

Governing AI Before It Governs the Numbers: A Practical Agenda for Finance Leaders

Introduction

A common misconception about AI in finance is that it's usually a catastrophic failure of the machines. It is usually an undocumented tool, an unchecked answer, a private spreadsheet that is pasted into a public system, or an answer that is based on assumptions that are uncheckable. The role and influence of AI on decisions, controls and trust are therefore shifting from user to stewardship by finance leaders.

Artificial intelligence is not usually implemented in an organization as a singular, formal transformation initiative. It arrives quietly. A team outfits a generative AI assistant to create a variance commentary. A procurement analyst instructs a model to make comparisons on bids. A controller tries out anomaly detection. A business unit purchases software with an AI component, but the contract calls it "automation". When senior management talks about an AI policy, the technology could be impacting reports, estimates, and decisions.

This results in an administrative void. The organization might be aware of the enterprise systems that it has, but not the types of AI systems which employees are using, the data fed into those AI systems, or how their output influences decision-making. That is the reason why the riskiest AI in the financial sector might not be the most sophisticated. Maybe the unrecorded one.

The need is not hypothetical. According to the World Economic Forum's Future of Jobs Report 2025, 86% of employers believe that AI and information-processing technologies will change the nature of their business by 2030. This figure is even higher in financial services, at 97 percent. Meanwhile, international advice is gaining in specificity. NIST categorizes AI risk management into four functions: Govern, Map, Measure, and Manage. ISO/IEC 42001 specifies requirements for an AI management system for an organization. Transparency, robustness, and accountability are key principles in the OECD AI Principles. Some high-risk obligations are given a later transition date, but the main parts of the European Union's AI Act entered into general application on 2 August 2026. In fact, these changes demonstrate the direction of travel: the use of AI must be made visible, risk-based, documented and reviewable, even for organizations outside Europe.

Why Finance Cannot Delegate AI Governance

Sometimes AI governance is perceived as an IT or legal project. Both functions are necessary, but neither can be the sole driver of the impact of AI on enterprise value. Finance has an understanding of materiality, internal controls, audit trails, decision rights, and cost-benefit trade-offs, as well as the difference between an estimate and a fact. Finance is a logical co-owner of AI governance because of those capabilities.

Try to think of a model that makes a cash flow prediction. The technical team may perform accuracy and cybersecurity assessments. Legal could review contractual and privacy issues. Finance has yet to decide on the right model for the planning cycle, whether any unusual events are included, who could be responsible for overriding the model, and how it will be explained if there is a material variance. The same holds true for credit scoring, pricing, expense review, tax analysis, fraud alerts, and management reporting. If you have a technically sound system but have the wrong context, incentives, or controls, then you will get a poor business decision.

Finance should not try to be the owner of all the technical decisions. It is designed to guarantee accountability in the use of AI in financial processes. This involves asking the "why" questions that ensure the integrity of the numbers: What decision is the system supporting? What are the chances of something going wrong? What might the repercussions be? Which aspects need to be preserved? Who is responsible if the product is disputed?

Govern the Use Case, Not Just the Tool

A common mistake is to approve or prohibit AI tools as a whole. Risk usually depends less on the brand of the tool than on the use case.



Dr. Salman Ahmed Khatani
Associate Professor & AI Literacy
Practitioner, IQRA University

Table 1: Risk-Based Governance Framework for AI Use in Finance

Risk tier	Finance examples	Primary concern	Minimum governance response
Low	Drafting public text; summarizing non-sensitive material	Poor wording or minor inaccuracy	User review; approved tool; no confidential data
Moderate	Variance commentary; spend classification; internal research	Misinterpretation or data leakage	Named owner; source checks; access controls; sample testing
High	Forecasts; credit or pricing recommendations; fraud alerts	Material financial or customer impact	Independent validation; documented limits; human approval; monitoring
Critical	Automated postings, payments, disclosures or consequential decisions	Direct loss, misstatement, unfair outcome or non-compliance	Executive approval; segregation of duties; audit trail; fallback and incident plan

Source: Author's synthesis based on NIST AI RMF and ISO/IEC 42001.

The same generative AI service may be low risk when it improves the wording of a public presentation, and high risk when it recommends provisions, evaluates employees, or processes confidential customer data. Governance must therefore begin with the purpose, data, decision, and possible consequence.

A Finance-Led Operating Model

Create an AI use-case register

Don't wait for a day of perfection. Use a practical register that includes columns for business purpose, owner, users, data sources, vendor or model, output, affected decision, risk tier and risk review date. Don't just use stand-alone chatbots, but integrate AI into current software applications. The register turns invisible experimentation into controlled work. It also avoids duplicate purchases and facilitates the comparison of the financing's value with the control cost.

Assign accountability before approval

Each use case must have a business owner, who will be held accountable for results; a technical owner, who has technical knowledge of the system; and a control owner, who is knowledgeable about the safeguards and will ensure they are working. Approval should be cross-functional for material applications. An AI committee can coordinate, but accountability should not be forgotten in a committee. Whether and under what circumstances the use is acceptable is up to the decision of a named executive.

