



Risk Taxonomy as Governance Infrastructure: Adaptation, Traceability and Industry- led Use Cases for Fintech Innovation



Risk Taxonomy as Governance Infrastructure: Adaptation, Traceability and Industry-led Use Cases for Fintech Innovation

John Finch, Adam Smith Business School, University of Glasgow, and Financial Regulation Innovation Lab, 2 Discovery Place, Glasgow G11 6EY, UK – john.finch@glasgow.ac.uk

Selen Shah, School of Law, University of Glasgow, and Financial Regulation Innovation Lab, 5-9 the Square, Glasgow, G12 8QQ, UK – selen.shah@glasgow.ac.uk

31 March 2026

Acknowledgements: We are grateful to Dominic Chalmers and Alessio Azzutti of the University of Glasgow and the Financial Regulation Innovation Lab, Kirsty McKenna at FinTech Scotland and the Financial Regulation Innovation Lab, and to an industry Partner and senior assurance leaders for guidance on practical relevance and feasibility. All remaining interpretations and any errors are our own. We received funding from the Innovate UK Innovation Accelerator programme (grant numbers 10055559 and 10157846) and from the University of Glasgow College of Social Sciences PGR Opportunities Hub PhD internship programme.

Declaration of AI use: We are guided by Gatrell, Muzio, Post and Wickert (2024) in their editorial article in the *Journal of Management Studies*. The authors completed the initial analysis and first draft of this paper without using AI tools. ChatGPT 5.2 was then used in a supervised, iterative way to support a further review of the nine annual reports, the restructuring of the draft into white paper form, the development and refinement of the use-case presentation, and editing for clarity. All source materials, coding decisions, interpretations, and final editorial choices were reviewed and validated by the authors. Grammarly was used at the final proofreading stage for limited language editing.

Open Access. Some rights reserved:



The publishers, the University of Glasgow and FinTech Scotland, and the authors, John Finch and Selen Shah, want to encourage the widest possible circulation of our work while retaining copyright. We therefore have an open-access policy that enables anyone to access our content online at no charge. Anyone can download, save, perform or distribute this work in any format, including translation, without written permission. This is subject to the terms of the Creative Commons By Share Alike licence. The main conditions are:

- The University of Glasgow, FinTech Scotland, and the authors are credited, including our web addresses www.gla.ac.uk, and www.fintechscotland.com
- If you use our work, you share the results under a similar licence

A full copy of the licence can be found at <https://creativecommons.org/licenses/by/4.0/>

You are welcome to ask for permission to use this work for purposes other than those covered by the licence. We gratefully acknowledge the work of Creative Commons in inspiring our approach to copyright. To find out more, go to www.creativecommons.org

Financial Regulation Innovation Lab

Who are we?

The Financial Regulation Innovation Lab (FRIL) is an industry-led collaborative research and innovation programme focused on leveraging new technologies to respond to, shape, and help evolve the future regulatory landscape in the UK and globally, helping to create new employment and business opportunities, and enabling the future talent.

FRIL provides an environment for participants to engage and collaborate on the dynamic demands of financial regulation, explore, test and experiment with new technologies, build confidence in solutions and demonstrate their ability to meet regulatory standards worldwide.

FRIL is part of the Glasgow City Region Innovation Accelerator programme, funded through Innovate UK on behalf of UK Research and Innovation. The Innovation Accelerator programme is investing £130 million in 26 transformative R&D projects to accelerate the growth of three high-potential innovation clusters, including the Glasgow City Region.

What is Actionable Research?

FRIL will integrate academic research with an industry relevant agenda, focused on enabling knowledge on cutting-edge topics such as generative and explainable AI, advanced analytics, advanced computing, and earth-intelligent data as applied to financial regulation. The approach fosters cross sector learning to produce a series of papers, actionable recommendations and strategic plans that can be tested in the innovation environment, in collaboration across industry and regulators.

Locally-led Innovation Accelerators delivered in
partnership with DSIT, Innovate UK and City Regions



Innovate
UK



GLASGOW
CITY REGION

Risk Taxonomy as Governance Infrastructure: Adaptation, Traceability and Industry-led Use Cases for Fintech Innovation

John Finch and Selen Shah

University of Glasgow

31 March 2026

Abstract

Risk taxonomies in financial services are often presented as stable classification models: a hierarchy of principal risks that supports aggregation, governance oversight, and regulatory reporting. This paper argues that such a view is incomplete. In practice, a risk taxonomy is a governance infrastructure shared across stakeholder communities, including firms, supervisors, regulators, and the risk profession. Its value depends on traceability.

By traceability, we mean the practical ability to track and explain how risk categories, data, judgments, mapping rules, and mitigations are defined, applied, changed, and communicated within and beyond an organisation. Traceability matters because the quality of risk management is rarely observable in real time. Weaknesses may only become visible much later as business performance outcomes, operational incidents, or supervisory concerns, at which point feedback is costly, and remediation may be too late for firms and the wider system.

For CFO and related decision-makers, the central issue is not whether a taxonomy contains the right high-level categories (which tend to be stable across major firms), but whether the organisation can operate the taxonomy reliably: govern changes, manage mappings and aggregation, and produce assurance-grade evidence that connects categories to decisions and outcomes

Conceptually, we position risk taxonomy as a boundary object: a shared artefact that different communities can use for their own purposes while recognising it as “the same thing.” A taxonomy must be robust enough to maintain comparability across firms and reporting regimes, yet flexible enough to adapt to local realities: portfolio differences, organisational structure, data maturity, and changes in technology and competition. This predictable tension does not imply failure; it highlights that the taxonomy is doing coordination work across communities. The design challenge is to make this coordination auditable.

Empirically, we review risk management disclosures in the annual reports of three large UK banks over 2022–2024. We code for risk categories, implied hierarchies, grouping logic, and signals of year-on-year change. We observe strong stability at the top level (credit, market, liquidity/capital, operational, conduct, financial crime, model risk and climate-related risk as recurring themes), with change occurring mainly through shifting emphasis, greater granularity, and increased attention to non-financial and resilience-related risks. The implication is that innovation opportunities are less about reinventing categories and more about strengthening the socio-technical system around the taxonomy: data lineage, mapping logic, change control, governance workflows, assurance, and explainability, including for data-sparse or emerging risks where judgement and scenario processes play a larger role.

The paper’s contribution is a proposal for industry-led use cases, as a portfolio of well-specified problem statements that large financial services firms can publish to invite targeted innovation from fintechs. Each use case is designed to be procurement-ready and assurance-aware. The portfolio is organised by decision ownership, minimum data inputs, workflow controls (audit trail, segregation of duties, approvals), outputs (MI, reporting support, evidence packs), and success measures (time and cost saved, reduction in reconciliation burden, fewer classification disputes, improved supervisory confidence).

We conclude with implications for the UK supervisor. Rather than advocating new rules, we propose that supervisory engagement can be strengthened by encouraging firms to set expectations for traceability-by-design around taxonomies. By this, we mean controlled change, transparent mapping between internal granularity and external reporting, auditable lineage from source data to risk decisions, and explicit governance of data-sparse risks., including scenarios, assumptions logs, and review cycles. These steps can reduce reporting friction while improving the credibility of risk disclosures and the resilience of firms and the system.

