mumbai issued unauthorised Letters of Undertaking via the SWIFT messaging system to companies controlled by diamond merchant Nirav Modi. The Letters of Undertaking — essentially guarantees to overseas banks — created \$2 billion in contingent liabilities for PNB.¹⁰ The fraud continued for seven years.

The mechanism was devastatingly simple. PNB's SWIFT messaging system operated independently of its Core Banking System. SWIFT messages creating liabilities were sent without corresponding entries appearing in the CBS. The contingent liabilities existed in one system but not the other. Periodic reconciliations between SWIFT messages and CBS records were either not performed or were ineffective. The result was that \$2 billion in guarantees were issued through one technology platform without appearing on the institution's balance sheet as recorded by another.

This is a **data integration** failure of the kind that the methodology's data management infrastructure requirements are designed to prevent. The seventh component in the table above — data auditing — specifically requires regular audits of data governance compliance. A reconciliation between SWIFT messages and CBS records is precisely the kind of cross-system integrity check that data auditing mandates. More broadly, the **master data management** requirement — a single authoritative source for key entities and risk events — would have identified the SWIFT/CBS disconnect as a structural gap in the institution's data architecture.

The aftermath was dramatic: the largest fraud in Indian banking history; Nirav Modi fled to the United Kingdom (later extradited); the employee was convicted; and the Reserve Bank of India mandated SWIFT-CBS integration across the entire Indian banking system.¹¹ The regulatory response — forced system integration — addressed the technology gap. But the underlying risk identification failure was that no one had asked the question that the methodology's ICT specialist sub-process (Chapter 7) requires: are all systems that create financial commitments integrated with the systems that record them?

Ethical and Regulatory Guardrails

The deployment of AI and ML in risk identification introduces risks of its own, and the methodology must account for them.

Bias. Algorithms trained on historical data inherit the biases embedded in that data. If historical loss data underrepresents certain geographies, product types, or customer segments — because those segments were not monitored, not because they were risk-free — ML models will systematically under-identify risks in those areas. Models must be tested for bias across portfolios, geographies, and segments, with documented results and remediation of identified biases.

Transparency and explainability. Models used in risk identification must be interpretable. When an ML system flags a pattern as a potential emerging risk signal, the Risk Identification Lead and the CRO must be able to understand *why* the system flagged it. Black-box models that produce outputs without interpretable logic are not acceptable for regulatory-facing risk assessments without appropriate model risk governance. The model risk specialist sub-process described in Chapter 7 applies to AI/ML models used for risk identification, not just models used for pricing or capital calculation.

Data privacy. All AI and ML applications must comply with applicable data protection regulations — GDPR in Europe, local equivalents elsewhere. NLP systems scanning social media or external sources must operate within legal boundaries for data collection and processing.

Model risk. This is the guardrail that institutions most frequently overlook. AI and ML models used for risk identification are themselves subject to model risk. They must be included in the institution's model risk management framework (Chapter 7), with independent validation, performance monitoring, and revalidation when conditions change. An ML model trained on pre-pandemic data may be actively misleading in a post-pandemic environment. The model risk framework must ensure that AI/ML tools used for risk identification are subject to the same governance as any other model — because the consequences of a risk identification model producing false reassurance are at least as severe as the consequences of a pricing model producing an incorrect valuation.

The Implementation Pathway

For institutions that have not yet integrated technology into their risk identification process, The methodology prescribes a structured pathway:

- 1. Feasibility assessment.** Evaluate organisational readiness, data availability, and expected return on investment. An institution with poor data quality (multiple source systems, no master data management, no data governance) should invest in data infrastructure before AI/ML deployment. Deploying ML on unreliable data produces unreliable outputs with a veneer of analytical sophistication — the technology equivalent of compliance theatre.
- 2. Data preparation.** Ensure high-quality, relevant, and unbiased data is available for model training. This may require a data remediation programme before any technology deployment. The Data Quality Ratings from the risk inventory (Chapter 9) provide an honest assessment of the institution's data readiness.
- 3. Algorithm selection.** Choose or develop models suited to the specific risk types being targeted. NLP for media scanning is a different technical challenge from anomaly detection in transaction data. Each application should be scoped, specified, and selected against the specific phase and activity it will support.
- 4. Pilot testing.** Conduct small-scale testing to validate effectiveness before enterprise rollout. Pilot with a single business unit or a single phase of the methodology. Measure both detection effectiveness (does it find signals that human analysts confirm as relevant?) and false positive rate (does it generate so many irrelevant flags that analysts ignore the system?).
- 5. Full deployment.** Integrate validated models into the existing risk management framework. The outputs must flow into the methodology's existing processes — the starting universe, the briefing packs, the KRI dashboards, the event-driven trigger framework — not into a parallel technology-driven process that operates independently.
- 6. Ongoing monitoring.** Continuously monitor model performance. Retrain models as conditions change. Technology that worked in a low-interest-rate environment may produce different results in a rising-rate environment. The ongoing cycle (Chapter 13) applies to technology as much as it applies to the rest of the methodology.

What Technology Does Not Replace

Throughout this chapter, the distinction between what technology can do and what it cannot has been deliberate. This distinction matters because the most dangerous outcome of technology investment in risk identification is not failure — it is misplaced confidence.

Technology does not replace the SWIFT workshop. No algorithm can replicate the moment when a structured prompt causes a senior executive to consider a risk they had not previously articulated. Technology does not replace the Delphi panel. No ML model trained on historical data can identify a risk that has never occurred — that requires human imagination constrained by structured methodology. Technology does not replace the reconciliation process. No automated system can navigate the political dynamics of assigning risk ownership between business units or challenging a CRO's assessment of materiality.

Technology does not replace the judgement that the methodology is designed to structure and discipline. It extends the reach of that judgement. It provides better information, faster. It automates surveillance that would otherwise depend on human attention spans. It scales processes across institutions too large and complex for purely manual approaches. But it does so in service of the methodology, not as a substitute for it.

The institutions that deploy technology most effectively are those that have a functioning methodology first. The institutions that deploy technology least effectively are those that hope technology will compensate for the absence of a methodology. The first group uses AI to enhance human judgement. The second group uses AI to avoid exercising it.

The Bridge

The technology enablers described in this chapter operate within a regulatory environment that defines minimum standards for risk identification, data management, and model governance. Chapter 15 maps the regulatory landscape — the sixteen frameworks across jurisdictions that mandate, constrain, and shape how institutions identify their risks. From the Fed's SR 15-18 and the PRA's SS31/15 to the ECB's SREP expectations and the EBA's Guidelines on ICT risk, each regulatory framework creates specific

requirements that the methodology must satisfy. Chapter 15 (The Regulatory Landscape) examines those requirements, maps them to the methodology's six phases, and provides the regulatory traceability that every supervisor will demand.

1. Dan McCrum, "Wirecard's Suspect Accounting Practices Revealed," *Financial Times*, 2015. The FT's investigative series on Wirecard began in early 2015 and continued through 2020, culminating in the company's insolvency filing on 25 June 2020.
2. KPMG, *Independent Special Investigation: Wirecard AG*, April 2020; EY subsequently confirmed that €1.9 billion in cash balances purportedly held at Philippine trustee banks could not be verified. See also: German Parliamentary Inquiry Committee on the Wirecard scandal, *Bundestag Drucksache* 19/30900, 2021.
3. Basel Committee on Banking Supervision, *Principles for Effective Risk Data Aggregation and Risk Reporting*, BCBS 239, January 2013. The January 2016 compliance deadline applied to global systemically important banks (G-SIBs); domestic systemically important banks were given an additional three years.
4. U.S. House of Representatives, Committee on Oversight and Government Reform, *The Equifax Data Breach*, Staff Report, 115th Congress, December 2018. The Apache Struts vulnerability (CVE-2017-5638) was publicly disclosed on 7 March 2017; Equifax did not patch the affected system until after the breach was discovered in late July 2017.
5. Equifax Inc., "Equifax Announces Cybersecurity Incident Involving Consumer Information," press release, 7 September 2017. The company initially reported approximately 143 million affected consumers; the figure was subsequently revised upward to 147 million.
6. Federal Trade Commission, "Equifax to Pay \$575 Million as Part of Settlement with FTC, CFPB, and States Related to 2017 Data Breach," press release, 22 July 2019. The settlement included up to \$425 million in a consumer restitution fund (later raised to a potential \$700 million cap), \$175 million in payments to states, and additional compliance and remediation costs. Equifax Inc., SEC filings (Form 10-K, fiscal years 2017–2019), report cumulative breach-related costs exceeding \$1.4 billion.
7. U.S. Securities and Exchange Commission, *In the Matter of Knight Capital Americas LLC*, Administrative Proceeding File No. 3-15570, 16 October 2013. The SEC found that Knight deployed untested code to production servers, which reactivated dormant functionality and generated approximately 4 million erroneous executions in 154 stocks over 45 minutes on 1 August 2012, resulting in a pre-tax loss of approximately \$440 million.
8. Knight Capital Group Inc., Form 10-Q filed with the SEC for the quarter ended 30 June 2012, reporting stockholders' equity of approximately \$365 million prior to the 1 August 2012 incident.
9. Knight Capital Group Inc., Form 8-K filed with the SEC, 19 December 2012, announcing the merger agreement with GETCO Holding Company, LLC. The merger was completed on 1 July 2013. GETCO subsequently rebranded as KCG Holdings Inc. and was acquired by Virtu Financial Inc. in 2017.
10. Central Bureau of Investigation (India), First Information Report and charge sheet filed in the Punjab National Bank fraud case, February 2018. PNB disclosed fraudulent Letters of Undertaking totalling approximately ₹14,000 crore (approximately \$2 billion at the time) issued through the SWIFT system between 2011 and 2018. See also: Reserve Bank of India, inspection findings and enforcement actions relating to Punjab National Bank, 2018.
11. Reserve Bank of India, circular on "SWIFT-Related Operational Controls," RBI/2017-18/133, 20 February 2018, directing all banks to integrate SWIFT infrastructure with their Core Banking Systems. Nirav Modi was arrested in London in March 2019 and extradited to India in April 2024. The Westminster Magistrates' Court approved the extradition in February 2021, with subsequent appeals exhausted by 2024.

