ADAPTIVE RECOVERABILITY GOVERNANCE

A Supervisory Layer Above VaR, Stress Testing, and Catastrophe Modelling

Adaptive Recoverability Systems

Overview Paper — Governance Framework Introduction

Overview

Modern financial risk systems — Value at Risk, Stressed VaR, and catastrophe models — are highly effective at estimating loss exposure under defined scenarios. They answer the question: what losses may occur?

They do not answer a prior and more fundamental question: can this institution still recover using its own internal correction mechanisms?

Adaptive Recoverability Governance (ARG) is a supervisory framework designed to sit above existing risk infrastructure and answer that question directly. Rather than replacing existing systems, ARG acts as a state-aware governance layer that responds to early signals of declining institutional resilience — before visible instability emerges. The framework is sector-agnostic and has been developed and tested across banking, reinsurance, and sovereign risk contexts.

1. Existing Risk Frameworks and Their Shared Limitation

1.1 Value at Risk

VaR estimates the maximum expected loss over a defined period at a given confidence interval under normal market conditions. It is effective for market exposure measurement, capital allocation, and day-to-day portfolio monitoring.

VaR assumes relatively stable market structure, functioning liquidity, and — critically — persistent institutional recoverability. The capacity of the institution to correct itself is treated as fixed and given.

1.2 Stressed Value at Risk

SVaR extends VaR by calibrating to a historically stressed period, improving tail-risk estimation and regulatory resilience. However, SVaR still assumes that institutions remain governable and that correction mechanisms continue functioning under stress. It tests what the exposure would be during a historical crisis; it does not test whether the institution retains the structural capacity to respond.

1.3 Catastrophe Models

CAT models simulate low-frequency, high-severity events across natural, cyber, pandemic, and systemic scenarios. They are highly effective for exposure aggregation, scenario modelling, and loss estimation. However, they primarily evaluate event severity and financial loss transmission. They do not model endogenous governance degradation or an institution's evolving capacity to adapt under stress.

1.4 The Shared Limitation

Every tool above shares a foundational assumption: the system being measured is structurally stable and capable of self-correction. The question being asked is always some form of how bad could a bad event be — never whether the institution is still capable of recovering when it does get bad.

This gap is not theoretical. A recognisable pattern runs through some of the most significant institutional failures of the past three decades:

- Japanese banking system, 1990s–2000s: banks passed contemporaneous solvency metrics while underlying correction capacity had been structurally impaired for years, a pattern documented in the academic literature on “zombie lending.”
- Lehman Brothers, 2008: internal risk metrics and disclosed risk models did not flag the firm's structural fragility in the period immediately preceding its collapse.
- Silicon Valley Bank, 2023: capital ratios were compliant and exceeded regulatory minimums in the years before failure, even as the underlying risks that ultimately drove the collapse went uncorrected.

In each case, the institution continued to operate and produce outputs while having lost much of its practical capacity to self-correct. The failure moment was not the visible collapse — it was the earlier, undetected erosion of that capacity.

2. The Missing Layer: Recoverability

The central distinction this framework introduces is between structural persistence and regulatory functionality.

An institution may remain solvent, operational, profitable, and externally stable while internally:

- governance quality deteriorates,
- optionality shrinks,
- intervention costs escalate,
- and the capacity for endogenous recovery weakens.

This creates a condition of operational persistence with declining recoverability. Existing systems will not detect it, because they are not designed to measure it. ARG is.

> *The most important moment in an institution’s risk lifecycle is not collapse. It is the moment the institution enters a state where recovery is no longer endogenous — where continued functioning depends on external support rather than internal correction capacity.*

ARG measures recoverability directly through two integrated layers, described conceptually below.

3. How ARG Is Structured

3.1 The Recoverability Monitoring Layer

This layer continuously assesses an institution's underlying capacity to self-correct. Rather than relying on a single solvency or liquidity metric, it draws on a small number of structural indicators of institutional health — governance independence, the diversity of analytical perspective feeding into risk decisions, the degree of genuine external challenge an institution is subject to, and the openness of the institution to corrective signals from outside its own incentive structure.

These indicators are combined into a single measure of correction capacity, which is then adjusted for a structural condition the framework treats as particularly significant: the degree to which the bodies responsible for detecting early warning signs are themselves dependent on the institutions they oversee. This condition — distinct from a simple policy choice to delay action — has been present, in varying degrees, in each of the historical cases referenced above.

The framework also monitors how aligned an institution's internal and external observers have become over time. When risk officers, regulators, rating agencies, and boards converge toward the same assumptions and the same models, the independent error-checking that genuine correction capacity depends on quietly erodes — often well before any external shock makes the underlying fragility visible.