Table of Contents

1. Introduction	1
1.1 Purpose and audience	1
1.2 Risk taxonomies as governance infrastructure	1
1.3 The golden thread is traceability	1
1.4 Why the boundary object concept is useful	2
1.5 What this paper contributes	2
2. Conceptual framing: boundary object and operating model	2
2.1 Taxonomy is a controlled artefact	2
2.2 Boundary object in one practical sentence	3
2.3 Three domains of interpretation or discretion - and why traceability is the answer	3
3. Why risk taxonomy is now an innovation focus	4
3.1 Why now: stable categories, rising operational pressure	4
3.2 Three drivers of increasing expectations for risk taxonomies	4
3.3 How fintechs can contribute with credibility	5
4. Data and Findings	6
4.1 Data, approach, and coding frame	6
4.2 Summary findings from UK Banks reporting risk management	6
4.3 Eight patterns across UK banks' risk management reporting	9
5. Findings - From concepts to problems as four challenges that define a use-case portfolio	9
5.1 Challenge 1 - Adaptation: stable identity, controlled change	9
5.2 Challenge 2 - Discretion: unavoidable choices that must be governable	10
5.3 Challenge 3 - Distribution: one taxonomy across communities and systems	10
5.4 Challenge 4 - Traceability to regulatory rule sets and governance decisions	10
6. Findings - Industry-led use cases for fintech innovation	11
Use Case 1 - Versioned taxonomy governance featuring controlled evolution	11
Use Case 2 - Cross-walk and reconciliation across risk artefacts	12
Use Case 3 - Evidence packs for risk categories, emphasising traceability-by-design	12
Use Case 4 - Data lineage and explainable aggregation	12
Use Case 5 - Scenario governance for data-sparse risks	13

Use Case 6 - Model risk integration with taxonomy.....	13
Use Case 7 - Translation layer internal granularity interacting with external reporting	13
Use Case 8 - Evidenced learning loops as taxonomy development that you can prove	14
7. Findings summary: Prioritising where to focus industry-led use cases	14
Tier 1 - Highest priority	14
Tier 2 - Strong priority	15
Tier 3 - Foundational, to maintain.....	15
8. Discussion and conclusions	15
8.1 What this paper shows	15
8.2 What actionable research adds	15
8.3 Industry actions: publish use cases, with assurance expectations built in.....	15
8.4 Implications for the UK prudential supervisor	16
9. Appendices	17
9.1 Appendix A. Coding scheme for annual reports	17
A1. Purpose	17
A2. Data sources	17
A3. Unit of analysis.....	17
A4. Coding questions (as our core framework)	17
A5. Year-on-year change coding	18
A6. Reliability and transparency.....	19
9.2 Appendix B. Annual report evidence base and extraction notes.....	19
B1. Document list (2022–2024).....	19
B2. Extraction approach.....	19
B3. Notes on interpretation	19
9.3 Appendix C. A use-case template for industry publication	19
Use Case ID and Title	20
9.4 Appendix D. Glossary and working definitions	21
10. References	23
11. About the Authors.....	26

1. Introduction

1.1 Purpose and audience

The purpose of this paper is to support industry-led innovation in risk taxonomy operations. Specifically, we make the case for large financial services firms to articulate a portfolio of use cases - grounded in real governance and reporting needs - that smaller and often start-up or scale-up fintechs can address with targeted tools and analytics.

The intended audience is practitioners and professionals, as senior leaders connected to CFO and CRO responsibilities, assurance and audit leaders, risk professionals (including professional bodies such as Global Association of Risk Professionals), and supervisory stakeholders in the UK, such as the Bank of England, Prudential Regulatory Authority, and Financial Conduct Authority. The paper is written as a white paper, reporting actionable research directed towards application in procurement and assurance.

1.2 Risk taxonomies as governance infrastructure

Risk taxonomies are foundational to enterprise risk management in financial services. They provide a structured language for classifying risks, enabling aggregation, governance oversight, and regulatory reporting. Yet the apparent stability of top-level categories can hide operational complexity: (i) definitional choices and thresholds; (ii) mapping between multiple artefacts, including risk registers, incident taxonomies, control libraries, model inventories, stress and scenario processes, disclosures; (iii) aggregation rules supporting the governance of hierarchies; and (iv) the governance of taxonomy and organisational change.

This operational reality matters because risk management practices and culture do not automatically translate into observable outcomes in real time. When outcomes become visible, it may be too late for corrective action. From the perspectives of senior management, assurance functions, and supervisors, early indicators of capability and credible evidence of how we know what we know are essential (Bank of England, PRA, 2023). Treating taxonomy as infrastructure also raises practical questions of ownership, maintenance, coordination, failure, and accountability across the organisations that build, use, and supervise it.

1.3 The golden thread is traceability

We argue that the value of a risk taxonomy depends on traceability: the ability to track and explain how risk categories, data, interpretation, and mitigation are defined, applied, and changed, and how that is communicated internally and externally. Traceability serves both corporate and supervisory objectives. For firms, it reduces reconciliation costs, strengthens governance, improves decision-making speed, and supports credible disclosure. For assurance functions, it increases testability and

repeatability. For supervisors, it increases visibility into capability and control, supports comparability, and reduces uncertainty when risks are emerging or data are sparse (ICAEW, 2018).

1.4 Why the boundary object concept is useful

A taxonomy sits in between key stakeholder communities: firms, supervisors and regulators, and professional bodies (e.g., Global Association of Risk Professionals). It must support coordination without requiring full agreement on meanings, priorities, or incentives. This is the practical reason why we frame a taxonomy as a boundary object: it is shared, it travels across communities, and it holds its identity while being used differently (Star and Griesemer, 1989). Recent work shows that boundary objects remains a live, productive lens for understanding coordination and knowledge integration in innovation settings (Caccamo, Pittino and Tell, 2023; Hummel, Berends and Tuertscher, 2025).

We use the concept lightly, as a discipline that keeps our analysis focused on where tension is normal (robustness versus local adaptation) and where design improvements are realistic (traceability, change control, mapping, assurance).

1.5 What this paper contributes

We offer three contributions:

1. An operational account of where interpretation, discretion, and adaptation occur within risk taxonomies, and why this is governance work rather than error or noise in risk mapping and analysis.
2. A comparative review of public risk disclosures across three major UK banks over 2022–2024, showing stability at the core and change through emphasis, granularity and non-financial and resilience focus.
3. A procurement-focused portfolio of industry-led use cases that invite fintech innovation in taxonomy operations, designed from the outset to satisfy assurance requirements and to improve supervisory expectations.

2. Conceptual framing: boundary object and operating model

2.1 Taxonomy is a controlled artefact

A risk taxonomy is a hierarchical classification system for categorising and defining the risks an organisation may encounter. It guides risk identification and supports aggregation and reporting. However, in financial services, the taxonomy is never just a classification. It is embedded in an operating model: accountabilities, controls, governance forums, reporting routines, and assurance and testing practices. It must

function across multiple domains (finance, operations, technology, legal and compliance, customer outcomes), across levels (from business units to the enterprise), and across different data environments (from data-rich to data-sparse).

2.2 Boundary object in one practical sentence

A boundary object is a shared artefact that different groups can use for their own purposes while still recognising it as the same thing (Star and Griesemer, 1989). In financial services, the risk taxonomy must remain robust enough for comparability and supervisory oversight, yet flexible enough to adapt to firm-specific realities and evolving risk landscapes. Recent work shows boundary objects also form part of wider boundary infrastructures sustaining collaboration across organisational settings (Finch and Geiger, 2010; Caccamo, Pittino and Tell, 2023; Hummel, Berends and Tuertscher, 2025). This is the tension of robust identity against local adaptation that makes traceability non-negotiable.