The Regulatory Landscape

Show Me Your Regulatory Traceability

A PRA supervisor opens a thick examination file and asks a single question: “Walk me through how your risk identification process satisfies SS31/15.”

The question is simple. The answer, for many institutions, is not.

The supervisor is not asking whether the institution identifies risks. Every bank identifies risks. She is asking whether the process that produced those identifications can be traced — requirement by requirement — to the specific provisions of the regulatory framework under which the institution is supervised. Can the Risk Identification Lead point to the section of the process that delivers what SS31/15 requires? Can they demonstrate that the output meets the standard’s expectations for comprehensiveness, for forward-looking identification, for stress testing as an identification tool?

A methodology designed with regulatory traceability built in from the start can answer this immediately. When every major regulatory requirement maps to a specific phase and activity within the methodology, the answer does not need to be constructed — it requires only opening the regulatory mapping appendix.

Most institutions cannot do this. They have a risk identification process. They have a long list of regulations they are subject to. What they lack is the bridge between the two — a systematic mapping that demonstrates, for each applicable framework, exactly how the process satisfies its requirements.

This gap matters. It matters because the regulatory landscape for bank risk identification is not a single framework. It is sixteen distinct frameworks across multiple jurisdictions — global standards from the Basel Committee, EU directives and EBA guidelines, PRA supervisory statements, US supervisory letters, and FCA conduct guidance — each imposing specific requirements that the identification process must satisfy. For an institution operating across the European Union, the United Kingdom, and the United States,

these requirements overlap, sometimes reinforce each other, and occasionally pull in different directions. A methodology that satisfies one framework but ignores another is a methodology that will fail its next regulatory examination.

This chapter maps the sixteen regulatory frameworks that govern bank risk identification, examines what each requires, and demonstrates how the methodology presented in this book satisfies all of them simultaneously. It is the chapter that answers the supervisor’s question before she asks it.

The Sixteen Frameworks

The sixteen frameworks fall into four jurisdictional tiers. The first tier — global standards issued by the Basel Committee on Banking Supervision — sets the baseline that all jurisdictions adopt and extend. The second tier — European Union regulations and EBA guidelines — transposes and elaborates these standards into a prescriptive legal architecture of extraordinary depth. The third tier — United Kingdom PRA supervisory statements and FCA guidance — applies a principles-based but increasingly detailed approach with personal accountability. The fourth tier — US Federal Reserve supervisory letters and OCC guidelines — imposes arguably the most operationally prescriptive requirements through the capital planning process.

| # | Framework | Jurisdiction | Core Risk Identification Requirement |
|---|--|--------------|--|
| 1 | BCBS Corporate Governance Principles — Principle 7 | Global | Bank-wide, ongoing risk identification |
| 2 | BCBS 239 — Risk Data Aggregation | Global | Data completeness, accuracy, aggregation capability |
| 3 | BCBS PSMOR — Operational Risk | Global | RCSA, KRIs, external loss data, near-miss identification |
| 4 | CRR III / CRD VI — The Banking Package | EU | ICAAP, ESG integration, consolidated identification |

| # | Framework | Jurisdiction | Core Risk Identification Requirement |
|----|---|--------------|--|
| 5 | EBA Internal Governance (GL/2021/05) | EU | Holistic risk view, M&A identification, AML/CFT integration |
| 6 | ECB SREP Methodology | EU/SSM | Comprehensive risk inventory, materiality determination |
| 7 | DORA (EU 2022/2554) | EU | ICT asset identification, dependency mapping, third-party register |
| 8 | EBA Outsourcing Guidelines (GL/2019/02) | EU | Pre-outsourcing risk assessment, concentration risk |
| 9 | EBA ESG Risk Guidelines (GL/2024/01) | EU | Transmission channels, multi-horizon identification |
| 10 | EBA ICT and Security Risk (GL/2019/04) | EU | ICT risk classification, information asset inventory |
| 11 | EU AMLD6 — Anti-Money Laundering | EU | Enterprise-wide ML/TF risk assessment |
| 12 | PRA SS31/15 — ICAAP | UK | Comprehensive identification, stress testing, reverse stress testing |
| 13 | PRA Step-In Risk (PS5/25) | UK | Non-contractual support obligations identification |
| 14 | Fed SR 15-18 — CCAR | US | Quarterly Material Risk Inventory, capital linkage |
| 15 | OCC Heightened Standards (12 CFR Part 30, App. D) | US | Front-line accountability, independent challenge |

| # | Framework | Jurisdiction | Core Risk Identification Requirement |
|----|---|--------------|--|
| 16 | FCA Conduct Risk — Five Conduct Questions | UK | Business-led conduct risk identification |

Chapter 2 (The Foundations: Standards and Frameworks) examined the three foundational standards — ISO 31000, ISO 31010, and COSO ERM — that provide the architecture, techniques, and enterprise lens upon which the methodology is built. These sixteen frameworks are the regulatory instruments that mandate how that architecture must be applied in practice. They are the frameworks that mandate, constrain, and shape how institutions identify their risks.

The Global Baseline: What the Basel Committee Requires

The Basel Committee on Banking Supervision does not directly supervise banks. Its standards are transposed into national legislation — through the CRR in Europe, the PRA Rulebook in the United Kingdom, and the Dodd-Frank Act in the United States. But the definitions and identification requirements set by the BCBS form the genetic code of global risk management. Every jurisdiction starts here.

BCBS Corporate Governance Principles — Principle 7 states the foundational mandate: “Risks should be identified, monitored and controlled on an ongoing bank-wide and individual entity basis.”¹ Four words in that sentence do significant work. *Ongoing* means not annual — it requires continuous or at minimum quarterly identification. *Bank-wide* means consolidated, across every legal entity, subsidiary, and branch. *Individual entity basis* means the enterprise view does not excuse the absence of entity-level identification. And *identified* is listed first, before monitoring and controlling, because identification is the prerequisite for everything that follows.

Principle 7 also requires identification of risks arising from mergers and acquisitions, new products, and changes to organisational structure.² These are the event-driven triggers that Chapter 13 (The Ongoing Cycle: Refresh, Events, and Audit) built into Phase 6. They are not discretionary additions to the methodology — they are regulatory requirements under the global baseline.

BCBS 239 — Principles for Effective Risk Data Aggregation and Risk Reporting connects risk identification to data infrastructure. Published in 2013 with a compliance deadline of January 2016 for G-SIBs,³ it remains — a decade later — one of the most frequently cited areas of supervisory concern. The reason is straightforward: if data cannot be aggregated accurately and quickly, risks remain invisible.

Principle 4 (Completeness) requires banks to capture and aggregate risk data across business line, legal entity, asset type, industry, region, and other groupings.⁴ This means the risk identification process must tag every risk with these metadata attributes — a requirement that Chapter 11 (Documentation: The Living Risk Inventory)'s fourteen-field risk inventory was designed to satisfy. Principle 7 (Accuracy) requires data reliability, which maps directly to Chapter 9 (Assessment — Scoring, Multi-Dimensional Impact, and Data Quality)'s Data Quality Rating system. And the standard's requirement for ad-hoc reporting capability — generating aggregate risk information during stress periods — demands the on-demand identification capability built into Phase 6's event-driven triggers.

BCBS Principles for the Sound Management of Operational Risk (PSMOR), revised in 2021,⁵ provides the toolkit for operational risk identification. It mandates specific instruments: Risk and Control Self-Assessments (RCSA), Key Risk Indicators (KRIs), external loss data analysis, business process mapping, and event management.⁶ Critically, the 2021 revision places heavy emphasis on ICT and third-party risk identification, and introduces a **near-miss requirement** — identifying events that did not result in financial loss but exposed a vulnerability.⁷ This shifts identification from lagging indicators to leading indicators. Chapter 7 (Bottom-Up Identification: Templates, RCSA, and Specialist Processes)'s ten specialist sub-processes — particularly the RCSA, ICT, and third-party outsourcing processes — were designed to deliver exactly what PSMOR requires.

The European Union: The Layered Architecture

The EU's approach to risk identification regulation is distinctive in its depth and prescriptiveness. Where the Basel Committee sets principles, the EU builds a legislative architecture of layered complexity: directly applicable regulations (CRR), directives requiring national transposition (CRD), and detailed guidelines from the European Banking Authority that supervisors use to assess compliance. For a bank operating under the Single Supervisory Mechanism, these eight frameworks create multiple overlapping layers of identification requirements.

CRR III / CRD VI

The 2024 Banking Package represents the EU's transposition of the final Basel III reforms.⁸ For risk identification, two innovations are significant.

First, the **Output Floor** limits how much internal models can reduce capital requirements below the standardised approach, effectively mandating a dual view of risk. Banks must identify risk through their own models and simultaneously through standardised regulatory metrics. Discrepancies between these two views serve as identification signals for model risk or data anomalies — a principle that aligns with the methodology's Data Quality Rating and independent challenge mechanisms.

Second, **CRD VI explicitly mandates that ESG risks be identified not as a standalone category but as drivers of traditional risk types** — credit, market, operational, liquidity.⁹ This is precisely the approach Chapter 5 (Setting the Context: External, Internal, and Risk Culture) established through the six transmission channels framework: physical risks and transition risks transmitting through existing risk categories rather than existing in isolation.

