RoboCFO.ai

a G3 Consulting company

The AI Risk Register

Bringing SOX Discipline to AI Governance in Finance

A practical framework for extending COSO, SOX, and ICFR controls to AI systems in the finance function, drawn from *The AI-Ready CFO* by Glenn Hopper (Wiley Finance, Sept 2026).

Glenn Hopper | G3 Consulting | 2026

Executive Summary

AI has entered the finance function. Seventy-eight percent of organizations reported using AI in 2024, and 87% of CFOs say the technology will be critical to finance operations in 2026. But governance has not kept pace. Only one in five companies has a mature model for governing autonomous AI agents, and most organizations have not yet integrated AI into their SOX control documentation.

This gap is a ticking compliance risk. When AI tools touch forecasting, close processes, disclosures, or approvals, they become part of the control environment. If they are not documented, tested, and governed with the same rigor applied to any other financial control, they represent an uncontrolled variable in the reporting process. One flawed AI-generated disclosure could trigger a regulatory restatement.

The AI Risk Register solves this problem by extending the frameworks finance teams already use. It is not a new bureaucracy. It is COSO, SOX 404, ICFR, and NIST AI RMF applied to a new class of tools. This whitepaper, drawn from the governance framework in Glenn Hopper's forthcoming *The AI-Ready CFO* (Wiley Finance, September 2026), provides the structure, templates, and implementation roadmap that finance leaders need to close the governance gap.

95%	1 in 5	1,100+
of GenAI pilots fail to deliver P&L results	companies have mature AI agent governance	AI-related bills introduced in U.S. states in 2025

Sources: MIT 2025 AI Report; Deloitte State of AI in the Enterprise 2026 (n=3,235); Jade Global AI Governance Report 2025.

Key Findings

- **AI governance is not a separate discipline.** It is the extension of existing SOX, ICFR, and COSO frameworks to a new class of tools. Finance teams that treat AI governance as a bolt-on will create silos. Those that integrate it into existing controls will scale safely.

- **The AI Risk Register is the operational unit.** A single document that inventories every AI touchpoint in finance, links each to its risks, controls, evidence, and accountable owner, and updates on the same quarterly cadence as other SOX documentation.

- **The CFO is the natural owner.** Finance sits at the intersection of data and controls. The CFO, as the "Chief Accountability Officer," cannot delegate the integrity of financial information to an algorithm without governing the algorithm itself.

- **The regulatory clock is ticking.** EU AI Act high-risk enforcement begins August 2026. Over 1,100 AI bills were introduced in U.S. states in 2025. The SEC's 2026 examination priorities now explicitly include AI and cybersecurity.

The Governance Gap

The pattern across organizations is consistent: AI adoption is running far ahead of AI governance. Teams are deploying tools in forecasting, close processes, and reporting before establishing the controls, documentation, and oversight that auditors and regulators expect. The result is a growing population of uncontrolled AI systems operating inside the most scrutinized function in the enterprise.

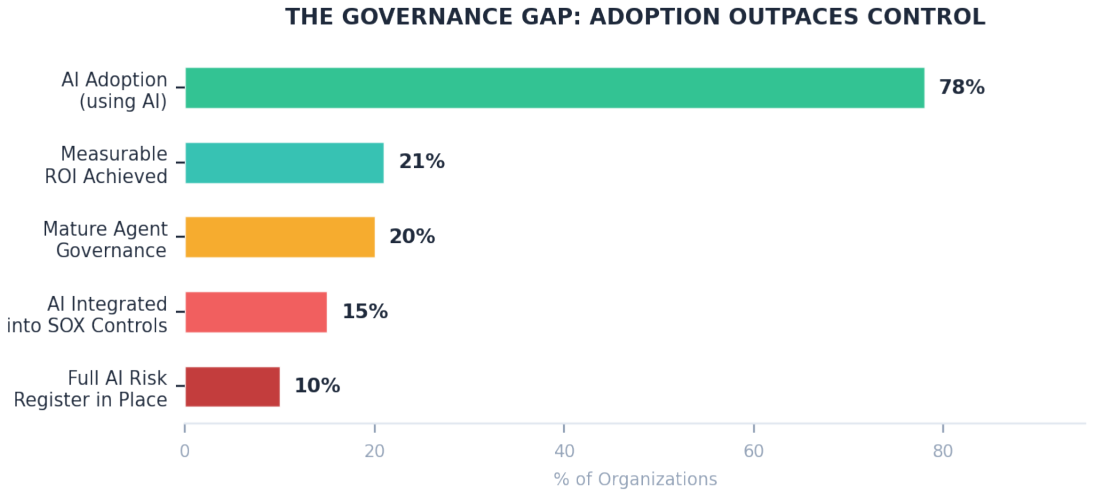

Exhibit 1: The governance gap. AI adoption significantly outpaces governance maturity. Sources: Stanford HAI 2025; Deloitte 2026 State of AI; G3 Consulting estimates.