2.3 Three domains of interpretation or discretion - and why traceability is the answer

Where we observe a taxonomy's stability in Section 4 (below), this can mask three predictable zones of interpretation or discretion:

(A) Content interpretation of what is identified, included and how it is defined. Which risks are deemed material, how categories are defined, how thresholds are set, and how measurement approaches vary. This includes how a risk is seen and made visible when data are thin.

(B) Structural discretion of granularity, hierarchy, and aggregation rules (Azzutti, 2026). How categories are grouped, what sits above/below what, how cross-cutting risks are handled, how aggregation is performed across business lines, and how the taxonomy relates to risk appetite frameworks, KRIs, incidents and controls.

(C) Adaptive discretion of how change is governed and justified. How the taxonomy evolves with external change (technology, competitive environment, supervisory focus, major events), how risks seem new to the taxonomy are introduced, how categories are re-weighted or decomposed, and how change is communicated, implemented and assured across the organisation.

We use interpretation (in A, above) to refer to decisions made in connection with supervisory reporting and compliance rules and guidance, and discretion (in B and C, above) to refer to business and enterprise decisions. Across all three zones, the design requirement is the same. Traceability requires explanations of what changed, why, by whom, based on what evidence, under what controls, and with what impacts (ICAEW, 2018; Bank of England, PRA, 2023).

3. Why risk taxonomy is now an innovation focus

3.1 Why now: stable categories, rising operational pressure

Risk taxonomies in large financial services firms appear mature and relatively stable. This stability supports governance continuity, comparability over time, and communication with boards, auditors, and supervisors. However, the operational load of maintaining a taxonomy has increased sharply. The cost is not primarily in writing the list of risks, but in governing definitions, managing mappings across systems and artefacts, ensuring consistent aggregation and reporting, and producing evidence that senior stakeholders can rely upon (Boulwood and Switzer, 2025).

Taxonomy has become an innovation focus because the demands on taxonomy operation have intensified while organisations simultaneously face pressures to simplify reporting, accelerate decision-making, and demonstrate credible controls (Bank of England, PRA, 2023; Financial Conduct Authority, 2018, 2025). For CFO-adjacent leaders, taxonomy operations sit on the critical path for governance sign-off, regulatory and public disclosures, assurance testing and audit comfort, and investment prioritisation in risk and resilience capabilities.

In banking, a significant part of the prudential backbone of risk taxonomy is drawn from the Basel framework and its domestic implementation. Categories such as credit risk, market risk, operational risk, counterparty credit risk, credit valuation adjustment (CVA) risk, liquidity, leverage, and interest rate risk in the banking book (IRRBB) are embedded in prudential rules, supervisory reporting and disclosure frameworks. The Basel Committee's finalisation of Basel III reforms was intended in part to restore credibility and reduce excessive variability in the calculation of risk-weighted assets, including through revised approaches to credit, market and operational risk and the introduction of the output floor. Firms therefore do not construct taxonomies from a blank slate. Rather, they interpret and adapt a regulatory framework within their own governance, reporting and assurance systems. This strengthens the case for understanding risk taxonomy both as an internal classificatory device, and as a rule-linked governance infrastructure that must remain traceable back to prudential categories and reporting requirements (Basel Committee on Banking Supervision, 2017; Bank of England, PRA, 2026).

3.2 Three drivers of increasing expectations for risk taxonomies

Risk taxonomies are gaining prominence among risk professionals, senior managers, and supervisors and regulators, and we identify three trends.

(A) Growth in cross-cutting and non-financial risk salience. Operational resilience, technology risk, economic and financial crime, customer outcomes, and climate-related issues increasingly cut across business lines and traditional organisational groups. These areas generate complex interdependencies and multiple classification and

reporting demands at once, such as incidents, controls, KRIs, scenarios, third-party assessments, and model inventories. The boundary-object tension becomes more visible as both stable identity and local adaptation (Star and Griesemer, 1989; Hummel, Berends and Tuertscher, 2025). The taxonomy must coordinate these domains without losing its stable comparability.

(B) Evidence requirements are moving upstream. Boards, assurance functions, and supervisors increasingly want credible evidence of capability and control, not only category names. If taxonomy is governance infrastructure, then evidence must travel with it, to include definitions, thresholds, mapping principles, lineage from data to metrics, decision logs, and controlled change histories (ICAEW, 2018; Bank of England, PRA, 2023). Evidence assembled late is expensive, harder to defend, and signals a missed opportunity for sharing and coordinating understandings across the organisation.

(C) Translation and reconciliation burdens are growing. Firms rarely operate a single taxonomy in a single system. They maintain multiple connected artefacts such as risk registers, incident taxonomies, control libraries, model inventories, stress and scenario libraries, policy and process repositories, and reporting and disclosure templates. These artefacts evolve at different speeds with different ownership. The hidden cost sits in translation zones between these, to include mapping rules, aggregation logic, and exception handling. Translation failures produce inconsistent MI, late reporting breaks, unresolved classification disputes, duplicated work, and avoidable assurance findings. All these challenges are visible to CFO leaders.

3.3 How fintechs can contribute with credibility

The challenges with risk taxonomies are both operational and socio-technical, to include data management, workflow, audit trail, controlled change, and explainable aggregation. This creates an opportunity for specialist providers, including fintechs, to suggest scaled solutions, particularly to traceability challenges around taxonomy operations. For solutions to be credible in financial services, they should be designed with assurance in mind from the outset.

Credible contributions require:

1. Procurement-ready use cases: decision ownership, defined data inputs, integration expectations, outputs, and success measures.
2. Assurance-aware design: immutable audit trails, access controls, segregation of duties, change control, testability, and retention (ICAEW, 2018).
3. Supervisory comfort: improved comparability and explainability without requiring identical taxonomies across firms (Bank of England, PRA, 2023).

4. Data and Findings

4.1 Data, approach, and coding frame

We present a comparative case-study design using annual-report risk management statements from three large UK banks across 2022–2024 (Barclays PLC, Lloyds Banking Group, and NatWest Group). Our approach follows the established logic of case-based theory development, while reflecting guidance on how cases and observations are constructed and made comparable in practice (Eisenhardt, 1989; Eisenhardt and Graebner, 2007; Ragin, 1992; Collins, Neely and Khan, 2024).

Demonstrating change is important for this paper as it supports the case for taxonomy as an innovation focus. Using annual reports as the primary source, we code: (1) the main risks identified, (2) implied hierarchy, (3) grouping logic (financial and non-financial), and (4) interpreting how disclosures could map to an underlying taxonomy. Our coding schema is in Appendix A towards the end of this paper. We then assess year-on-year change in emphasis. Annual reports do not present technical taxonomies per se, but they enable a consistent reading of risk categories, their relationships, and shifting emphasis over time. We set out our data collation in Appendix B towards the end of this paper.

Overall, the findings show strong stability in core risk categories, with change primarily driven by increased granularity, narrative emphasis, and the prominence of non-financial, cross-cutting, and resilience-related risks. Furthermore, some of the most stable categories identified across the annual reports are likely to reflect not only firm-level continuity, but also the institutional influence of Basel-derived prudential classifications. This is especially relevant for credit, market, operational, counterparty and liquidity and capital-related categories, which sit close to regulatory capital, disclosure and supervisory reporting frameworks. Stability in these categories should not be read only as organisational conservatism. It may also indicate the durability of externally anchored prudential categories that firms must continuously interpret, operationalise and reconcile with broader internal governance and public reporting needs (Basel Committee on Banking Supervision, 2017; Bank of England, PRA, 2026).