EBA Internal Governance Guidelines (GL/2021/05)

These guidelines impose three requirements that directly shape the methodology.¹⁰ First, the **holistic view** requirement — Section 17 requires risk identification to aggregate across legal entities and risk types to identify cross-cutting themes.¹¹ This is the enterprise portfolio view of Chapter 8 (Reconciliation and the Enterprise Portfolio View)'s reconciliation process. Second, the detection of **unapproved exposures** — Section 20.4 places a specific duty on the risk management function to identify instances where the business has taken risk outside the agreed appetite.¹² This requires the active monitoring capability built into Phase 6's KRI framework. Third, the 2021 revision explicitly integrates **ML/TF risk** into the general risk management framework, requiring financial crime to appear in the enterprise-wide risk inventory — not isolated in a Compliance silo.

The guidelines also mandate risk identification in the context of mergers and acquisitions and new product approvals — requirements that Chapter 13's event-driven triggers and Chapter 7's specialist sub-processes were designed to deliver.

ECB SREP Methodology

The ECB's Supervisory Review and Evaluation Process assesses the approximately 110 significant institutions under its direct supervision.¹³ Built on CRR/CRD and EBA guidelines, the SREP is the mechanism through which those requirements are tested. The ECB examines not only what risks have been identified but how the institution determined which are material, what data quality underpins the assessments, and whether the identification process captures all relevant risk categories.

For a European G-SIB with London, Swiss, and US operations, the methodology must simultaneously satisfy FINMA's SREP-equivalent assessment for the Swiss entity, the Fed's CCAR requirements for the US operations, and the PRA's expectations for the London branch. The ECB SREP requires a comprehensive risk inventory with explicit materiality determination — the process described in Chapter 9's four-dimensional assessment and materiality framework. Institutions that cannot demonstrate compliance face direct supervisory consequences: increased capital requirements, restrictions on activities, and mandatory remediation programmes.

DORA — Digital Operational Resilience Act (EU 2022/2554)

DORA entered into force in January 2025¹⁴ and represents the most prescriptive regulatory instrument for ICT risk identification. Article 8 requires financial entities to “identify, classify and adequately document all ICT-supported business functions, roles and responsibilities... and identify all ICT assets.”¹⁵ This is not risk-level identification — it is **asset-level identification**.

DORA requires banks to map the links between ICT assets and Critical or Important Functions, creating a dependency map that exposes single points of failure. And the **Register of Information** requirement — a complete register of all contractual arrangements with ICT third-party service providers — demands that banks identify concentration risks in their technology supply chain. If all core banking systems sit on one cloud provider in one region, that is a concentration risk that DORA requires to be identified. Chapter 7's ICT specialist sub-process, Chapter 14 (Technology: AI, ML, and Data Analytics)'s data management infrastructure requirements, and the enterprise portfolio view's concentration analysis collectively address these requirements.

EBA Outsourcing Guidelines (GL/2019/02)

The guidelines require a **pre-outsourcing risk assessment** before any contract is signed, including a determination of whether the arrangement involves a critical or important function.¹⁶ This classification is itself a risk identification step — it determines the level of ongoing monitoring and the restrictions on the arrangement. The guidelines also require identification of concentration risks at both the individual provider level and at the sector level. Chapter 7's third-party and outsourcing specialist sub-process was designed to deliver these requirements.

EBA ESG Risk Guidelines (GL/2024/01)

Applicable from 2026/2027,¹⁷ these guidelines operationalise the integration of environmental, social, and governance risks into the risk management framework. The core identification requirement is the mapping of **transmission channels**: exactly how physical risks (floods, storms) and transition risks (policy changes, technology shifts, market preferences) impact counterparties' ability to repay or the institution's own operations. This requires new data — geolocation data for flood risk, carbon intensity data for transition risk — and new time horizons. Risk identification must span short-term (1-3 years), medium-term (3-10 years), and long-term (10-30 years) horizons.

Chapter 5's climate and ESG risk identification framework, with its six transmission channels and three time horizons, was built on these requirements and on the earlier **ECB Guide on Climate-Related and Environmental Risks (2020)**,¹⁸ which established the supervisory expectations that the EBA guidelines subsequently codified into detailed implementation standards.

EBA ICT and Security Risk Guidelines (GL/2019/04)

These guidelines preceded DORA and remain relevant for the identification methodology — specifically the requirement that ICT risks be classified using the same taxonomy and scoring methodology applied to all other risk types.¹⁹ They mandate an **information asset inventory**, annual ICT risk assessments, and the integration of ICT risk into the enterprise risk framework. Chapter 7's ICT specialist sub-process applies the same four-dimensional assessment, the same Data Quality Rating, and the same reconciliation process to ICT risks as to credit, market, or operational risks. This prevents ICT risk from remaining isolated in a technology silo disconnected from enterprise risk management.

EU AMLD6

The Sixth Anti-Money Laundering Directive²⁰ mandates an **enterprise-wide ML/TF risk assessment** across four vectors: customers, countries, products, and channels. The identification process must analyse transaction data against national risk assessments published by member state governments. Chapter 7's AML/CFT specialist sub-process addresses this by requiring financial crime risks to be identified within the enterprise risk taxonomy — not confined to a Compliance-only register — and assessed using the same four-dimensional scoring methodology. The AML/CFT sub-process also requires that new product approvals include an explicit ML/TF risk assessment before launch, delivering on the EBA Internal Governance Guidelines' requirement for new product risk identification.

The United Kingdom: Accountability and Proportionality

The UK's post-Brexit regulatory architecture is principles-based but increasingly detailed, with three characteristics that distinguish it from the EU approach: personal accountability through the Senior Managers regime, a proportionality principle that tailors requirements to firm size and complexity, and a willingness to lead on emerging risk types — particularly model risk, operational resilience, and step-in risk.

PRA SS31/15 is the supervisory statement that governs the Internal Capital Adequacy Assessment.²¹ For risk identification, its requirements are the most consequential in the UK framework. The PRA views stress testing not merely as a capital calculation tool but as a **primary method of risk identification**. By subjecting the business model to severe but plausible scenarios, banks identify latent risks invisible in benign conditions. **Reverse stress testing** — identifying the scenarios that would render the business model non-viable — is explicitly expected. The methodology addresses this through Chapter 12 (Integration: Capital Planning, Strategy, and the Board)'s capital planning integration, which links the risk inventory to stress scenario design and includes reverse stress testing drawing on Chapter 10 (Risk Interaction: Bow-Ties, Matrices, and Concentration)'s interaction analysis and cascade pathways.

The consequences of deficient risk identification under SS31/15 are direct and quantifiable. As Chapter 12 described, the PRA can impose **Pillar 2A capital add-ons with scalars up to 40%** of the Pillar 2A requirement for firms whose risk identification is

assessed as inadequate. This is not a theoretical penalty. Institutions regularly receive supervisory feedback citing specific gaps in their identification process — risks that should have been in the inventory but were not, risk categories that lacked adequate coverage, enterprise-level concentration risks that were identified at entity level but not aggregated. The methodology's enterprise portfolio view, regulatory traceability mapping, and documentation standards are designed to prevent exactly these findings.

PRA Step-In Risk — introduced through PS5/25 in 2025²² — requires the explicit identification of financial support obligations that extend beyond contractual requirements. These are the reputational compulsions to support unconsolidated entities — sponsored funds, structured vehicles, affiliated entities — during stress. Identifying these non-contractual obligations is difficult precisely because they do not appear in legal agreements or financial statements. They exist in the relationship between the institution's brand and the entity's investors. Chapter 4 (The Risk Taxonomy)'s risk taxonomy includes step-in risk as a Level 1 category, and Chapter 5's internal context assessment examines the institution's relationship with unconsolidated entities as part of the organisational structure review.

FCA Conduct Risk — the Five Conduct Questions represents a fundamentally different approach to risk identification.²³ Where prudential frameworks ask “what could go wrong financially?”, the FCA asks “what incentives could lead to poor customer outcomes?” The Five Conduct Questions framework requires business-level identification of conduct drivers — sales incentives, power dynamics, conflicts of interest, product design that profits from customer inertia. Product governance rules require identification at the design stage: who is the target market, and what are the risks if the product is sold outside that market? Chapter 7's conduct risk specialist sub-process, with its desk-by-desk granularity, was designed to deliver this requirement.

The FCA's approach is complemented by the **Senior Managers and Certification Regime (SM&CR)**,²⁴ which personalises responsibility for risk identification. Every Senior Manager's Statement of Responsibilities specifies what they are accountable for. If a risk emerges in a business unit that was not identified, the regulator reviews the Statement of Responsibilities to determine who is liable. The defence is proving “reasonable steps” — and a robust, documented risk identification process is the primary evidence of reasonable steps.

The United States: Capital Planning as Risk Identification

The US regulatory framework imposes arguably the most operationally prescriptive risk identification requirements globally, driven by the capital planning process.

Fed SR 15-18 governs the Comprehensive Capital Analysis and Review (CCAR) for the largest US banking organisations.²⁵ Its requirements transform risk identification from a periodic exercise into a continuous discipline with direct capital consequences. Unlike many other jurisdictions where annual identification cycles are the norm, SR 15-18 expects a dynamic, quarterly risk identification process that feeds directly into capital planning.

The **Material Risk Inventory** is the centrepiece. For every risk in the inventory, the bank must map it to the stress test: is it captured in the scenario? Is it captured in the P&L model? Or does it require a separate capital add-on? This “map to capital” requirement ensures no identified risk is left unfunded. Integrating the risk identification process with CCAR means the Material Risk Inventory becomes the foundation for stress scenario design — if a risk is identified as material, the scenario must address it. If the scenario reveals a risk not in the inventory, the inventory must be updated. This bidirectional linkage between identification and capital planning is what Chapter 12 described as the methodology’s integration architecture.

