

**“48th RSEP International Conference on Economics, Finance and Business” 1-2 October
2026, University of Washington Rome Center, Rome, Italy**

Operational Resilience and Technology Risk Management in Regulated Financial Institutions

Jayanth M. Prasanna

Senior Software QA Analyst, Lloyds Bank, New York, USA

E-mail: jayanth.prasanna3@lbusa.com

DOI: <https://doi.org/10.19275/RSEPCONFERENCES403>

Abstract

Technology failure has matured into a material source of operational risk for banks and other regulated financial institutions. Payment outages, defective software releases, and disruption at third-party service providers now occupy the agendas of boards and prudential supervisors alongside credit and market exposure. This paper examines how a regulated financial institution can govern technology risk as an embedded component of its operational risk profile and what contribution software assurance mechanisms make to that task. The study rests on a structured review of 3 supervisory frameworks, the Basel principles for operational resilience, the European Digital Operational Resilience Act, and the operational resilience regime of the United Kingdom, combined with a synthesis of peer-reviewed empirical evidence on operational and cyber losses in the financial sector. The analysis yields a managerial framework in which 4 assurance mechanisms function as controls over technology risk. Risk-based test prioritization allocates scarce verification capacity toward exposures with the greatest regulatory and financial materiality. Performance verification confronts critical services with realistic demand and measures their behavior against disruption tolerances. Reusable automation sustains continuous control execution and preserves an audit trail. Release governance converts the decision to deploy a change into a formal act of risk acceptance by business, technology, and compliance stakeholders. The paper concludes that engineering assurance and prudential regulation converge on a single object, the capacity of an institution to deliver critical services under stress, and that operational resilience should be treated as an economic characteristic of the firm.

Keywords: operational resilience, technology risk, operational risk management, regulated financial institutions, release governance

Jel codes: G21, G28, G32, M15

1. Introduction

Contemporary banking operates as a software enterprise with a balance sheet attached. A retail payment that fails to settle, an onboarding queue frozen by a defective release, or a sanctions screening engine that silently drops transactions can, within hours, convert into customer harm, supervisory scrutiny, and financial loss. Loss event data confirm the scale of this exposure. Operational value-at-risk in banks ranges from 6% to 12% of total gross income, depending on the estimation method, and an average operational loss takes more than 1 year to be discovered and recognized in the books (Aldasoro et al., 2023). The same body of evidence shows that stronger supervision is associated with lower operational losses, suggesting that the quality of governance shapes the loss distribution itself.



The articles on the RSEP Conferences website are bear Creative Commons Licenses either CC BY or CC BY-NC-ND licenses that allow the articles to be immediately, freely, and permanently available on-line for everyone to read, download, and share.

The technology component of this exposure carries a systemic dimension. A model of the wholesale payments network of the United States estimates that the impairment of any 1 of the 5 most active banks would spill over to 38% of the network on average, and that liquidity hoarding by uncertain counterparties could push forgone payment activity above 2.5 times the daily gross domestic product (Eisenbach et al., 2022). A reverse stress test in the same study shows that outages originating in banks with less than 10 billion US dollars in assets are sufficient to impair a material share of the system. Technology failure within a single institution falls within the vocabulary of financial stability, and the private cost of an outage understates its social cost.

Prudential regulation has internalized this lesson. The Basel Committee on Banking Supervision issued principles for operational resilience in March 2021, which ask banks to map critical operations and to define a tolerance for disruption. The European Union adopted the Digital Operational Resilience Act, known as DORA, which entered into force on January 16, 2023, and applies from January 17, 2025, harmonizing ICT risk management, incident reporting, resilience testing, and oversight of third-party providers across the financial sector (Clausmeier, 2023). The United Kingdom requires firms to identify key business services, set impact tolerances, and demonstrate the capacity to remain within those tolerances under severe but plausible scenarios, with full compliance expected by March 2025. Questions that once lived inside test plans and deployment checklists now occupy board packs and supervisory dialogues.

Technology risk in this paper refers to the component of operational risk that arises from the design, change, and operation of information systems, as well as from services procured from technology providers. The definition places change at the center of the category. A banking platform in a steady state rarely generates new failure modes, whereas every release, migration, configuration update, and vendor substitution injects candidate defects into an environment whose complexity conceals them until activation. Managing technology risk centers on managing the flow of change, and the instruments that discipline change, verification, measurement, and gated approval, become the natural carriers of that management.