4.2 Summary findings from UK Banks reporting risk management

[Barclays](#) (2022–2024): Stable set of principal risks, hierarchical presentation, consistent financial/non-financial distinction, climate and resilience treated as cross-cutting, year-on-year change mainly through granularity and interdependencies rather than category redesign.

[NatWest Group](#) (2022–2024): Stable split between financial and non-financial categories, disclosures signal a structured “risk directory” logic, change expressed through embedding risk appetite, stress testing, and resilience themes rather than redesign.

[Lloyds Banking Group](#) (2022–2024): Stable financial and non-financial categories, hierarchical presentation supported by drivers, controls and governance, movement largely in narrative emphasis and governance simplification, with rising resilience/economic crime focus.

We present our findings in Table 1 and Table 2 (below). Table 1 aggregates the set of reports and identifies how banks present common, seemingly stable risk categories. Table 2 identifies areas of discretion, indicating where traceability is multi-layered, for example, across data, interpretation, and review. We include a glossary of terms, mostly as set out in the tables below and in our findings, in Appendix D towards the end of this paper.

Table 1 Cross-bank stability, evolution, and traceability burden (2022–2024)

Risk family (top level)	Stability	Change	Traceability burden	Hotspots	Typical challenges
Credit	●●●	▲▲○	■□	M, L	Portfolio mapping and threshold choices create friction.
Market	●●●	▲○○	■□	L, S	Stress/scenario linkages; reproducibility.
Liquidity / Capital	●●●	▲○○	■□	L, G	Aggregation rules; interfaces with capital processes/models.
Pension (where applicable)	●●○	▲○○	■□	L, G	Assumptions and governance documentation.
Operational	●●●	▲▲○	■■■	M, L, G	Incidents, controls, key reporting incidents reconciliation – all labour-intensive.
Conduct / customer outcomes	●●●	▲▲○	■□	M, G	Definitions, accountability trails, MI consistency.
Compliance / legal / regulatory	●●○	▲▲○	■□	M, G	Cross-framework mapping questions; need for repackaging for reporting.
Financial/economic crime	●●●	▲▲▲	■■■	L, M, G	Typologies, lineage for decisions and reporting consistency.
Model risk	●●●	▲▲▲	■■■	L, G	Inventory/tiering/validation evidence, and usage context links.
Climate (cross-cutting)	●●●	▲▲▲	■■■	S, L, G	Data-sparse, so scenario-heavy; assumptions need audit trail.
Reputational	●●●	▲○○	■□	M, G	Causal attribution and linkage to mitigants is hard to evidence.
Technology / operational resilience	●●○	▲▲▲	■■■	M, L, G, S	Interdependencies, testing evidence and change logs.

Legend:

Stability ●●● very stable | ●●○ stable | ●○○ variable

Evolution ▲▲▲ rising strongly | ▲▲○ rising | ▲○○ modest

Traceability burden ■■■ high | ■■□ medium | ■□□ lower

Hotspots: M mapping/cross-walk | L lineage/reproducibility | G governance/change control | S scenarios/assumptions

The interpretation of Table 1 (above) is that the innovation focus sits where both Change and Traceability burdens are high. Notably, this is at operational risk, financial and economic crime, model risk, climate (by virtue of being a cross-cutting issue), technology, and resilience. This also helps explain why some apparently mature categories may still carry a high traceability burden as they are linked to prudential rule sets, supervisory review, and enforcement-sensitive reporting obligations (Basel Committee on Banking Supervision, 2017; Bank of England, PRA, 2026).

Table 2 Where interpretation and discretion sit and what good evidence looks like.

Interpretation/ discretion domain	Typical decision	Minimum control expectation	Evidence artefact
Content interpretation	Define category scope and thresholds	Named ownership and approval	Definition register, rationale and effective date
Structural discretion	Change hierarchy / introduce sub-categories	Change control, impact assessment	Versioning, cross-walk and impact summary
Grouping logic	Financial vs non-financial; cross-cutting placement	Governance sign-off	Mapping principles and exception/outliers log
Adaptive discretion	Re-weight emphasis / introduce new focus	Review schedule and documentation	Review minutes, decision log and MI/disclosure link
Data-sparse governance	Scenarios / expert judgement	Assumptions governance and challenge	Assumptions register, scenario library and review notes

4.3 Eight patterns across UK banks' risk management reporting

Drawing further on our cross-case analysis presented in Tables 1 and 2 (above) (Eisenhardt, 1989; Collins, Neely and Khan, 2024), we identify eight patterns from our coding of the banks' risk management sections in the annual reports:

1. Core category stability dominates, with change being mostly below the surface.
2. Year-on-year change appears as shifts in emphasis, rather than category redesign, and provides insights into changes that are otherwise multi-layered for any particular year.
3. Granularity seems to increase where operating and supervisory pressures rise.
4. Disclosures increasingly blend risk phenomena with signals of governance maturity.
5. Interdependencies across the stable categorisations, such as climate, technology and resilience, are increasingly explicit.
6. Translation zones across the artefacts of risk management carry a significant operational and assurance burden.
7. Data-rich and data-sparse risks require different governance and evidence.
8. Traceability is the unifying factor across CFO, assurance, and supervisory requirements.

To summarise, the implication of the comparison is that the potential value of innovation lies in taxonomy operations, including traceability, mapping, versioning, lineage and scenario governance, more than in redefining top-level risk categories.

5. Findings - From concepts to problems as four challenges that define a use-case portfolio

We present these findings in two steps. First, in this section, we identify four challenges faced across the sector by financial services firms, the fintechs that may offer supporting services, and regulators and supervisors. We then in Section 6 (below) draw out eight use cases, framed as potential opportunities for fintechs offering services to financial services companies.

5.1 Challenge 1 - Adaptation: stable identity, controlled change

Problem: contexts change faster than taxonomies can be redesigned without breaking comparability.

Risk: category drift, inconsistent definitions, and costly reconciliations to establish and share understandings of what changed.

Good practice: versioned governance, including cross-walks that can be read both prospectively and retrospectively, together with clear mapping principles for cross-cutting risks (Bank of England, PRA, 2023).

Use-case links: UC1, UC2, and UC7, set out in Section 6 (below).

5.2 Challenge 2 - Discretion: unavoidable choices that must be governable

Problem: category boundaries, thresholds, hierarchy and aggregation choices involve interpretation and discretion; it must be governable and auditable

Risk: disputes, inconsistent MI, unclear accountability, limited testability in different domains (Star and Griesemer, 1989; Caccamo, Pittino and Tell, 2023).

Good practice: explicit ownership, documented aggregation rules, evidence packs linking judgement to data and actions (ICAEW, 2018).

Use-case links: UC3, UC4, UC5.

5.3 Challenge 3 - Distribution: one taxonomy across communities and systems

Problem: different artefacts in risk management change at different speeds, often stimulated by a need for operational responses and lines of defence.

Risk: duplication, late-stage reporting risk, inconsistent disclosures, and translation deficits among those not directly involved in any specific operational challenge.

Good practice: controlled cross-walk across artefacts; a clear internal-to-external translation layer; and governance and access controls (ICAEW, 2018).

Use-case links: UC2, UC6, UC7.

5.4 Challenge 4 - Traceability to regulatory rule sets and governance decisions

Problem: without traceability, firms cannot explain changes, reproduce metrics, map decisions to regulatory rule sets and supervisory expectations, or evidence governance under scrutiny.

Risk: slow sign-off, higher assurance cost, and weaker supervisory comfort.