OCC Heightened Standards (12 CFR Part 30, Appendix D)²⁶ apply to large national banks and are distinctive in two respects. First, they codify the Three Lines Model with explicit risk identification duties for each line. **Front-line units — the business — are responsible for assessing material risks in their activities.** They cannot delegate identification to the risk function. This is the same principle that underlies the methodology’s bottom-up identification track (Chapter 7), where business units actively identify risks in their own operations rather than merely receiving and approving risk lists produced by the centre.

Second, the OCC requires **effective challenge** — Independent Risk Management must identify and assess material aggregate risks independently. If its assessment diverges from the front line’s, the divergence must be reported. This dual identification system is structurally equivalent to the methodology’s top-down/bottom-up reconciliation (Chapter 8), which was designed to create exactly the productive tension that the OCC requires.

The Regulatory Traceability Matrix

The purpose of the preceding analysis is not academic. It is to demonstrate that every requirement imposed by these sixteen frameworks maps to a specific phase and activity within the methodology. The following matrix provides this mapping — the regulatory traceability that every supervisor will demand.

| Framework | Phase 1 | Phase 2 | Phase 3 | Phase 4 | Phase 5 |
|--------------------------------|---------------------------|------------------------------|---------------------------|------------------------------|------------------------------|
| BCBS Principle 7 | Bank-wide scope | Entity/portfolio/BU coverage | — | — | Board reporting |
| BCBS 239 | Data infrastructure | — | DQ Rating | 14-field metadata | Aggregated reporting |
| BCBS PSMOR | — | RCSA, KRI design | — | Event management | — |
| CRR III / CRD VI | ESG transmission channels | Dual-track identification | Output Floor comparison | — | ICAAP/ILAAP |
| EBA Internal Governance | Internal context | M&A trigger | — | AML in inventory | Holistic enterprise view |
| ECB SREP | Context establishment | Comprehensive coverage | Materiality determination | Full documentation | Capital/strategy integration |
| DORA | ICT asset identification | ICT sub-process | — | ICT register, dependency map | — |
| | — | | | | — |

| Framework | Phase 1 | Phase 2 | Phase 3 | Phase 4 | Phase 5 |
|---------------------------------|-----------------------------------|------------------------------|-------------------------------|--------------------------|----------------------------------|
| EBA Outsourcing | | Third-party sub-process | Concentration assessment | Outsourcing register | |
| EBA ESG | Transmission channels, 3 horizons | Climate in workshops | Multi-horizon scoring | — | Risk appetite integration |
| EBA ICT | Information asset inventory | ICT risk assessment | — | Annual ICT documentation | — |
| EU AMLD6 | — | AML/CFT sub-process | — | EWRA documentation | — |
| PRA SS31/15 | Context for stress testing | Comprehensive identification | Materiality, stress scenarios | — | Capital planning, reverse stress |
| PRA Step-In Risk | Unconsolidated entity review | Support obligations ID | — | Step-in documentation | Capital treatment |
| Fed SR 15-18 | — | Quarterly re-identification | — | Material Risk Inventory | Map to capital/scenarios |
| OCC Heightened Standards | — | Front-line + IRM dual ID | — | — | Effective challenge reporting |
| FCA Conduct Risk | — | Conduct sub-process | Conduct impact scoring | — | — |

Chapter 2 established a phase-to-standards traceability table mapping the six phases to ISO 31000, ISO 31010, COSO ERM, and BCBS Corporate Governance Principles. This matrix extends that traceability to all sixteen regulatory frameworks. Together, the two

tables demonstrate that the methodology was designed not as a theoretical best-practice framework that happens to overlap with regulation. It was designed to satisfy all applicable frameworks simultaneously — because that is the reality of operating a bank in any major jurisdiction.

The matrix also serves the **regulatory vs economic risk gap analysis** introduced in Chapter 12. The sixteen frameworks define the “regulatory” half of that analysis — what each regulator requires in terms of risk identification, capital treatment, and supervisory reporting. The methodology’s four-dimensional assessment defines the “economic” half — the institution’s own view of each risk’s severity, likelihood, speed, and interconnectedness. Where these two views diverge — a sovereign exposure treated as zero risk weight under Basel but assessed as a material concentration risk through the methodology — the gap analysis escalates to the Board. The regulatory mapping table translates between the two views, and the traceability matrix demonstrates that the methodology covers every requirement the gap analysis must test against.

The Regulatory Taxonomy Divergence Problem

Chapter 4 introduced a challenge that this chapter must address directly: regulators do not agree on how to categorise risks.

The Basel Committee defines three Pillar 1 risk categories — credit, market, operational — with a broad Pillar 2 requirement to identify additional risks. The EBA uses a more granular categorisation including ICT risk as a distinct category, ESG as a cross-cutting driver, and ML/TF as integrated within operational and compliance risk. The PRA recognises step-in risk as a category that neither the EBA nor the Fed explicitly names. The Fed’s CCAR process uses its own risk typology aligned with stress scenario design — credit losses, trading and counterparty losses, operational risk losses, pre-provision net revenue effects — which does not map one-to-one onto any prudential taxonomy. And the FCA’s conduct risk framework uses a behavioural lens — drivers of poor customer outcomes — that has no counterpart in any prudential taxonomy.

For an institution operating across these jurisdictions, this divergence creates a practical problem. The internal taxonomy cannot be four different things simultaneously. It must be one coherent classification structure that can be translated into the language of each applicable regulator.

Chapter 4's **regulatory mapping table** is the solution. The institution maintains a separate document — active, not aspirational — that maps every node of the internal taxonomy to the corresponding category used by each regulator. When the PRA asks about step-in risk, the mapping table shows which L1/L2/L3 categories contain those exposures. When the Fed asks about the Material Risk Inventory's coverage of operational risk, the mapping table translates the internal taxonomy's more granular classification into the Fed's categories. When the EBA examiner asks how ICT risks are identified, the mapping table points to the ICT entries in the same enterprise risk inventory that contains credit and market risks.

The Risk Identification Lead is responsible for maintaining the currency of this mapping table. Every time a regulator introduces a new category — as the PRA did with step-in risk in 2025, as the EBA has done with ESG risks — the mapping table must be updated to show how the internal taxonomy captures it. If the internal taxonomy lacks coverage, that gap is itself a risk identification finding requiring taxonomy amendment.

For a G-SIB operating across jurisdictions, this divergence is acute. The London operation is supervised by the PRA. The Swiss entity by FINMA. The US operations by the Fed. Different divisions maintain different numbers of risk categories — different classifications for the same exposures. The regulatory mapping table is not an optional administrative exercise — it is the mechanism that makes multi-jurisdictional compliance possible without maintaining separate risk identification processes for each jurisdiction.

One Process, Multiple Jurisdictions

Consider the practical reality. A European bank headquartered in the eurozone, with a PRA-regulated subsidiary in London and a Fed-regulated branch in New York. The parent is subject to the ECB SREP, CRR III/CRD VI, EBA guidelines on internal governance, ESG, ICT, outsourcing, and AML/CFT. The London subsidiary is subject to PRA SS31/15, step-in risk requirements, SM&CR, and FCA conduct risk. The New York branch is subject to Fed SR 15-18 and OCC Heightened Standards.

This is not a hypothetical. This is the operating reality for every major European banking group.

The methodology addresses this through three mechanisms.

First, **a single process with jurisdictional overlays**. The six-phase methodology is the same everywhere. Phase 1 establishes context. Phase 2 identifies risks through top-down and bottom-up tracks. Phase 3 assesses. Phase 4 documents. Phase 5 integrates with capital planning and strategy. Phase 6 maintains the ongoing cycle. What varies by jurisdiction is the specific content within each phase — the PRA subsidiary runs reverse stress testing in Phase 5 that the ECB parent may not require in the same form; the US branch runs quarterly re-identification in Phase 6 that the annual EU cycle does not mandate at the same frequency. But the process architecture is identical.

Second, **the regulatory mapping table** translates between taxonomies. The internal taxonomy is the institution's own — designed for the process, not for any single regulator. The mapping table converts this internal taxonomy into the language each regulator expects. One identification process, one taxonomy, multiple regulatory translations.

Third, **the enterprise portfolio view** satisfies the consolidation requirement that every framework shares. BCBS Principle 7 requires bank-wide identification. The ECB SREP requires a comprehensive risk inventory. PRA SS31/15 requires consolidated identification for ICAAP. Fed SR 15-18 requires a Material Risk Inventory. The enterprise portfolio view — Chapter 8's reconciled, consolidated picture of the institution's risk landscape — satisfies all four simultaneously.

The cost of getting multi-jurisdictional risk identification wrong was illustrated by Standard Chartered's experience with sanctions enforcement. Operating across multiple jurisdictions with different regulatory expectations for financial crime risk identification, the institution's failure to maintain consistent identification standards across its global operations contributed to a \$667 million fine. The methodology's insistence on one process with jurisdictional overlays — rather than separate processes for each jurisdiction — exists precisely to prevent this kind of fragmentation.²⁷

The Consequences of Failure

Regulatory frameworks are not advisory. They carry consequences.

Under **PRA SS31/15**, a supervisory assessment that the risk identification process is inadequate triggers Pillar 2A capital add-ons. As Chapter 12 described, the PRA can apply scalars of up to 40% of the Pillar 2A requirement. For a large UK bank, this can represent billions of pounds in additional capital that must be held — capital that cannot be

deployed, distributed, or used for lending. The cost of deficient risk identification is not a fine. It is a permanent drag on the institution's economics until the deficiency is remediated and the supervisor is satisfied.

Under **Fed SR 15-18**, the consequence is the CCAR qualitative objection — a determination that the institution's capital planning process, including its risk identification inputs, is inadequate. A qualitative objection restricts capital distributions: share buybacks and dividend increases are blocked. For a publicly traded US bank, this is a market event.

