RISK MANAGEMENT AND COMPLIANCE

How AI Can Transform Risk Self-Assessments in Banking

By Matteo Coppola, Lorenzo Fantini, Brian O'Malley, Giovanni Lucini, Michele Rigoni, and Benedetta Testino

ARTICLE AUGUST 10, 2026 12 MIN READ

Financial institutions often struggle to manage non-financial risks (NFRs) with the same rigor they apply to financial risks. Unlike financial risks, NFRs tend to rely on unstructured data that is not always systematically captured, organized, or actively managed.

Given the lack of a consistent data foundation, the primary tool for measuring NFRs—the risk and control self-assessment (RCSA)—is often considered a box-checking compliance exercise with limited effectiveness. However, thanks to advances in data observability, analytics, and AI,

banks have an opportunity to reinvent the RCSA as a more accurate, forward-looking tool for management decision-making.

The timing is propitious, given that many regulators have higher expectations for how banks use analytics to manage risk and the growing financial and reputational impact of operational, conduct, and compliance failures. In the face of these challenges, we believe financial institutions should begin managing NFRs with the same rigor and discipline they apply to financial risks.

Closing the Divide Between Financial and Non-Financial Risk Management

Financial and non-financial risk management address fundamentally different exposures and therefore differ markedly in the details of their processes. (See Exhibit 1.) The former measures financial exposure and volatility with a high degree of accuracy based on frequent, data-rich metrics. NFR management, in contrast, focuses on whether processes and controls operate as intended and on the risks inherent in human behavior. These risks are harder to quantify and less predictable, so measuring and managing them tends to rely more on judgment and preventive controls.

EXHIBIT 1

Financial and Nonfinancial Risk Processes Differ Markedly

Dimension	Financial risk management	VS	Non-financial risk management
Timing	High-frequency—daily/intraday monitoring of exposures and limits		Periodic assessments at rigidly predefined intervals, but increasingly trigger-based
Data sources	Reliance on structured, numerical data (prices, contractual cash flows, counterparty metrics, transaction history)		Mix of structured and unstructured data and "signals" (e.g., incidents, audit/control outcomes, complaints)
Data availability	Typically, more abundant data, enabling repeatable aggregation and revaluation		Data heterogeneous and harder to obtain, making analysis more difficult
Approach to measurement	Easier to quantify, producing exposure/loss measures (e.g., "How much could we lose?")		Inherently harder to quantify
	Model-driven measurement		Control- and framework-driven
Exposure	Frequent but smaller amounts		Less frequent but larger amounts
Control vs. exposure	"Manage the exposure"—reduce/reshape risk via pricing, hedging, collateral, portfolio actions		"Manage the system"—strengthen control frameworks, operating effectiveness, end-to-end processes
Operating model	Centralized "risk engine" with stronger second-line-of-defense role		Stronger emphasis on first-line-of-defense ownership

Source: BCG analysis.

Today, however, the “observability divide” between the two types of risk management is narrowing, providing an opportunity to apply many of the strengths of financial risk management to NFR management. For example, banks are increasingly using AI-enabled code scanners to identify redundant controls, along with near-duplicates, enabling them to modify or decommission these controls. AI is also being used to review web behavior, mobile app usage, and call-center interactions to model fraud risk in real time.

As observability increases, inconsistencies become easier to detect, risk patterns more predictable, and controls more targeted and preventive. This shift fundamentally changes what is possible in NFR management and paves the way for a reinvention of the risk self-assessment process.

RCSA: A Solid Foundation, with Some Constraints

RCSA has always had limitations, but it has endured because it is conceptually robust and well structured. It covers multiple NFR types simultaneously, including compliance, conduct, operational resilience, IT and cyber risk, financial crime, and fraud. It also builds on a structured view of inherent risk, controls, and residual risk and could potentially draw on a broad set of contextual data, including internal and external losses, regulatory findings, control testing, and audit outcomes. However, while the theoretical framework is sound, constraints such as limited data availability, significant manual effort, and implementation complexity lead to gaps between the design intentions and real-world application. RCSA is prone to the following flaws:

- **Backward-Looking—by Design.** RCSA remains a point-in-time exercise, anchored in a limited set of historical data and blind to how the risk landscape is evolving.
- **High Friction, Low Value.** Manual data collection, spreadsheets, and complex coordination make RCSA time-consuming and resource-heavy, and therefore can disengage users.
- **Subjective and Inconsistent Outcomes.** Lacking sufficient hard data, RCSAs rely heavily on judgment, and responses tend to skew toward the over-positive based on uneven rationale. Comparability across assessment units is also limited, implying that results are often not actionable and difficult to apply in a coherent manner across an organization.