Good practice: immutable audit trails, reproducible calculations with lineage, transparent mapping to regulatory rule sets and supervisory expectations, controlled scenario or assumption governance for data-sparse risks, and evidenced learning loops (ICAEW, 2018; Bank of England, PRA, 2023). This is the golden thread that turns taxonomy into an assured capability

Use-case links: UC1, UC3, UC 4, UC5, and UC8 (portfolio-wide).

In summary, these four challenges show the ways in which taxonomy is an actionable innovation question. The opportunity for innovation is not endless reclassifications of risks. Rather, it is in the development through the interactions of financial services and fintechs of an evidence infrastructure at scale that makes taxonomies into reliable boundary objects under assurance and supervisory scrutiny.

6. Findings - Industry-led use cases for fintech innovation

Drawing on our first set of findings as sector-wide challenges in Section 5 (above), we develop these further in this section by presenting eight use cases. These are framed from the perspective of services that fintechs could offer to financial services companies in supporting the development and implementation of risk taxonomies, recognising the critical importance of supervision and regulation.

We present eight use-cases using a common template and emphasising the qualities of procurement-readiness from the perspective of financial services buying from fintechs and an awareness of assurance, regulation and supervision requirements. Each use case includes: decision owner and users, problem identification, minimum inputs, workflow and controls to be logged, outputs, assurance requirements, and success measures. We include a template for the uses cases as Appendix C towards the end of this paper.

The use cases below are designed to be: (i) procurement-ready for financial services companies, (ii) achievable for fintechs, and (iii) assurance-aware by construction, consistent with the expectations of a skilled-person review and auditability principles (ICAEW, 2018) and addressing the expectations of model risk governance where relevant (Bank of England, PRA, 2023). The innovation logic is to connect firm-level operational and commercial practice with institutional expectations in supervision and professional standards. In terms of our methodology, we draw on a pattern also analysed in research on the linking of commercial and institutional entrepreneurship for grand challenges (Volmar and Eisenhardt, 2025).

Use Case 1 - Versioned taxonomy governance featuring controlled evolution

Decision owner and users: enterprise risk; taxonomy owner; assurance; finance disclosure.

Problem: taxonomy changes are hard to govern and explain across time.

Minimum inputs: taxonomy repository; approvals; mapping tables; reporting categories.

Workflow and controls: change request, rationale, impact assessment, approvals, and effective date. Supported by version control, Segregation of Duties (SoD), and access logs.

Outputs: version history, cross-walk old to new, impact summary and evidence pack.

Assurance: immutable audit trail, reproducibility, and periodic testing.

Success: fewer late reporting breaks, faster sign-off, and fewer audit findings.

Use Case 2 - Cross-walk and reconciliation across risk artefacts

Decision owner and users: operational risk, compliance, finance reporting, assurance.

Problem: parallel artefacts drift, reconciliation is manual and inconsistent.

Minimum inputs: risk register, incident system, control library, KRI catalogue, mapping rules.

Workflow and controls: mapping governance, exceptions, automated reconciliation with review, approvals.

Outputs: mapping tables, reconciliation dashboard, exceptions log, lineage view.

Assurance: mapping change control, review evidence, access and SoD.

Success: reduced reconciliation effort, shorter reporting cycles, improved MI consistency.

Use Case 3 - Evidence packs for risk categories, emphasising traceability-by-design

Decision owner and users: assurance and audit, risk owners, and board reporting.

Problem: evidence is scattered and assembled late, making it hard to answer the question of ‘how do we know?’

Minimum inputs: definitions, metrics, data sources, control tests, incidents, actions.

Workflow and controls: automated collection; tagging; approvals; refresh cadence; retention.

Outputs: per-category evidence pack, linking data to judgement, actions and outcomes.

Assurance: provenance, time-stamping, retention, independent review workflow.

Success: faster responses; reduced searches of historic presentations to specified audiences, and improved board confidence.

Use Case 4 - Data lineage and explainable aggregation

Decision owner and users: finance reporting, data governance, risk analytics and assurance.

Problem: aggregation rules and metric definitions are implicit; results are not reproducible.

Minimum inputs: data catalogue, data transformations, metrics registry, hierarchy and aggregation logic.

Workflow and controls: lineage capture, approvals, checks, and exception handling.

Outputs: lineage maps, metric registry, reproducible calculation logs, diagnostics of why an item moved in a taxonomy hierarchy.

Assurance: reproducibility, change control, test scripts/results.

Success: fewer breaks, fewer disputes, faster sign-off.

Use Case 5 - Scenario governance for data-sparse risks

Decision owner and users: enterprise risk, resilience, board risk committee, and assurance.

Problem: scenario work relies on judgement, weak governance undermines confidence.

Minimum inputs: scenario library, assumptions, expert-led inputs, outputs and mitigations.

Workflow and controls: assumptions register; challenge process; approvals; refresh calendar.

Outputs: scenario evidence pack; assumptions log; change history; decision link.

Assurance: documented challenge, independence where required, versioning.

Success: strengthening governance credibility, clearer supervisory dialogue, reduced black box risk.

Use Case 6 - Model risk integration with taxonomy

Decision owner and users: model risk, finance, assurance and enterprise risk.

Problem: model governance is siloed, taxonomy reporting needs clear linkage to model reliance.

Minimum inputs: model inventory, tiering, validation, usage contexts, and taxonomy mapping.

Workflow and controls: mapping model risk as a category, approvals, and validation cycle integration.

Outputs: category-level model exposure view, limitations, validation status, and mitigations

Assurance: validation evidence linkage, access controls, and change logs.

Success: better senior decisions on model reliance, fewer surprises, improved audit comfort.

Use Case 7 - Translation layer internal granularity interacting with external reporting

Decision owner and users: enterprise risk, finance disclosure, regulatory reporting, and assurance.

Problem: internal detail and external comparability pull in opposite directions.

Minimum inputs: internal taxonomy, external categories, mapping rules, MI/disclosure outputs.

Workflow and controls: mapping principles, exceptions, periodic review, narrative alignment checks.

Outputs: stable external set and transparent mapping, drill-down to internal views.

Assurance: mapping control, review evidence, change impact summaries.

Success: reduced disclosure friction, improved comparability, and fewer escalations.

Use Case 8 - Evidenced learning loops as taxonomy development that you can prove

Decision owner and users: enterprise risk, assurance, transformation, and strategy.

Problem: learning is claimed but not evidenced, changes are not linked to outcomes.

Minimum inputs: incidents and near misses, root cause, taxonomy change log, and control improvements.

Workflow and controls: learning events, insight, decision, taxonomy control change, and impact check.

Outputs: learning register, linkage to taxonomy changes, and closure evidence dashboards.

Assurance: auditability of decisions, independent challenge, and closure evidence.

Success: fewer repeat events, clearer investment prioritisation, demonstrable improvement cycle.

7. Findings summary: Prioritising where to focus industry-led use cases

This section sets out an evaluation and scoring approach to compare and rank the risks presented in Table 1 (above). From the perspective of financial services firms, risks are grouped as highest priority, strong priority, and foundational to maintain. We highlight where these risks feature in the use cases presented in Section 6 (above). Prioritisation is derived from the analysis set out in Table 1 and informed by assurance and supervisory logic (ICAEW, 2018; Bank of England, PRA, 2023).

Tier 1 - Highest priority

Operational risk (including resilience, as this connects across taxonomy categories), featuring in Use Cases 1, 2, 3 and 4.

1. Financial/economic crime, featuring in Use Cases 1, 3, 4 and 7.
2. Model risk, featuring in Use Cases 1, 3 and 6.
3. Climate as cross-cutting with data-sparse governance, featuring in Use Cases 3, 5 and 7.
4. Technology and operational resilience, featuring in Use Cases 1, 2, 4 and 5.