Under **SM&CR**, the consequence is personal. If a risk materialises in a business unit and the Senior Manager responsible cannot demonstrate reasonable steps to identify it, they face individual enforcement action — fines, prohibition from the industry, and public censure.

These consequences share a common structure. The regulator does not penalise the crystallisation of risk — risk is inherent in banking. The regulator penalises the failure to identify it. An institution that identifies a risk, assesses it accurately, holds appropriate capital, and monitors it through the ongoing cycle has discharged its regulatory obligation — even if the risk eventually materialises. An institution that fails to identify that same risk has not.

The methodology's regulatory traceability matrix is therefore not a compliance artefact. It is evidence. It is the documentation that demonstrates — framework by framework, requirement by requirement — that the institution's risk identification process satisfies what every applicable regulator requires. When the supervisor opens the examination file, the traceability matrix is the answer.

The Bridge

The sixteen frameworks analysed in this chapter exist for a reason. They were developed — most of them — in response to bank failures. Basel III emerged from the 2008 financial crisis. DORA from the recognition that digital disruption could produce systemic instability. SM&CR from the realisation that no individual was accountable for the failures at HBOS, RBS, and elsewhere.²⁸ The regulations are the codified lessons of institutional catastrophe.

Chapter 16 (Lessons from Bank Failures) examines those catastrophes directly. Drawing on the industry loss database of 179 bank failures spanning six decades and thirty-five countries, it analyses what went wrong, identifies the recurring failure modes, and demonstrates — case by case — what a structured risk identification methodology would have caught. The regulations tell institutions what they must do. The evidence tells them why.

1. Basel Committee on Banking Supervision, *Corporate Governance Principles for Banks* (BCBS 328), Principle 7, paragraph 100, July 2015. The quoted text appears in the risk management guidance section requiring ongoing, bank-wide risk identification.
2. Basel Committee on Banking Supervision, *Corporate Governance Principles for Banks* (BCBS 328), Principle 7, paragraphs 101-103, July 2015. These paragraphs specify risk identification expectations for M&A activities, new products, and organisational changes.
3. Basel Committee on Banking Supervision, *Principles for Effective Risk Data Aggregation and Risk Reporting* (BCBS 239), January 2013. The standard set a compliance deadline of January 2016 for global systemically important banks (G-SIBs), with national supervisors encouraged to apply the principles to domestic systemically important banks (D-SIBs) within three years of their designation.
4. Basel Committee on Banking Supervision, *Principles for Effective Risk Data Aggregation and Risk Reporting* (BCBS 239), Principle 4 (Completeness) and Principle 7 (Accuracy), January 2013. Principle 4 requires that data aggregation capabilities cover all material risk exposures across the banking group, and Principle 7 requires that risk management reports accurately and precisely convey aggregated risk data.
5. Basel Committee on Banking Supervision, *Revisions to the Principles for the Sound Management of Operational Risk* (BCBS d515), March 2021. This revision replaced the original 2011 version (BCBS 195).
6. Basel Committee on Banking Supervision, *Revisions to the Principles for the Sound Management of Operational Risk* (BCBS d515), Principles 5-7, March 2021. These principles mandate the use of RCSA, KRIs, external loss data analysis, business process mapping, and internal event management as core operational risk identification tools.
7. Basel Committee on Banking Supervision, *Revisions to the Principles for the Sound Management of Operational Risk* (BCBS d515), Principle 6, paragraph 6.3, March 2021. The 2021 revision introduced the requirement to identify near-miss events and strengthened expectations for ICT and third-party risk identification.
8. Regulation (EU) 2024/1623 of the European Parliament and of the Council (CRR III), 31 May 2024, and Directive (EU) 2024/1619 (CRD VI), 31 May 2024. Together these constitute the EU Banking Package implementing the final Basel III reforms, with phased application beginning 1 January 2025.
9. Directive (EU) 2024/1619 (CRD VI), Article 87a, 31 May 2024. This article requires institutions to identify, measure, manage, and monitor ESG risks as drivers of existing prudential risk categories (credit, market, operational, liquidity, and concentration risk).
10. European Banking Authority, *Guidelines on Internal Governance* (EBA/GL/2021/05), 2 July 2021, effective 31 December 2021. These guidelines repeal and replace the earlier EBA/GL/2017/11.
11. European Banking Authority, *Guidelines on Internal Governance* (EBA/GL/2021/05), Section 17 (paragraphs 72-74), 2 July 2021. Section 17 requires institutions to establish a holistic view of risks across legal entities and risk types.
12. European Banking Authority, *Guidelines on Internal Governance* (EBA/GL/2021/05), Section 20.4 (paragraph 119), 2 July 2021. This paragraph places a duty on the risk management function to detect and report unapproved exposures.

13. European Central Bank, *SSM Supervisory Priorities 2024-2026*, December 2023. The ECB Banking Supervision directly supervises approximately 110 significant institutions within the Single Supervisory Mechanism (SSM). The precise number fluctuates as institutions are designated or de-designated. See also ECB, *SREP Methodology Booklet — 2023 Edition*.
14. Regulation (EU) 2022/2554 of the European Parliament and of the Council on Digital Operational Resilience for the Financial Sector (DORA), 14 December 2022. DORA entered into force on 16 January 2023, with the application date of 17 January 2025.
15. Regulation (EU) 2022/2554 (DORA), Article 8(1), 14 December 2022. The article requires financial entities to identify, classify, and document all ICT-supported business functions, information assets, and ICT assets.
16. European Banking Authority, *Guidelines on Outsourcing Arrangements* (EBA/GL/2019/02), 25 February 2019, effective 30 September 2019. Section 12 specifies the pre-outsourcing risk assessment requirements, including the classification of functions as critical or important.
17. European Banking Authority, *Guidelines on the Management of ESG Risks* (EBA/GL/2024/01), 9 January 2025. The guidelines apply to large institutions from 11 January 2026 and to small and non-complex institutions from 11 January 2027.
18. European Central Bank, *Guide on Climate-Related and Environmental Risks: Supervisory Expectations Relating to Risk Management and Disclosure*, November 2020. This guide set out 13 supervisory expectations for how banks should integrate climate and environmental risks into their business strategy, governance, risk management, and disclosure frameworks.
19. European Banking Authority, *Guidelines on ICT and Security Risk Management* (EBA/GL/2019/04), 28 November 2019, effective 30 June 2020. These guidelines require institutions to classify ICT risks using the same taxonomy applied to other risk types and to maintain an information asset inventory.
20. Directive (EU) 2018/1673 of the European Parliament and of the Council on Combating Money Laundering by Criminal Law (AMLD6 / the Sixth Anti-Money Laundering Directive), 23 October 2018, transposition deadline 3 December 2020. The enterprise-wide ML/TF risk assessment requirement is complemented by the EBA's *Guidelines on ML/TF Risk Factors* (EBA/GL/2021/02), revised March 2021.
21. Prudential Regulation Authority, *Supervisory Statement SS31/15: The Internal Capital Adequacy Assessment Process (ICAAP) and the Supervisory Review and Evaluation Process (SREP)*, originally issued November 2015, most recently updated December 2023. SS31/15 sets out the PRA's expectations for how firms should assess the adequacy of their capital, including through stress testing and reverse stress testing as risk identification tools.
22. Prudential Regulation Authority, *Policy Statement PS5/25: Step-In Risk*, 2025. This policy statement introduces requirements for the identification of non-contractual financial support obligations to unconsolidated entities, implementing the BCBS framework on step-in risk (BCBS d349, October 2017).
23. Financial Conduct Authority, *Five Conduct Questions*, introduced as part of the FCA's supervisory approach from 2015. The framework requires firms to assess conduct risk through five questions covering culture, products, governance, remuneration, and treatment of customers. See also FCA, *Business Plan 2023/24*, for the current application of the conduct questions framework.
24. Financial Conduct Authority and Prudential Regulation Authority, *Senior Managers and Certification Regime (SM&CR)*, implemented for banks and PRA-designated investment firms from 7 March 2016, extended to all FCA solo-regulated firms from 9 December 2019. Established under the Financial Services (Banking Reform) Act 2013, Part 4.
25. Board of Governors of the Federal Reserve System, *SR 15-18: Federal Reserve Supervisory Assessment of Capital Planning and Positions for LISC Firms and Large and Complex Firms*, 18 December 2015. This supervisory letter establishes the qualitative and quantitative expectations for capital planning, including the requirement for a comprehensive Material Risk Inventory.

26. Office of the Comptroller of the Currency, *OCC Guidelines Establishing Heightened Standards for Certain Large Insured National Banks, Insured Federal Savings Associations, and Insured Federal Branches* (12 CFR Part 30, Appendix D), final rule effective 1 January 2014, with amendments. The guidelines establish minimum standards for risk governance, including front-line unit risk identification responsibilities and independent risk management challenge functions.
27. Standard Chartered Bank was fined a total of approximately \$667 million by US and UK authorities in 2012 (US: \$340 million OFAC, \$100 million Fed, \$227 million to New York DFS), with additional enforcement actions in subsequent years relating to sanctions compliance deficiencies. See US Department of the Treasury, OFAC Enforcement Action, 10 December 2012, and New York State Department of Financial Services, Consent Order, 6 August 2012.
28. Parliamentary Commission on Banking Standards, *Changing Banking for Good* (HL Paper 27-I / HC 175-I), June 2013. The Commission's findings on the failures at HBOS and RBS, and the lack of individual accountability under the then-existing Approved Persons Regime, led directly to the creation of SM&CR under the Financial Services (Banking Reform) Act 2013.

