

SERRAQ HQ, OPERATING UNDER THE WORKING NAME RISK ACCOUNTING STANDARDS BOARD (RASB), STEVE BAILEY

Proposal and Comment Information

Title: Regulatory Capital Rule: Category I and II Banking Organizations, Banking Organizations with Significant Trading Activity, and Optional Adoption for Other Banking Organizations, R-1887

Comment ID: FR-2026-0007-01-C20

Submitter Information

Organization Name: SERRAQ HQ, operating under the working name Risk Accounting Standards Board (RASB)

Organization Type: Organization

Name: Steve Bailey

Submitted Date: 06/12/2026

June 12, 2026

Via: Federal Reserve website

Docket/RIN: Board Docket No. R-1887; RIN 7100-

AH20 Comment due date: June 18, 2026

Subject: Regulatory Capital Rule: “Category I and II Banking Organizations, Banking Organizations with Significant Trading Activity, and Optional Adoption for Other Banking Organizations”

Comment on Operational Risk Measurement, Supervisory Evidence, and Forward-Looking Risk Sensitivity under the Proposed Expanded Risk-Based Approach

Questions addressed: Questions 5, 84, 86-92, and 195-199,
with principal emphasis on Questions 5, 84, 89, 92, 195, and 199.

Submitting organisation: SERRAQ HQ, operating under the working name Risk Accounting Standards Board ([RASB](#))

Contact: Steve Bailey, Chairman, steve.bailey@rasb.org, +44 7909 991086,
Bishopstone, 36 Crescent Road, Worthing, West Sussex, England, BN11 1RL, rasb.org.

Who We Are

The Risk Accounting Standards Board (RASB) is a standards-oriented body concerned with the disciplined development, review and application of methods for evidencing, measuring and reporting non-financial and operational risk. RASB's interest in this proposal is limited to operational-risk measurement, supervisory evidence, and the question of how forward-looking operational information can be made more standardised, reviewable and empirically testable.

Executive Summary

This comment addresses the operational-risk aspects of the agencies' proposal to modernise the regulatory capital framework for Category I and II banking organisations and other covered banking organisations under the proposed expanded risk-based approach. It focuses principally on Questions 5, 84, 89, 92, 195 and 199, with supporting observations relevant to Questions 87 and 198.¹

We support the agencies' proposal to remove bank-owned internal models for operational risk as determinants of minimum regulatory capital requirements. The agencies' concerns are well founded: operational-risk tail modelling is vulnerable to limited data, model and parameter uncertainty, judgement-sensitive assumptions, opacity, inconsistent implementation across firms and potential model-optimisation incentives that may not align with supervisory comparability.²

Removing internal operational-risk capital models should not be understood as requiring the supervisory framework to rely only on realised losses, income measures, business indicators or asset-based proxies. Realised operational loss data are valuable because they provide objective evidence of outcomes and have demonstrated informational value for future operational loss experience. Business indicator, income and asset measures are useful scale and activity proxies, but they are not designed to measure the current condition of a firm's operating environment.^{3 4}

The agencies' own research provides the starting point. Federal Reserve staff analysis has found that historical operational losses remain informative about future operational losses even after controlling for a wide range of financial characteristics, likely because past losses proxy for conditions that are otherwise hard to observe, such as control quality, risk culture and risk appetite. If that is so, the natural next question is empirical: whether those underlying conditions can be observed more directly, in a standardised and evidence-backed way, rather than only inferred after losses have crystallised. This comment recommends that the agencies test that question, while preserving the comparability and simplicity of the proposed standardised framework.

We therefore recommend that the agencies consider a distinct, standardised supervisory evidence layer based on operational-condition indicators: a limited set of standardised, evidence-backed information fields describing observable features of a banking organisation's operational environment, such as control testing results, unresolved audit or assurance findings, remediation ageing, incident and near-miss classifications, operational resilience test results, material technology or third-party concentration indicators, and documented review, challenge, approval and versioning. These indicators would not be bank-specific capital models, would not generate capital requirements through an internal formula, and would not provide automatic capital relief.

The central recommendation is a bounded pilot or supervisory data study. Such a study would test whether standardised operational-condition information adds explanatory or supervisory value beyond business indicator measures, historical losses and other scale proxies. It would also allow the agencies to assess feasibility, burden, comparability, susceptibility to model optimisation or reporting incentives, predictive value and potential usefulness for supervisory analysis before considering any broader reporting or calibration implications.

¹ NPR, pp. 30, 198-199, 208, 214-216, 529, 659-660.

² NPR, pp. 25-29.

³ Filippo Curti and Marco Migueis, The Information Value of Past Losses in Operational Risk, Finance and Economics Discussion Series 2023-003, Board of Governors of the Federal Reserve System, 2023, abstract and pp. 1-2, 16.

⁴ NPR, pp. 197-199.

I. Introduction and Scope of Comment

This comment is submitted in response to the agencies' notice of proposed rulemaking on the proposed expanded risk-based approach. It addresses the operational-risk elements of the proposal and, in particular, the agencies' questions concerning the removal of internal models, the use of realised operational losses, alternative standardised approaches, calibration, impact analysis, related reporting and the possible need for additional data. The proposal identifies risk sensitivity, consistency, transparency and simplification as central objectives of the revised capital framework.⁵

The comment is directed principally to Questions 5, 84, 89, 92, 195 and 199. It also includes supporting observations relevant to Questions 87 and 198, and limited observations relevant to Questions 86, 88, 90, 91, 196 and 197 where those questions bear on operational-risk measurement and supervisory information. It does not address Question 85 except to the extent that accounting classification issues may indirectly affect operational-risk exposure.⁶

The core position is deliberately limited. We do not recommend restoring the advanced measurement approaches or permitting banking organisations to use internal operational-risk models to determine minimum regulatory capital requirements. We also do not recommend that any new operational-risk indicator should directly reduce operational-risk capital, override the proposed business indicator framework, or substitute for realised operational loss data.