5. Compliance, legal and regulatory, featuring in Use Cases 1, 2 and 7.

Tier 2 - Strong priority

6. Credit, featuring in Use Cases 1 and 4.
7. Market, featuring in Use Cases 4 and 5.
8. Liquidity and capital, featuring in Use Cases 1 and 4.
9. Conduct and customer outcomes, featuring in Use Cases 1, 3 and 7.

Tier 3 - Foundational, to maintain

10. Reputational, featuring in Use Cases 3 and 8.
11. Pension (where applicable), featuring in Use Cases 1 and 4.
12. Pension (where applicable), featuring in Use Cases 1 and 4

The implication of prioritising is that the highest-return fintech opportunities can focus on scaling up the core quality of traceability within Tier-1 risk categories. Mapping these into the use cases, the qualities to address in any solutions could include versioned governance, cross-walk/reconciliation, evidence packs, lineage, and the governance of scenarios and assumptions.

8. Discussion and conclusions

8.1 What this paper shows

Through reviewing annual reports, major UK banks portray risk taxonomy as a stable governance infrastructure. Change tends to be expressed through emphasis, granularity, and cross-cutting interdependencies rather than top-level redesign (Barclays PLC, 2022–2024; Lloyds Banking Group, 2022–2024; NatWest Group, 2022–2024).

8.2 What actionable research adds

The contribution to actionable research is to provide a more precise account of what practitioners are already doing and why establishing traceability is so hard (Azzutti, 2026). Treating taxonomy as infrastructure brings into view questions of ownership, maintenance, coordination, failure, and accountability. Boundary-object framing explains coordination across different domains without full consensus (Star and Griesemer, 1989).

8.3 Industry actions: publish use cases, with assurance expectations built in

We recommend that large firms publish an industry-led portfolio of procurement-ready use cases, each specifying decision ownership, minimum data inputs, workflow controls (audit trail, approvals, SoD), outputs, and success measures. This creates a structured interface between industry needs and fintech innovation, while ensuring solutions produce assurance-grade evidence (ICAEW, 2018).

8.4 Implications for the UK prudential supervisor

For the Prudential Regulation Authority, the practical relevance is that firms' taxonomy operations should support supervisory confidence through the principle of traceability-by-design. This point is sharpened by the Basel framework and its UK implementation. A substantial part of the prudential core of banking risk taxonomy is inherited from Basel categories and reporting logic rather than being designed by firms. The challenge for firms is to translate Basel-linked categories into operating, reporting and assurance architectures that remain intelligible across business functions, boards, assurance functions and supervisors. This makes interpretation of rule sets, alongside managerial discretion, central to taxonomy governance. It also reinforces why mature categories such as credit, market, liquidity and capital, and compliance-related risks may continue to carry high traceability burden, particularly as the UK moves toward Basel 3.1 implementation from 1 January 2027, with the market-risk internal-model component following from 1 January 2028 (Basel Committee on Banking Supervision, 2017; Bank of England, PRA, 2026).

Without prescribing firm-specific structures, supervisory engagement can reward controlled taxonomy change (versioning, approvals, impact assessment), transparent internal-to-external mapping, auditable lineage from source data to risk metrics and governance decisions, and explicit governance for data-sparse risks (Bank of England, PRA, 2023; Bank of England, PRA, 2024a; Bank of England, PRA, 2024b). Despite top-level stability and similarity, we would not expect firms to adopt or present identical taxonomies.

A secondary implication concerns the broader regulatory ecosystem, including the Financial Conduct Authority, where customer outcomes and conduct-related themes intensify demands for consistent categorisation, evidence, and accountability trails (Financial Conduct Authority, 2018, 2025). Industry-led use cases can support both prudential and conduct priorities without introducing a single taxonomy standard, by strengthening the evidence infrastructure that makes taxonomies reliable boundary objects under scrutiny (Star and Griesemer, 1989; Hummel, Berends and Tuertscher, 2025).

9. Appendices

9.1 Appendix A. Coding scheme for annual reports

A1. Purpose

This appendix documents the coding approach used to analyse how risk taxonomies are signalled through public risk disclosures in annual reports. The goal is to identify stable categories, shifts in salience, and implied hierarchy/grouping logic in a way that is transparent and replicable. There is no expectation that banks will disclose their taxonomies in annual reports.

A2. Data sources

We analysed risk disclosures in annual reports for three large UK banks across 2022–2024 (see Appendix B for document list).

A3. Unit of analysis

- Primary unit: each bank-year annual report (e.g., “Bank A - 2023”).
- Coding unit: risk-related sections describing principal risks, including narrative descriptions, risk factor groupings, and supporting governance and controls where disclosed.

A4. Coding questions (as our core framework)

We used four guiding questions (with standard prompts) for each bank-year document:

Question 1 - What risks are identified?

Record:

- List of disclosed “principal risks” (or equivalent).
- Presence or absence of recurring top-level risk families (credit, market, liquidity/capital, operational, conduct/customer outcomes, financial/economic crime, model risk, climate, technology/resilience, reputational, legal/regulatory, and pensions where applicable).
- Any new or newly elevated risks.

Prompt: Are risks presented as a stable list or as an evolving set? Are some framed as risk categories or drivers?

Question 2 - What hierarchy is implied?

Record:

- Any explicit levels (e.g., principal risks relative to sub-risks/drivers).

- Whether cross-cutting risks are treated as standalone categories or as overlays (e.g., climate or resilience influencing other categories).
- Whether the annual report implies a taxonomy-like directory or classification reference.

Prompt: Where do risk categories appear to sit relative to controls, governance and decision cycles?

Question 3 - What grouping logic is used?

Record:

- Explicit or implied grouping (financial vs non-financial; strategic vs operational; prudential vs conduct; etc.).
- Where “hybrid” or cross-cutting risks are placed (climate, technology/resilience, financial crime).
- Any alignment cues to supervisory or professional classifications.

Prompt: What boundary work is being done? How does the disclosure help different stakeholder communities interpret the same risk set?

Question 4 - How does it map to a taxonomy in practice?

- Record, as signals only, given public disclosure limitations in annual reports:
- Evidence of mapping/translation between internal granularity and external reporting categories.
- Signs of controlled change (year-on-year continuity language, revisions, re-labelling, changes in number of categories).
- References to governance routines that imply traceability (committees, risk appetite alignment, stress/scenario processes, model governance).

Prompt: Does the disclosure allow an informed reader to see how categories connect to measurement, decision and accountability?

A5. Year-on-year change coding

For each bank, we recorded:

- Category stability: unchanged top-level categories (yes/no).
- Change type: (i) re-labelling; (ii) re-grouping; (iii) increased granularity; (iv) change in narrative emphasis; (v) introduction of new cross-cutting theme.
- Traceability burden signals: mentions of mapping, reconciliation, governance strengthening, resilience/testing, model validation, scenario processes.

A6. Reliability and transparency

- Coding was undertaken using a consistent template per bank annual report year.
- Outputs were summarised into cross-bank comparison tables and pattern narratives.
- Limitations are acknowledged: annual reports are disclosure artefacts, not internal working taxonomies; the analysis reads them as signals of taxonomy operation rather than as a full organisational description.

9.2 Appendix B. Annual report evidence base and extraction notes

B1. Document list (2022–2024)

The following annual reports were used as the empirical evidence base:

- Barclays PLC - Annual Report 2022; 2023; 2024.
- Lloyds Banking Group - Annual Report and Accounts 2022; 2023; 2024.
- NatWest Group - Annual Report and Accounts 2022; 2023; 2024.