Lessons from Bank Failures

The \$2.3 Trillion Evidence Base

Two point three trillion dollars. That is the aggregate cost of the 179 bank failures, frauds, and near-collapses in the industry loss database that underpins this book. Spanning six decades, thirty-five countries, and every major risk category in the taxonomy, this database represents the most comprehensive collection of risk identification failures ever assembled for a single methodology.¹

This database was built not as an academic exercise but as a forensic tool. For each of the 179 entries, the same question posed in Chapter 1 was applied: *Was the risk identifiable before the loss materialised?* The answer, in every case, was yes. Not with the benefit of hindsight. Not with information that only became available after the fact. The risks were identifiable at the time, with the information available at the time, using techniques available at the time. What was missing was the process.

Chapter 15 (The Regulatory Landscape) described the sixteen regulatory frameworks that mandate risk identification across jurisdictions. Those regulations tell institutions what they must do. This chapter presents the evidence that tells them why.

The database reveals patterns that should alarm any practitioner. Credit risk accounts for 42% of entries and \$1.37 trillion in aggregate losses — the largest category by dollar amount. Operational risk accounts for 43% of entries and \$626 billion — the largest by frequency. Conduct risk, a category that barely existed in institutional taxonomies before 2010, accounts for nearly 7% of entries and \$110 billion. The temporal distribution is equally revealing: 44% of entries cluster in the 2000s, driven by the Global Financial Crisis's 63 events and \$1.5 trillion in aggregate losses, while 47% fall in the 2010s — demonstrating that the post-crisis regulatory reforms, however extensive, did not eliminate risk identification failures. They changed the character of the failures.²

Geographically, Europe accounts for 54% of entries and North America for 30%, but the database spans Africa, Asia-Pacific, Latin America, and the Middle East. Risk identification failure is not a problem of any single regulatory regime. It is a structural problem of how institutions approach the identification step.

The COSO objective category distribution is instructive: 41% of failures map to the Strategic category, 22% to Operations, 19% to Compliance, and 18% to Reporting. The dominance of Strategic means that the most frequent risk identification failures occur at the level of the institution's fundamental business decisions — the risks that only senior management and the Board can identify and own. This is the evidence basis for the methodology's dual-track approach: the top-down workshops described in Chapter 6 (Top-Down Identification: Workshops, SWIFT, and Delphi) are not optional. They are where the most consequential identification occurs.

Across all 179 events, ten failure modes recur with a consistency that makes them effectively predictable. Those ten — Concentration Blindness, Model Overreliance, Governance Bypass, Silo Thinking, Cultural Suppression, Emerging Risk Blindness, Control Environment Failure, Information Asymmetry, Regulatory Arbitrage Masking, and Complacency — were introduced in Chapter 1 as recurring themes. This chapter examines them through the full weight of the evidence.

The Concentration Trap

Concentration Blindness is the single most destructive failure mode in the database. It appears in every decade, every jurisdiction, and across every risk type. The pattern is always the same: an institution accumulates exposure to a single counterparty, sector, geography, or funding source, and the risk identification process either fails to detect the concentration or detects it but classifies it as acceptable.

Alpha Bank illustrates the most insidious form: concentration that regulatory frameworks actively disguise. As a major Greek bank, Alpha held disproportionate sovereign debt exposure to the Greek state. Under the Capital Requirements Regulation, eurozone sovereign debt carried a zero risk weight — a regulatory treatment that, in effect, told risk identification processes there was nothing to identify. Alpha did not flag domestic sovereign debt as a concentration risk because the regulatory framework explicitly stated it

was not one. When the eurozone debt crisis triggered Greek sovereign restructuring, Alpha required approximately €4 billion in recapitalisation through the Hellenic Financial Stability Fund, part of the €30.9 billion injected across all four Greek systemic banks.³

The same regulatory blind spot destroyed the Bank of Cyprus, where uninsured depositors lost up to 47.5% in the largest European bail-in of its time.⁴ Both institutions demonstrate a failure mode that the methodology's regulatory versus economic risk gap analysis — described in Chapter 12 (Integration: Capital Planning, Strategy, and the Board) — was designed to catch. The four-step gap analysis requires institutions to compare the regulatory treatment of every material risk against their own economic assessment. A risk identification process that simply accepts zero risk weights for home-sovereign debt is not identifying risks — it is transcribing regulatory assumptions.

Countrywide Financial demonstrates a different flavour of concentration blindness: the originate-to-distribute illusion. Countrywide became America's largest mortgage originator by aggressively expanding into subprime and Alt-A lending. The stated business model was that securitisation transferred the risk to investors. The risk identification failure was in not recognising that retained tranches, warehouse lines, representations-and-warranties obligations, and reputational attachment to the product meant that securitisation transferred the asset but not the risk. Countrywide's acquisition by Bank of America for \$4 billion ultimately cost Bank of America over \$40 billion in settlements and litigation — a ten-to-one ratio that demonstrates just how badly the residual exposure was underestimated.⁵

The methodology's enterprise portfolio view, described in Chapter 8 (Reconciliation and the Enterprise Portfolio View), requires explicit assessment of aggregate position against appetite, including concentration analysis across all three forms defined in Chapter 10 (Risk Interaction: Bow-Ties, Matrices, and Concentration): single-name, structural, and systemic. A PESTLE assessment conducted under Phase 1 would have classified the rapid growth in subprime origination as a macro-economic risk factor. A SWIFT workshop would have asked: *What are we assuming about the securitisation market that might not be true?* The originate-to-distribute assumption was precisely the kind of unexamined premise that the methodology's assumption register is designed to capture and challenge.

The same concentration pattern, with local variation, appears in Northern Rock's 75% wholesale funding dependence, AIG's \$527 billion CDS notional on mortgage-linked CDOs, Citigroup's off-balance-sheet SIV exposure, the Icelandic banks' collective

growth to ten times their sovereign's GDP, and the Spanish cajas whose concentrated property lending culminated in Bankia's €22.4 billion bailout.⁶ In each case, the concentration was visible. In each case, the process for identifying it was absent or inadequate.

Models, Ratings, and False Precision

Model Overreliance is the failure mode that produces the most intellectually sophisticated losses. The institutions that fall to this pattern are not unsophisticated — they are often the most analytically advanced in the industry. The failure is not the absence of models but the absence of independent challenge to those models.

Long-Term Capital Management remains the defining illustration. LTCM's founders included two Nobel Prize-winning economists. Their Value-at-Risk models, calibrated to historical market conditions, indicated that the fund's convergence trades across global fixed income markets were well-diversified. The models assumed that correlations between different markets and instruments were stable and that liquidity would always be available to unwind positions. When the Russian crisis of 1998 caused correlations to spike to one across virtually all markets simultaneously, LTCM's entire portfolio moved against it. The positions that VaR showed as diversified turned out to be a single concentrated bet on market normality. The \$3.6 billion bailout, brokered by the Federal Reserve with fourteen counterparty banks,⁷ demonstrated that the model's failure was not a calculation error — it was a structural inability to capture the regime shift that the methodology's scenario analysis and speed-of-onset dimension are designed to identify.

Merrill Lynch demonstrates the corporate variant: the AAA rating as a substitute for independent analysis. Merrill accumulated approximately \$40 billion in gross subprime exposure, of which some \$32 billion consisted of super-senior CDO tranches retained on its balance sheet⁸ — the most senior tranches that would only suffer losses after all junior tranches were wiped out. The risk identification process accepted the AAA rating as sufficient evidence that these positions were near-risk-free. No independent analysis examined the scenario in which a systemic housing market collapse would impair even super-senior tranches. The \$51.8 billion in write-downs and emergency sale to Bank of America⁹ demonstrated that an external credit rating is not a risk assessment. It is a single input to a risk assessment.

The methodology's response to Model Overreliance is the independent challenge requirement described in Chapter 9 (Assessment — Scoring, Multi-Dimensional Impact, and Data Quality): models are one input, not a substitute for judgement. The four-dimensional assessment framework requires explicit documentation of the model's assumptions, identification of the scenarios under which the model would break down, and a recorded management judgement on the model's applicability. The Data Quality Rating forces disclosure of the evidence basis — and positions assessed primarily through models calibrated to benign conditions should carry a Data Quality Rating of Low, triggering the conservatism adjustment that prevents low data quality from supporting low impact or likelihood ratings.

When Governance Fails

Governance Bypass, Cultural Suppression, and Information Asymmetry share a common root: the risk was identified or identifiable, but the institutional mechanisms for acting on that identification were broken. These three failure modes account for many of the most preventable losses in the database.

Standard Chartered's \$667 million sanctions fine¹⁰ illustrates governance bypass in its purest form. The bank's compliance function identified the sanctions risk at the operational level. Compliance officers flagged the processing of approximately \$250 billion in transactions for Iranian clients. Senior management in London overruled the concerns, treating the Iranian business as commercially important. This is not a case where risk identification failed — the risk was identified. The governance framework failed to protect the identification output from being suppressed by the people generating the revenue.

The methodology addresses this through the three forms of independence defined in Chapter 3 — structural, operational, and intellectual — and through the principal risk report that flows directly from the inventory to the Board Risk Committee without being filtered by the business lines that generated the risk. The Risk Identification Lead's mandate, as described in Chapter 3 (Governance: Who Owns What), must include the authority to record risks in the inventory regardless of business line objections, with escalation to the CRO and ultimately the Board.

Credit Suisse's Mozambique hidden debt scandal demonstrates a different governance failure: one where the risk identification framework never had the chance to engage. Credit Suisse bankers arranged \$2 billion in hidden loans to Mozambican state entities¹¹ — loans ostensibly for maritime security and fishing but involving military-connected borrowers, kickback payments to bankers and officials, and secrecy requirements that concealed the debt from the IMF and international donors. The sovereign lending risk identification process focused on credit analysis of the Mozambican state. It did not identify the corruption risk embedded in the transaction structure, despite red flags that a structured new-product approval process would have surfaced: the involvement of intermediaries, the military nature of the borrowers, the secrecy requirements, and the fee structures.