Instead, we recommend that the agencies consider whether a narrow, standardised supervisory evidence layer could complement the proposed framework. Such a layer would be designed to provide forward-looking information about operational condition without reintroducing bank-specific capital modelling. It would support supervisory insight, benchmarking, impact analysis, validation and future empirical study.

In this comment, "operational-condition indicators" means standardised information about the current state of material operational processes, controls, remediation, resilience, technology dependency and related governance evidence. Examples include control testing or assurance results, unresolved audit or compliance findings, remediation status, incident and near-miss classifications, operational resilience test outcomes, material third-party or technology concentration indicators, and evidence of independent review, challenge, approval, audit trail and version history.

These indicators are not intended to produce a capital number. They are intended to make relevant operational evidence more comparable, reviewable and capable of validation against subsequent outcomes.

It is also important to distinguish this capital NPR from the Board's separate supervisory stress-test operational-risk model and related stress-test transparency proposals. This comment refers to stress testing only insofar as operational-condition evidence may, in the future, assist supervisory benchmarking, impact analysis or model validation. It does not propose replacing the Board's operational-risk stress-test model, nor does it recommend immediate use of operational-condition indicators in stress capital buffer calculations. The Board's current operational-risk stress-test documentation describes a separate distributional modelling framework for projecting operational losses under stress; that framework should remain analytically distinct from the minimum capital design questions addressed in this NPR.^{7 8}

II. Removal of Internal Operational-Risk Models as Minimum Capital Determinants

Questions addressed: Q5 prominently; Q89 and Q92 indirectly.

⁵ NPR, pp. 1, 13, 16-18.

⁶ NPR, pp. 30, 198-199, 208, 214-216, 529, 659-660.

⁷ Board of Governors of the Federal Reserve System, Supervisory Stress Test Model Documentation: Operational Risk Model, October 2025, pp. 4-9.

⁸ Board of Governors of the Federal Reserve System, Enhanced Transparency and Public Accountability of the Supervisory Stress Test Models and Scenarios; Modifications to the Capital Planning and Stress Capital Buffer Requirement Rule, Enhanced Prudential Standards Rule, and Regulation LL, notice of proposed rulemaking, 2025, pp. 1, 5, 14.

We support the agencies' proposal to remove internal models for operational risk as determinants of minimum regulatory capital requirements. This is an appropriate response to the experience of the advanced approaches and to the particular difficulties of operational-risk tail modelling.⁹

Operational risk differs from market risk in important respects. Market-risk models can often be tested against frequent price observations and daily feedback from traded positions. Operational-risk models, by contrast, frequently rely on sparse severe-loss observations, heterogeneous event histories, judgement-sensitive scenario assumptions and difficult-to-verify distributional choices. The agencies' concern that such models can produce inconsistent minimum capital outcomes across firms with similar risk profiles is therefore justified. The proposal recognises that the advanced approaches have limitations as a means of setting consistent minimum risk-based capital requirements, even while internal models may remain useful for internal risk management, stress testing and planning.¹⁰

The removal of internal operational-risk capital models has several advantages.

First, it should improve comparability. Minimum regulatory capital requirements should not depend materially on firm-specific modelling choices that supervisors and market participants cannot readily compare across institutions. This is especially important for operational risk, where assumptions about frequency, severity, dependence, scenario design, boundary events and tail behaviour can have material effects on model output.

Second, it should improve transparency. A standardised approach allows supervisors, firms and external stakeholders to understand more clearly the main drivers of minimum operational-risk capital requirements. This does not eliminate judgement from the broader supervisory process, but it reduces the role of bank-specific modelling judgement in the calculation of minimum capital.

Third, it should reduce incentives for model optimisation. Even where a firm has strong internal governance, any capital framework that allows internal modelling choices to determine minimum requirements creates a structural risk that modelling approaches will be evaluated partly by their capital effects. Removing internal models from the minimum requirement calculation mitigates this concern.

Fourth, it should simplify the framework. Complexity may be justified where it produces demonstrably better risk measurement and can be validated. For operational-risk tail modelling, the combination of limited data, low-frequency severe events and judgement-sensitive assumptions makes that standard difficult to satisfy.

These advantages should be acknowledged without overstating the case. Internal operational-risk models, scenario analysis, internal stress testing and operational-risk analytics may still provide useful information for a banking organisation's own risk management and capital planning. The issue is not whether firms should understand and model operational risk internally. They should. The issue is whether such models should determine minimum regulatory capital requirements. For that purpose, the agencies' concerns about opacity, inconsistency, subjectivity and gaming are sufficiently strong to support removal.

However, the conclusion that internal capital models should not determine minimum regulatory capital requirements does not imply that all operationally derived information should be excluded from supervisory use. There is an important distinction between, on the one hand, allowing each firm to estimate its own operational-risk capital requirement through a proprietary tail model, and, on the other hand, collecting standardised evidence about operational condition for supervisory review, benchmarking, validation and empirical study.

That distinction should be preserved in the final rule and in related supervisory development. A standardised supervisory evidence layer would not ask firms to select loss distributions, calibrate tail percentiles, estimate firm-specific capital offsets, or translate internal assessments directly into capital relief. Instead, it would ask whether specified operational facts and evidence can be collected in a consistent way across firms and tested against subsequent outcomes.

Such information may be useful precisely because the proposed minimum capital framework will necessarily rely on proxies. Business indicator, income and asset measures can support simplicity and comparability, but they do not directly measure current control condition, remediation discipline,

⁹ NPR, pp. 18-20, 25-30.

¹⁰ NPR, pp. 20, 25-30.

operational resilience, technology dependency or outsourcing concentration. Historical operational losses are empirically important, but they reflect outcomes that have already occurred. A standardised supervisory evidence layer could help supervisors observe whether current operational conditions appear to be improving, deteriorating or diverging from what loss history and scale proxies alone would suggest.