The economics of the subject extend beyond regulatory obligation. Disruption destroys revenue during the outage window, triggers remediation and compensation expenditure, invites enforcement action, and erodes the customer trust on which deposit franchises and payment volumes rest. An institution that can evidence its resilience gains a lower expected cost of disruption and a stronger standing in supervisory dialogue, both of which carry measurable value. Resilience, on this reading, behaves as an asset that management can build through deliberate investment in control capability.

A gap between 2 professional languages nevertheless persists. Engineering literature discusses testing, performance verification, and release control as matters of product quality, while prudential literature treats operational risk as a distribution of losses to be measured, capitalized, and reported. Few studies translate the engineering toolkit into the vocabulary of institutional risk management. This paper asks how a regulated financial institution can embed software assurance into its operational risk architecture so that test prioritization, performance engineering, and release governance function as documented controls over technology risk. The contribution is threefold. First, the paper assembles the empirical picture of operational and technology losses from recent peer-reviewed studies. Second, it compares 3 supervisory frameworks for operational resilience along a common set of dimensions. Third, it develops a managerial framework in which 4 assurance mechanisms operate as risk controls, and the release decision constitutes a formal act of risk acceptance.

The remainder of the paper proceeds as follows. Section 2 reviews 3 strands of literature and compares the regulatory frameworks. Section 3 describes the materials and the method. Section 4 presents the results in 2 steps: the evidence drawn from published data and the assurance-based control framework. Section 5 discusses the implications for practice and policy, as well as the limitations of the study. Section 6 concludes.

2. Literature Review

The banking operational risk literature has reached a critical mass, developing its own data sources, research methods, and discourses. A bibliometric analysis of 676 operational risk articles published between 2010 and March 2023 outlines the growth in research and the intellectual structure of the field using topical citation and coupling networks and identifies lagging areas of research (digitization, newer risk types, and governance) (Jadwani et al., 2024). Measurement issues, as this strand has highlighted, remain a central problem, with a lack of common language and meaning around technology-related loss events making comparisons across firms and jurisdictions difficult. Originally created as a data collection tool, a common space for industry cyber event recording is provided by the operational risk event taxonomy (Curti et al., 2023), complemented by loss event studies. Cross-country event-level evidence establishes the magnitude of operational value-at-risk, the multi-year lag between occurrence and appreciation, the sensitivity of these losses to the economic cycle, and the association between supervisory intensity and the containment of operational losses (Aldasoro et al., 2023). The

strand also records a shift in attention from capital measurement to management practice, a movement that this paper continues.

A second strand examines how technology reshapes the distribution of operational losses. Analysis of more than 100,000 cyber incidents across sectors shows that costs grow with firm size, that events touching several organizations at once carry a heavier price, and that the financial sector attracts more attacks than other industries yet suffers lower average losses thanks to proportionate investment in information technology security (Aldasoro et al., 2022). The same study finds that reliance on cloud services reduces the cost of smaller incidents while concentration among a handful of providers builds tail risk for the system as a whole. Firm-level evidence points in a compatible direction. A panel of Chinese banks observed between 2008 and 2019 links the adoption of cloud computing to higher profit efficiency and, at the same time, to greater operational risk, with synergy effects between cloud and other emerging technologies in the containment of that risk (Cheng et al., 2022). Network analysis adds a systemic layer, since third-party providers connect banks that share no direct financial relationship, thereby opening channels for the propagation of operational shocks across the system (Eisenbach et al., 2022).

A third strand concerns the regulation of digital operational resilience. DORA is a cross-sectoral regulation that is divided into five pillars: ICT risk management, incident classification and reporting, digital operational resilience testing, including threat-led penetration testing for firms in scope, management of ICT third-party risk, and voluntary information sharing (Clausmeier, 2023). The ability to manage convergence, coordination, and fragmentation in supervisory architecture for DORA has been analyzed, with the oversight of critical ICT third-party providers in 27 member states seen as the sharpest case (Buttigieg & Zimmermann, 2024). Cross mapping between DORA and ISO 27001 revealed 9 gaps across the 6 requirement areas. The weakest academic evidence for DORA was observed over threat-led penetration testing, major incident management, and third-party risk (Neumannová et al. 2023). Table 1 compares the 3 supervisory frameworks that anchor the present study with respect to their legal nature, scope, core requirements, testing expectations, and timeline.

Table 1. Regulatory frameworks for operational resilience in the financial sector.