The emergence of more thorough data observability addresses the foregoing pain points, paving the way for a fundamental reset of RCSA. Leading financial institutions are focusing on three actions: First, they are anchoring RCSA in risk prioritization, employing scenario analysis to target the most salient and significant risks. Second, they are leveraging AI to build near-real-time views of critical processes. And third, they are using clear language to express current outcome-oriented metrics, making results more actionable for business leaders.

> “ Leading financial institutions are anchoring RCSA in risk prioritization, employing scenario analysis to target the most salient and significant risks.

Three Steps to Reinventing RCSA

RCSA’s weaknesses do not invalidate its worth as a guide to capturing emergent risks and decision making. The task, therefore, is to use technology to correct its faults and make it a more robust tool for risk owners and business leaders. Three key moves can jumpstart the journey.

Rethink RCSA from Top to Bottom

One of the RCSA’s major shortcomings is its point-in-time perspective, which rapidly becomes dated and less relevant as time passes. Today, however, with improved data observability, leading institutions are adapting threat-based scenario methodologies—long used in cyber risk and operational resilience—to design a forward-looking RCSA that becomes a top-down scenario-based risk prioritization framework supported by a continuous flow of bottom-up control evidence. (See Exhibit 2.)

EXHIBIT 2

A Forward-Looking RCSA Becomes a Top-Down Scenario-Based Risk Prioritization Framework

Source: BCG analysis.

Scenario analysis combines internal loss databases, external loss databases, and forward-looking indicators to model the potential impact of specific NFR events. And rather than producing single-point scores, the approach estimates loss distributions by simulating outcomes based on event frequency and severity assumptions. The results are easily interpretable metrics such as percentile losses or zero-loss frequency.

Leading institutions also systematically map severe risks to the processes that generate them. From there, processes are top-down prioritized based on the number and severity of risks they pose—highlighting true operational hotspots for management consideration and attention.

Once the task of top-down risk prioritization is complete, the bottom-up RCSA focuses exclusively on control monitoring (by first line of defense) and control testing (mostly by the second line of defense). Control monitoring focuses on control effectiveness, to confirm actual control execution and detection capability. Control testing entails a mix of trigger-based control design—based on significant changes, as opposed to calendar-driven—and control effectiveness, leveraging multiple methods ranging from process walkthroughs and data analytics to sample-based reperforming.

The scenario-based approach is gaining momentum across the industry and increasingly resonates with regulators, who are signaling their openness to more risk-based, efficient RCSA approaches and pushing for a step-up in risk quantification through scenario analysis.

Building a Dynamic, Data-Driven Process with AI

With the right data foundations (for example, code, call/app logs, system logs, transactions), AI can be leveraged to deliver a real-time view of key processes, identifying exceptions, control

failures, and emerging patterns that reveal systemic weaknesses often missed by traditional RCSA approaches. AI tools boost effectiveness across the full process.

Risk Prioritization. AI can be used to design an RCSA that is accurate and robust enough to be reused across business units. The technology can analyze a bank's archive of past RCSAs to uncover previously hidden connections and patterns, map and clarify process boundaries, and diagram consistent taxonomies for risks and process controls, so all parties are using the same terms for the same risks. These steps enable institutions to focus on the scenarios that warrant deeper analysis and avoid wasted effort on nonmaterial risks.

> AI can analyze past RCSAs to uncover previously hidden connections and patterns, map process boundaries, and diagram consistent taxonomies for risks and process controls.

Scenarios and Control Assessment Execution. AI can run scenarios and execute RCSAs that are grounded in data, less subjective, and more consistent. It then automates the drafting of scenario narratives based on historical losses, emerging risk signals, tail-risk classifications, and business context.

In risk and control assessment, AI can propose risk and control ratings based on observable data —such as control performance, incidents, losses, and risk events—rather than self-assessment alone. AI-assisted calibration of expert assessments further improves consistency and reliability by enabling risk owners to anchor their judgments to benchmarks, reducing bias and overconfidence.