The main risk in this proposal is that it could be misunderstood as an attempt to reintroduce the advanced measurement approaches under another name. That is not the recommendation. Any such standardised supervisory evidence layer should be subject to clear limitations:

- it should not determine minimum regulatory capital requirements through a bank-owned model;
- it should not provide automatic capital relief;
- it should not replace realised operational loss data;
- it should not rely on unconstrained firm-specific taxonomies;
- it should not produce falsely precise estimates of operational-risk exposure;
- it should be validated against subsequent outcomes before any broader supervisory or calibration use is considered.

The agencies' concerns about subjectivity should therefore shape the design of any future data work. Operational-condition indicators would inevitably include some judgement-based elements. The relevant question is not whether judgement can be eliminated. It cannot. The question is whether judgement can be bounded, evidenced, challenged, versioned, benchmarked and validated over time. If those safeguards cannot be met, the information should not be used for any consequential regulatory purpose. If they can be met, the information may provide useful supervisory insight without undermining the agencies' decision to remove internal models from minimum capital determination.

III. Operational-Risk Measurement: Standardisation, Risk Sensitivity and Forward-Looking Evidence

Questions addressed: Q84, Q89 and Q92 prominently; Q86, Q87, Q88, Q90 and Q91 where relevant.

The agencies' proposed standardised approach for operational risk represents a defensible response to the limitations of internal operational-risk capital models. A standardised framework based on a business indicator has important advantages over bank-specific modelled estimates: it is more transparent, more comparable across banking organisations, less dependent on firm-specific modelling judgement, and more administrable for purposes of minimum regulatory capital requirements.¹¹

The business indicator also improves on simpler gross-income or asset-based approaches by seeking to capture a broader range of banking activities through defined components of financial activity. The proposed business indicator component also increases at higher marginal rates as the business indicator rises, reflecting the agencies' view that operational risk generally increases more than proportionately with business volume and complexity.¹²

For these reasons, we do not recommend replacing the proposed standardised operational-risk approach with a bank-owned model or with a purely judgement-based assessment of operational condition. Nor do we recommend that forward-looking operational indicators be used at this stage to provide automatic capital relief. The proper question is narrower: whether the agencies' standardised approach can be complemented by operational evidence that improves supervisory insight, benchmarking, validation and future empirical study without undermining the simplicity and comparability of minimum regulatory capital requirements.

¹¹ NPR, pp. 197-199.

¹² NPR, pp. 209-211.

A. The residual forward-looking gap in business indicator, income, asset and realised-loss measures

The principal limitation of the business indicator is not that it is inappropriate, but that it is a proxy. It measures scale, activity and broad business volume. It does not directly measure the current condition of a banking organisation's operating environment.

This limitation is also present in the alternative approaches discussed by the agencies. A gross-income approach would be simpler and more transparent but would remain a scale-based measure. An asset-based approach would also be administrable and stable but may not appropriately capture operational risks associated with activities that have a small on-balance-sheet footprint and few quantifiable off-balance-sheet exposures. This is particularly relevant for activities such as payments, custody, asset servicing, technology-enabled delivery, outsourcing-intensive operating models, shared services, and other transaction-intensive functions where operational exposure may not be well represented by assets alone.¹³

The same point applies to noninterest income and expense measures. These may be useful indicators of activity, but they do not distinguish between a high-volume activity supported by stable, well-tested processes and the same activity supported by ageing technology, unresolved audit issues, weak remediation discipline, third-party concentration, or repeated control failures. Likewise, exclusions for recharge income or non-financial affiliate activity should not depend solely on accounting label. Some non-financial activities may be low-risk pass-throughs; others, such as technology, processing, human resources, infrastructure or shared-service functions, may carry material operational exposure.

For example, two high-volume payment processing platforms may present similar noninterest income profiles and transaction intensity, yet materially different operational conditions. One may operate on ageing core infrastructure with unresolved critical audit findings, manual reconciliation dependencies and delayed remediation; the other may operate on a more resilient, automated control framework with recent assurance testing and timely issue closure. A business indicator or income-based proxy may treat these profiles similarly, while a standardised evidence layer could allow supervisors to observe and test whether the operational conditions are materially different.

Realised operational losses occupy a different position. They are not merely proxies for business volume. They are objective evidence that operational-risk outcomes have crystallised. Federal Reserve staff analysis has found that past operational losses are informative of future operational loss experience even after controlling for a wide range of financial characteristics, and proposes that this informational value may arise because past losses capture hard-to-quantify factors such as operational-risk control quality, risk culture and risk appetite.¹⁴ The same analysis finds that loss frequency is generally more predictive of future operational-risk outcomes than loss severity, likely because frequency is a more stable metric.¹⁵

This evidence supports the continued use of realised loss data in operational-risk analysis. However, it does not eliminate the forward-looking gap. Historical losses may be informative in part because they proxy for operational conditions that are otherwise difficult to observe.¹⁶ They are valuable evidence of realised outcomes, but they remain lagging indicators. They may not capture current remediation, deteriorating control conditions, recent technology migration, new third-party dependencies, business model change, or improvements following past control failures.

The agencies should therefore avoid two extremes. Historical losses should not be treated as irrelevant merely because they are backward-looking. Equally, historical losses should not be treated as the only defensible source of information about operational-risk condition. A balanced framework would preserve realised losses as essential calibration and validation evidence while studying whether standardised operational-condition indicators can provide incremental, forward-looking supervisory information.

¹³ NPR, pp. 210-215.

¹⁴ Curti and Migueis 2023, abstract and pp. 1-2, 16.

¹⁵ Curti and Migueis 2023, abstract and pp. 1-2, 16.

¹⁶ Curti and Migueis 2023, abstract and pp. 1-2, 16.

