



Why We Call It Risk Accounting

Measuring the Missing Present

Revisiting Peter J. Hughes's vision for integrating risk measurement and financial reporting

RASB WHITE PAPER

August 2026

Based on the published Risk Accounting work of Peter J. Hughes, the 2016 JRMFI expert commentary, the 2021 published test of RU inherent predictiveness, the 2024 Butler-Brooks academic endorsement and extension, and primary accounting, regulatory and supervisory sources.



Risk management tells us what might go wrong under different scenarios. Risk Accounting asks what the risk we have already accepted may be costing us - before anything actually goes wrong.

Author's Note

This paper grew out of a return to Peter J. Hughes's life's work and to the question that sat behind his Risk Accounting method: how should organizations understand the economic significance of non-financial risk before that risk becomes a realized loss? Revisiting his book, papers and correspondence makes clear that his ambition reached beyond improving risk assessment. He wanted risk information to become measurable, aggregable and capable of being connected with the accounting information used to describe financial performance and condition.

Peter approached that problem from an unusual combination of perspectives. He was a chartered accountant who had also led finance, audit, operations, treasury, trading and risk functions in a global bank. That career allowed him to see the same transaction through several professional lenses: the revenue it generated, the operating environment that processed it, the risks accepted with it, the controls intended to mitigate those risks and the accounting consequences that could emerge much later.

A private email written shortly before his death captured the proposition with characteristic directness:

'...we account for the future statistical loss probability of one class of risk (credit) but not the other (non-financial risk). This is because nobody has figured out a common measurement framework to quantify an exposure to non-financial risk... except us!'

Peter J. Hughes, private correspondence

The pages that follow revisit that proposition in contemporary terms. They examine the measurement problem Peter was trying to solve, connect it with developments in expected-loss accounting and operational-risk regulation, and explore the management and reporting possibilities that follow when accepted non-financial risk can be measured credibly. In that sense, the paper is both an analysis and a tribute to Peter's vision for a closer integration of risk management, operations and accounting.

Author perspective and disclosure

The author worked closely with Peter Hughes during the later development of Risk Accounting and remains directly engaged with the Risk Accounting Standards Board and the continuing development of the discipline. The perspective is therefore engaged rather than detached. To keep that perspective transparent, this paper distinguishes among three kinds of statement: propositions published by Hughes; external accounting, regulatory or empirical evidence; and implications developed in this paper. The author also participated directly in one of the empirical research exercises discussed later; that role is disclosed alongside the evidence in Section 10.1.

Terminology Used in This Paper

Risk terminology varies across institutions and regulatory frameworks. This paper uses established banking definitions where a widely recognized definition exists and states a working convention where industry usage remains broader. The full expression non-financial risk is used throughout for readability.

Term	Working definition in this paper
Operational risk	The Basel definition: the risk of loss resulting from inadequate or failed internal processes, people and systems, or from external events. The Basel definition includes legal risk and excludes strategic and reputational risk.[5]
Non-financial risk	A broad working term in this paper for material risks outside the conventional financial-risk categories. Operational risk is central to the discussion; exact institutional boundaries vary, so the term is used as an umbrella rather than as a claim of universal taxonomy.
Inherent risk	The risk an activity would pose before the effect of controls or other mitigating factors is taken into account.[6] In Hughes's banking method, inherent exposure is operationalized through activity uncertainty and daily-new-business scale.
Residual risk	The risk remaining after the effect of mitigation is taken into account.[6]
Risk Unit (RU)	Hughes's common quantitative unit for expressing non-financial-risk exposure. The published method uses RUs to quantify inherent and residual exposure and to aggregate risk information through reporting hierarchies.[2, pp. 136-141]
Exposure Uncertainty Factor (EUF)	A factor assigned to a product or activity to represent the relative uncertainty or operational complexity associated with a principal risk type. EUFs are method-specific inputs to inherent RU.[2, pp. 138, 148, 156-160]
Value Band Weighting (VBW)	A weighting selected from Hughes's value table by reference to daily new business for a product. It introduces scale into inherent-risk measurement.[2, pp. 139, 149, 161-163]
E-RCSA	The Enhanced Risk and Control Self-Assessment. Hughes's framework replaces a purely qualitative traffic-light assessment with benchmark questions, numeric scores and category weightings that feed the RMI.[2, pp. 140, 153-155]
Risk Mitigation Index (RMI)	A 0-100 measure of mitigation effectiveness derived through Actual Score Aggregate and Maximum Score Aggregate calculations and recalculated through the reporting hierarchy.[2, pp. 140, 164-169]
RU_m	Hughes's notation for the monetary value of one RU, proposed as a longer-term empirical calibration derived through statistical correlation and back-testing of residual RUs and associated loss experience.[2, p. 141]

The professional definitions establish the vocabulary; the Risk Accounting terms describe Hughes's particular measurement architecture.

Executive Summary

Modern risk management has become increasingly sophisticated at looking in two directions. Regulators and institutions collect detailed evidence about what has already happened: losses, incidents, causes and historical outcomes. Risk practitioners use scenarios, stress tests, forecasts, resilience testing and models to explore multiple possible futures. Risk Accounting focuses first on the condition between them: the exposure that already exists because an organization has chosen to conduct business.

A transaction can create risk before an error, fraud, default, outage or conduct failure becomes visible. Business volume can increase faster than mitigation capability. Sales incentives can alter the quality of decisions. A control environment can deteriorate while recent loss history remains reassuring. Peter Hughes's central proposition was that this accepted exposure should become measurable before an adverse event turns it into accounting data.[2]

Risk Accounting separates inherent exposure from mitigation effectiveness and expresses the remaining exposure through a common quantitative architecture based on Risk Units, the E-RCSA and the Risk Mitigation Index. In the published banking method, daily new-business values determine a Value Band Weighting, product characteristics determine Exposure Uncertainty Factors, and the combination produces inherent RUs. Benchmark-based E-RCSA assessments are then weighted into Actual and Maximum Score Aggregates, from which RMI and residual RUs are calculated. Appendix A works one of Hughes's own published examples through that chain.[2, pp. 136-169]

The method's claim of additivity deserves precision. Inherent RUs and the ASA/MSA control totals aggregate through the hierarchy; RMI and residual exposure are recalculated at each higher level. The common-unit architecture therefore supports aggregation by treating parent residual exposure as an exposure-weighted recomputation rather than as the simple arithmetic sum of pre-calculated child residual figures. A further question remains empirical: whether an RU carries sufficiently stable monetary meaning across different risk classes and institutions. Hughes's proposed RU_m calibration is the test that ultimately connects the index to expected monetary loss.[2, p. 141]

Expected-credit-loss accounting provides the important precedent. IFRS 9 and the US CECL framework recognize that an existing credit exposure can contain economically meaningful expected loss before the exact future default is known.[14][16] For non-financial risk, current IFRS follows a different recognition architecture: IAS 37 recognizes provisions when a present obligation arising from a past event satisfies its requirements; future operating losses that depend on future operations fall outside that recognition basis.[15] A broad expected-non-financial-risk treatment would therefore require its own standard-setting basis rather than a simple extension of current provision rules.

The regulatory context also makes the measurement question timely. Basel removed the Advanced Measurement Approaches for operational-risk capital after finding excessive complexity and variability in internal modelling, and the Basel III reforms moved to a standardized approach.[10][11] In the European Union, CRR3 makes the Pillar 1 own-funds requirement for operational risk equal to the Business Indicator Component, while still requiring substantial institutions to maintain detailed loss data and operational-risk governance.[12] DORA, applicable since January 2025, adds a parallel regulatory emphasis on digital operational resilience.[13] These

frameworks strengthen capital, loss-data and resilience disciplines; they leave room for a separate management question about current control-adjusted exposure.

The practical value of Risk Accounting begins before any change to accounting standards. A common exposure measure can support risk-adjusted profitability, product pricing, control-investment decisions, management accountability, quantitative risk appetite and a longitudinal record linking the conditions under which risk was accepted with later outcomes. That last idea is developed here as a risk-provenance record: a structured history joining origination, changing risk conditions, event occurrence, discovery, accounting recognition and remediation.

Risk Accounting also enters this discussion with a record of published scrutiny. Six invited experts commented on the 2016 JRMFI papers and surfaced questions about additivity, interpretation, comparability, governance and implementation.[4] Hughes and Marzouk then published an initial empirical test in 2021 using data from fifteen large US banks; the study reported a 58.1% increase in inherent operational-risk RUs over the eight-year period leading into the financial crisis and described the result as an initial indication rather than a conclusive validation.[3] Butler and Brooks, who were unaffiliated with the original development of the method, subsequently published an academic endorsement and extension within a broader critique of prevailing operational-risk practice, highlighting granular measurement, aggregation, transparency and auditability while proposing a wider digital architecture around the method.[25]

The paper therefore treats Risk Accounting as both a management proposition and a continuing research programme. Its claims become stronger to the extent that the measure responds sensibly to changes in exposure and mitigation, discriminates between materially different risk states, survives independent validation and manipulation-resistant governance, and eventually supports robust monetary calibration. Evidence that challenges those properties provides a basis for refinement and, where necessary, narrower claims. That explicit burden of proof is one of the defining features of a measurement discipline.