The \$475 million coordinated global settlement across the SEC, DOJ, and FCA¹² was one of several scandals that contributed to Credit Suisse's eventual demise. The methodology's Phase 1 internal environment assessment, using the seven COSO elements, would have examined whether the institution's ethical values and risk management philosophy were genuinely embedded — or whether, as in this case, individual bankers could structure transactions that bypassed every control designed to prevent exactly this kind of exposure. The event-driven trigger framework from Chapter 13 (The Ongoing Cycle: Refresh, Events, and Audit) requires that new products and new business structures undergo risk identification assessment before launch. The Mozambique loans were a new structure. They were never assessed.

ICBC's experience adds a dimension rarely examined in Western risk management literature: what happens when the state itself directs the risk-taking. As the world's largest bank by assets, ICBC accumulated significant non-performing loan exposure through policy-directed lending to state-owned enterprises and local government financing vehicles. Credit decisions directed by state industrial policy could not be challenged by the risk function. NPL recognition was delayed by regulatory forbearance and ever-greening practices. Risk identification was structurally constrained — the dual role of Chinese state banks as both commercial lenders and policy instruments meant that the governance framework required by BCBS Principle 7 and the independence requirements of every Western regulatory framework were fundamentally compromised.

What the Risk Framework Could Not See

Silo Thinking and Emerging Risk Blindness share a structural characteristic: the risk exists but falls outside the boundaries of how the institution has organised its identification process. Silo Thinking means the risk is visible within individual units but invisible at the enterprise level. Emerging Risk Blindness means the risk is visible in the external environment but invisible within the institution's taxonomy.

Lloyd's of London after September 11 illustrates silo thinking at the market level. The \$5.6 billion in claims¹³ arose because individual syndicates covered different aspects of the same underlying event — property, liability, aviation, business interruption — without identifying the aggregate market-level exposure. Catastrophe risk identification models did not include a scenario combining coordinated terrorist attacks on multiple high-value insured properties in a single city. The aggregation risk across syndicates covering different aspects of the same underlying event was not identified at the market level. This is the insurance equivalent of the Citigroup SIV pattern: individually managed exposures that were collectively catastrophic because no enterprise-level — or in Lloyd's case, market-level — portfolio view existed.

The methodology's reconciliation process and enterprise portfolio view, described in Chapters 8 (Reconciliation and the Enterprise Portfolio View) and 10 (Risk Interaction: Bow-Ties, Matrices, and Concentration), directly address this pattern. The four assessments of the enterprise portfolio view — common exposures, simultaneous crystallisation, aggregate position against appetite, and diversification and correlation — are designed to identify precisely the kind of hidden aggregation that Lloyd's experienced. The SWIFT guide word "*Where are the gaps between...*" applied across business areas and risk categories would have prompted examination of cross-syndicate exposure to correlated catastrophe events.

UBS's initial €4.5 billion French tax evasion fine in 2019 — the largest tax evasion penalty in French history at the time, later reduced to €1.8 billion on appeal and ultimately settled for €835 million in 2025¹⁴ — demonstrates emerging risk blindness. UBS systematically sent Swiss-based relationship managers to France to solicit undeclared accounts. The cross-border wealth management model was treated as a normal distribu-

tion strategy. Risk identification classified it as business-as-usual. The conduct risk embedded in helping clients maintain undeclared offshore accounts was not identified as a distinct risk category requiring assessment and controls.

The same pattern — a profitable business practice that is reclassified as misconduct when regulatory expectations shift — appears across the post-GFC conduct era. **Lloyds' PPI mis-selling** is the most expensive example: over £50 billion in industry-wide redress, with Lloyds paying the largest share at approximately £22 billion.¹⁵ PPI sales were treated as a profitable cross-selling activity. Sales incentive structures that rewarded PPI attachment rates were not identified as conduct risk drivers. Customer suitability was not a risk category in the taxonomy. When the FCA forced retrospective review, the scale of customer detriment became apparent — but the detriment had accumulated over years during which risk identification frameworks could not see it because the taxonomy did not include a place to look.

This is where Chapter 4 (The Risk Taxonomy)'s Taxonomy Test becomes critical. Take any historical loss and map it to the institution's taxonomy at the time. If it does not map cleanly, the taxonomy was deficient. PPI does not map to a taxonomy that lacks conduct risk. UBS's cross-border solicitation does not map to a taxonomy that classifies it as sales and marketing rather than compliance risk. Lloyd's aggregation risk does not map to a framework that assesses individual syndicates without an enterprise-level concentration view. The Taxonomy Test, applied retrospectively, reveals the structural gap that enabled the loss.

The methodology's Delphi Method, described in Chapter 6, is specifically designed for emerging risks that fall outside current taxonomy boundaries. External panellists, protected by anonymity, provide perspectives that internal processes are structurally incapable of generating. The Wirecard case — where the Financial Times and short-sellers identified the €1.9 billion fraud that BaFin could not see¹⁶ — demonstrates what happens when no mechanism exists for incorporating external dissent into internal risk identification.

The Complacency Cycle

Complacency and Control Environment Failure share a temporal characteristic: they develop over extended periods of apparent success. The longer an institution operates without a major loss, the more convinced it becomes that its risk profile is well-man-

aged. Risk identification processes degrade incrementally. The quarterly re-identification cycle becomes a roll-forward exercise. Event-driven triggers are not processed. New risks are not identified because the old risks have not crystallised.

Bear Stearns provides the most stark illustration of a complacency failure that included a direct, explicit warning. In June 2007, two Bear Stearns hedge funds collapsed from CDO exposure. The parent bank treated this as an isolated subsidiary event — a contained loss that did not require re-examination of the bank's own retained structured credit positions. Nine months later, Bear Stearns itself collapsed from the same concentration risk.¹⁷ The methodology's event-driven trigger framework, described in Chapter 13, defines a material loss or near-miss as the first of six triggers that require immediate re-identification. The collapse of two hedge funds managed by the parent bank, holding positions similar to those on the parent's own balance sheet, is not an ambiguous signal. It is a direct warning. Bear Stearns did not process it.

American Savings & Loan illustrates the longer-duration form: complacency bred by structural change. The Depository Institutions Deregulation and Monetary Control Act of 1980 and the Garn-St. Germain Act of 1982 removed interest rate ceilings and expanded the permissible activities for savings institutions. American Savings responded with aggressive high-risk mortgage origination and speculative securities investments. The risk identification failure was in treating deregulation as an opportunity without identifying that interest rate liberalisation had fundamentally altered the asset-liability risk profile. The institution was seized by FSLIC in 1988 with \$5.4 billion in losses.¹⁸

The same complacency pattern appears with remarkable consistency across eras and jurisdictions. Silicon Valley Bank failed to re-identify interest rate risk as the Federal Reserve raised rates by 500 basis points over eighteen months — a rate-rising cycle that was publicly signalled, widely discussed, and directly relevant to a bank holding a \$91 billion HTM securities portfolio. The CRO position was vacant for eight months during this period.¹⁹ Signature Bank failed to process the TerraUSD collapse in May 2022, the Celsius bankruptcy in June, and the FTX collapse in November as event-driven triggers for its own crypto deposit concentration — three separate warnings in six months, each of which should have forced immediate re-identification. NAB Australia failed to learn from CBA's AUD 700 million AUSTRAC fine in 2018²⁰ despite operating with analogous AML monitoring gaps — a peer institution event that the methodology explicitly defines as an event-driven trigger requiring internal re-identification of analogous risks.

The pattern is always the same: the ongoing cycle degrades because the institution equates the absence of loss with the absence of risk. In each case, the mechanisms described in Chapter 13 — quarterly re-identification, event-driven updates, monthly KRI monitoring, and internal audit assurance — would have caught the degradation before it became catastrophic. The process performance indicator that Chapter 13 defined — “an institution where no new risks are identified across three consecutive quarterly cycles is either operating in a static environment or running a process that has succumbed to compliance theatre” — is the diagnostic for this failure mode.

What the Methodology Catches

The evidence from 179 failures maps directly to the six-phase methodology described in Chapters 3 (Governance: Who Owns What) through 14 (Technology: AI, ML, and Data Analytics). This is not coincidental. The methodology was designed from the evidence.

Phase 1 — Foundation Setting catches the failures that begin with unexamined assumptions. The PESTLE assessment would have identified the macro-economic conditions that enabled the S&L crisis, the eurozone sovereign debt crisis, and the GFC. The internal environment assessment using the seven COSO elements would have surfaced the cultural dysfunction at Anglo Irish Bank, the governance compromise at HBOS, and the ethical failures at Credit Suisse. The risk criteria — particularly the multi-dimensional impact scales covering financial, regulatory, reputational, and customer dimensions — would have prevented the single-dimensional assessments that made AIG’s CDS portfolio and Merrill Lynch’s super-senior tranches appear low-risk.

Phase 2 — Dual-Track Identification catches the failures that arise from incomplete coverage. Top-down SWIFT workshops with structured guide words would have surfaced the Archegos cross-counterparty concentration, the Countrywide originate-to-distribute assumption, and the Bear Stearns parent-subsidiary correlation. Bottom-up standardised templates would have captured the front-line knowledge that Wells Fargo’s account-opening practices were fraudulent, that Standard Chartered’s sanctions processing was non-compliant, and that AIB’s Allfirst subsidiary had inadequate trade confirmation processes. The Delphi Method would have incorporated the external perspectives that identified Wirecard’s fraud years before the regulator acted.