Aspect	BCBS principles of 2021	EU DORA	UK regime
Legal nature	Supervisory principles without direct binding force	Regulation 2022/2554, binding in full across the European Union	Rules and supervisory statements of the PRA and the FCA
Scope	Internationally active banks	More than 20 categories of financial entities, together with critical ICT third-party providers	Banks, insurers, financial market infrastructures, and larger investment firms
Core requirements	Mapping of critical operations, tolerance for disruption, ICT, and cyber risk management	ICT risk management, incident classification and reporting, third-party risk, information sharing	Identification of important business services, impact tolerances, and mapping of dependencies
Resilience testing	Scenario-based testing of the ability to remain within tolerance	Digital operational resilience testing program, including threat-led penetration testing for selected entities	Scenario testing against impact tolerances under severe but plausible scenarios
Key dates	Published in March 2021	In force since January 16, 2023, applicable from January 17, 2025	In force since March 2022, full compliance is required by March 2025

Source: compiled by the author on the basis of Clausmeier (2022), Neumannová et al. (2023), and Buttigieg and Zimmermann (2024)

The comparison reveals a common architecture beneath divergent legal forms. Each framework asks an institution to know which of its services matter, to quantify how much disruption those services can absorb, and to prove through testing that the tolerance holds under stress. Each framework also extends the perimeter of responsibility beyond the corporate boundary to third-party providers, and each ties accountability to the firm's management body. The demand for assurance capability is therefore rising across jurisdictions at the same time, while the supply of that capability continues to depend on engineering practices whose risk content the frameworks presuppose without describing.

The taxonomic position of technology risk within these frameworks carries practical weight. Loss data providers and supervisors record technology failure under the operational risk umbrella alongside internal fraud, process error, and conduct events, and the classification template for cyber events extends this convention by anchoring each event type to the business process it disrupts (Curti et al., 2023). The convention has consequences for management. Once technology failure enters the operational risk register, it competes for capital, attention, and control investment on the same terms as every other loss category, and the case for any specific control must be argued in the currency of expected loss reduction. The framework developed in Section 4 accepts those terms.

A fourth strand supplies the methodological toolbox. A survey of operational research and artificial intelligence methods in banking covers applications ranging from credit scoring to operational risk measurement and outlines a research agenda in which data-driven techniques permeate every risk management function of the firm (Doumpos et al., 2023). What remains scarce at the intersection of these strands is work that connects the engineering assurance apparatus of a bank, its test portfolios, performance laboratories, automation estates, and release gates, with the prudential architecture of operational resilience at the level of the institution. The present paper addresses that gap by translating assurance mechanisms into the vocabulary of risk control and by anchoring the translation in the published empirical record.

3. Data & Methodology

The study follows an integrative review design with framework synthesis. It draws on 3 categories of material. The first category contains the supervisory frameworks themselves, the Basel principles of 2021, DORA, and the operational resilience regime of the United Kingdom, as read through the peer-reviewed legal and comparative analyses cited in Section 2. The second category contains published empirical research on operational and cyber losses in the financial sector, from which the study extracts stylized facts without collecting any new loss data. The third category contains generalized professional practice accumulated across multinational quality engineering and performance engineering programs delivered within large regulated banking groups over more than 2 decades of delivery across domains, including customer due diligence, financial crime compliance, payment messaging, and high-volume transaction processing. Confidentiality obligations govern this third input, so the paper states it at the level of principles and refrains from naming institutions, projects, or proprietary figures. The 3 frameworks were selected because they span the global standard setter, the largest supranational legal order, and a major national regime, which together cover the regulatory environment of most internationally active institutions.

The procedure comprises 3 steps. First, resilience requirements were extracted from the 3 frameworks along 5 dimensions: legal nature, scope, core requirements, testing expectations, and timeline, which produced the comparison in Table 1. Second, the empirical literature was distilled into stylized facts on the magnitude, latency, technology dependence, and systemic propagation of operational losses, with each fact traced to the study that documents it. Third, 4 assurance mechanisms drawn from professional practice were mapped onto the requirements and the facts, and the mapping was formalized in Figure 1 for the overall architecture, in Figure 2 for the release decision, and in Figure 3 for the full chain from loss evidence to the resilience outcome. The mapping proceeded in both directions, from regulatory demand toward the mechanism that satisfies it and from the mechanism toward the risk property that justifies its cost.

3 boundaries delimit the analysis. All quantitative statements are based on the cited publications, and no numbers in this paper describe a specific implementation within any single institution. The expert component is qualitative and serves as an interpretive frame, so the resulting framework carries the status of a structured synthesis open to empirical testing. Jurisdictional coverage focuses on the Basel, European, and United Kingdom regimes, with comparable regimes in other regions outside the scope of the comparison.

