AI Governance Belongs in the Finance Operating Model

Kashif Zafar
Independent researcher, Dubai, United Arab Emirates
11 October 2026

Abstract

This paper argues that artificial intelligence governance in corporate settings is best treated as an extension of existing financial control and investment discipline rather than as a standalone policy regime. It proposes a six-part operating framework grounded in finance practice: inventorying decisions rather than models, assigning named business and technical owners with an explicit approval route, scaling controls to the materiality, data sensitivity, reversibility and reliance of each use case, making human review meaningful by supplying evidence, expertise and authority, monitoring benefit and failure in a single record, and carrying validation, monitoring and exit costs into the investment business case. The framework is offered as a management perspective for finance leaders and does not constitute legal or regulatory advice.

Keywords: AI governance; model risk; internal control; accountability; proportionate control; finance operating model

1. Introduction

Here's what AI governance looks like at most companies: a policy document. Some principles. An approval form. A committee that meets once a quarter to feel important. And then everyone goes back to using ChatGPT however they want.

I'll be honest — that's not governance. That's paperwork. Real AI governance is a management job, and here's the thing: finance already has most of the machinery for it. We know how to ask where a number came from. Who signed off on it. What evidence supports it. Those exact same questions apply when an AI tool is shaping a decision.

2. Start with decisions, not technology

First step: don't inventory your AI models or your vendors. Inventory the decisions. Where does AI touch the work? And what does it actually influence?

An AI drafting a meeting summary — low risk. Someone reads it, fixes it, moves on. But an AI suggesting payment approvals? Adjusting a forecast? Drafting disclosure language? Ranking capital projects? That's different. The question isn't 'is this AI?' — it's 'what decision does this shape, and what happens if it gets it wrong?'

3. Assign accountable owners

Every AI application needs a business owner who answers for the outcome. A technical owner who answers for how the thing works. And a clear approval route. If

nobody can tell you who'd explain a bad result to the audit committee, that tool is not ready for anything material.

This is how we treat any key control. Same thing. Ownership is what turns a policy into actual behaviour.

4. Scale controls to consequences

Not every use case needs the same scrutiny. I look at four things: how material is the decision, how sensitive is the data, how easily can you reverse a mistake, and how much are people going to rely on the output without questioning it.

A low-risk drafting assistant? Basic acceptable-use guidance is probably fine. A forecasting model feeding guidance to the market? That needs validation, documentation, change control, periodic review — the works. And here's the important bit — if you over-control the small stuff, people learn to route around the whole process. Proportionality keeps governance credible.

5. Make human review meaningful

'Human in the loop' is my favourite corporate comfort blanket. It only protects you if the human can actually push back. They need the evidence behind the output. The expertise to judge it. And the authority to reject it without someone asking why they're slowing things down.

A reviewer who approves 200 AI suggestions per hour, with no context and no time? That's a rubber stamp, not a control. Design the review step the same way you'd design a segregation-of-duties check.

6. Monitor value and failure together

Once a tool is live, track benefits and problems in the same place. Time saved. Accuracy gained. But also: errors, overrides, incidents, drift in performance. A rising override rate is often the earliest sign that something changed — in the data, the model, or the business.

Reporting both sides together stops the conversation from being either blind enthusiasm or paranoid caution. You want honesty, not vibes.

7. Connect governance to investment decisions

And finally — put governance into the business case. The cost of validation, monitoring, and controls is part of the cost of the tool. So is vendor dependency. So is the exit plan when the provider changes terms, pricing, or capability.

When those costs are visible up front, leaders make better choices about where AI actually pays back. And governance stops looking like overhead someone bolted on after the fact.

8. Conclusion

None of this needs a new bureaucracy. None of this needs a new department. It needs finance leaders to apply the discipline they already have — ownership, evidence, proportionate control, honest measurement — to a new class of tools.

This is a practical management perspective, not legal advice. Your specific regulatory obligations depend on your jurisdiction and sector — check those separately.

References

1. National Institute of Standards and Technology (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1. Gaithersburg: NIST.
2. International Organization for Standardization and International Electrotechnical Commission (2023). ISO/IEC 42001:2023 Information technology — Artificial intelligence — Management system. Geneva: ISO/IEC.
3. European Union (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L series.
4. Organisation for Economic Co-operation and Development (2019). Recommendation of the Council on Artificial Intelligence, OECD/LEGAL/0449. Paris: OECD. (Adopted 2019; revised 2024.)