3.2 Trajectory, Not Just Position

Once an institution's correction capacity is measured, ARG tracks how it is moving — not just where it currently stands. A static score answers “where is this institution today?” ARG is additionally designed to answer “which direction is this institution heading, and how quickly?”

This distinction matters because, in each of the historical cases ARG is calibrated against, the deterioration was gradual and measurable in hindsight, followed by a collapse that conventional metrics did not anticipate. ARG's trajectory tracking is designed to surface that gradual deterioration while it is still happening, rather than only after the fact.

3.3 The Decision Robustness Layer

The second layer translates recoverability signals into a governance response. As measured recoverability declines, this layer is designed to progressively constrain new exposure, leverage, and irreversible commitments — tightening governance in response to structural health signals, rather than waiting for a crisis threshold to be breached.

This produces a small number of qualitative governance postures, ranging from near-unconstrained authority under healthy conditions through to maximum constraint activation when structural signals indicate severe and compounding stress. A small number of hard governance floors — limits on single-exposure concentration, minimum post-deployment liquidity, and staged commitment structures under elevated stress — apply regardless of the calculated posture, providing a consistent baseline of institutional discipline.

4. Relationship to Existing Risk Systems

ARG is not intended to replace VaR, SVaR, CAT models, liquidity systems, or regulatory capital frameworks. These systems continue to perform the functions they are designed for: loss estimation, capital allocation, and scenario analysis.

ARG acts as a supervisory governance and recoverability overlay. The relationship is additive:

Existing Risk Systems	Adaptive Recoverability Governance
What is the exposure?	Can the institution still correct itself?
How bad could a bad event be?	Is the institution structurally equipped to respond when it does?
Event-driven: responds after a threshold is breached	State-aware: responds to a structural decay trajectory
Contemporaneous or scenario-based signals	Forward-looking trajectory signals
Correction capacity assumed constant and implicit	Correction capacity measured and tracked explicitly

Institutions already running VaR, SVaR, and CAT systems can integrate ARG as an additional governance layer without modifying existing infrastructure.

5. Illustrative Validation

ARG's underlying methodology has been implemented as a working computational system and tested across a range of institutional profiles spanning strong, moderate, and severely degraded governance conditions, under consistent shock conditions designed to isolate the effect of structural health differences.

Institutions modelled with strong, independent governance and low analytical convergence consistently demonstrated sustained resilience under simulated stress. Institutions modelled with weak oversight and high analytical convergence showed measurable deterioration well in advance of eventual failure, with the framework's governance layer constraining new exposure significantly before that failure occurred.

Profiles calibrated to resemble the structural conditions present in the Silicon Valley Bank case were particularly instructive: the framework's measures of correction capacity registered as critically impaired before any simulated external shock, and the governance layer halted new commitments immediately as a result. In that scenario, simulated capital remained substantially intact at the point of failure, because new exposure had already been curtailed.

The operational proposition is not that the framework prevents external shocks from occurring. It is that governance posture reflects structural reality before those shocks arrive — rather than only adjusting after losses have already begun to materialise.

6. Sector Applications

The framework is sector-agnostic. Its underlying indicators and structural logic can be calibrated to any risk-bearing institution or system. Priority sectors for initial deployment include:

- Reinsurance and specialty insurance — governance over catastrophe exposure, contracting portfolio exposure ahead of trigger conditions being met.
- Banking and prudential supervision — a structured, quantitative complement to the qualitative assessment of institutional correction capacity currently conducted under frameworks such as Basel Pillar 2.
- Sovereign and emerging market risk — early structural indicators of deterioration ahead of ratings downgrades and debt distress events.
- Infrastructure and project finance — correction capacity monitoring across long-duration counterparty exposures, where conventional credit ratings may not reflect early-stage structural fragility.
- Regulatory bodies and central banks — a structured framework for assessing institutional governance quality across supervised sectors.

7. Conclusion

Modern financial systems operate in environments of interconnected fragility, rapid feedback loops, and nonlinear systemic stress. Traditional risk systems remain essential for exposure estimation and capital management.

However, a recognisable pattern across major institutional failures is not sudden collapse. It is operational persistence alongside undetected declining recoverability — institutions that continued to function, absorb resources, and pass conventional metrics while having already lost much of their capacity to correct themselves.

Adaptive Recoverability Governance is designed to monitor that transition and adapt governance behaviour before visible instability emerges. It does not compete with existing risk infrastructure. It is intended to provide the supervisory layer above it that the industry has not yet had.

Adaptive Recoverability Systems

This overview paper introduces the Adaptive Recoverability Governance framework at a conceptual level. The underlying scoring methodology, indicator weighting, calibration procedures, and implementation architecture are proprietary and are made available under licence. For licensing or research enquiries, see adaptiverecoverability.com.