(Full citations are in Appendix E.)

B2. Extraction approach

For each bank-year report, we extracted and reviewed:

- The section labelled “Principal risks”, “Risk management”, “Risk factors” or similar.
- Any adjacent material describing risk governance, risk appetite linkage, operational resilience, scenario/stress testing, model governance, and financial crime frameworks (where present).
- Where risk category tables were presented, we recorded the categories and any ordering/grouping logic.

B3. Notes on interpretation

- Public disclosures are written for multiple audiences (investors, regulators, counterparties, analysts). They are, therefore, useful evidence of how taxonomies are communicated across boundaries.
- However, disclosures are not internal operating documents, such that the absence of evidence is not evidence of absence.
- Our comparative focus is on stable category sets and how change is expressed (typically via emphasis and granularity rather than wholesale redesign).

9.3 Appendix C. A use-case template for industry publication

This template is designed so that a large firm can publish a credible problem statement that a fintech/scale-up can build against, while remaining procurement-ready and assurance-aware.

Use Case ID and Title

- Use Case X - [short, action-led name].

Decision owner and primary users

- Accountable owner: [e.g., enterprise risk, finance reporting, model risk, financial crime].
- Primary users: [roles/functions].
- Assurance stakeholders: [audit/assurance, compliance, second line, etc.].

Problem statement

- What fails today: [one paragraph].
- Why it matters: [cost/time, control risk, sign-off risk, supervisory risk].
- What “good” looks like: [3–5 bullets].

Scope and boundaries

- In scope: [processes, artefacts, systems, risk domains]
- Out of scope: [explicit exclusions].
- Dependencies: [data governance, identity and access, existing platforms].

Minimum inputs (data and artefacts)

- Data sources required (minimum viable set).
- Artefacts required (taxonomy, mapping tables, metrics registry, scenario library, etc.).

Workflow and controls to log (assurance-by-design)

- Required workflow steps (create/review/approve/change/retire).
- Required controls (audit trail, role-based access, segregation of duties, retention).
- Exception handling process and evidence.

Outputs

- Operational outputs (dashboards, alerts, MI).
- Reporting outputs (disclosure support, regulatory reporting support).
- Evidence outputs (evidence packs, change logs, lineage maps).

Success measures

- Time/cost measures (e.g., reduction in reconciliation hours).
- Quality measures (e.g., fewer classification disputes; fewer late-cycle breaks).
- Assurance measures (e.g., fewer findings; faster evidence retrieval).
- Supervisory dialogue measures (e.g., faster responses; clearer mapping explanation).

Implementation notes

- Integration expectations (APIs, data formats).
- Minimal viable deployment (pilot scope).
- Scale-up considerations.

9.4 Appendix D. Glossary and working definitions

This glossary uses what we could term white-paper definitions, intended to be operationally clear and good enough for shared understanding.

Aggregation: Combining risk information from multiple sources or business areas into a single view (e.g., enterprise risk profile), using defined rules.

Assumptions register: A controlled log of key assumptions used in scenario analysis, judgement-based assessments, or model limitations—ideally versioned, reviewable, and linked to decisions.

Audit trail: A record of who did what, when, and why, sufficient to reconstruct decisions and changes during assurance or supervisory review.

Boundary object: A shared artefact used across different communities for different purposes while still recognised as “the same thing.” In this white paper, the risk taxonomy functions as a boundary object across firms, assurance communities, and supervisors.

Boundary infrastructure: A more durable set of linked boundary objects, routines, and systems that sustains collaboration across organisational boundaries over time (useful language for describing mature traceability ecosystems).

Cross-walk: A mapping table that links categories in one taxonomy/artefact to categories in another (e.g., internal category set to external reporting categories), ideally versioned and governed.

Data lineage: The end-to-end record of how data moves and transforms from source systems through calculations into reported metrics and decisions.

Data-sparse risk: A risk domain where historical data is limited or non-representative (e.g., emerging risks), requiring structured judgement, scenario analysis, and controlled assumptions.

Evidence pack: A structured bundle of documentation and data outputs that demonstrates how a risk category is defined, monitored, governed, and acted upon—built for rapid retrieval and assurance review.

Exception handling: The controlled process for handling cases that do not fit standard taxonomy mapping or aggregation rules, including documentation and approvals.

Granularity: The level of detail in risk categorisation (e.g., a single “operational risk” category versus multiple sub-categories).

Internal to external translation layer: The governed mapping between the firm’s internal, decision-useful risk categorisation and the more stable categories used for public/regulatory reporting.

Learning loop (evidenced): A documented cycle linking events (incidents/near misses), insights, governance decisions, taxonomy/control changes, and verified impacts.

Mapping principles: Explicit rules guiding how categories are aligned across artefacts or reporting regimes (including how cross-cutting risks are treated).

Model risk: Risk arising from the use of models (including ML/AI) due to limitations, misuse, poor governance, or incorrect assumptions, requiring inventory, tiering, validation and usage controls.

Operational resilience: The ability to prevent, adapt to, respond to, recover from and learn from operational disruptions, often spanning technology, third parties, processes and governance.

Segregation of duties (SoD): Control principle ensuring no single individual can both initiate and approve high-impact changes, reducing fraud/error risk.

Taxonomy (risk taxonomy): A structured classification of risk categories used to support identification, governance, aggregation, reporting and accountability.

Traceability (golden thread): The ability to reconstruct what happened across the taxonomy lifecycle: definitions, mapping rules, data and metric lineage, changes, approvals, scenario assumptions, and resulting decisions/actions.

Versioning: Controlled management of successive taxonomy (or mapping/metric) versions, enabling backward comparison and reproducibility.

10. References

- Azzutti, A. (2026) ‘AI governance after MiFID II: beyond (mere) technological neutrality?’, *ERA Forum*. [doi:10.1007/s12027-026-00871-1](https://doi.org/10.1007/s12027-026-00871-1).
- Barclays PLC (2022) *Barclays PLC Annual Report 2022*. Available at: <https://home.barclays/content/dam/home-barclays/documents/investor-relations/reports-and-events/annual-reports/2022/AR/Barclays-PLC-Annual-Report-2022.pdf> (Accessed: 1 March 2026).
- Barclays PLC (2023) *Barclays PLC Annual Report 2023*. Available at: <https://home.barclays/content/dam/home-barclays/documents/investor-relations/reports-and-events/annual-reports/2023/Barclays-PLC-Annual-Report-2023.pdf> (Accessed: 1 March 2026).
- Barclays PLC (2024) *Barclays PLC Annual Report 2024*. Available at: <https://home.barclays/content/dam/home-barclays/documents/investor-relations/reports-and-events/annual-reports/2024/Barclays-PLC-Annual-Report-2024.pdf> (Accessed: 1 March 2026).
- Bank of England, Prudential Regulation Authority (PRA) (2023) *SS1/23: Model risk management principles for banks*. London: Bank of England. Available at: <https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/supervisory-statement/2023/ss123.pdf> (Accessed: 1 March 2026).
- Bank of England, Prudential Regulation Authority (PRA) (2024a) *SS7/14: Reports by skilled persons* (current version published 15 November 2024). London: Bank of England. Available at: <https://www.bankofengland.co.uk/prudential-regulation/publication/2014/reports-by-skilled-persons-ss> (Accessed: 1 March 2026).
- Bank of England, Prudential Regulation Authority (PRA) (2024b) *SS34/15: Guidelines for completing regulatory reports* (future version; near-final versions include September 2024). London: Bank of England. Available at: <https://www.bankofengland.co.uk/prudential-regulation/publication/2015/guidelines-for-completing-regulatory-reports-ss> (Accessed: 1 March 2026).
- Bank of England, Prudential Regulation Authority (PRA) (2026) *PS1/26: Implementation of Basel 3.1: Final rules*. London: Bank of England. Available at: <https://www.bankofengland.co.uk/prudential-regulation/publication/2026/january/implementation-of-the-basel-3-1-final-rules-policy-statement> (Accessed: 9 March 2026).
- Boulwood, B. and Switzer, J. (2025) ‘ERM: The importance of building a risk taxonomy’, *GARP Risk Intelligence*. 17 April. Available at: <https://www.garp.org/risk-intelligence/culture-governance/erm-importance-building-250417> (Accessed: 1 March 2026).