The past gives us observations. The future gives us possibilities. Risk Accounting seeks to measure the exposure that exists between them - now.

Contents

Author's Note	1
Author perspective and disclosure.....	1
Terminology Used in This Paper	2
Executive Summary	3
Contents	5
1. The Accounting Gap	7
1.1 Expected credit loss as a precedent	7
1.2 The missing non-financial-risk denominator	7
1.3 Measurement, valuation and recognition.....	8
2. Risk Has a Time Problem	8
2.1 The past: what happened?	8
2.2 The future: what might happen?	8
2.3 The missing present: what have we already accepted?.....	8
3. Using Loss History in Context	9
3.1 Detection and reporting shape the observed record.....	9
3.2 Remediation and adaptation change the system	9
3.3 Clean history and emerging exposure	10
3.4 How loss data enters Risk Accounting.....	10
4. From Loss Categories to Risk Origins.....	10
4.1 Credit losses can contain operational causes.....	10
4.2 The global financial crisis as a causal-layer example.....	10
4.3 Sales incentives: a clearer organizational case	11
5. Connecting Decisions, Exposures and Outcomes.....	11
5.1 The transaction as a carrier of risk information	11
5.2 A risk-provenance record	11
5.3 Turning outcomes into tests of the measurement	12
6. Measuring the Missing Present.....	13
6.1 Inherent exposure: uncertainty and scale.....	13
6.2 Mitigation effectiveness: E-RCSA to RMI.....	13
6.3 Residual exposure	13
6.4 What additivity means	14
6.5 From residual RU to monetary expected loss	14
6.6 Professional scrutiny and refinement of the method	14
7. From Loss Modelling to Exposure Measurement	15
7.1 A different measurement target	15
7.2 BEICFs: important prior art.....	16

7.3 The current standardized capital context	16
7.4 Richer event chronology and operational resilience.....	16
7.5 Relationship to the existing operational-risk toolkit.....	17
8. What Measurable Non-Financial Risk Changes for Management.....	17
8.1 Risk-adjusted profitability and product pricing.....	17
8.2 Measuring the return on mitigation.....	17
8.3 Accountability, incentives and behavioural change.....	18
8.4 Risk appetite, tolerance and capacity	19
8.5 Risk budgeting and strategic planning	19
9. From Management Accounting to External Reporting	19
9.1 A management risk ledger	19
9.2 Risk-adjusted management performance.....	20
9.3 Disclosure can precede recognition	20
9.4 The recognition boundary under current IFRS.....	20
9.5 An evidence-based progression to standard setting.....	20
10. Evidence So Far: Validation, Governance and Falsifiability.....	21
10.1 Initial empirical evidence: Hughes and Marzouk (2021).....	21
10.2 External academic endorsement and extension: Butler and Brooks (2024). 21	
10.3 Mechanical reproducibility.....	22
10.4 Responsiveness	22
10.5 Discrimination	22
10.6 Monetary calibration and stability	23
10.7 Comparability and economic commensurability	23
10.8 Manipulation resistance and independent challenge.....	23
10.9 Implementation evidence and proportionality.....	24
10.10 A research programme with explicit success conditions	24
11. Why We Call It Risk Accounting.....	24
Appendix A. A Published-Method Worked Example.....	26
A.1 Product-scale inputs and inherent exposure	26
A.2 One business-component RMI calculation	26
A.3 Hierarchy aggregation.....	27
A.4 The monetary step remains symbolic until calibrated.....	27
A.5 What this example establishes - and what the research programme must test	27
References.....	28

1. The Accounting Gap

The word Accounting in Risk Accounting is literal. Accounting creates common representations of economic activity. Products, customers, transactions and business units can differ radically, yet their financial effects can be classified, aggregated, reconciled, controlled and audited through a common system. Peter Hughes believed non-financial risk lacked an equivalent common measurement language.[2]

The gap becomes visible when profit and risk move on different timelines. A bank can originate a loan, execute a derivative, sell a financial product or process a payment and recognize the commercial benefit of the activity immediately. The risk accepted through the same decision may remain outside financial performance until an adverse event later produces a recognizable loss. Hughes summarized this asymmetry as 'risk and profit taking now... losses later'.[2, pp. 49-50]



Figure 1. Visible performance, invisible risk.

1.1 Expected credit loss as a precedent

Credit risk offers a useful comparison. IFRS 9 includes expected-credit-loss impairment requirements for financial assets and commitments within its scope, while US GAAP applies the CECL framework under Topic 326.[14][16] A performing financial asset can therefore carry an expected-loss allowance while the identity and timing of the particular future default remain uncertain.

The conceptual lesson is significant: uncertain future loss can become current accounting information when a recognized exposure and an accepted measurement architecture exist. Credit risk also has a comparatively visible denominator. The bank knows the financial asset exists, knows its contractual amount and cash flows, and can associate the exposure with a counterparty and credit profile.

1.2 The missing non-financial-risk denominator

Non-financial risk is more heterogeneous. Cyber exposure, process complexity, conduct risk, model risk, sales practices, vendor dependence and operational resilience are often expressed through different measures. A traffic-light rating, an incident count, a fraud scenario and a vendor criticality score use different dimensions and therefore resist meaningful aggregation into one portfolio quantity.

Hughes's diagnosis was that measurement comes first. The Risk Unit was developed as a common quantitative metric for heterogeneous non-financial-risk exposures, with the longer-term objective of relating residual RUs to monetary loss

experience.[2, pp. 136-141] This creates a progression from exposure measurement to economic valuation and then to the reporting question.

The accounting gap is preceded by a measurement gap: a credible reporting debate begins with a credible way to describe and measure the exposure.

1.3 Measurement, valuation and recognition

These stages carry different professional requirements. Risk Accounting provides a structure for measuring inherent exposure, mitigation effectiveness and residual risk. Economic valuation asks what monetary expected cost can be associated with that measured residual exposure. Accounting recognition asks whether a resulting amount meets the definition, scope and recognition requirements of an accounting standard.

The distinction is central to the paper. A quantified residual exposure can already matter to management even while statutory accounting leaves it outside recognized profit or liabilities. A statistically calibrated expected economic loss can become an internal management measure even while the recognition question remains open. Financial reporting then becomes a separate standard-setting problem involving recognition, presentation, disclosure, comparability and assurance.

2. Risk Has a Time Problem

Risk management sees the organization through several clocks. Loss databases look backward. Scenario analysis looks forward. Accounting records economic consequences when the relevant recognition conditions are met. The operating environment, however, is continuously accepting and carrying exposure in the present.

2.1 The past: what happened?

A realized loss is valuable evidence. It can be dated, classified, investigated and assigned a monetary amount. It can expose a weak process, an ineffective control, an inappropriate incentive or a faulty model. The evidence is strongest when interpreted as an observation about the conditions that produced the event.

2.2 The future: what might happen?

Scenario analysis, stress testing, emerging-risk processes and resilience exercises address another legitimate question: what could happen under alternative conditions? They help institutions explore severe outcomes, challenge assumptions and prepare responses. Many of these tools begin from current portfolios and current systems, but their output is conditional on a hypothetical future state.

2.3 The missing present: what have we already accepted?

Between historical observation and future possibility sits a present-state question. Customers have already been onboarded. Products have already been sold. Transactions have already been processed. Systems and third parties are already in use. Controls are already strong or weak. Business volumes are already imposing a processing burden on the operating environment.

Risk Accounting focuses on this existing condition. Its proposition is structural rather than event-specific: high inherent exposure combined with weak mitigation represents a materially different present risk state from low inherent exposure protected by strong mitigation. That distinction can be measured before anyone knows which particular loss event, if any, will crystallize next.



Figure 2. Risk Accounting measures the exposure that exists now.

Temporal view	Primary question	Typical evidence or tool
Past	What happened?	Loss events, incidents, root-cause analysis, audit findings
Present	What exposure is being carried now?	Inherent RU, RMI, residual RU, control and activity state
Future	What might happen?	Scenarios, stress tests, resilience exercises, forecasts

3. Using Loss History in Context

Historical loss data remains essential, but the information content of a loss depends on the context preserved around it. A monetary amount alone describes an outcome. A loss linked to the exposure and mitigation state under which the relevant activity existed can also test the measurement system.

3.1 Detection and reporting shape the observed record

Fraud illustrates the problem clearly. Successful fraud is designed to remain undiscovered. The observable sequence is therefore longer than 'fraud occurs, fraud enters the dataset': occurrence is followed by detection, investigation, quantification, classification and accounting. Events can be absent, incomplete or delayed at several points in that chain.

Similar effects arise with misconduct, cyber events and control failures where discovery and reporting are imperfect. A dataset of discovered and recorded events is valuable evidence, but it represents discovered and recorded outcomes rather than a complete census of the underlying exposure.

