

FINANCIAL STABILITY BOARD · FRONTIER AI

The FSB's Frontier AI Warning Exposes *a Second Systemic Risk Channel*

MAUREEN DOYLE-SPARE

© 2026 Maureen Doyle-Spare.

SWIPE TO READ →

The FSB's Frontier AI Warning Exposes *a Second Systemic Risk Channel*

The Financial Stability Board is warning about frontier AI, common technology providers and cross-border disruption.

MAUREEN DOYLE-SPARE

Frontier AI has moved squarely into the Financial Stability Board[1]'s financial-stability agenda. In a letter dated 28 August and published on 31 August, FSB Chair Andrew Bailey warned that frontier AI models are developing increasingly sophisticated autonomy, problem-solving abilities and threat capabilities. Cyber risk is the immediate concern, especially the possibility that advanced AI changes the speed, scale and economics of disruption across a financial system increasingly dependent on common technology providers and shared infrastructure.

In June, the FSB had already proposed 12 Sound Practices for Responsible Adoption of Artificial Intelligence[2] and explicitly asked whether the practices appropriately balance risks across existing and emerging forms of AI, including agentic AI. Together, the consultation and the August warning move the discussion beyond individual models toward autonomy, concentration and resilience.

Shared infrastructure is only one common dependency

Financial stability supervision already knows how to think about concentration. Institutions can depend on the same cloud platforms, service providers, data utilities or foundation models. A compromise or failure in one dependency can affect many firms at once.

Agentic AI adds another shared surface because autonomous systems resolve operational meaning before they act. A customer-onboarding agent may need to determine which beneficial-owner definition applies to a jurisdiction. A payment agent may need to resolve an exception policy or authorization threshold before committing to an action.

Those decisions can depend on common foundation models, orchestration platforms, regulatory-interpretation services or policy-translation infrastructure. If one of those surfaces shifts, multiple firms could arrive at the same unauthorized Operational Interpretation[3] while their systems continue to function normally.

© 2026 Maureen Doyle-Spare.

My research frames that cross-institutional possibility through Systemic Semantic Contagion and Semantic Concentration Risk. Both are proposed research directions for financial-stability supervision. The immediate control still sits inside each institution, where the task is to determine whether an autonomous system is acting on meaning the institution actually authorized.

Figure 1 shows the propagation surface across shared AI infrastructure.

The control problem begins inside each institution

Take the customer-onboarding agent resolving a beneficial-owner definition for a particular jurisdiction. Its Operational Interpretation is the definition it is preparing to act on. Runtime Semantic State captures that resolution in governance-visible form. For that decision, the Reasoning Baseline supplies the institutionally authorized definition for the jurisdiction.

Semantic Deviation Index measures whether the agent's resolved definition has drifted from that Baseline. Measurement then reaches the Deterministic Gate, which turns the result into an authorization decision. A materially divergent definition can hold the onboarding workflow before execution authority is emitted.

Measurement becomes governance when it changes what happens next and leaves evidence of the decision.

Move the evidence point before execution

For that onboarding decision, the Pre-Execution Assurance Protocol establishes the applicable Baseline, captures what the agent resolved, measures divergence, applies the Gate, and generates a Runtime Semantic State Record. Possible Gate outcomes are permit, permit-with-flag, hold, or mandatory hold.

Timing is the important change. Evidence is created before execution authority is emitted. A permitted decision still leaves a record. A held decision preserves the runtime state and the reason the onboarding action did not proceed.

If the agent has applied the wrong beneficial-owner definition, investigators do not have to infer that later from an account outcome or downstream control failure. That pre-execution record shows the definition the agent resolved, the authorized reference it was tested against, and the decision that followed.

Across many decisions, those records begin to populate the second map by showing how operational meaning resolves alongside the infrastructure that carries it.

© 2026 Maureen Doyle-Spare.



Figure 1. *Systemic Semantic Propagation and Correlated Exposure.*
Source: Doyle-Spare (2026).

Prevention is only half of the supervisory problem

Pre-execution assurance addresses whether an individual autonomous action should proceed. Financial-stability oversight adds another need. When an incident crosses workflows, providers or institutions, investigators need to reconstruct the semantic and authorization state that existed when execution was permitted or withheld.

An activity log alone cannot do that. Verification of Runtime Semantic Resolution (V-RSR) is the verification activity at the point of decision. Its evidentiary output is the Runtime Semantic State Record (RSSR), which preserves the active Baseline, the agent's resolved state, the divergence result and Gate decision, together with enough operational context to reconstruct the execution decision.