- Caccamo, M., Pittino, D. and Tell, F. (2023) ‘Boundary objects, knowledge integration, and innovation management: A systematic review of the literature’, *Technovation*, 122, 102645. [doi:10.1016/j.technovation.2022.102645](https://doi.org/10.1016/j.technovation.2022.102645).
- Collins, C., Neely, M.T. and Khan, S. (2024) “Which cases do I need?” Constructing cases and observations in qualitative research’, *Annual Review of Sociology*, 50, pp. 391–414. [doi:10.1146/annurev-soc-031021-035000](https://doi.org/10.1146/annurev-soc-031021-035000).
- Eisenhardt, K.M. (1989) ‘Building theories from case study research’, *Academy of Management Review*, 14(4), pp. 532–550. [doi:10.5465/AMR.1989.4308385](https://doi.org/10.5465/AMR.1989.4308385).
- Eisenhardt, K.M. and Graebner, M.E. (2007) ‘Theory building from cases: opportunities and challenges’, *Academy of Management Journal*, 50(1), pp. 25–32. [doi:10.5465/amj.2007.24160888](https://doi.org/10.5465/amj.2007.24160888).
- Financial Conduct Authority (2018) *FCA Handbook* (version as of 10 December 2018). Available at: <https://handbook.fca.org.uk/handbook?date=10-12-2018&timeline=True> (Accessed: 1 March 2026).
- Financial Conduct Authority (2025) *FCA Handbook* (version as of 6 May 2025). Available at: <https://handbook.fca.org.uk/handbook?date=06-05-2025&timeline=True> (Accessed: 1 March 2026).
- Finch, J.H. and Geiger, S. (2010) ‘Positioning and relating: Market boundaries and the slippery identity of the marketing object’, *Marketing Theory*, 10(3), pp. 237-251, <https://doi.org/10.1177/1470593110373188>
- Gatrell, C., Muzio, D., Post, C. and Wickert, C. (2024) ‘Here, there and everywhere: On the responsible use of artificial intelligence (AI) in management research and the peer-review process’, *Journal of Management Studies*, 61(3), pp. 1-13, <https://doi.org/10.1111/joms.13045>
- Hummel, J.T., Berends, H. and Tuertscher, P. (2025) ‘From boundary objects to boundary infrastructure: A process study of collaboration between big science and big business’, *Journal of Management Studies*, 62(4), pp. 1644–1679. [doi:10.1111/joms.13118](https://doi.org/10.1111/joms.13118).
- Institute of Chartered Accountants in England and Wales (ICAEW) (2018) *Guidance for skilled person reviews (TECH 01/18 FSF)*. London: ICAEW. Available at: <https://www.icaew.com/-/media/corporate/files/technical/technical-releases/financial-services/tech-01-18-fsf-guidance-for-skilled-persons-reviews.ashx> (Accessed: 1 March 2026).
- Lloyds Banking Group (2022) *Lloyds Banking Group plc Annual Report and Accounts 2022*. Available at: <https://www.lloydsbankinggroup.com/assets/pdfs/investors/financial-performance/lloyds-banking-group-plc/2022/full-year/2022-lbg-annual-report.pdf> (Accessed: 1 March 2026).

- Lloyds Banking Group (2023) *Lloyds Banking Group plc Annual Report and Accounts 2023*. Available at: <https://www.lloydsbankinggroup.com/assets/pdfs/investors/financial-performance/lloyds-banking-group-plc/2023/q4/2023-lbg-annual-report.pdf> (Accessed: 1 March 2026).
- Lloyds Banking Group (2024) *Lloyds Banking Group plc Annual Report and Accounts 2024*. Available at: <https://www.lloydsbankinggroup.com/assets/pdfs/investors/financial-performance/lloyds-banking-group-plc/2024/q4/2024-lbg-annual-report.pdf> (Accessed: 1 March 2026).
- NatWest Group (2022) *NatWest Group plc Annual Report and Accounts 2022*. Available at: <https://investors.natwestgroup.com/~media/Files/R/RBS-IR-V2/results-center/17022023/nwg-annual-report-and-accounts.pdf> (Accessed: 1 March 2026).
- NatWest Group (2023) *NatWest Group plc Annual Report and Accounts 2023*. Available at: <https://investors.natwestgroup.com/~media/Files/R/RBS-IR-V2/results-center/16022024/nwg-annual-report-and-accounts-2023.pdf> (Accessed: 1 March 2026).
- NatWest Group (2024) *NatWest Group plc Annual Report and Accounts 2024*. Available at: <https://www.investors.rbs.com/~media/Files/R/RBS-IR-V2/results-center/14022025/nwg-annual-report-and-accounts-2024.pdf> (Accessed: 1 March 2026).
- Ragin, C.C. (1992) “Casing” and the process of social inquiry’, in Ragin, C.C. and Becker, H.S. (eds.) *What is a Case? Exploring the Foundations of Social Inquiry*. Cambridge: Cambridge University Press, pp. 217–226.
- Slawinski, N., Brito, B., Brenton, J. and Smith, W.K. (2025) ‘Reflections on deep academic–practitioner partnering for generative societal impact’, *Strategic Organization*, 23(1), pp. 134–146. [doi:10.1177/14761270241274590](https://doi.org/10.1177/14761270241274590).
- Star, S.L. and Griesemer, J.R. (1989) ‘Institutional ecology, “translations” and boundary objects: amateurs and professionals in Berkeley’s Museum of Vertebrate Zoology, 1907–39’, *Social Studies of Science*, 19(3), pp. 387–420. [doi:10.1177/030631289019003001](https://doi.org/10.1177/030631289019003001).
- Van de Ven, A.H. (2007) *Engaged Scholarship: A Guide for Organizational and Social Research*. Oxford: Oxford University Press.
- Volmar, E. and Eisenhardt, K.M. (2025) ‘Mavericks and diplomats: Bridging commercial and institutional entrepreneurship for society’s grand challenges’, *Organization Science*, 36(2), pp. 572–600. [doi:10.1287/orsc.2020.13810](https://doi.org/10.1287/orsc.2020.13810).

11. About the Authors



John Finch (PhD, Lancaster University) is Professor of Marketing at the University of Glasgow's Adam Smith Business School. His research examines innovation systems, evidence use and policy implementation, with particular emphasis on how organisations translate analysis into actionable decisions in complex public-sector settings. He leads multi-partner, externally funded programmes combining interdisciplinary research with sustained engagement with government, industry and society, including on evidence-informed infrastructure investment and innovation policy. His publications include *Research Policy*, *Marketing Theory*, *Industrial Marketing Management*, and *Journal of Business Research*. He brings policy experience through a UK Government secondment on procurement-led innovation.



Get in touch



FRIL@FinTechScotland.com



University
of Glasgow



University of
Strathclyde
Glasgow