B. Historical loss data should remain essential for calibration and validation

With respect to Q84 and Q92, historical loss data should remain central to operational-risk calibration and validation, whether or not the agencies include an internal loss multiplier in the calculation of operational risk-weighted assets.¹⁷

The internal loss multiplier raises difficult trade-offs. On the one hand, incorporating historical loss experience may improve risk sensitivity because firms with persistently higher operational losses may have higher underlying operational-risk outcomes. On the other hand, a mechanical loss multiplier may introduce volatility, may overemphasise losses from operating conditions that have materially changed, and may be difficult to reconcile with the objective of a simple and stable minimum capital requirement. Federal Reserve staff analysis has recognised this tension: past losses and income metrics may be predictive of future operational loss experience, but reliance on backward-looking windows can raise questions about forward-lookingness and capital stability.¹⁸

For that reason, we do not propose that operational-condition indicators be used to offset a loss multiplier or to reduce minimum regulatory capital requirements. Nor do we propose that they replace historical losses in calibration. Historical losses should remain an essential empirical anchor because they record actual outcomes, are capable of consistent monetary measurement, and allow validation against realised experience.

The additional question is whether the agencies should test, over time, whether standardised operational-condition indicators add explanatory value beyond historical losses and business-volume proxies. If past losses are informative because they proxy for control quality, risk culture, risk appetite and related operational conditions, then a standardised and evidence-backed measure of those conditions may help supervisors understand whether a firm's current operational condition is consistent with, better than, or worse than its historical loss profile.¹⁹

To make this concrete, and responsive to the agencies' question of which methodology could be used to estimate operational-risk exposure, one candidate construct could be specified for testing without reintroducing internal capital modelling. The construct would express inherent operational exposure as a function of two supervisor-defined inputs: the relative complexity of an activity and its operational throughput, moderated by evidenced mitigation-condition bands describing the current state of controls, remediation and operational resilience. Critically, the complexity and throughput parameters, and the definitions of the mitigation-condition bands, would be defined and held by the agencies, not by individual firms; firms would supply only standardised, evidenced inputs. The construct would carry no capital consequence and would function solely as a research and validation object, used to test whether it explains subsequent operational-risk outcomes more fully than scale, income and loss proxies alone. It would not constitute a differentiated capital scalar, a firm-specific model or a basis for capital relief.

Such information should be used first for supervisory review, benchmarking and validation. It should not be used to produce immediate capital relief. Before any broader regulatory use could be considered, the agencies would need evidence that the indicators are comparable across firms, stable over time, resistant to gaming, and validated against subsequent operational-risk outcomes.

C. Standardised operational-condition indicators as a supervisory complement

The appropriate complement to the proposed standardised approach is not an internal capital model, but the standardised supervisory evidence layer described above: a limited, standardised and evidence-backed view of current operational condition.

The purpose of these indicators would be to make operational evidence more comparable and reviewable. They would not ask a banking organisation to estimate a 99.9th percentile loss distribution, assign its own capital scalar, or translate an internal score directly into a capital requirement. They would

¹⁷ NPR, pp. 197-199, 215-216.

¹⁸ Marco Migueis, Is Operational Risk Regulation Forward-looking and Sensitive to Current Risks?, FEDS Notes, Board of Governors of the Federal Reserve System, May 21, 2018.

¹⁹ Curti and Migueis 2023, abstract and pp. 1-2, 16.

instead provide a structured way to observe operational conditions that are relevant to risk but not captured directly by BI, gross income, total assets or realised losses.

This approach is consistent with the agencies' concern for simplicity. The minimum regulatory capital requirement can remain standardised, while supervisors separately collect or review a bounded set of operational-condition evidence. Such a structure would preserve the advantages of a standardised capital framework while allowing the agencies to assess whether additional information improves supervisory discrimination and validation.

The approach is also relevant to Q87, although business-line differentiation should be treated cautiously. The agencies' consideration of downscaling certain activities recognises that operational-risk outcomes may differ by activity. However, differentiated scalars may introduce classification risk and supervisory burden if they are not empirically supported. Operational-condition indicators could help test whether differences across activities are persistent, observable and evidenced before any expansion of business-line differentiation is considered.²⁰

D. Mitigating subjectivity, inconsistency and gaming

The principal objection to operational-condition indicators is that they may be subjective. That objection is valid and should shape the design of any future supervisory evidence layer.

Operational-condition information should not be treated as fully objective in the same sense as realised loss amounts recorded in financial systems. Many relevant indicators will involve judgement: the severity of a control weakness, the adequacy of remediation, the effectiveness of a resilience test, or the materiality of a third-party dependency. Such judgement should not determine minimum regulatory capital requirements through a firm-controlled process.

However, judgement-based information can still be made more reliable if it is standardised, evidenced, challenged and validated. The relevant safeguards would include:

- a fixed operational-risk taxonomy and business/activity taxonomy;
- defined evidence requirements for each indicator;
- banded indicators rather than falsely precise numerical scores;
- documented ownership and approval;
- independent challenge by risk, compliance or audit functions;
- clear version history and audit trail;
- periodic benchmarking across comparable activities and firms;
- validation against subsequent realised losses, near misses, audit findings, supervisory findings, resilience incidents and remediation outcomes.

These safeguards would not eliminate subjectivity. They would make it visible, reviewable and capable of challenge. This is the essential distinction between an unconstrained internal model and a standardised supervisory evidence layer. The former allows firm-specific assumptions to determine capital outcomes. The latter would require firms to produce comparable evidence under a common taxonomy, with no automatic capital consequence.

E. GRC systems as systems of record for operational evidence

Many banking organisations already maintain governance, risk and compliance systems, issue-management platforms, audit-management systems, operational-risk event databases, control libraries, resilience testing records and third-party risk management tools. Where sufficiently mature, these systems may provide practical systems of record for the evidence supporting operational-condition indicators.