Protect data and preserve evidence

The finance data is particularly confidential as it may show performance, pricing, payroll, customers, suppliers and future strategy. Rules of approved use should include what kinds of information can be placed in each system, where it is processed, whether data is allowed to be retained or processed on that system,

and mechanisms for access to the data. Equally crucial, there should be evidence of material outputs. Keep source data, relevant prompt or configuration, model version (if applicable), human reviewer, approval, and final decision. If there is no traceability, there's only guesswork in the review.

Keep human judgment at the decision point

Human oversight is NOT being done by putting a human at the end of an automated process. The reviewer should have the authorization, skills, time, and data to critique the product. It is important for controls to recognize situations where AI might make a decision, decisions AI might suggest, and decisions AI should never make on its own. A reviewer who is accustomed to clicking 'approve' adds ceremony; it doesn't add governance.

Test for business fitness, not only accuracy

Accurate is not enough; it is necessary. Finance leaders should also question the stability, explain ability of the system (to the decision), security, and fairness of the system when people are involved, and resilience of the system in the case of changes in data or conditions. The test must not be based on average performance but on exceptional cases. This could be supply interruption, currency fluctuation, and/or an unexpected shift in demand for a forecasting model. In the case of generative AI, it could be fabricated citations, inaccurate calculations, prompt injection, or leakage of confidential data. There need to be thresholds on material use cases that, on review, suspension, or fallback to manual process.

Govern vendors and model change

The purchase of AI doesn't shift responsibility. Data use, security, audit rights, service continuity, model updates and notification, and exit arrangements should all be covered in contracts.

A tool which passed validation six months ago may act differently after an undetected update. The register should therefore be linked to vendor review and monitoring for changes. To understand what the organization relies on, there must be a shared understanding between procurement, technology, legal, and finance.

Treat AI literacy as an internal control

No policies can cover all prompts or features. Individuals should be aware when an AI answer isn't certain, when data isn't suitable to be shared, and when escalation is necessary. Training should be job specific. The financial analyst requires other examples than does the accounts-payable officer, the internal auditor or a board member. Short scenario-based exercises are more effective than a generic annual presentation. It is not about becoming a data scientist. Its purpose is to develop educated doubt and sensible faith.

What finance leaders can do in the next 90 days?

Day 1-30: Assess what AI is being used. Make a set of use cases, review software features and procurement records, and conduct team surveys. Release an interim regulation about acceptable tools, confidential information, and required human review. Choose 2-3 material use cases to explore in greater detail.

Days 31-60: Categorize those use cases based on the consequences and assign owners, and document some minimum controls. Verify reproduction and challenge of outputs. Check terms with vendors and determine areas of need for logging, access, data management, and incident reporting.

Days 61-90: Approve, amend, or suspend priority use cases. Set up regular reviews of high-risk applications, define escalation and fallback plans, and start delivering AI literacy training to various groups. Report to senior management in business language: value created, risk accepted, controls operating, and decisions required.

This way you do not fall into two undesirable traps. One is an embargo that breaks the cycle of experimentation. The other is that it's an unrestricted adoption that takes speed for progress. Good governance provides a regulated innovation journey. Low-risk uses are fast-moving. Strong evidence, oversight, and accountability for high-impact uses.

Conclusion: The Advantages of Trust

The role of AI will grow more significant in predicting, planning for resource allocation, anomaly identification, and communicating performance. The finance function may react to the problems as they start to arise, or be proactive in defining the rules under which AI can enter the enterprise. The second one is more valuable.



Good governance does not come with guarantees that AI will never get it wrong. It ensures that the organization is aware of where AI is being used, what failure would entail, adds controls in accordance with consequences, and is able to explain the final decision-maker. That ability to factor in decisions becomes a marketable skill in an age of fake content and AI-driven recommendations.

The question finance leaders are now faced with is no longer, **'Should we use AI?'** It is, **'What shall we do when we must?'** The answer needs to be accurate numbers, safeguarded information, transparency and accountability, evidence that can be checked, and human judgment in the right places. Regulating those conditions does not impede innovation. That's the path to winning trust in innovation.

References

- European Commission. AI Act: regulatory framework and application timeline, updated 2026. Available at: <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- IFAC. IFAC Technology Matrix: Artificial Intelligence and Accounting resources. Available at: <https://www.ifac.org/knowledge-gateway/discussion/ifac-technology-matrix>
- ISO. ISO/IEC 42001:2023, Artificial intelligence management system. Available at: <https://cdn.standards.iteh.ai/samples/81230/4c1911ebc9a641fcb6ee21aa09c28ad3/ISO-IEC-42001-2023.pdf>
- NIST. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, 2024. Available at: <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- OECD.AI. OECD AI Principles, updated May 2024. Available at: <https://oecd.ai/en/ai-principles>
- World Economic Forum. Future of Jobs Report 2025. Available at: <https://www.weforum.org/publications/the-future-of-jobs-report-2025/>

About the Author: Dr. Salman Ahmed Khatani is an Associate Professor and Cluster Head for HRM and Social Sciences at IQRA University. A futurist and AI literacy practitioner, he works at the intersection of strategic foresight, responsible AI and organizational transformation.