3.2 Remediation and adaptation change the system

Loss events frequently trigger intervention. The organization changes a control, process, limit, procedure or system. In adversarial areas such as fraud and cyber risk, the threat actor may adapt as well. Historical observations can therefore come from several materially different operating environments.

This is a form of non-stationarity. The useful question becomes less 'how often did exactly this event happen before?' and more 'what measurable risk conditions preceded it, what changed afterwards, and how comparable are those conditions with the present state?' That framing preserves the value of history while recognizing that the mechanism generating future events evolves.

3.3 Clean history and emerging exposure

A clean loss record also permits several explanations. Strong controls are one possibility. Luck, non-detection, limited exposure, recent business growth or a changing threat environment are others. Emerging technologies and new products create the same challenge in reverse: the absence of a long internal loss history can coexist with a material current exposure.

No-loss history is evidence about observed outcomes. Exposure measurement asks a different question: what risk conditions exist today?

3.4 How loss data enters Risk Accounting

Hughes's longer-term valuation proposal uses historical outcomes differently from a conventional loss-distribution model. Residual RUs and RMI provide independently measured context about the operating state; associated loss experience then becomes evidence for statistical correlation, back-testing and eventual derivation of RU_m . [2, pp. 136-141]

The sequence matters. Current exposure is measured independently from activity, scale and mitigation. Historical losses then test whether those measured risk states discriminate meaningfully between different outcomes and, if that relationship proves stable enough, help estimate the monetary meaning of residual RU. This resolves the apparent tension: loss history can be incomplete as a stand-alone exposure measure while remaining indispensable as validation and calibration evidence.

That calibration will itself need to respond to structural change. Different products, risk classes, jurisdictions and periods may require segmentation or re-estimation. The empirical programme therefore needs to test both stability and drift rather than assume a universal conversion in advance.

4. From Loss Categories to Risk Origins

The category in which a loss is ultimately reported describes the financial form of the consequence. Risk taxonomies are useful for classification; causal analysis can then travel further upstream to identify the operating conditions that created the exposure.

4.1 Credit losses can contain operational causes

A loan default is appropriately recorded as a credit loss, yet the quality of the exposure can be shaped by decisions made at origination. Growth targets can intensify sales pressure. Underwriting standards can be relaxed. Independent challenge can weaken. Model overrides can become routine. Documentation quality can fall. The eventual loss retains its credit classification while part of the causal chain lies in governance, policy, process, people and control conditions.

This distinction matters for historical modelling. Defaults observed in one period were produced under the lending standards, approval authorities, incentives, models and economic environment of that period. Later portfolios may operate under materially different conditions. Credit-loss history remains relevant, but the institution gains additional information when it preserves the operating state under which the exposure was accepted.

4.2 The global financial crisis as a causal-layer example

Hughes used the global financial crisis to argue that losses conventionally seen through credit, market and liquidity categories also had substantial non-financial-risk causes upstream. His point is strongest when read as causal layering rather than

reclassification. Mortgage defaults remained credit events; market and liquidity stress remained financial risks. The preceding chain also involved underwriting standards, securitization incentives, model assumptions, product complexity, ratings, due diligence, governance and operational capacity.[2, pp. 44-49][19][22]

The published RU study by Hughes and Marzouk restated public data for fifteen large US banks and reported a pronounced increase in measured operational exposure in the years surrounding the crisis, associated with rising throughput and concentration in complex subprime products.[3] The study establishes prior empirical engagement with the RU proposition; broader validation remains part of the research agenda developed later in this paper.

4.3 Sales incentives: a clearer organizational case

Wells Fargo offers a simpler causal chain. US enforcement findings linked unauthorized accounts to sales targets and compensation incentives.[20] The financial consequences arrived after the operating environment had already created the pressure. Hughes's own Selling Operations EUF framework explicitly treats products linked to sales incentive schemes and bundled products such as Loan + Insurance or Loan + Interest Rate Swap as additional sources of uncertainty.[2, p. 160]

The lesson is general: commercial success can increase exposure when the mechanisms used to generate it alter product complexity, decision quality or control effectiveness. Measuring that condition when business is written preserves information that a later loss classification by itself leaves unavailable.

A more recent illustration comes from Credit Suisse. Butler and Brooks use the bank's 2023 rescue to frame a sequence of governance, misconduct and operational-risk events that preceded the final financial consequences.[25] The point here is causal layering rather than reclassification: the ultimate financial outcome can remain correctly classified while earlier control, conduct and governance conditions help explain how the exposure developed.

5. Connecting Decisions, Exposures and Outcomes

The timing correction is crucial. The risk state that matters most may exist when the organization accepts the exposure rather than when the loss finally crystallizes. A loan can be originated under relaxed underwriting standards and default years later. A mis-sold product can generate revenue immediately and redress costs much later. A fraud can continue for months before discovery.

5.1 The transaction as a carrier of risk information

Hughes designed Risk Accounting around transaction-level and product-level information. His method envisages Risk Units being associated with transactions at capture and financial and risk values travelling through common processing and reporting structures.[2, pp. 127-128, 136-142] The conceptual result is a persistent link between the economic activity and the risk conditions under which it entered the organization.

5.2 A risk-provenance record

This paper uses the term risk-provenance record for the chronology that follows. The expression is introduced here as a practical implication of Hughes's transaction-oriented architecture rather than as terminology attributed to him.

Risk-Provenance Record

Tracking a single exposure from origination to remediation – with the original risk state preserved.




Figure 3. A risk-provenance record preserves the original risk state from acceptance to remediation.

Stage	Information preserved
1. Exposure creation	When was the transaction, product, relationship or activity accepted? What inherent RU, RMI, residual RU, control state and appetite position applied?
2. Exposure evolution	How did volume, product characteristics, controls, policies, systems or counterparties change while the exposure remained live?
3. Event occurrence	When did the adverse event happen or begin?
4. Discovery	When did the organization become aware of it?
5. Accounting recognition	When did the event produce a loss, reserve or provision in the accounts?
6. Remediation	What changed after the event, and how did the measured risk state respond?

This chronology now has an interesting regulatory parallel. CRR3 requires relevant EU institutions' operational-loss datasets to capture the date of occurrence, date of discovery and date of accounting, together with descriptive information on drivers or causes.[12, Art. 317] Risk provenance adds an earlier layer: the state at exposure creation, plus subsequent changes in the risk state before the event becomes visible.

5.3 Turning outcomes into tests of the measurement

Suppose two transactions eventually produce the same monetary loss. One was originated under high inherent exposure and weak mitigation; the other under a strong measured control environment and low residual exposure. The accounting consequence can be identical while the measurement implication differs sharply. The first observation may support the ranking; the second should trigger investigation of evidence quality, omitted factors, control circumvention, changed conditions or tail behavior.

This turns historical outcomes into model-risk information. The cycle becomes: measure when exposure is accepted; preserve changes; observe outcomes; compare outcomes with prior measured states; improve mitigation or the measurement framework; measure again. Past outcomes then test whether earlier risk conditions were understood and represented meaningfully.

Accounting tells us when the financial consequence becomes recognizable. Risk provenance can also preserve the conditions under which the organization accepted the exposure that eventually produced it.

6. Measuring the Missing Present

Risk Accounting's central proposition becomes testable only when the measurement mechanics are visible. Hughes's published banking method uses a compact chain: activity uncertainty and scale create inherent RUs; benchmark-based assessments create an RMI; the remaining portion becomes residual RUs. Appendix A reproduces a published Commercial Loans example so readers can inspect the arithmetic directly.[2, pp. 136-169]

6.1 Inherent exposure: uncertainty and scale

For each product p and principal risk type r , the published method assigns an Exposure Uncertainty Factor and one Value Band Weighting selected by reference to the product's daily new business. The risk-type-specific inherent exposure is:

$$I(p,r) = EUF(p,r) \times VBW(p)$$

The total inherent exposure for a product is the sum of its applicable principal-risk-type components:

$$I(p) = \sum_r I(p,r) = VBW(p) \times \sum_r EUF(p,r)$$

The value table deliberately converts monetary throughput into a weighting rather than multiplying RUs linearly by transaction value. In the worked example, Commercial Loans carry daily new business of \$147 million, a VBW of 33.1 and total applicable EUFs of 32, producing 1,059.2 inherent RUs before display rounding.[2, pp. 148-163]

6.2 Mitigation effectiveness: E-RCSA to RMI

The E-RCSA scores business components against defined best-practice benchmarks. Category scores are combined with category weightings and the relevant inherent RU to produce an Actual Score Aggregate (ASA). A Maximum Score Aggregate (MSA) uses the same structure while assuming a score of 100 for every applicable category. ASA and MSA are therefore weighted score aggregates; Hughes also uses their additive properties as control totals for checking the arithmetic integrity of Risk Accounting reports.[24, p. 132][25, pp. 241-242] At a reporting scope h :

$$ASA(h) = \sum_j \sum_k [s(j,k) \times w(j,k) \times I(j)]$$

$$MSA(h) = \sum_j \sum_k [100 \times w(j,k) \times I(j)]$$

$$RMI(h) = 100 \times ASA(h) / MSA(h)$$

Here $s(j,k)$ is the observed score for category k on exposure component j , $w(j,k)$ is the category weighting, and $I(j)$ is the associated inherent RU. Hughes's published calculations aggregate the weighted score aggregates through the reporting hierarchy and recalculate RMI at each reportable level.[25, pp. 241-242]

6.3 Residual exposure

Residual exposure is the portion of inherent RU left after the measured effect of mitigation:

$$R(h) = I(h) \times [1 - RMI(h) / 100]$$

The calculation therefore distinguishes growth in business from deterioration in controls. A business can become more exposed because daily-new-business scale increases while RMI remains unchanged. It can also reduce residual exposure by strengthening mitigation while the underlying activity remains similar.