Independent Risk Challenger. AI can also act as a "guardian" within the RCSA process, reviewing and questioning entries. By comparing current assessments against internal patterns, historical events, peer benchmarks, and external data, AI flags outlier assumptions, inconsistencies, and unusually optimistic ratings before outputs are finalized.

RCSA Agentic Support. Agentic RCSA models can deploy dialogue-based AI assistants to help first-line risk owners refine inputs, test assumptions, tune assessment parameters, and prepare for formal challenge sessions. Agentic can also act as an RCSA workflow orchestrator, triggering each RCSA step as needed.

RCSA as a Tool for Risk-Reward Decision Making

With risk prioritization and AI-driven processes in place, RCSA outputs can now play a decisive role in supporting decision making by business leaders. The new RCSA, no longer a box-checking exercise with limited actionable outputs, becomes more effective in two ways: it can be applied in a proportionate and selective manner, as opposed to a one-size-fits-all approach; and the outputs are more focused and relevant for business.

A Proportionate RCSA Process. RCSA can now be scaled to fit each legal entity's size and complexity—where allowed by regulation. Large "flagship" entities run the full-fledged process; smaller units could leverage the group results and run the full RCSA only for those risks that have been deemed severe or are expected to deviate from group results. This right-sizing ensures that risk owners are required to assess only those risks that have been prioritized, saving both time and attention.

Clear, Relevant Outputs. RCSA outputs can now be expressed in clear, outcome-oriented terms such as "expected loss," as opposed to abstract risk scores. These outputs are mapped to P&L ownership and the organization's structure, so that each P&L owner sees the expected loss for their risk. Risk figures are clear and easy to understand. Countries, business units, and legal entities see views that map to the institution's organizational hierarchy. Recommendations are specific, actionable, and tightened to risk—that is, they describe which key controls are essential to mitigate a risk event and align root causes with operational losses.

The result of these pragmatic improvements is to foster a no-excuses culture across the organization and to give teams time back to focus on value-adding tasks. In our work with clients, we have seen additional tangible benefits such as:

- A significant reduction in risk management execution costs through automation and risk-based prioritization
- More effective deployment of control investments, achieved by focusing on the controls and processes that matter most
- Better, more data-based business and strategic decisions through quantified, forward-looking risk insights linked to P&L impact

The New RCSA: Making It Happen

As data observability continues to improve, enabling NFR management to edge closer to financial risk management in its accuracy and comprehensiveness, RCSA comes into its own as a

decision-making tool. No longer a periodic standalone certification exercise, it becomes a future-oriented, decision-enabling living platform that is embedded into day-to-day management

For financial institutions, arriving at this point does not entail a Big Bang redesign, but rather a sequenced transformation. Successful organizations will anchor the transformation in the materiality or severity of risk, starting with a small number of NFRs that could meaningfully move the P&L. As with any transformation, proving value early is key to sustained success; in this case, banks should build a minimum viable observability layer (system logs, transactions, control testing results, incident data, audit findings) around four or five critical business processes. When it comes time to scale, discipline is crucial: pilots should be run in a flagship or complex entity, and only when successful should they be scaled to the rest of the organization with explicit thresholds that indicate which entities must adopt the full approach.

With this stepwise, measured approach, financial institutions can access the benefits of a revitalized RCSA within roughly 24 months. What was once an imperfect process riddled with shortcomings becomes a tool that both risk owners and business leaders can turn to with confidence and high expectations.

Authors

Matteo Coppola

Managing Director & Senior Partner
Milan

Lorenzo Fantini

Managing Director & Senior Partner
Milan

Brian O'Malley

Managing Director & Partner
Minneapolis

Giovanni Lucini

Partner & Director
Milan

Michele Rigoni

Partner
Milan

Benedetta Testino

Partner and Associate Director; Risk & Regulatory Compliance in FI
Milan

ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group bridges the gap between ambition and outcomes for the world's leading companies and organizations. We are built for this era of unprecedented change — bringing strategic clarity rooted in over 60 years of deep domain knowledge, combined with applied AI shaped by our practitioners. BCG works shoulder-to-shoulder with CEOs across industries and geographies to deliver transformative impact at scale: stronger returns, transferred capabilities, and change that sticks. For more information, visit bcg.com.