4. Results

4.1. Evidence from published data on operational and technology risk

The published record supports 3 stylized facts. The first fact concerns magnitude and latency. Operational value at risk between 6% and 12% of total gross income places technology and process failure in the same order of relevance as classical financial exposures, and the average lag of more than 1 year between the occurrence of a

loss and its recognition means that the causes of future losses already sit inside current systems (Aldasoro et al., 2023). For technology risk, this latency has a direct engineering interpretation. Defects embedded during change surface long after deployment, so controls that prevent defective change deliver more value than controls that detect damage after materialization. The inverse association between supervisory intensity and loss levels documented in the same study reinforces the reading that governance, together with the assurance practices through which governance operates, shapes outcomes.

The second fact concerns the technology channel itself. The financial sector absorbs more attacks than any other industry while paying less per event on average, a pattern the data attributes to sustained investment in information technology security (Aldasoro et al., 2022). Investment in verification and protection capability could be measured in the loss distribution. It varies depending on the cloud. Cloud use lowers the cost of low-impact events but concentrates tail risk with service providers. There is also bank-level evidence that cloud adoption has increased operational risk, even while improving operational efficiency (Cheng et al., 2022). Beyond the bank's own infrastructure, modern bank technology is highly integrated across multiple providers. There are many 'failure surfaces' in a bank's technology, and an increasing share of those are exposed to externalities from other providers. The third-party pillar of DORA reacts to this concentration with a set of mandatory contractual provisions, information registers and oversight of critical third-party service providers (Clausmeier, 2023).

The third fact is that if 1 of the 5 most important banks in the United States payment system failed, on average 38% of the system is affected. When liquidity hoarding takes place, the missing payments can be more than 2.5 times daily gross domestic product. Moreover, the failure of a bank with less than 10 billion US dollars of assets can also affect a material fraction of the system (Eisenbach et al., 2022). Disruption tolerances protect more than the institution that sets them. When a supervisor asks a bank to demonstrate that a payment service will remain within tolerance under a severe but plausible scenario, the request expresses the financial system's claim on the firm's internal quality controls.

Taken together, the 3 facts impose design requirements on any control system for technology risk. The magnitude of losses justifies dedicated investment. The latency of discovery demands controls positioned before deployment, where causes remain observable and inexpensive to correct. The technology channel requires coverage of third-party and integration surfaces in addition to code owned by the firm. The systemic dimension elevates the evidentiary standard, since a control whose output cannot be shown to a supervisor protects the institution without discharging its obligation to the system.

The regulatory response converts these facts into testing obligations. DORA requires a digital operational resilience testing program and threat-led penetration testing for selected entities (Clausmeier, 2023). The gap analysis against ISO 27001 shows that advanced testing, major incident management, and third-party risk are the areas where established information security practices offer the least coverage of the new requirements (Neumannová et al., 2023). Institutions that already operate mature test prioritization, performance verification, and release governance hold the raw material for compliance. What such institutions lack is a shared representation that positions those mechanisms inside the operational risk architecture and makes their contribution legible to risk committees and supervisors. Section 4.2 supplies that representation.

4.2. An assurance-based control framework for technology risk

The evidence assembled above motivates a reframing of software assurance. Under this reframing, testing and performance engineering acquire the status of controls over a risk category that regulators measure and price, and their outputs become risk evidence with an audience beyond the delivery team. Figure 1 depicts the resulting architecture. Regulatory frameworks define expectations from above; technology risk sits within the institution's operational risk profile; 4 assurance mechanisms act on that risk as controls; and the outcome is the demonstrated capacity to remain within impact tolerances.




Figure 1. Technology risk and assurance mechanisms within the operational resilience architecture

Source: compiled by the author

The first mechanism is risk-based prioritization of assurance effort. Verification capacity is scarce, so its allocation should be based on expected loss. Professional practice in regulated banking classifies requirements and business processes into risk tiers using criteria of materiality, the presence of a regulatory obligation such as customer due diligence, anti-money laundering controls, sanctions screening, or payment messaging, the financial exposure of the transaction flow, the breadth of customer impact, the complexity of integrations with external platforms, and the history of production incidents and audit findings. Components in the highest tier receive earlier test cycles, deeper negative testing, and wider stakeholder review, while components in lower tiers receive coverage proportional to their contribution to the risk profile. In economic terms, the mechanism maximizes the reduction in expected disruption cost per unit of verification effort, mirroring the proportionality logic embedded in the supervisory frameworks in Table 1.