From there, the Semantic Audit Trail carries the evidence forward. V-RSR verifies the decision. RSSR preserves reconstruction-grade evidence for one execution. Across decisions and time, the Semantic Audit Trail preserves the institutional history. If the same beneficial-owner definition begins to drift across many onboarding decisions, that history gives risk teams, internal audit and supervisors evidence of the pattern instead of a collection of disconnected outcomes.

© 2026 Maureen Doyle-Spare.



Figure 2. *The Pre-Execution Assurance Protocol. Source: Doyle-Spare (2026).*

That evidence chain is a logical extension of the FSB's concern with common providers, simultaneous disruption and cross-border propagation. Operational records can show that several institutions used the same provider or executed similar actions. RSSR and the Semantic Audit Trail add evidence that can help determine whether those actions were also grounded in a common drifted interpretation. These are proposed Doyle-Spare mechanisms for the adjacent reasoning-layer problem, not FSB-prescribed controls.

Comparable per-execution evidence could help supervisors distinguish infrastructure correlation from semantic correlation.

What supervisors could see at scale

Comparable runtime evidence creates a possible supervisory signal without turning every correlation into a conclusion.

Anonymized distributions of Semantic Deviation Index results across similar regulated workflows could provide one signal. A synchronized shift among firms sharing a foundation model, orchestration platform or interpretation service would warrant investigation without being treated as proof of a common semantic failure.

Semantic Concentration Risk provides a second lens. Supervisors already examine concentration in common providers and operational infrastructure. The same discipline could be extended to common semantic infrastructure used for the same consequential decision class.

The test is whether shared semantic infrastructure is associated with common interpretation patterns.

A second lens on concentration risk

This distinction also appeared in an earlier research contribution. On April 18, 2026, I submitted

© 2026 Maureen Doyle-Spare.

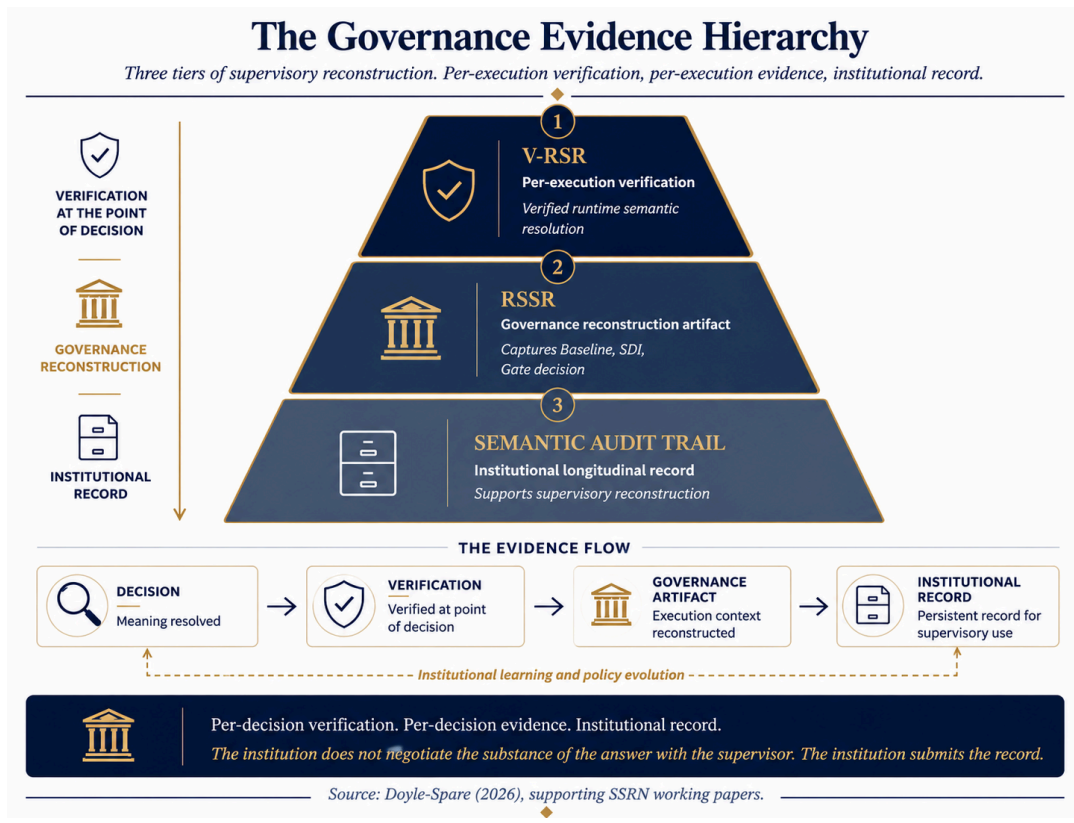


Figure 3. *The Governance Evidence Hierarchy. Source: Doyle-Spare (2026).*

practitioner research to the Financial Stability Board on agentic AI cross-border propagation risk and pre-execution measurement[4]. It described correlated divergence in how institutions may interpret and act on shared operational definitions and proposed consistent Semantic Deviation Index measurement as one route toward system-level observability.