Phase 3 — Assessment catches the failures that arise from inadequate analytical frameworks. The four-dimensional scoring methodology would have required AIG to assess its CDS portfolio across financial, regulatory, reputational, and customer dimensions — not just the single credit dimension that produced the near-zero score. The Data Quality Rating would have flagged any assessment based primarily on rating agency opinions or models calibrated to benign conditions. The independent challenge requirement would have prevented the self-referencing model reliance that destroyed LTCM and UBS's structured credit book.

Phase 4 — Documentation catches the failures that arise from fragmentation. The fourteen-field risk inventory with Risk Interaction Summary would have connected the individually assessed risks at Lehman Brothers into visible causal chains. The enterprise-level inventory would have aggregated Danske Bank's AML exposure, AMP's complaints data, and SNS Reaal's subsidiary concentration into single, visible entries. The audit trail requirement would have prevented the compliance theatre that characterises institutions where risk registers exist but risk identification does not.

Phase 5 — Integration catches the failures that arise from disconnection. The capital planning integration mechanisms would have linked Washington Mutual's and Dexia's risk identification outputs to capital adequacy requirements. The regulatory versus economic risk gap analysis would have identified the zero-risk-weight treatment of Greek sovereign debt as a regulatory assumption requiring independent economic assessment. The strategic planning integration would have subjected Bankia's merger of seven struggling cajas to enterprise portfolio view analysis before the aggregation of concentrated property exposures without recapitalisation.

Phase 6 — The Ongoing Cycle catches the failures that arise from degradation over time. The quarterly re-identification would have prevented Silicon Valley Bank's failure to re-identify interest rate risk across eight consecutive rate increases. The event-driven triggers would have forced Bear Stearns to re-examine its own balance sheet when its hedge funds collapsed, Signature Bank to reassess crypto deposit concentration when three major crypto firms failed, and NAB to review its AML monitoring when CBA paid a AUD 700 million fine for analogous gaps. The internal audit assurance, testing seven specific areas of the methodology, would have identified the process degradation before it produced catastrophic outcomes.

The Common Thread

Across 179 events, thirty-five countries, and six decades of banking history, the same structural gap appears: **the institution did not have a process designed to find the risks it was exposed to.**

Some had risk registers but no identification methodology. Some had identification processes that were undermined by the governance structures that should have protected them. Some had sophisticated analytical capabilities that were never applied to the identification step. And some had processes that worked once but degraded over time as complacency replaced rigour.

The losses were not inevitable. They were preventable. The risk was there. The information was there. What was missing was the process.

This book has presented that process. Chapter 2 (The Foundations: Standards and Frameworks) established the standards architecture — ISO 31000's risk management principles, ISO 31010's technique toolkit, COSO ERM's enterprise lens, and the BCBS banking mandate. Chapters 3 through 14 described the six-phase methodology in full practitioner detail: the governance that protects the process, the taxonomy that structures it, the foundation setting that contextualises it, the dual-track identification that populates it, the assessment that prioritises it, the risk interaction analysis that connects it, the documentation that records it, the integration that activates it, the ongoing cycle that sustains it, and the technology that enables it. Chapter 15 mapped the sixteen regulatory frameworks that mandate it.

This chapter has presented the evidence that validates it.

The methodology is built on three foundations: standards, regulation, and evidence. The standards provide the architecture. The regulation provides the mandate. The evidence — 179 failures, \$2.3 trillion, thirty-five countries, six decades — provides the reason.

Every institution in this database had talented risk professionals. Many had sophisticated measurement systems, advanced models, and comprehensive regulatory reporting. What they did not have was a systematic, end-to-end process for identifying the risks they were exposed to before those risks destroyed value, reputations, and — in too many cases — the institution itself.

The methodology described in this book exists because the evidence demands it. The process is comprehensive. It is detailed. And when implemented with the governance, rigour, and institutional commitment it requires, it works.

The evidence is clear. The methodology is available. The only remaining question is whether the institution has the will to implement it.

-
1. All statistics from the Industry Loss Database cited in this chapter are calculated from the complete 179-entry dataset. See Appendix A: Industry Loss Database Methodology for inclusion criteria, data sources, loss definitions, USD standardisation methodology, and the full event list.
 2. Distribution statistics calculated from the Industry Loss Database. The temporal clustering in the 2000s reflects the concentration of Global Financial Crisis events (2007-2010). See Appendix A for methodology.
 3. Hellenic Financial Stability Fund, *Annual Report*, 2013; European Commission, State Aid Decisions SA.34823, SA.34825, SA.34826, SA.34827 (2012-2013). The €30.9 billion represents aggregate HFSF capital injections into Alpha Bank, Eurobank, National Bank of Greece, and Piraeus Bank.
 4. Central Bank of Cyprus, Decree of the Resolution Authority, 29 March 2013. Uninsured deposits above €100,000 at Bank of Cyprus were converted to equity at rates resulting in losses of up to 47.5%.
 5. Bank of America Corporation, Annual Reports and SEC filings, 2008-2014. The \$4 billion acquisition was announced 11 January 2008. Subsequent settlements include the \$16.65 billion DOJ National Mortgage Settlement (2014) and the \$8.5 billion private-label RMBS settlement (2011), among others. See also Financial Crisis Inquiry Commission, *The Financial Crisis Inquiry Report*, January 2011, Chapter 11.
 6. Northern Rock: House of Commons Treasury Committee, *The Run on the Rock*, Fifth Report of Session 2007-08, HC 56-I, January 2008. AIG: Form 10-K for fiscal year 2007, filed 28 February 2008. Icelandic banks: Rannsóknarnefnd Alþingis (Special Investigation Commission), *Report to the Althingi*, 2010. Bankia: Fondo de Reestructuración Ordenada Bancaria (FROB) and European Commission State Aid Decision SA.35253, 28 November 2012.
 7. Lowenstein, R. *When Genius Failed: The Rise and Fall of Long-Term Capital Management*. Random House, 2000. Federal Reserve Bank of New York, press statement, 23 September 1998. The consortium of fourteen institutions included Goldman Sachs, Merrill Lynch, J.P. Morgan, Morgan Stanley, and ten others.
 8. Financial Crisis Inquiry Commission, *The Financial Crisis Inquiry Report*, January 2011, pp. 256-261. See also Merrill Lynch & Co., Form 10-K, fiscal year 2007.
 9. Merrill Lynch & Co., SEC filings, 2007-2008. The cumulative write-downs on mortgage-related positions were disclosed across quarterly earnings announcements. The sale to Bank of America was announced 15 September 2008.
 10. New York State Department of Financial Services, Consent Order Under New York Banking Law §39, 6 August 2012. The \$667 million figure represents cumulative penalties from the NYDFS, OFAC, DOJ, FCA, and the Federal Reserve, accumulated 2012-2019. The \$250 billion transaction figure is from the NYDFS order.
 11. US Department of Justice, Press Release, "Three Former Credit Suisse Investment Bankers Indicted for Fraud Scheme Involving Billions of Dollars in Loans to Mozambique," 3 January 2019. SEC, Litigation Release No. 25252, 19 October 2021.
 12. DOJ Press Release, "Credit Suisse Group AG Agrees to Pay More Than \$475 Million," 19 October 2021. SEC, Order Instituting Cease-and-Desist Proceedings, File No. 3-20625. FCA Final Notice to Credit Suisse Securities (Europe) Limited, 2021. The DOJ component was attributed to Credit Suisse Group AG.

13. Lloyd's of London, *Annual Report*, 2001. The figure represents net claims to the Lloyd's market from the September 11, 2001 attacks.
14. Tribunal de Grande Instance de Paris, Judgment of 20 February 2019 (€4.5 billion); Cour d'appel de Paris, Judgment of 13 December 2021 (reduced to €1.8 billion); settlement at €835 million reported 2025.
15. FCA, *Payment Protection Insurance Complaints: Aggregate Data*, final reporting period ending August 2019. Industry-wide PPI redress exceeded £50 billion across all UK providers. Lloyds Banking Group, Annual Reports, 2011-2019.
16. EY Special Audit of Wirecard AG, 2020. Wirecard AG insolvency filing, Amtsgericht München, 25 June 2020. German Parliamentary Committee of Inquiry (*Wirecard-Untersuchungsausschuss*), 2021.
17. Financial Crisis Inquiry Commission, *The Financial Crisis Inquiry Report*, January 2011, pp. 238-242. The Bear Stearns High-Grade Structured Credit Fund and Enhanced Leverage Fund collapsed in June-July 2007. Bear Stearns Companies Inc. was acquired by JPMorgan Chase on 16 March 2008.
18. FDIC, *Managing the Crisis: The FDIC and RTC Experience*, 1998. American Savings was seized by the Federal Savings and Loan Insurance Corporation in 1988. The Depository Institutions Deregulation and Monetary Control Act of 1980 (Pub.L. 96-221) and the Garn-St. Germain Depository Institutions Act of 1982 (Pub.L. 97-320) are the referenced legislation.
19. SVB Financial Group, Form 10-K, fiscal year 2022 (\$91 billion HTM securities portfolio). Federal Reserve Board, *Review of the Federal Reserve's Supervision and Regulation of Silicon Valley Bank*, 28 April 2023, pp. 6, 10 (CRO vacancy and supervisory findings). The Federal Open Market Committee raised the federal funds rate from near-zero to 5.00-5.25% between March 2022 and May 2023.
20. AUSTRAC v Commonwealth Bank of Australia [2018] FCA 830. AUSTRAC, Statement of Claim, filed 3 August 2017. CBA paid AUD 700 million in civil penalties for 53,506 contraventions of the Anti-Money Laundering and Counter-Terrorism Financing Act 2006.



EON Risk Services Ltd
Dublin, Ireland

© 2025 Rory Roberts. All rights reserved.
No part of this publication may be reproduced
without the prior written permission of the author.