The second mechanism is performance verification against disruption tolerances. Load models constructed from business volumes, peak usage patterns, transaction mixes, concurrency assumptions, and service level objectives confront critical services with realistic and adverse demand before customers do. When a payment or onboarding platform sustains projected peak volumes within its response and error thresholds, the test result constitutes documented evidence that the institution can remain within the impact tolerance of that service, and the margin between demonstrated capacity and projected demand quantifies the platform's resilience headroom. Root cause discipline strengthens the mechanism, since bottlenecks cluster at identifiable layers, application code, database access, infrastructure capacity, and integration paths, and remediation aimed at causes outlasts remediation aimed at symptoms. Verification of behavior under degraded network conditions and service dependency failures extends the same logic to severe but plausible scenarios that regulatory frameworks address.

The third mechanism is reusable automation as a continuous control. A layered framework that separates test logic, business rules, object repositories, test data, environment configuration, and reporting makes each automated validation an asset that survives application changes and travels across releases. Automation of compliance-critical regression testing converts verification from a periodic exercise into a standing control and produces the timestamped evidence trail that auditors and supervisors expect from any control that carries prudential weight. Selection of automation candidates follows business criticality, execution frequency, regression exposure, and return on investment, which keeps the automation estate aligned with the risk tiers of the first mechanism and prevents the accumulation of maintenance burden around low-value scenarios.

The fourth mechanism is release governance, the point at which technology risk becomes a formal decision of the institution. Figure 2 details this component.




Figure 2. Release governance as a formal risk acceptance procedure

Source: compiled by the author

The Go/No-Go gate aggregates 4 categories of release readiness evidence. Quality metrics cover test coverage, pass rates, requirement traceability, and defect leakage trends. Compliance validation confirms the behavior of customer due diligence, anti-money laundering, sanctions screening, and payment workflows. Performance verification demonstrates the capacity to handle expected volumes. Operational readiness certifies monitoring, rollback procedures, and disaster recovery arrangements. A Go decision documents the acceptance of residual risk by named business, technology, and compliance owners, while a No-Go decision returns the change to remediation and expresses the institution's risk appetite in operational terms. Grounds for a No-Go include unresolved compliance defects, material performance shortfalls, incomplete coverage of critical flows, and gaps in operational readiness, each of which corresponds to a documented loss category in the empirical record. Since defects introduced through change dominate many operational loss narratives, the release gate marks the last moment at which technology risk remains governable before it is exposed to production.

The 4 mechanisms operate as a system, mutually reinforcing one another. Prioritization decides where the performance and automation investment lands; automation feeds the release gate with current evidence at negligible marginal cost; performance results calibrate the tolerance margins that the gate enforces; and the record of Go/No-Go outcomes flows back into the incident and audit history on which the next prioritization cycle draws. The loop confers a learning property on the institution. Every release, whatever its outcome, improves the information set that governs the following release, and the quality of technology risk management compounds across cycles.

Figure 3 consolidates the argument of Sections 4.1 and 4.2 and anticipates the practical implications developed in Section 5. It traces the chain from the published loss evidence to the design requirements that the evidence imposes, through the closed loop of the 4 mechanisms with their respective cost structures, to the 3 evidence artefacts that the loop produces, the audiences that consume those artefacts, and the resilience outcome in which the chain terminates.




Figure 3. From loss evidence to resilience outcome: the assurance control loop, its evidence artefacts, and their consumers

Source: compiled by the author

The 4 mechanisms map onto the supervisory architecture without translation loss. Risk-based prioritization operationalizes the proportionality principle. Performance verification provides the basis for scenario testing against tolerances and complements the advanced testing agenda of DORA. Reusable automation delivers the continuous control execution and the audit evidence that examinations require. Release governance provides management bodies with the accountable decision-making point that all 3 frameworks presuppose. Assurance, in

this reading, is the machinery through which an institution manufactures its own resilience, and the maturity of that machinery is a leading indicator of the operational risk profile.

4.3. Correspondence between mechanisms and regulatory expectations

The correspondence between the 4 mechanisms and the supervisory frameworks of Table 1 can be stated with precision. Risk-based prioritization addresses the demand of all 3 frameworks for institutions to identify their critical operations and important business services, because the tiering and service identification exercises use the same inputs: regulatory obligations, customer reach, financial exposure, and interdependence. An institution that maintains current risk tiers for its application estate holds a mapping of critical operations as a byproduct of routine work.