Current FSB work and the April submission address different parts of the same landscape. FSB attention is focused on autonomy, frontier-model capability, common providers, cyber resilience and shared technological dependencies. The earlier research considers whether common AI infrastructure can also produce correlated patterns in operational interpretation.

Financial stability may need two maps

One map describes technology concentration and the channels through which outages, compromise or cyber disruption can spread. A second describes semantic concentration and where institutions depend on common mechanisms for resolving regulated meaning.

Both maps can overlap. A foundation-model provider, orchestration platform or regulatory-interpretation service can create both forms of concentration. As autonomy increases, more institutional authority is exercised through runtime interpretation. As concentration increases, more institutions may depend on the same upstream surfaces for that interpretation.

Financial stability oversight of agentic AI may ultimately need both containment and evidence. Containment stops unauthorized interpretation before execution authority is emitted. Evidence establishes what the

© 2026 Maureen Doyle-Spare.

agent resolved, which Reasoning Baseline governed the decision, what divergence was measured, why execution was permitted or withheld, and which workflows were exposed.

Pre-execution assurance addresses containment. V-RSR, RSSR and the Semantic Audit Trail address reconstructability. At population scale, those institutional records could support future supervisory work on aggregated audit evidence, concentration disclosure and propagation mapping.

Institutions can make autonomous reasoning more governable and reconstructable now. As systems become more capable and the infrastructure beneath them more concentrated, financial-stability oversight may need to trace two things: where a disruption propagated, and what meaning propagated with it. ◆

References

1 *Financial Stability Board. FSB Chair's letter to G20 Finance Ministers and Central Bank Governors. Dated August 28, 2026; published August 31, 2026. fsb.org/2026/08/fsb-chairs-letter-to-g20-finance-ministers-and-central-bank-governors-august-2026*

2 *Financial Stability Board. Sound Practices for Responsible Adoption of Artificial Intelligence (consultation report). June 10, 2026. fsb.org/2026/06/sound-practices-for-responsible-adoption-of-artificial-intelligence-ai-consultation-report*

3 *Doyle-Spare, Maureen. Agentic AI Systems Governance: A Runtime Reference Architecture for the Reasoning Layer and the Semantic Control Plane in Regulated Financial Institutions. Zenodo, 2026. doi.org/10.5281/zenodo.20749051*

4 *Doyle-Spare, Maureen. FSB practitioner research submission, April 18, 2026 (provenance record). maureendoylespare.substack.com/p/fsb-cross-border-propagation-2026-04-18*

5 *Financial Stability Board. The Financial Stability Implications of Artificial Intelligence. November 14, 2024. fsb.org/2024/11/the-financial-stability-implications-of-artificial-intelligence*

6 *Doyle-Spare, Maureen. Agentic Workflow Drift and Agentic Workflow Subversion: A New Risk Taxonomy for the Governance of Agentic AI in Financial Services. SSRN No. 6459612. doi.org/10.2139/ssrn.6459612*

7 *Doyle-Spare, Maureen. Agentic AI Drift Measurement and the Semantic Deviation Index. SSRN No. 6531238. doi.org/10.2139/ssrn.6531238*

Originally published in *The Agentic AI Governance Playbook* on Substack, September 2, 2026. maureendoylespare.substack.com

© 2026 Maureen Doyle-Spare.



Financial-stability oversight may need to trace two things: where a disruption propagated, and what meaning propagated with it.

THE AGENTIC AI GOVERNANCE PLAYBOOK

Download the research

Papers, briefs and reference architecture at maureendoylespare.com

Visit the Playbook

WEEKLY NEWSLETTER

Follow the newsletter

The Agentic AI Governance Playbook Weekly Newsletter

Follow Newsletter

INTELLECTUAL PROPERTY: © 2026 Maureen Doyle-Spare. All rights reserved. No part of this article, including its text, figures, tables, and images, may be reproduced or distributed without the author's prior written permission.