Their role should be carefully limited. GRC systems should not be treated as alternative capital models, substitutes for the agencies' operational-risk framework, or direct capital calculation engines. They should instead preserve the underlying evidence and workflow records needed to make operational-condition information reviewable. This includes evidence attachments, control test results, issue

²⁰ NPR, pp. 206-208.

ownership, remediation dates, challenge comments, approval history, audit trail, versioning and status changes over time.

This distinction is important because GRC data may themselves contain judgement-based records. The value of such systems lies not in assuming that their contents are automatically objective, but in improving evidential discipline. A judgement recorded in a controlled workflow, supported by evidence, independently challenged, versioned and later tested against outcomes is materially more useful than an informal or unstructured assessment.

In this respect, existing operational-risk infrastructures may reduce implementation burden and support consistency if the agencies later decide to study standardised operational-condition indicators. The relevant supervisory question is not whether every banking organisation already has perfect data. It is whether a limited, standardised and evidence-backed set of indicators can be defined in a way that builds on existing infrastructure while remaining comparable across firms.

F. Conclusion on operational-risk measurement

The agencies' proposed standardised approach is a defensible foundation for minimum regulatory capital requirements. It improves transparency and comparability and avoids the weaknesses of internal operational-risk capital models. Realised loss data should remain essential because they provide objective evidence of outcomes and are critical for calibration and validation.

The residual limitation is that BI, income, assets and realised losses do not directly measure current operational condition. A standardised supervisory evidence layer could help address that limitation without reintroducing internal capital models. Its purpose should be supervisory insight, benchmarking, validation and future empirical study, not immediate capital relief.

The appropriate standard for such information should be deliberately demanding: the safeguards set out in Section III.D above. If these safeguards are not met, operational-condition indicators should not be used for any consequential regulatory purpose. If they are met, they may provide a disciplined way to add forward-looking operational information to a framework that otherwise relies primarily on scale proxies and realised outcomes.

IV. Proposed Pilot or Supervisory Data Study

Questions addressed: Q195 and Q199 prominently; Q92 and Q198 in support; Q87, Q196 and Q197 where relevant.

The agencies should consider a limited pilot or supervisory data study to assess whether standardised operational-condition indicators can provide incremental supervisory value without being incorporated into minimum regulatory capital requirements. Such a study would be consistent with the agencies' broader objectives of improving risk sensitivity, transparency and consistency, while preserving the simplicity of the proposed standardised operational-risk capital framework.²¹

The purpose of the study would not be to create a new internal model, to calculate a firm-specific capital adjustment, or to provide automatic capital relief. Its purpose would be to test whether standardised, evidence-backed operational information can improve supervisory understanding of current operational condition, support peer benchmarking, assist impact analysis, and validate loss- and proxy-based operational-risk measures.

A. Purpose and supervisory use

A pilot or data study should be designed to answer a limited empirical question: whether operational-condition indicators add useful information beyond business indicator measures, income or asset proxies, and realised operational loss data.

The study could support five supervisory objectives: impact analysis; benchmarking; validation against subsequent outcomes; analysis of business-line differentiation; and the coherence of related supervisory frameworks. It could also allow the agencies to assess whether operational-condition evidence has any

²¹ NPR, pp. 1, 13, 16-18, 529, 660.

value for interpreting or validating operational-risk stress-test outputs, while keeping the capital NPR and the Board's supervisory stress-test model analytically distinct.

B. Scope and sequencing

The study should be limited in scope at the outset. The agencies could begin with Category I and II banking organisations, or with a voluntary or supervisory sample of firms with sufficiently mature operational-risk, control, audit, resilience and GRC infrastructures. A limited scope would allow the agencies to test feasibility, burden, data quality and comparability before considering any broader reporting requirement.

The study should also be sequenced in stages. An initial design stage could define common taxonomies, evidence expectations, indicator bands, governance requirements and data submission formats. A collection stage could gather a limited number of indicators across selected activities, legal entities or business lines. An assessment stage could test data completeness, interpretation consistency, burden, evidence quality, auditability and susceptibility to gaming. A validation stage could compare the indicators with subsequent outcomes over time.

C. Illustrative data elements

The pilot should use a narrow set of fields. The following categories are illustrative and should be refined through supervisory design before any collection begins:

- business line, activity and process taxonomy;
- operational-risk event and control taxonomy;
- control testing or assurance result, expressed in standardised bands;
- unresolved audit, compliance or operational-risk findings;
- remediation ageing, overdue status and closure evidence;
- incident and near-miss classifications;
- operational resilience or business continuity test results;
- material technology, platform, data, outsourcing or third-party concentration indicators;
- evidence source, date, owner and approver;
- independent challenge status;
- audit trail and version history.

The indicators should be banded rather than falsely precise. For example, a remediation indicator may distinguish between current, moderately aged, materially overdue and long-overdue remediation items. A control evidence indicator may distinguish between tested-effective, tested-with-deficiencies, untested, or subject to unresolved challenge. Such bands are more defensible than numerical scores that imply a precision that the underlying evidence may not support.

D. Evidence and governance expectations

The study should test not only reported indicator values, but also the evidence and governance supporting them. A reported indicator should be accompanied by evidence sufficient for supervisory review. This evidence might include control test records, audit workpapers or finding summaries, issue-management records, remediation approvals, resilience test reports, third-party risk assessments, incident records, management attestations, or independent challenge documentation.

Where sufficiently mature, existing GRC systems and related operational-risk infrastructures may serve as systems of record for this evidence. Their role should be evidential and procedural: preserving records of evidence, workflow, ownership, review, challenge, remediation, approval, version history and audit trail. They should not be used as alternative capital models or direct regulatory measurement engines.

The study should also test whether the evidence can be reviewed consistently across firms. If similar evidence produces materially different indicator classifications across institutions, the taxonomy or evidence hierarchy would need refinement before the information could be considered for any consequential regulatory purpose.

