ALDEN ADVISORY

RESEARCH PAPER

AI Governance in Financial Services

Why most AI governance frameworks are model-risk frameworks relabeled

Ravi Sharma — Principal, Regulatory Strategy

Research Paper — Alden Advisory

Alden Advisory — London — aldenadvisory.example (concept site)

Executive Summary

Financial services firms are rapidly building AI governance frameworks in response to emerging supervisory expectations. This paper argues that most of these frameworks are existing model-risk management frameworks with AI-specific language layered on top — and that this substitution misses the genuine governance question AI-assisted decision-making raises: who is accountable when an AI-assisted decision is challenged, and can that accountability actually be evidenced.

Model Risk Is Necessary but Not Sufficient

Traditional model-risk governance focuses on model validation, performance monitoring, and documented model limitations — all necessary, none sufficient for AI-assisted decisions embedded in a live decision workflow. The governance question that model-risk frameworks were not built to answer is: when a human reviewer accepts an AI-generated recommendation without meaningfully interrogating it, who owns the resulting decision?

The Decision-Rights Gap

Our engagements consistently surface the same gap: firms can describe, in detail, how a model was validated, but cannot describe, with the same precision, what a human reviewer is actually expected to do with the model's output before a decision is finalised. Supervisors are increasingly probing this gap directly, and firms without a defensible answer face a genuine governance exposure that model-risk documentation alone will not resolve.

Building Defensible AI-Assisted Decision Governance

We recommend firms map, decision by decision, exactly what independent judgment a human reviewer is expected to exercise before an AI-assisted recommendation becomes a final decision — and build evidence, not just policy, that this judgment is genuinely being exercised rather than rubber-stamped. This is a decision-rights exercise, not a technology exercise, and firms treating it as the latter are building governance frameworks that will not hold up under scrutiny.