Performance verification addresses the need for scenario testing within tolerances. The Basel principles and the United Kingdom regime ask for demonstration under severe but plausible scenarios, and a governed load test that pushes a critical service beyond projected peak demand delivers such a demonstration in executable form. Under DORA, the same capability supports the digital operational resilience testing program, and the discipline of environment realism, production-like data volumes, representative integrations, and adverse network conditions, determines whether the demonstration convinces (Clausmeier, 2023).

Reusable automation answers the demand for continuous control and for evidence. The gap analysis against ISO 27001 locates the weakest coverage of DORA requirements in advanced testing, incident management, and third-party risk (Neumannová et al., 2023). An automation estate that exercises compliance-critical flows on every release provides recurring verification that closes part of this gap while generating the artifacts that examinations consume.

Release governance answers the demand for accountability of the management body. Each framework in Table 1 ties resilience obligations to named executive responsibility, and the Go/No-Go register gives that responsibility a documentary form, a sequence of dated decisions with attached evidence and named owners. The register converts an abstract accountability into an auditable trail that survives staff rotation and organizational change.

5. Discussion

The central finding is a convergence. Prudential regulation and engineering assurance address the same object from opposite directions: the capacity of an institution to deliver critical services under stress. The stylized facts of Section 4.1 explain why the convergence carries economic weight. Losses of the documented magnitude, discovered with the documented latency and propagated through the documented network channels, make prevention more valuable than detection after the fact, and the 4 mechanisms of Section 4.2 constitute the prevention apparatus that a regulated institution already funds. Recognition of that apparatus as a portfolio of operational risk controls raises its visibility, its budget rationale, and its governance standing at once, and aligns the incentives of delivery organizations with those of risk functions that have historically focused on loss data.

The convergence also resolves an incentive problem inside the firm. Delivery organizations are held accountable for speed, while risk functions are held accountable for losses, and the 2 mandates collide at every release date. A shared representation of assurance as risk control gives both sides a common objective function, the expected cost of disruption, against which acceleration and caution can be traded in explicit terms. The Go/No-Go record then serves as the ledger for those trades, converting a recurring negotiation into a governed process with organizational memory.

A cost lens sharpens the practical argument. The 4 mechanisms have different cost structures: prioritization is close to free once the risk information exists; performance verification requires environments and tooling; automation requires upfront construction and ongoing maintenance; and governance requires senior time at every release. Against these costs stand the avoided losses whose scale the empirical record documents, and the comparison favors the mechanisms wherever the value at risk of the protected service is material. The allocation question dominates the volume question, and the loss distribution supplies the allocation rule.

For practice, the framework implies integrating reporting lines. Quality indicators such as defect leakage, coverage of compliance-critical flows, and performance headroom against tolerances belong in operational risk dashboards alongside loss events and established key risk indicators, because these quality indicators drive the losses the dashboards record, with the latency documented in the loss event literature. Change advisory boards gain the character of risk committees once the Go/No-Go record is understood as a register of risk acceptance decisions, and the register itself becomes an audit object with evidentiary value in supervisory reviews. Contracts with third-party providers should carry assurance clauses that extend the 4 mechanisms across the corporate boundary, since the evidence on provider concentration shows that the perimeter of technology risk exceeds the perimeter of ownership.

For policy, the analysis suggests that supervisors can request release governance artifacts, prioritization rationales, and performance test results as resilience evidence with little additional burden on well-managed firms, since these artifacts are already produced in the ordinary course of delivery. The concerns about supervisory convergence and fragmentation raised in the legal literature apply here with force, as heterogeneous evidence standards across 27 national authorities would erode the comparability that the harmonized frameworks aim to achieve (Buttigieg & Zimmermann, 2024). Methodological progress offers a complementary path. The spread of data-driven techniques across banking risk functions opens the possibility of predictive quality analytics, in which assurance telemetry, defect histories, and incident data feed models of technology risk, sharpening both the prioritization and the release decision (Douplos et al., 2023).

4 limitations qualify the results. The framework rests on a synthesis of published evidence and generalized professional practice, so its causal claims await empirical validation. The quantitative statements describe the datasets of the cited studies and may not transfer to every jurisdiction or size class. The expert component reflects experience in large international banking groups, and smaller institutions may face different constraint sets and different economic conditions for automation. The regulatory comparison freezes a moving target, since implementing measures under DORA continue to evolve, and the supervisory practice around them is young. Future research can test the association between assurance maturity and the frequency and severity of technology incidents on a cross-bank panel, extend the mapping to insurance undertakings and market infrastructures, and quantify the cost of the 4 mechanisms against the expected loss they avert, a direction consistent with the research agenda identified in the bibliometric record (Jadwani et al., 2024).