E. Validation against subsequent outcomes

Validation should be central to the study. Operational-condition indicators should be assessed against subsequent observable outcomes, including:

- realised operational losses;
- loss frequency and severity;
- near misses;
- internal audit findings;
- compliance or operational-risk findings;
- supervisory findings, where available;
- remediation delays and repeated issue reopening;
- operational resilience failures;
- material incidents involving technology, outsourcing, data, cyber or processing failures.

This validation should be longitudinal. A single collection would be insufficient to determine whether the indicators are stable, comparable or informative. Repeated observations over time would allow the agencies to test whether changes in operational condition precede or coincide with changes in realised outcomes.

The study should also assess whether indicators add value after controlling for business indicator measures, firm size, income measures, assets, business mix and historical operational losses. This would directly test whether the proposed evidence layer provides incremental information rather than simply restating information already captured through scale proxies or loss history.

The validation stage should therefore include out-of-sample empirical testing, so that indicators collected at one point in time can be assessed against subsequent operational-risk outcomes rather than only against contemporaneous judgement or historical fit.

As a concrete and falsifiable starting hypothesis, the agencies could test whether a standardised control-condition band predicts a banking organisation's next-period operational-loss frequency incrementally, after controlling for the business indicator, firm size, income measures and three-year loss history. Loss frequency is a suitable dependent variable because Federal Reserve staff analysis has found it to be more stable, and more predictive of future operational-risk outcomes, than loss severity. Such a test could in principle be conducted on data the agencies already hold, including the special supervisory data collection from large bank holding companies referenced in the impact analysis, before any new or permanent reporting requirement is contemplated. A null result would itself be informative, indicating that condition indicators add little beyond existing scale and loss measures and should not be pursued for any consequential supervisory use.

F. Treatment of subjectivity and gaming risk

The pilot should be designed with the agencies' concerns about subjectivity and potential incentive effects at its centre. Operational-condition indicators will inevitably include judgement-based elements. The study should therefore test whether those judgements can be standardised, evidenced, challenged and validated.

The agencies could mitigate subjectivity and gaming risk through common definitions and taxonomies, prescribed evidence hierarchies, limited and banded indicator values, documented ownership and approval, independent challenge by control, risk, compliance or audit functions, supervisory review of selected evidence packs, peer comparison, tracking of version changes and validation against subsequent outcomes.

The absence of automatic capital consequences during the pilot would itself reduce gaming incentives. If firms know that indicators are being studied for supervisory understanding rather than used immediately to reduce capital, the agencies may obtain a clearer view of data quality and behavioural incentives before considering any broader use.

G. Reporting burden and proportionality

A pilot should be burden-sensitive. The agencies should avoid creating a comprehensive new operational-risk reporting universe before testing whether the information is useful. The initial study

should use a small number of indicators, clear definitions, limited frequency and a defined sample of firms or activities.

The agencies should also evaluate whether the data can be sourced from existing systems. Many firms already maintain operational-risk event databases, issue-management tools, control libraries, audit-management systems, resilience testing records and third-party risk systems. These infrastructures may reduce incremental burden where they are sufficiently mature, although they should not be assumed to produce comparable data without supervisory standardisation.

The study should explicitly measure burden, including mapping effort, data extraction effort, evidence preparation, review time, reconciliation requirements and supervisory follow-up. Burden should be assessed alongside informational value. Indicators that are costly to produce but weakly validated should not be retained.

H. Non-use for minimum regulatory capital requirements during the study

The agencies should make clear that participation in the study would not affect minimum regulatory capital requirements during the pilot period. Operational-condition indicators should not be used to reduce operational-risk capital, increase operational-risk capital, override the business indicator framework, or create firm-specific capital add-ons through an unpublished formula.

This limitation is important for two reasons. First, it preserves the integrity and comparability of the proposed standardised capital framework. Second, it allows the agencies to study the data without creating immediate incentives for firms to optimise reported indicators for capital purposes.

If the study later shows that certain indicators are comparable, evidenced, stable, resistant to gaming and validated against subsequent outcomes, the agencies could then consider whether they have a role in supervisory review, benchmarking, reporting, impact analysis or model validation. Any consequential use should follow further notice, explanation and opportunity for comment.

Conclusion on the pilot or data study

A limited pilot or supervisory data study would provide a measured way to test standardised operational-condition indicators without reintroducing internal capital models or altering minimum regulatory capital requirements. It would allow the agencies to assess whether forward-looking operational evidence can be made sufficiently standardised, evidenced, challengeable, benchmarkable and validated to be useful.

This approach recognises both sides of the operational-risk measurement problem. The agencies are right to avoid opaque internal capital models and to preserve the comparability of minimum regulatory capital requirements. At the same time, operational risk is affected by current control condition, remediation discipline, operational resilience, technology dependency and third-party concentration. These conditions are not fully captured by business indicator, income, assets or realised losses alone.

A bounded study would allow the agencies to test that residual gap empirically before considering any broader use of operational-condition evidence.

V. Responses to Selected Remaining NPR Questions

Questions addressed: Q87 and Q198 clearly; Q195–Q199 where relevant to impact analysis, benefits, costs, feasibility and empirical testing; Q86, Q88, Q90, Q91, Q196 and Q197 briefly where they support the main argument.

This section addresses selected remaining questions relevant to operational-risk measurement and supervisory information. It should be read together with Sections II–IV. The central recommendation remains limited: the agencies should not adopt operational-condition indicators as determinants of minimum regulatory capital requirements at this stage. Rather, the agencies should consider whether a standardised supervisory evidence layer could be tested through a bounded pilot or data study to support supervisory benchmarking, impact analysis, model validation and future empirical evaluation.

A. Q87 — Business-line differentiation and differentiated scalars

Q87 asks whether the agencies should adopt differentiated scalars for certain operational-risk business lines, including whether the proposed treatment for investment management, investment services and treasury services is appropriate.²²

We agree that operational-risk outcomes may differ materially by activity. A custody, asset-servicing, payments, investment management or treasury services business may present a different operational-risk profile from traditional balance-sheet lending. We also agree that some degree of business-line differentiation may improve risk sensitivity where there is robust evidence that operational loss outcomes differ persistently across activities.