The upper boundary of the RMI also has a specific meaning in Hughes's architecture. For transaction processing, he describes the theoretical reference state as 100%

straight-through processing (STP), totally reliable and secure core applications, and flawless data. If the interaction of automated processing, manual processes and data were flawless, measured processing exposure would reach the model's zero-residual boundary. Hughes explicitly describes this benchmark as hypothetical and unattainable in practice; its purpose is to provide an ideal reference point against which mitigation effectiveness can be measured.[25, p. 235]

The same interpretation applies to an RMI of 100 in the residual-risk equation: it is the measurement architecture's ideal benchmark of complete mitigation against the defined best-practice standard. Real operating environments remain exposed to imperfections, change and failure, which is why actual observations and subsequent losses remain essential to validation and calibration.

6.4 What additivity means

Hughes describes the RU as a common, additive metric.[2, p. 136] The published hierarchy mechanics make the intended meaning more precise. Inherent RUs are additive. ASA and MSA are weighted score aggregates whose additive properties also allow them to function as control totals. At a higher reporting level, those aggregates are combined, RMI is recalculated from aggregate ASA and MSA, and residual exposure is then recalculated from aggregate inherent RU and aggregate RMI.[24, p. 132][25, pp. 241-242]

$$I(H) = \sum_j I(j) \quad ASA(H) = \sum_j ASA(j) \quad MSA(H) = \sum_j MSA(j) \\ RMI(H) = 100 \times ASA(H)/MSA(H) \quad R(H) = I(H) \times [1 - RMI(H)/100]$$

Accordingly, additivity describes a shared aggregation architecture rather than a rule that every pre-calculated child residual figure is simply summed. The distinction matters because different child exposures can carry different RMIs, and the aggregate RMI is exposure-weighted through ASA/MSA.

A second, stronger question concerns economic commensurability. The method establishes a common RU index by construction. Whether one residual RU ultimately carries sufficiently stable monetary meaning across processing, conduct, cyber, lending-related operational exposures and different institutions is an empirical claim. Hughes's RU_m programme is therefore more than a valuation convenience: it is a direct test of the economic meaning of the common unit.

6.5 From residual RU to monetary expected loss

Hughes's longer-term proposal is to derive the monetary value of one RU, denoted RU_m, through statistical correlation and back-testing of product-level residual RUs and associated loss history.[2, p. 141] Once an empirical RU_m is available for an appropriate calibration domain, expected economic loss can be expressed as:

$$\text{Expected non-financial-risk loss}(h) = R(h) \times RU_m$$

The formula is conceptually simple; the calibration burden is substantial. Stable monetary valuation requires longitudinal data, clear linkage between exposure states and outcomes, segmentation where relationships differ, out-of-sample testing and periodic recalibration. The accounting ambition therefore rests on evidence developed after the exposure measure itself has been established.

6.6 Professional scrutiny and refinement of the method

The first published Risk Accounting papers were followed in 2016 by an unusual form of external scrutiny. After interest during peer review, JRMFI invited six senior academics and practitioners to comment specifically on whether the RU could provide a sufficiently risk-sensitive common metric for expressing and aggregating

risk.[4] Their responses were broadly supportive, but they also identified questions that remain useful for interpreting the method.

Julian Williams viewed the RU as a practical abstraction that could aggregate more atomic organizational risk features and described the overall approach as a sizeable step toward a workable enterprise-risk framework. His quantitative discussion also placed an important boundary around additivity: a common unit can be useful for management aggregation without implying that every interaction relevant to asset pricing or portfolio dynamics disappears.[4, pp. 413-416]

Roger Chen welcomed the common measurement architecture but challenged early wording that mixed the ideas of exposure and probability. His example suggested that ten residual RUs are more naturally interpreted as materially greater loss potential or exposure than as a probability ten times larger. Later published descriptions of the method use the cleaner formulation adopted in this paper: residual RUs represent the amount of non-financial-risk exposure remaining after mitigation - the actual exposure carried by the organization.[4, p. 417][2, p. 136]

Madelyn Antoncic emphasized the value of standardized comparability while calling for a clearer path from expert judgement toward increasingly objective calibration. She also stressed that metrics operate alongside strong governance and risk culture rather than replacing them. Adam Litke similarly clarified the scope of the RU: it measures the non-financial or operational exposure associated with underlying financial activities, rather than substituting for the financial-risk measure itself. Mark Abbott and Robert Mark highlighted the potential gains in comparability and transparent risk-adjusted returns.[4, pp. 418-420]

These comments are useful because they show the method developing under professional challenge rather than emerging fully formed. The present paper follows the later Hughes formulation where RUs express exposure, keeps financial-risk and associated non-financial-risk quantities distinct, and treats calibration, governance and economic commensurability as empirical questions rather than assumptions.

7. From Loss Modelling to Exposure Measurement

Any new quantitative operational-risk framework enters a field with important history. Basel II's Advanced Measurement Approaches allowed banks to use internal models to determine operational-risk capital. The Basel Committee later concluded that the AMA's complexity and the lack of comparability across internal modelling practices had amplified variability in risk-weighted assets, and it proposed removal of the AMA in 2016. The 2017 Basel III reforms replaced the earlier operational-risk approaches with a single standardized framework.[10][11]

7.1 A different measurement target

The most useful comparison begins with the quantity being estimated. AMA models were designed for prudential capital and, under Basel II, were associated with a 99.9% confidence standard for annual operational loss.[8] Risk Accounting begins with a different object: current accepted exposure expressed in RUs. Monetary expected-loss estimation enters later through RU_m . This difference changes the data, assumptions and validation questions, while leaving a substantial burden of proof on both approaches.

Dimension	Advanced Measurement Approaches	Risk Accounting
Primary purpose	Regulatory capital requirement	Present-state management measurement; later economic valuation
Primary quantity	Extreme annual operational-loss / capital measure	Inherent RU, RMI and residual RU; later expected monetary loss

Dimension	Advanced Measurement Approaches	Risk Accounting
Core information	Internal/external loss data, scenarios, BEICFs and modelling choices	Current activity characteristics, daily-new-business scale, product EUFs and assessed mitigation state
Role of control environment	Included through BEICFs and related adjustments	Measured explicitly through E-RCSA, ASA/MSA and RMI
Aggregation challenge	Tail distributions, dependence and model-choice variability	Common-unit calibration, hierarchy rules and empirical commensurability
Current consequence	Historically determined Pillar 1 capital under AMA	Management/risk information; regulatory or accounting use remains an evidence-based future question

7.2 BEICFs: important prior art

Basel's AMA framework already recognized that operational-risk assessment requires information about the current business environment and internal controls. The Committee described Business Environment and Internal Control Factors as forward-looking factors capable of reflecting growth, employee turnover, new products, audit results, system downtime and control effectiveness.[9]

That lineage matters. BEICFs established an important supervisory precedent: operating condition and control effectiveness belong in forward-looking operational-risk analysis. Risk Accounting gives those conditions a more central and standardized measurement role. In AMA practice, BEICFs were combined with other elements in widely varying ways, and the Basel Committee documented significant differences across institutions.[9] Hughes's framework separates exposure from mitigation and places the mitigation measure directly into the residual-risk calculation.

7.3 The current standardized capital context

The standardized capital framework answers a prudential question with a deliberately simpler architecture. In the European Union, CRR3 - applicable from 1 January 2025 - makes the operational-risk Pillar 1 own-funds requirement equal to the Business Indicator Component.[12, Arts. 312-315] The regulation's recitals explain that historical operational losses are disregarded in the minimum own-funds calculation in the EU implementation.

Loss information remains important elsewhere in the framework. Institutions above the relevant business-indicator threshold must calculate annual operational losses and maintain detailed loss datasets, subject to quality controls, independent review and supervisory scrutiny.[12, Arts. 316-323] This distinction is important: the EU capital formula emphasizes comparability and simplicity, while management and supervisory processes continue to require detailed operational-risk information.