6. Conclusion

This paper set out to connect 2 bodies of knowledge that describe the same institutional capability in different languages. The review of 3 supervisory frameworks showed that regulators across jurisdictions now expect a bank to know its critical services, quantify the disruption those services can absorb, and demonstrate through testing that the quantification holds. The synthesis of published loss evidence showed why the expectation is justified, since operational losses are large, slow to surface, sensitive to technology choices, and capable of propagating through payment networks and third-party dependencies beyond the originating firm. The comparison further showed that the 3 regimes, despite their differences in legal form and scope, converge on mapping, tolerance setting, and testing as the operative verbs of resilience.

The answer developed here places 4 assurance mechanisms at the center of technology risk management. Risk-based prioritization directs scarce verification capacity toward material exposures. Performance verification measures critical services against disruption tolerances. Reusable automation maintains continuous control execution and audit evidence. Release governance converts deployment into an accountable act of risk acceptance by the institution. Together, these mechanisms make operational resilience a manufactured property of the firm, one with consequences for regulatory standing, customer trust, and the cost of disruption. Operational resilience, on this account, is an economic characteristic of a financial institution, and the engineering practices that produce it deserve a permanent seat in the firm's risk management architecture.

3 recommendations follow for institutions preparing for the current regulatory cycle. Establish a single register that links application components to risk tiers and to the important business services they support. Route every material release through a governed gate whose evidence categories cover quality, compliance, performance, and operational readiness. Report the resulting indicators into the operational risk function on the same cadence as loss data. None of these steps requires new science, and together they convert an existing engineering practice into a demonstrable resilience capability.

References

- Aldasoro, I., Gambacorta, L., Giudici, P., & Leach, T. (2022). The drivers of cyber risk. *Journal of Financial Stability*, 60, Article 100989. <https://doi.org/10.1016/j.jfs.2022.100989>
- Aldasoro, I., Gambacorta, L., Giudici, P., & Leach, T. (2023). Operational and cyber risks in the financial sector. *International Journal of Central Banking*, 19(5), 340–402.
- Buttigieg, C. P., & Zimmermann, B. B. (2024). The digital operational resilience act: Challenges and some reflections on the adequacy of Europe's architecture for financial supervision. *ERA Forum*, 25, 11–28. <https://doi.org/10.1007/s12027-024-00793-w>
- Cheng, M., Qu, Y., Jiang, C., & Zhao, C. (2022). Is cloud computing the digital solution to the future of banking? *Journal of Financial Stability*, 63, Article 101073. <https://doi.org/10.1016/j.jfs.2022.101073>