The residual issue is that business-line differentiation can create classification risk. If scalars are tied to activity categories without sufficient evidence discipline, firms may have incentives to classify activities in ways that reduce capital requirements. Differentiated scalars can also increase supervisory burden if business-line boundaries are not clear, consistently applied and empirically supported.

The proposed evidence layer could help address this problem, not by replacing business-line scalars, but by allowing the agencies to test whether observed differences across activities are supported by current operational evidence. This approach would be bounded and defensible because it would not require the agencies to adopt new business-line scalars immediately. Any future use of differentiated scalars should be preceded by empirical testing, common taxonomy, supervisory review of evidence and further consultation.

Recommended treatment: pilot/data study and, if warranted by evidence, further consultation. No immediate expansion of differentiated scalars should be adopted solely on the basis of unevidenced business-line judgement.

B. Q198 — Interaction with related proposals, including stress testing and reporting changes

Q198 asks what modifications, if any, the agencies should consider to this proposal due to related proposals. In this comment, that question is relevant principally to interactions with the Board's stress testing framework, stress capital buffer requirements, GSIB surcharge proposal and related reporting implications.²³

We agree that the operational-risk capital framework should be considered coherently with related supervisory and reporting frameworks. The proposed expanded risk-based approach concerns minimum regulatory capital requirements. The Board's operational-risk stress-test model and related transparency proposals serve a distinct function: projecting operational losses under supervisory stress scenarios and informing the stress capital buffer framework. These frameworks should remain analytically distinct.^{24 25}

The residual issue is that both frameworks rely, in different ways, on realised losses, firm scale, business activity and supervisory data. These inputs are valuable, but they do not directly measure current operational condition. As discussed in Sections III and IV, current operational condition includes control effectiveness, issue remediation, resilience testing, technology dependency, outsourcing concentration and other operational features that may be relevant before losses crystallise.

The proposed evidence layer could help the agencies evaluate whether additional evidence improves interpretation or validation of loss- and proxy-based approaches. Such indicators should not replace the Board's operational-risk stress-test model, should not be used automatically in stress capital buffer

²² NPR, pp. 206-208.

²³ NPR, pp. 17, 30, 659.

²⁴ Operational Risk Model Documentation, 2025.

²⁵ Stress-Test Transparency Proposal, 2025.

calculations, and should not be used to create firm-specific capital adjustments without further analysis and consultation.^{26 27}

Recommended treatment: pilot/data study, coordinated with related reporting and stress-test transparency work. No immediate capital or stress capital buffer change should be made on the basis of operational-condition indicators.

C. Q195 — Impact analysis, data and estimation methodology

Q195 asks whether the agencies' impact analysis and estimation methodologies are appropriate and whether additional data or methodologies should be considered.²⁸

We agree that the agencies' impact analysis should remain transparent, empirical and administrable. The impact analysis should not be used to introduce a new internal operational-risk model or to bypass the proposed standardised framework.

The residual issue is that impact analysis based on business indicator, income, assets and realised losses may not fully identify whether capital outcomes vary appropriately with current operational condition. Two firms with similar business indicators and similar loss histories may have different current control states, remediation backlogs, resilience profiles, technology dependencies or third-party concentrations.

A standardised operational-condition data study could help the agencies test whether these differences are material and observable. The agencies could compare outcomes under BI-only measures, BI plus realised losses, gross income, asset-based measures and BI plus selected indicators. The purpose would be to assess incremental explanatory and supervisory value, not to create immediate capital effects.

Recommended treatment: pilot/data study for impact analysis and empirical testing. No immediate rule change to capital calibration should be made on this basis.

D. Q199 — Related reporting, implementation feasibility and burden

Q199 asks what modifications, if any, the Board should consider to this proposal or to other Board rules indirectly affected by the proposal, including possible reporting and burden implications.²⁹

We agree that any additional reporting should be limited, proportionate and justified by demonstrated supervisory value. Operational-risk reporting should not become a broad new reporting universe before feasibility, comparability and burden have been tested.

The residual issue is that existing regulatory data do not appear to provide a common, evidence-backed view of current operational condition across firms. Without such information, the agencies may have limited ability to evaluate whether differences in operational-risk outcomes are being captured by BI, income, assets and realised losses alone.

A bounded reporting pilot could test feasibility before any permanent reporting change. The pilot could be limited to Category I and II firms or to a voluntary/supervisory sample. It could use a small number of fields, such as activity taxonomy, control evidence bands, unresolved findings, remediation ageing, incident and near-miss classifications, resilience test results, material technology or third-party concentration indicators, evidence source, challenge status and version history.

GRC systems may be relevant here as systems of record. Where sufficiently mature, they may reduce the incremental burden of evidence collection by preserving issue-management records, control testing evidence, audit trails, approvals, challenge history and remediation status. However, GRC data should not be assumed to be objective or comparable without supervisory standardisation and validation.

²⁶ Board of Governors of the Federal Reserve System, Supervisory Stress Test Model Documentation: Operational Risk Model, October 2025.

²⁷ Board of Governors of the Federal Reserve System, Enhanced Transparency and Public Accountability of the Supervisory Stress Test Models and Scenarios, notice of proposed rulemaking, 2025.

²⁸ NPR, p. 529.

²⁹ NPR, p. 660.

Recommended treatment: pilot/data study followed by further consultation if the agencies later consider permanent reporting changes. No broad immediate reporting requirement should be adopted without feasibility and burden analysis.