7.4 Richer event chronology and operational resilience

CRR3's loss-data requirements are particularly relevant to the temporal argument in this paper. Article 317 requires the capture, where available, of the date of occurrence, date of discovery and date of accounting, as well as drivers or causes of operational-risk events. The regulation also addresses operational events related to credit risk and market risk, showing that classification boundaries and causal relationships require explicit treatment.[12]

DORA adds another contemporary layer. Applicable since 17 January 2025, it requires financial entities to build, assure and review digital operational resilience and to manage ICT risk across systems and third-party dependencies.[13] DORA provides a rigorous resilience framework rather than a common enterprise RU-like exposure denominator. Its relevance here is the direction of travel: regulators increasingly expect institutions to understand the operating conditions and dependencies from which loss can emerge before a major event occurs.

7.5 Relationship to the existing operational-risk toolkit

Approach	Primary role	Risk Accounting contribution
Loss database	Records realized events and monetary outcomes	Adds exposure and mitigation context to the event history
Scenario analysis	Explores severe but plausible futures	Provides a present-state baseline for scenario selection and challenge
KRIs	Detects movement in selected indicators	Connects indicators with measured exposure or mitigation where appropriate
Conventional RCSA	Assesses risks and controls, often qualitatively or ordinally	Uses benchmark scoring and a separate exposure measure to derive RMI and residual RU
Operational resilience	Tests continuity of important services under disruption	Provides a complementary measure of conditions that can make disruption more or less likely or severe
Standardized operational-risk capital	Provides a comparable prudential capital measure	Addresses a different management question: current control-adjusted accepted exposure

8. What Measurable Non-Financial Risk Changes for Management

The first benefits of a common exposure measure belong to management rather than statutory reporting. Once residual risk can be expressed consistently, organizations can bring risk into the same decisions that already use cost, revenue, capital and volume.

8.1 Risk-adjusted profitability and product pricing

Two business units can report the same accounting profit while carrying very different residual non-financial-risk exposures. Appendix A provides a concrete baseline from Hughes's published Commercial Loans example: approximately 1,059 inherent RUs, an RMI of 68.0% and 339 residual RUs.[2, pp. 168-169][23] An internal economic view can therefore show accounting contribution alongside measured residual RU and, once RU_m is empirically established, expected non-financial-risk cost. This is management information rather than an IFRS profit measure, but it changes how growth and product economics can be assessed.

Product pricing follows the same logic. Management accounting already allocates direct costs, overhead, funding and capital. Hughes argued that a monetary risk cost could eventually become another input to economic pricing.[2] A product generating attractive accounting margin while consuming unusually high residual RU may be economically less attractive than the income statement alone suggests.

8.2 Measuring the return on mitigation

Control expenditure is often defended through qualitative statements about compliance, resilience or reduced risk. RUs create an intermediate quantitative output: management can compare residual exposure before and after a control intervention. Using the published Commercial Loans baseline as an illustration, holding inherent exposure at 1,059 RUs while a hypothetical control programme raises RMI from 68% to 78% would reduce residual exposure from about 339 RUs to about 233 RUs - a measured reduction of roughly 106 RUs. The 78% scenario is developed here for illustration; the 1,059/68%/339 baseline comes from Hughes's published example.

Once a validated RU_m exists, that 106-RU reduction could also be translated into an estimated economic benefit and compared with implementation cost. Until then, the RU reduction remains useful on its own because it makes the measured effect of mitigation visible in the same unit used to describe exposure.

8.3 Accountability, incentives and behavioural change

Risk ownership becomes more meaningful when exposure can be attributed to the activities that create it. Business units receive revenue today while risk, compliance, operations, legal and technology functions can absorb much of the cost of controlling that business. A common exposure measure can bring the risk consumed by a commercial decision into performance discussions alongside the financial output and make accountability more closely follow the decision that created the exposure.

Risk culture is commonly pursued through governance, policies, training, codes of conduct, tone from the top and supervisory expectations. Risk Accounting adds a measurement-based control layer to those mechanisms. Hughes used expense management as an analogy: budgeting influences behaviour because expenditure is measured, reported and compared with approved limits; he argued that risk culture benefits from the same discipline of explicit measurement and accountability.[2, pp. 131-135] In the 2016 Risk Accounting papers, Grody and Hughes went further and proposed benchmarking RMIs within and between banks as a basis for monitoring and proactively promoting improvements in risk culture.[25, pp. 246-247]

Commercial incentive systems already communicate powerful signals about what an organization values: revenue, volume, margin, market share and growth. When those rewards are visible immediately while the associated non-financial-risk exposure remains diffuse or deferred, increased risk-taking can continue to look economically attractive. Placing inherent RU, RMI and residual RU alongside the financial return allows management to see both sides of the decision at the same time. This is the practical bridge from risk culture as an aspiration to risk-aware behaviour as an observable management choice.

The same information can shape product strategy. Management can grow products whose returns remain attractive relative to residual exposure; strengthen controls where mitigation materially improves the risk-adjusted economics; reprice products whose risk burden is inadequately reflected in margin; redesign products whose structure creates excessive uncertainty; constrain growth where volume is outrunning mitigation capability; and retire or replace products whose risk-adjusted contribution no longer supports their continued use. The objective is better risk selection: take risk where the expected return justifies the exposure, improve mitigation where it changes the economics, and direct resources toward products and activities that create more durable value.

From risk culture aspirations to measurable incentives.




Figure 4. From risk culture aspirations to measurable incentives.

Risk culture becomes materially more actionable when the economic reward for taking risk and the measurable exposure created in earning that reward are visible in the same decision framework.

The same incentive architecture requires strong measurement governance. If management evaluation depends on residual RU while RMI relies partly on assessment evidence supplied from the organization, favourable scoring can itself become economically valuable. A credible framework therefore separates assessment ownership from validation, requires evidence for material scores, provides independent challenge and internal-audit access, and back-tests measured states against observable incidents, control-test results and loss outcomes. The accounting analogy is useful here: credibility comes from the control environment around the number as well as from the formula.

8.4 Risk appetite, tolerance and capacity

Risk appetite becomes more operational when approved limits and actual exposure can be expressed on a common scale. Hughes explicitly proposed using RUs to set and monitor risk appetite.[2] The paper uses the contemporary governance concepts cautiously: risk capacity describes the maximum risk the organization can bear within its financial, regulatory and control constraints; appetite describes the risk it deliberately chooses to assume; tolerance or operating limits can govern local deviations and escalation.

RUs can provide the quantitative non-financial-risk spine across those concepts. For example, a board can approve an enterprise residual-RU appetite and allocate portions to business lines or products, while retaining separate financial and qualitative constraints. Capacity remains broader than an RU number because capital, liquidity, regulatory constraints and management capability also matter.

8.5 Risk budgeting and strategic planning

A business plan typically projects revenue, cost, capital and volume. Risk Accounting creates the possibility of adding a risk budget: planned business volume implies expected inherent RUs; planned control investments imply expected mitigation; residual RU can then be compared with approved appetite. Strategy becomes a linked financial-and-risk plan rather than two documents prepared in parallel.

How much revenue do we expect to generate - and how much residual risk do we expect to carry in generating it?

9. From Management Accounting to External Reporting

The accounting ambition can progress in stages. Management accounting can use relevant internal measures before those measures qualify for recognition in statutory accounts. External disclosure can also develop before a general balance-sheet treatment exists.

9.1 A management risk ledger

A residual-RU balance can be reconciled over time in a manner analogous to a management ledger. Opening exposure is adjusted for new activity, changing scale or uncertainty, mitigation improvements or deterioration, retired exposures and controlled methodology changes. The objective is movement analysis: management sees both the closing position and the reasons it changed. A simple one-factor illustration can use the same Commercial Loans baseline developed above:

Illustrative management reconciliation	Residual RU
Opening position: published Commercial Loans example (RMI 68%)	339

Mitigation improvement: hypothetical RMI 68% → 78%, inherent RU unchanged	-106
Closing position after mitigation improvement	233

Illustrative movement only. The opening state is drawn from Hughes's published Commercial Loans example; the 78% RMI closing state is a paper-developed management scenario.

Once such reconciliations become material to decisions, the decomposition convention itself becomes part of the accounting design. Interacting volume, product-mix and mitigation effects can be allocated using a disclosed convention, reported with an interaction residual, or handled through another stable variance-analysis approach. The important property is consistency and auditability rather than a claim that one decomposition is naturally unique.

9.2 Risk-adjusted management performance

If RU_m calibration becomes sufficiently robust, an internal performance statement could show accounting contribution and a separate expected non-financial-risk economic cost. The two measures would serve different purposes: statutory profit follows the accounting standards; risk-adjusted economic performance provides management with an additional view of the risk cost accepted to generate that profit.

9.3 Disclosure can precede recognition

External reporting has a wide space between silence and balance-sheet recognition. Companies already disclose management measures, principal risks, segment information and narrative explanations about future prospects. IFRS 8's management approach to operating segments illustrates how information used internally by the chief operating decision maker can become relevant to external segment reporting.[17] The revised IFRS Practice Statement 1 on Management Commentary, issued in June 2025, similarly emphasizes management's perspective, company-specific metrics, risks, prospects and connected information.[18]

A mature Risk Accounting framework could therefore first support governed disclosure of inherent RU, RMI, residual RU, major concentrations and movement in exposure, provided the methodology, limitations and assurance are explained clearly. Such disclosure would keep management measures distinct from recognized IFRS amounts while giving investors information about the quality and sustainability of performance.