- Clausmeier, D. (2023). Regulation of the European Parliament and the Council on digital operational resilience for the financial sector (DORA). *International Cybersecurity Law Review*, 4, 79–90. <https://doi.org/10.1365/s43439-022-00076-5>
- Curti, F., Gerlach, J., Kazinnik, S., Lee, M., & Mihov, A. (2023). Cyber risk definition and classification for financial risk management. *Journal of Operational Risk*, 18(2), 37–58. <https://doi.org/10.21314/JOP.2022.036>
- Doumpos, M., Zopounidis, C., Gounopoulos, D., Platanakis, E., & Zhang, W. (2023). Operational research and artificial intelligence methods in banking. *European Journal of Operational Research*, 306(1), 1–16. <https://doi.org/10.1016/j.ejor.2022.04.027>
- Eisenbach, T. M., Kovner, A., & Lee, M. J. (2022). Cyber risk and the U.S. financial system: A pre-mortem analysis. *Journal of Financial Economics*, 145(3), 802–826. <https://doi.org/10.1016/j.jfineco.2021.10.007>
- Jadwani, B., Parkhi, S., & Mitra, P. K. (2024). Operational risk management in banks: A bibliometric analysis and opportunities for future research. *Journal of Risk and Financial Management*, 17(3), Article 95. <https://doi.org/10.3390/jrfm17030095>
- Neumannová, A., Bernroider, E. W. N., & Elshuber, C. (2023). The Digital Operational Resilience Act for financial services: A comparative gap analysis and literature review. In M. Papadaki, P. Rupino da Cunha, M. Themistocleous, & K. Christodoulou (Eds.), *Information systems: EMCIS 2022 (Lecture Notes in Business Information Processing, Vol. 464, pp. 570–585)*. Springer. https://doi.org/10.1007/978-3-031-30694-5_40.
- Dunning, K. H., 2020. Building resilience to natural hazards through coastal governance: a case study of Hurricane Harvey recovery in Gulf of Mexico communities. *Ecological Economics*, 176, 106759, ISSN 0921-8009. <https://doi.org/10.1016/j.ecolecon.2020.106759>
- Google Scholar. Available from: <https://scholar.google.com> (accessed on 2026 Sep 25).
- Head Office of Geodesy and Cartography (GUGiK). Home page. Data. National Register of Boundaries. Available from: <https://www.geoportal.gov.pl/en/data/national-register-of-boundaries/> (accessed on 2026 Aug 02).
- Local Data Bank Statistics Poland. Available from: <https://bdl.stat.gov.pl/bdl/start> (accessed on 2026 Sep 25)
- Marto, R., Papageorgiou, C., & Klyuev, V. (2018). Building resilience to natural disasters: An application to small developing states. *Journal of Development Economics*, 135, 574–586. <https://doi.org/10.1016/j.jdeveco.2018.08.00>
- National Bank of Poland. Home page/Statistics/Rates/ Table no. 251/A/NBP/2025 of 2025-12-31. Available from: <https://nbp.pl/en/archiwum-kursow/table-no-251-a-nbp-2025-of-2025-12-31/> (accessed on 2026 Sep 29).
- Pieniążek, M., Szejgiec, B., Zych, M., Ajdyn, A. & Nowakowska, G. (2014). Graficzna prezentacja danych statystycznych. Wykresy, mapy, GIS. Warszawa: Główny Urząd Statystyczny. [in Polish]
- Piepiora, P. A., Piepiora, Z. N., Stackeová, D. and Bagińska, J. (2023). Editorial: Physical culture from an interdisciplinary perspective. *Frontiers in Psychology*, 14, 1254027. <https://doi.org/10.3389/fpsyg.2023.1254027>
- Piepiora, Z., Bagińska, J. & Piepiora, P., 2026, July. Financing the removal of the effects of natural disasters in physical culture facilities by local government units in Poland. In R. M. Kalina & J. Kalra (Eds.), *Human Factors in Sports, Preventive Medicine and Innovative Agonology. AHFE (2026) International Conference. AHFE Open Access (Vol. 226, pp. 32-43)*. New York: AHFE International. <https://doi.org/10.54941/ahfe1007820>
- Piepiora Z. N., 2012. Ekonomiczne aspekty lokalnej polityki przeciwdziałania skutkom katastrof naturalnych (Economic Aspects of Local Policy for Mitigating the Effects of Natural Disasters). Kowary. [in Polish]
- Piepiora Z. N, 2019. Ekonomiczna efektywność strukturalnej ochrony przeciwpowodziowej w Polsce (The Economic Efficiency of Structural Flood Protection in Poland). Wrocław: Wydawnictwo Uniwersytetu Przyrodniczego. [in Polish]
- QGIS. Available from: <https://qgis.org> (accessed on 2025 Aug 13).
- Robinson, L. & Jarvie, J. K. (2008), Post-disaster community tourism recovery: the tsunami and Arugam Bay, Sri Lanka. *Disasters*, 32, 631–645. <https://doi.org/10.1111/j.1467-7717.2008.01058.x>
- Scopus. Available from: <https://www.scopus.com> (accessed on 2026 Sep 25).
- Sovacool, B. K., Tan-Mullins, M. & Abrahamse, W. 2018. Bloated bodies and broken bricks: Power, ecology, and inequality in the political economy of natural disaster recovery. *World Development*, 110, 243–255. <https://doi.org/10.1016/j.worlddev.2018.05.028>
- Statistics Poland. Yearly price indices of consumer goods and services from 1950. Available from: <https://stat.gov.pl/en/topics/prices-trade/price-indices/price-indices-of-consumer-goods-and-services/yearly-price-indices-of-consumer-goods-and-services-from-1950/> (accessed on 2026 Oct 06)
- Udeorah, S. F., Amaefule, C., & Omolehinwa, L. C. (2026). Tourism Financing and Natural Disasters: A Case of Desertification Threats in Nigeria. In S. Udeorah, H. Okoro, G. Thinnakkakath, & A. Ogbo (Eds.), *Natural Disasters*,

the Environment, and Technological Pathways to Sustainable Recovery (pp. 403-432). Hershey, PA: IGI Global Scientific Publishing. <https://doi.org/10.4018/979-8-3373-2680-1.ch015>

UN Tourism. Available from: <https://www.untourism.int/glossary-tourism-term> (accessed on 2026 Oct 05)