E. Q86, Q88, Q90 and Q91 — Alternative proxies and specific components of the business indicator

Q86, Q88, Q90 and Q91 address the design of particular operational-risk proxies or alternatives, including the noninterest component, recharge or non-financial income, gross income and asset-based approaches.³⁰

We agree that these measures may have practical value. Noninterest income and expense can help capture fee, service, trading and other activity-based exposure. Gross income is simple and transparent. Asset-based measures are stable and administrable. Certain recharge or affiliate income may overstate operational-risk exposure where it reflects low-risk pass-through arrangements.

The residual issue is that all of these measures remain proxies. They measure scale, activity, income, accounting classification or balance-sheet footprint. They do not directly measure current operational condition. The proposed evidence layer could assist supervisors in understanding whether a given proxy is directionally consistent with observable operational conditions, without replacing the proxy or creating a firm-specific capital model.

Recommended treatment: no immediate rule change on this basis. These questions support the rationale for a pilot/data study, but do not require separate operational-condition capital treatment at this stage.

F. Q196 — Benefits, costs and macroeconomic considerations

Q196 asks what alternative approaches, if any, the agencies should consider for assessing the macroeconomic benefits and costs of the changes in the level of required capital that would result from the proposals.³¹

We agree that any assessment of macroeconomic benefits and costs should remain empirical and proportionate. Better operational-risk information may have potential supervisory value, but this comment does not assert quantified macroeconomic benefits from operational-condition indicators.³²

The residual issue is that those benefits should not be overstated. There is not yet sufficient evidence to quantify macroeconomic benefits from operational-condition indicators, nor to assert that such indicators would reduce capital requirements or improve lending capacity. The appropriate contribution of operational-condition indicators at this stage is empirical.

Recommended treatment: pilot/data study only. No quantified macroeconomic claim or immediate capital conclusion should be drawn at this stage.

G. Q197 — Competitiveness effects

Q197 asks what alternative approaches, if any, the agencies should consider for assessing the proposal's effect on competitiveness.³³

We agree that competitiveness arguments should be handled cautiously. Operational-risk capital should not be adjusted merely to favour particular firms, sectors or activities. However, measurement design can affect competitive neutrality if some business models are systematically over- or under-represented by scale proxies.

The residual issue is that operational-risk outcomes may vary across activities in ways not fully captured by income or assets. The proposed evidence layer could help the agencies test whether proxy-based capital measures produce systematic differences across business models that are not supported by

³⁰ NPR, pp. 207-215.

³¹ NPR, p. 594.

³² NPR, p. 594.

³³ NPR, p. 649.

operational evidence. The analysis should be symmetrical: it should identify both possible overstatement and possible understatement of operational risk.

Recommended treatment: pilot/data study and further empirical analysis. No immediate competitiveness-based capital adjustment should be made.

H. Q85 — Operating lease income and expense classification

Q85 concerns whether income from operating leases and depreciation and impairments relating to operating leased assets should be included in the interest component rather than the noninterest component. This issue appears primarily to concern accounting classification and Basel alignment rather than the forward-looking measurement of operational condition.³⁴

We do not recommend separate treatment of Q85 in this comment, except to note the general principle that accounting classification should not obscure material operational exposure where such exposure exists. That point is addressed more directly in relation to Q86, Q88, Q90 and Q91.

Recommended treatment: no action in this comment, unless a narrow conforming point is needed for completeness.

Summary of recommended treatment

The selected remaining questions support the same conclusion reached in Sections III and IV. The agencies should preserve the comparability and simplicity of the proposed standardised operational-risk capital framework. They should also consider whether a bounded pilot or supervisory data study can test operational-condition indicators as a complement for supervisory insight, benchmarking, validation and future impact analysis.

This approach avoids automatic capital relief, avoids reintroducing internal operational-risk models, and preserves realised loss data as essential evidence of outcomes. It also provides a practical way for the agencies to evaluate whether standardised, evidence-backed operational-condition information can improve forward-looking risk capture over time.

VI. Conclusion

The agencies' proposal seeks to improve risk sensitivity, consistency and transparency in minimum regulatory capital requirements while reducing complexity. We support those objectives. In operational risk, they are particularly important because severe loss events are difficult to model, available data are limited, and capital outcomes should not depend on opaque or inconsistent firm-specific assumptions.³⁵

For that reason, we support the removal of internal operational-risk models as determinants of minimum regulatory capital requirements. Internal modelling, scenario analysis and operational-risk analytics may continue to serve important internal risk management and capital planning purposes. However, they should not determine minimum regulatory capital requirements where the relevant modelling choices are difficult to compare, validate and supervise consistently across firms.

At the same time, a standardised framework based on realised losses, business indicator measures, income components or asset-based proxies may not, on its own, provide a sufficiently forward-looking view of current operational condition. Realised loss data remain essential because they provide objective evidence of outcomes and support calibration and validation. Business indicator and other proxy measures remain useful because they support comparability and administrability. Their limitation is that they do not directly observe the current condition of controls, remediation, resilience, technology dependency, outsourcing concentration or other operational features that may affect future operational-risk outcomes.

Accordingly, we recommend that the agencies consider a bounded pilot or supervisory data study on standardised, evidence-backed operational-condition indicators. Such a study would allow the agencies to test whether a limited set of indicators, supported by fixed taxonomy, evidence requirements,

³⁴ NPR, pp. 202-203.

³⁵ NPR, pp. 1, 13, 16-18, 25-30, 197-198.

independent challenge, audit trail, benchmarking and validation against subsequent outcomes, can improve supervisory insight, impact analysis, benchmarking and model validation.

This recommendation is not a request for immediate capital relief, not a return to the advanced measurement approaches, and not a substitute for realised operational loss data. It is a modest proposal for empirical supervisory exploration before any consequential regulatory use is considered. We would welcome further supervisory consultation on whether such a standardised supervisory evidence layer could assist the agencies in strengthening forward-looking operational-risk assessment while preserving the comparability and defensibility of the proposed expanded risk-based approach.

Respectfully submitted,

Steve Bailey

Chairman

SERRAQ HQ (working name | Risk Accounting Standards Board – RASB)