9.4 The recognition boundary under current IFRS

The statutory recognition question requires its own conceptual basis. IFRS 9 contains expected-credit-loss requirements within the financial-instruments framework.[14] IAS 37 recognizes provisions around present obligations arising from past events; expected future operating losses that depend on future operations remain outside that recognition basis.[15]

The accounting pathway for a general expected non-financial-risk measure therefore points toward new standard-setting work. Such work would need to define the relevant accounting object, recognition trigger, measurement basis, presentation, treatment of reversals and changes in estimates, disclosure requirements, interaction with regulatory capital and the assurance model. A monetary RU measure would provide an input to that debate rather than decide those questions by itself.

9.5 An evidence-based progression to standard setting

Stage	Evidence produced
1. Measurement	Define and reproduce inherent RU, RMI and residual RU.
2. Management use	Apply the measures to pricing, control investment, appetite, budgeting and performance.
3. Longitudinal evidence	Link risk states at exposure creation and through time with later outcomes.
4. Monetary calibration	Test and estimate RU_m within appropriate calibration domains.
5. Assurance	Establish governance, comparability, traceability and independent validation.

Stage	Evidence produced
6. External disclosure	Test decision-usefulness for investors and other stakeholders.
7. Standard-setting dialogue	Bring empirical measurement and reporting evidence to bodies such as the IASB and FASB.

Tax, dividend and remuneration effects sit downstream from this sequence. Their treatment depends on the accounting, corporate-law, contractual and tax rules that would apply to any future recognized amount; they are consequences of a reporting model rather than the foundation of the measurement argument.

10. Evidence So Far: Validation, Governance and Falsifiability

Risk Accounting is no longer an entirely untested proposition, but the different forms of evidence need to be separated carefully. Published expert commentary has tested the conceptual architecture; the 2021 Hughes-Marzouk study provided an initial empirical test of inherent RU predictiveness; Butler and Brooks provided external academic endorsement and an implementation-oriented extension of the method; and Hughes published transparent worked calculations that can be mechanically reproduced. The next phase is prospective, independently replicated evidence on residual exposure, cross-institution comparability and monetary calibration.

10.1 Initial empirical evidence: Hughes and Marzouk (2021)

Hughes and Marzouk published the first reported test of the RU's inherent predictiveness in 2021.[3] The study restated publicly available US bank holding-company financial data into inherent operational-risk RUs and compared the resulting trend with net income before taxes (NIBT) over the period leading into the 2007-08 financial crisis. The exceptional decline in NIBT was used as a proxy for unexpected losses. The test covered fifteen of the largest US banks; banks undergoing significant reorganizations during the test period were excluded to preserve time-series comparability.[3]

The study reported a 58.1% increase in operational-risk exposure over the eight-year period, attributing the buildup to sharply increased throughput concentrated in more complex products - the combination of higher VBWs and higher EUFs that the method is designed to capture.[3] The authors concluded that the result provided an initial indication of inherent predictiveness and argued that a buildup of this magnitude would have supplied a clear risk signal before the subsequent collapse in profitability.

The paper also states its limitations directly. Public accounting data required extensive proxies and assumptions; the test was retrospective; NIBT was a proxy for unexpected loss; and the exercise tested inherent RU rather than live residual exposure produced from contemporaneous E-RCSA data. Hughes and Marzouk therefore described the result as initial rather than conclusive and identified simulations using direct bank inputs as the next research phase.[3] The present author contributed to downloading, formatting and restating the FR Y-9C data in RUs and to producing the analytics used in that study, as acknowledged in the published paper.[3]

10.2 External academic endorsement and extension: Butler and Brooks (2024)

Tom Butler and Robert Brooks published a peer-reviewed critique of the prevailing operational-risk paradigm in Risk Analysis and presented Risk Accounting as a candidate architecture for the paradigm change they argued was required.[25] The

article first appeared online in 2023 and is published in the 2024 volume of the journal. Its authors were unaffiliated with the original development of Risk Accounting, but the paper is best understood as an external academic endorsement and extension rather than an arms-length validation study: it adopts the method as part of the solution it advances.

Butler and Brooks assessed Risk Accounting through publications, presentations and demonstrations during 2021-2023. They described the Calculation Engine and knowledge base as highly transparent and auditable, highlighted granular RU-based exposure information, aggregation, RMI, risk appetite in RUs and stronger first-line participation, and proposed a wider AI-enabled digital architecture around the method.[25] Their work therefore provides external support for the problem diagnosis, the transparency of the mechanics and the plausibility of an operational information architecture.

Their paper also uses stronger language on monetary valuation than the evidence threshold adopted here, describing financial valuation of granular risk exposure as an enabled capability.[25] This whitepaper follows the more conservative sequencing in Hughes's own published method: RU exposure measurement and residual-RU calculation are current mechanics, while a robust monetary value of the RU (RU_m) remains an empirical objective requiring longitudinal calibration against associated realized losses.[2, p. 141][3] The difference is worth stating because endorsement and empirical validation perform different roles in the evidence base.

Published evidence	Primary contribution	What it already supports	Next evidence that would add most value
2016 JRMFI expert commentary	Six invited academic/practitioner reviews	Conceptual scrutiny of additivity, RU interpretation, comparability, governance and implementation	Empirical validation and implementation evidence
2021 Hughes & Marzouk study	Retrospective test using fifteen large US banks	Initial evidence that inherent RU buildup preceded material deterioration; 58.1% increase over the eight-year test period	Prospective direct-input tests of residual RU, independent replication and richer outcome linkage
2023 Butler & Brooks assessment	Independent academic/practitioner analysis	Support for transparency, auditability, granular aggregation, first-line relevance and digital implementation	Live implementation evidence and independent longitudinal outcome testing

10.3 Mechanical reproducibility

The first test is basic but essential. Two suitably trained practitioners applying the same published rules to the same inputs should obtain materially the same result. EUF assignment, VBW selection, E-RCSA applicability, category weighting, ASA/MSA calculation, hierarchy aggregation and rounding conventions therefore need documented rules and controlled versions. Appendix A demonstrates the arithmetic for one published example; broader reproducibility requires a reference corpus across products and hierarchy levels.

10.4 Responsiveness

The measure should respond in an intelligible direction when the underlying condition changes. Higher daily-new-business scale should increase inherent exposure under the value table. Changes in product uncertainty should alter EUF-driven inherent RU. Stronger mitigation should increase RMI and reduce residual RU, while weaker mitigation should move in the opposite direction. Sensitivity testing can examine these properties before any monetary loss calibration exists.

10.5 Discrimination

Across sufficiently large groups of comparable exposures, materially different residual-risk states should be associated with meaningfully different subsequent outcome patterns or independently observable control-failure indicators. The 2021

study provides an initial result at the inherent-RU level: rising inherent exposure preceded the crisis-period deterioration in its fifteen-bank sample.[3] The stronger test now required is prospective discrimination using contemporaneous residual RU, mitigation evidence and event outcomes. Successful discrimination means that the metric rank-orders or separates risk states in a way that carries useful information, while individual high-RU exposures can still resolve without loss.

A persistent absence of useful separation across appropriately designed tests would require revision of the measurement claim, the factors, the scoring architecture or the scope within which comparability is asserted. This is a core falsifiability condition: a common unit earns credibility by surviving evidence that could have contradicted it.

10.6 Monetary calibration and stability

RU_m carries an additional burden. Hughes's published proposition is to derive monetary value through statistical correlation of residual RUs with associated realized losses and to use back-testing to refine the exposure-measurement system.[2, p. 141][3] A usable monetary conversion should therefore be estimated out of sample, tested for bias and drift, and segmented wherever materially different products or risk types show different relationships between residual RU and loss. The theoretical $RMI=100$ boundary remains the ideal reference state defined by the measurement architecture rather than an empirical assertion that a real operating environment will achieve it.

As a robustness diagnostic for the valuation programme - rather than as a modification of Hughes's residual-risk definition - empirical work can compare the published proportional relationship with an unconstrained specification such as $Loss = \alpha + \beta R$. A material intercept would signal a systematic loss component not explained by measured residual RU within the calibration domain and would trigger investigation of segmentation, omitted factors or model specification. Either result would inform calibration while leaving the published RU and RMI mechanics unchanged.

10.7 Comparability and economic commensurability

The RU can operate as a standardized index before full monetary commensurability has been proven. Cross-institution comparison, however, requires common definitions, factor tables, assessment content and governance. The strongest version of the Risk Accounting proposition would show that equivalent exposures measured by different institutions produce sufficiently consistent RUs and that residual RUs carry a stable economic relationship to outcomes. If only part of that claim holds, the method's use should be scoped accordingly.

10.8 Manipulation resistance and independent challenge

RMI incorporates assessment evidence and therefore deserves controls comparable to other decision-critical measures. A strong governance model would include evidence requirements for material answers, defined assessment ownership, independent second-line challenge, audit access, change control for questionnaires and weights, versioned historical states, validation against control testing and incidents, and clear escalation for overrides or unusual scores.

The same principle appears in contemporary regulation. CRR3 requires loss-data quality processes, independent review, an independent operational-risk management function, regular monitoring and audit review for the operational-risk framework.[12, Arts. 322-323] Risk Accounting can adopt comparable governance disciplines even where the measure is used purely for management purposes.

10.9 Implementation evidence and proportionality

Computability has already been demonstrated in software. Grody and Hughes reported in 2016 that the method had been replicated in software producing explicit, dynamic and aggregatable risk measurements.[24] Hughes later described the Calculation Engine as capable of operating standalone, integrating with existing information infrastructure, being embedded in third-party platforms or being developed in-house.[2, p. 142] Hughes and Marzouk reported prototype software supporting ongoing research and testing, while Butler and Brooks examined Calculation Engine demonstrations and proposed a wider digital architecture around the method.[3][25] These implementations establish that the mechanics can be operationalized; they do not by themselves establish enterprise-scale proportionality. Pilot deployments should therefore record data requirements, assessment effort, integration cost, operational burden, user consistency, audit effort, processing performance and management decisions changed by the measurement.

10.10 A research programme with explicit success conditions

Test	Evidence that would strengthen the claim
Responsiveness	Outputs move coherently when scale, uncertainty or mitigation changes.
Discrimination	Different measured risk states show useful differences in later outcomes or control-failure indicators.
Calibration	Monetary expected-loss estimates remain reasonably calibrated out of sample within defined domains.
Comparability	Equivalent exposures produce sufficiently consistent results across assessors and, ultimately, institutions.
Governance	Evidence, overrides, changes and calculations remain traceable and independently challengeable.
Proportionality	Decision value justifies the data, assessment and systems burden at the intended scope.

The result is an unusually useful position for an emerging discipline: the whitepaper can state what evidence would strengthen Risk Accounting and what observations would narrow or revise its claims. That is the difference between a measurement research programme and a purely persuasive narrative.

11. Why We Call It Risk Accounting

The word Accounting captures the method's central ambition: to give accepted risk a common measurement structure with the properties that make accounting useful - aggregation, reconciliation, comparability, traceability, control and auditability.

The risk profession already has sophisticated tools for the past and the future. Historical databases show what went wrong. Scenario, stress and resilience approaches explore what could happen under alternative conditions. Risk Accounting focuses first on the exposure created by decisions already made and the quality of the mitigation currently protecting those decisions.

That present-state measure strengthens both other perspectives. Historical outcomes become richer when linked to the conditions under which exposure was accepted. Future scenarios become more disciplined when they begin from a quantified current risk state. The regulatory framework itself increasingly records richer event chronology and operating resilience, while standardized capital rules deliberately prioritize simplicity and comparability. Risk Accounting addresses a different layer: current accepted exposure.

The accounting comparison explains why this matters. Expected-credit-loss accounting already accepts that a current financial exposure can contain an economically meaningful future loss before an individual default occurs. Peter's

challenge was to create a common exposure measure for non-financial risk and, through evidence, build a path toward expected economic loss measurement.[2][14]

The method begins by making exposure observable. It separates inherent risk from mitigation effectiveness, expresses residual exposure in a common unit, associates that information with activities and transactions, and creates the possibility of relating measured risk states to realized monetary loss. The published equations are simple enough to audit; the empirical claims surrounding calibration and commensurability are substantial enough to demand independent testing.

That capability already matters to management accounting, pricing, control investment, performance measurement, risk appetite and governance. As evidence accumulates, it also creates a reporting question for investors, regulators, auditors and standard-setters: how should reliable information about the expected economic cost of accepted non-financial risk be communicated and, where an appropriate standard-setting basis is developed, reflected in financial reporting?

Peter Hughes believed that the missing link between risk management and accounting was a common measurement framework for exposure. His life's work was an attempt to build that link. The idea has since moved through published methodology, expert commentary, an initial empirical predictiveness test and an independent academic-practitioner assessment.[2][3][4][25] The next stage is prospective implementation, replication, refinement and monetary calibration with the level of evidence that a genuine accounting discipline demands.

For readers interested in exploring the method beyond this paper, the most direct path is Hughes's detailed 2023 book and worked examples, the original two-part 2016 JRMFI description of the architecture, the 2016 expert commentary, the 2021 predictiveness study and the 2024 Butler-Brooks academic endorsement and extension.[2][23][24][4][3][25] Together they show the mechanics of Risk Accounting, the professional scrutiny applied to it and the research questions that remain open.

Measure risk before it becomes loss, and it can inform economic decisions before it becomes an accounting surprise. That is why we call it Risk Accounting.

Appendix A. A Published-Method Worked Example

This appendix uses Hughes's own fictitious Commercial Loans (Trade Receivables) example rather than a newly invented case. Its purpose is expository: to show the published mechanics from daily new business and EUFs through inherent RU, E-RCSA scoring, RMI and residual RU. Values below are drawn from the worked example in Risk Accounting and the companion Calculation Engine material.[2, pp. 144-169][23]

A.1 Product-scale inputs and inherent exposure

The published example assigns Commercial Loans daily new business of \$147 million, which maps to a Value Band Weighting of 33.1. The product carries the following EUFs:

Principal risk type	EUF	VBW	Inherent RU
Processing	10	33.1	331.0
Lending	8	33.1	264.8 (265 displayed)
Selling	2	33.1	66.2 (66 displayed)
Funding	6	33.1	198.6 (199 displayed)
Interest rate	6	33.1	198.6 (199 displayed)
Total	32	-	1,059.2 (1,059 displayed)

$$I(\text{Commercial Loans}) = 33.1 \times (10 + 8 + 2 + 6 + 6) = 1,059.2 \text{ RUs}$$

The table illustrates the method's separation of product uncertainty from scale. The same VBW applies to the product's applicable risk-type EUFs; the resulting inherent components can be aggregated to the product total.

A.2 One business-component RMI calculation

For the Processing risk associated with Commercial Loans, the published inherent RU is 331. The Product & Service Pricing business component has eight generic E-RCSA category scores and weightings. Hughes states that the worked examples are entirely fictitious and provided to illustrate the methodology; extreme scores such as 100 and 0 in this example should therefore be read as values chosen to exercise the calculation rather than as a claim about a typical real assessment.[2, disclaimer]

E-RCSA category	Score	Weight
Control (CTRL)	20	10
People (PEOPLE)	100	10
Execution (EXECTN)	0	10
Business continuity (BUSCON)	65	8
Risk control (RSKCTL)	15	6
Internal logical access (INTLAM)	20	6
Physical access (PHSACC)	40	4
Policies & procedures (POLPRO)	60	2

The Actual Score Aggregate uses each score, weight and the 331 inherent RUs:

$$ASA = (20 \times 10 \times 331) + (100 \times 10 \times 331) + (0 \times 10 \times 331) + (65 \times 8 \times 331) + (15 \times 6 \times 331) + (20 \times 6 \times 331) + (40 \times 4 \times 331) + (60 \times 2 \times 331)$$

$$ASA = 731,510$$

The Maximum Score Aggregate uses the same weights and inherent RU while setting every category score to 100:

$$MSA = 1,853,600$$

$$RMI = 100 \times 731,510 / 1,853,600 = 39.46\% \text{ (39.5 displayed)}$$

$$\text{Residual RU} = 331 \times [1 - 0.3946] = 200.37 \text{ (200 displayed)}$$

The result is intuitive in the framework's own terms: 331 inherent RUs enter this business component, an RMI of about 39.5% represents the measured mitigation effectiveness, and roughly 200 residual RUs remain.

A.3 Hierarchy aggregation

The same architecture applies upward through the reporting hierarchy. Hughes's published Commercial Loans summary shows approximately 1,059 inherent RUs, aggregate ASA of 28,956,542, aggregate MSA of 42,593,080, an RMI of 68.0% and residual exposure of 339 RUs.[2, pp. 168-169][23]

$$RMI(\text{product}) = 100 \times 28,956,542 / 42,593,080 = 67.98\% \text{ (68.0 displayed)}$$

$$\text{Residual RU}(\text{product}) \approx 1,059 \times [1 - 0.6798] \approx 339 \text{ RUs}$$

This is also why hierarchy treatment matters to the additivity discussion. The method aggregates inherent RU and the ASA/MSA weighted score aggregates, which Hughes also uses as control totals, and recalculates the higher-level RMI and residual exposure. A single business-component MSA is therefore not expected to reconcile directly to the full product MSA: the product result incorporates all applicable assessed business components and their category structures before the aggregate ratio is calculated. A higher-level residual result is an exposure-weighted recomputation rather than the simple sum of already-rounded component residual figures.[25, pp. 241-242]

A.4 The monetary step remains symbolic until calibrated

Hughes's longer-term accounting proposition begins only after an empirical monetary value per RU has been established. For the displayed Commercial Loans residual exposure, the symbolic relationship is:

$$\text{Expected non-financial-risk loss} = 339 \times RU_m$$

RU_m remains symbolic here because longitudinal evidence must establish its monetary value. The example therefore distinguishes what the published method already calculates from what the future valuation programme must establish.

A.5 What this example establishes - and what the research programme must test

The worked example establishes transparency of mechanics. A reader can reproduce the movement from product inputs through inherent RU, assessment scores, RMI and residual RU. It also shows how ASA/MSA weighted score aggregates support hierarchical aggregation, control checking and auditability.

The empirical questions remain larger. EUF and weighting calibration must discriminate meaningfully between risk states. Assessment governance must resist optimism and manipulation. Cross-product and cross-institution comparability need independent evidence. RU_m must be estimated and back-tested within domains where the relationship is sufficiently stable. Those questions belong to Section 10's validation programme rather than to the arithmetic of the example itself.

References

- [1] Hughes, Peter J. Private correspondence with the author, shortly before his death.
- [2] Hughes, Peter J. (2023). Risk Accounting: The Complete Guide to Quantifying and Accounting for Non-Financial Risks. Grosvenor House Publishing Ltd, 27 July 2023. ISBN 978-1-80381-510-7. Method overview and worked example especially pp. 28, 44–50, 123–142 and 144–169.
- [3] Hughes, Peter J. & Marzouk, Mahmoud (2021). 'A Test of the Inherent Predictiveness of the RU, a New Metric to Express All Forms of Operational Risk in Banks.' Journal of Risk Management in Financial Institutions, 14(2), 173–194.
- [4] Williams, Julian; Chen, Roger; Antoncic, Madelyn; Litke, Adam; Abbott, Mark; & Mark, Robert (2016). 'Comments on Risk Accounting.' Journal of Risk Management in Financial Institutions, 9(4), 413–420. DOI: 10.69554/CSWG1354. <https://doi.org/10.69554/CSWG1354>
- [5] Basel Committee on Banking Supervision. Basel Framework, OPE10 – Definitions and application. Operational-risk definition. Current framework accessed 14 August 2026. https://www.bis.org/basel_framework/chapter/OPE/10.htm
- [6] Office of the Comptroller of the Currency (2019). Comptroller's Handbook: Corporate and Risk Governance. Version 2.0, July 2019; current OCC booklet subsequently updated to remove references to reputation risk. See especially the risk-assessment and risk-governance sections. <https://www.occ.treas.gov/publications-and-resources/publications/comptrollers-handbook/files/corporate-risk-governance/pub-ch-corporate-risk.pdf>
- [7] Basel Committee on Banking Supervision (2021). Revisions to the Principles for the Sound Management of Operational Risk. March 2021, BCBS d515. <https://www.bis.org/bcbs/publ/d515.htm>
- [8] Basel Committee on Banking Supervision (2006). International Convergence of Capital Measurement and Capital Standards: A Revised Framework – Comprehensive Version. June 2006, BCBS 128. Includes the Advanced Measurement Approach requirements for operational risk. <https://www.bis.org/publ/bcbs128.htm>
- [9] Basel Committee on Banking Supervision (2009). Observed Range of Practice in Key Elements of Advanced Measurement Approaches (AMA). 28 July 2009, BCBS 160b. Includes internal loss data, external data, scenario analysis and Business Environment and Internal Control Factors. <https://www.bis.org/publ/bcbs160b.htm>
- [10] Basel Committee on Banking Supervision (2016). Standardised Measurement Approach for Operational Risk. Consultative Document, 4 March 2016, BCBS d355. Proposed withdrawal of the AMA, citing its inherent complexity and the lack of comparability arising from divergent internal modelling practices. <https://www.bis.org/bcbs/publ/d355.htm>
- [11] Basel Committee on Banking Supervision (2017). Basel III: Finalising Post-Crisis Reforms. December 2017, BCBS d424. Includes the operational-risk reforms replacing the earlier approaches with a single standardised framework. <https://www.bis.org/bcbs/publ/d424.htm>
- [12] European Parliament and Council (2013, as amended 2024). Regulation (EU) No 575/2013 on prudential requirements for credit institutions and investment firms, as amended by Regulation (EU) 2024/1623 of 31 May 2024 (CRR3), in particular Articles 312–323 concerning the operational-risk own-funds requirement, Business Indicator Component, operational-loss data and operational-risk governance. Consolidated text current as at 1 January 2026. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:02013R0575-20260101>
- CRR3 amending regulation: <https://eur-lex.europa.eu/eli/reg/2024/1623/oj/eng>
- [13] European Parliament and Council (2022). Regulation (EU) 2022/2554 of 14 December 2022 on digital operational resilience for the financial sector (Digital Operational Resilience Act – DORA). Applicable from 17 January 2025. <https://eur-lex.europa.eu/eli/reg/2022/2554/oj/eng>
- [14] International Accounting Standards Board. IFRS 9 Financial Instruments. IFRS Foundation. Issued July 2014 and subsequently amended; expected-credit-loss impairment requirements. Current standard accessed 14 August 2026. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-9-financial-instruments/>
- [15] International Accounting Standards Board. IAS 37 Provisions, Contingent Liabilities and Contingent Assets. IFRS Foundation. Current standard accessed 14 August 2026; includes the present-obligation recognition requirements and treatment of future operating losses. <https://www.ifrs.org/issued-standards/list-of-standards/ias-37-provisions-contingent-liabilities-and-contingent-assets/>
- [16] Financial Accounting Standards Board (2016). Accounting Standards Update No. 2016-13, Financial Instruments—Credit Losses (Topic 326): Measurement of Credit Losses on Financial Instruments. Issued 16 June 2016. Introduced the Current Expected Credit Losses (CECL) methodology and Topic 326. <https://www.fasb.org/page/document?pdf=ASU+2016-13.pdf>
- [17] International Accounting Standards Board. IFRS 8 Operating Segments. IFRS Foundation. Issued November 2006 and subsequently amended; management approach to operating-segment reporting. Current standard accessed 14 August 2026. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-8-operating-segments/>
- [18] International Accounting Standards Board (2025). IFRS Practice Statement 1 Management Commentary. Revised Practice Statement issued 23 June 2025. IFRS Foundation. <https://www.ifrs.org/supporting-implementation/supporting-materials-by-ifrs-standards/management-commentary/>
- [19] Financial Crisis Inquiry Commission (2011). The Financial Crisis Inquiry Report: Final Report of the National Commission on the Causes of the Financial and Economic Crisis in the United States. Washington, DC: U.S. Government Printing Office, January 2011. <https://www.govinfo.gov/app/details/GPO-FCIC>
- [20] Consumer Financial Protection Bureau (2016). 'Consumer Financial Protection Bureau Fines Wells Fargo \$100 Million for Widespread Illegal Practice of Secretly Opening Unauthorized Accounts.' 8 September 2016. The release identifies sales targets and compensation incentives as drivers of the unauthorized-account practices.

- <https://www.consumerfinance.gov/archive/newsroom/consumer-financial-protection-bureau-fines-wells-fargo-100-million-widespread-illegal-practice-secretly-opening-unauthorized-accounts/>
- [21] Financial Conduct Authority (2014). TR14/14: Redress for Payment Protection Insurance (PPI) Mis-Sales. Thematic Review, 28 August 2014. <https://www.fca.org.uk/publications/thematic-reviews/tr14-14-redress-payment-protection-insurance-ppi-mis-sales>
- [22] Bernanke, Ben S. (2010). 'Causes of the Recent Financial and Economic Crisis.' Statement before the Financial Crisis Inquiry Commission, Washington, DC, 2 September 2010. Board of Governors of the Federal Reserve System. <https://www.federalreserve.gov/newsevents/testimony/bernanke20100902a.htm>
- [23] Grody, Allan D. & Hughes, Peter J. (2016). 'Risk Accounting – Part 1: The Risk Data Aggregation and Risk Reporting (BCBS 239) Foundation of Enterprise Risk Management (ERM) and Risk Governance.' *Journal of Risk Management in Financial Institutions*, 9(2), 130–146. DOI: 10.69554/DWJF5684. <https://doi.org/10.69554/DWJF5684>
- [24] Grody, Allan D. & Hughes, Peter J. (2016). 'Risk Accounting – Part 2: The Risk Data Aggregation and Risk Reporting (BCBS 239) Foundation of Enterprise Risk Management (ERM) and Risk Governance.' *Journal of Risk Management in Financial Institutions*, 9(3), 224–248. DOI: 10.69554/VZHP9888. <https://doi.org/10.69554/VZHP9888>
- [25] Butler, Tom & Brooks, Robert (2024). 'Time for a Paradigm Change: Problems with the Financial Industry's Approach to Operational Risk.' *Risk Analysis*, 44(6), 1285–1304. First published online 26 October 2023. DOI: 10.1111/risa.14240. <https://doi.org/10.1111/risa.14240>