Deloitte's 2026 State of AI in the Enterprise survey of 3,235 leaders found that worker access to AI rose by 50% in 2025, and the number of companies with 40% or more AI projects in production is set to double within six months. Yet only one in five has mature governance for autonomous AI agents. The skills gap compounds the problem: education, not workflow redesign, was the number-one way companies adjusted their talent strategies for AI.

For publicly traded companies, this gap has direct regulatory implications. SOX Section 302 requires CEOs and CFOs to personally certify the accuracy of financial reports and the effectiveness of internal controls. Section 404 places the burden on management to prove those controls work. If AI tools influence the numbers that flow into certified filings and those tools sit outside the control framework, the certification rests on an undocumented foundation.

The Regulatory Landscape

Regulators are no longer issuing guidance and waiting. Enforcement has arrived. The regulatory environment for AI in financial services is tightening from multiple directions simultaneously, and finance leaders who have not yet documented their AI controls face increasing exposure.

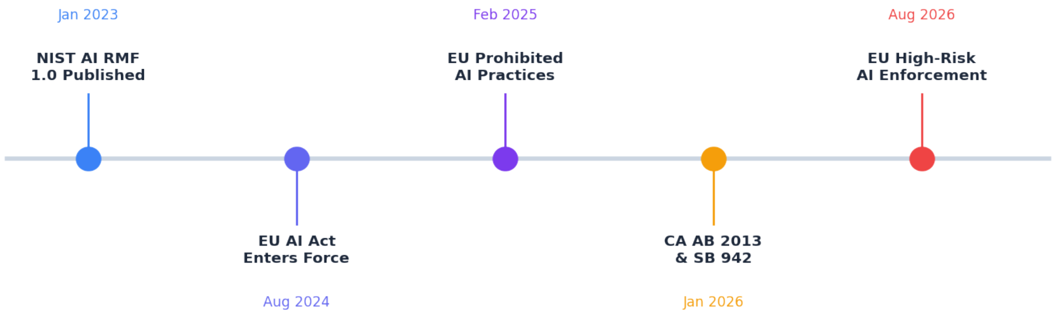

Exhibit 2: Key regulatory milestones for AI governance in finance.

FRAMEWORK	SCOPE	KEY REQUIREMENT FOR FINANCE
SOX 404 / ICFR	U.S. public companies	AI touching financial reporting must be in the control matrix with evidence
COSO ERM (2017)	Enterprise risk management	AI risks documented, assessed, and monitored like any operational risk
NIST AI RMF 1.0	U.S. voluntary standard	Govern, Map, Measure, Manage framework for AI risk lifecycle
EU AI Act	All EU-deployed AI systems	High-risk AI requires documentation, human oversight, logging by Aug 2026
SEC 2026 Priorities	SEC-regulated entities	AI and cybersecurity now in examination scope

Exhibit 3: Regulatory frameworks applicable to AI in finance.

The AI Risk Register

Most finance teams within publicly traded companies already maintain a risk register under COSO's Enterprise Risk Management framework and SOX 404/ICFR documentation. The AI Risk Register is simply an extension of those existing processes, bringing AI into the same structure used for operational, financial, and compliance risks. For private firms, the register serves the same purpose of mitigating AI-related risks, even when it is not required by regulation.

Purpose

The AI Risk Register has two goals: inventory every place AI touches finance workflows, and link each AI use case to its risks, controls, evidence, and accountable owner. This single source of truth lets the CFO demonstrate to auditors, the board, and regulators that AI adoption is governed, documented, and testable within the company's established control environment.

Structure of a Register Entry

FIELD	EXAMPLE
AI Use Case	Generative tool drafts MD&A disclosure commentary
Risk Category	Reporting accuracy
Potential Impact	Material misstatement in external filings
Mitigation / Control	Human review + tie-out to ledgers + GAAP checklist
Evidence / Testing	Review notes, signed checklist, versioned drafts
Owner	Director of External Reporting
Status / Next Review	Operating effectively; Q4 retest scheduled

Exhibit 4: AI Risk Register entry template. Adapted from The AI-Ready CFO, Chapter 8.

Mapping to COSO and ICFR

Auditors will expect to see AI controls expressed through the frameworks they already use. Mapping AI assurance activities to COSO's five components demonstrates that these systems sit squarely inside the existing governance structure. The goal is not to create a parallel control environment. It is to extend the one that is already in place.

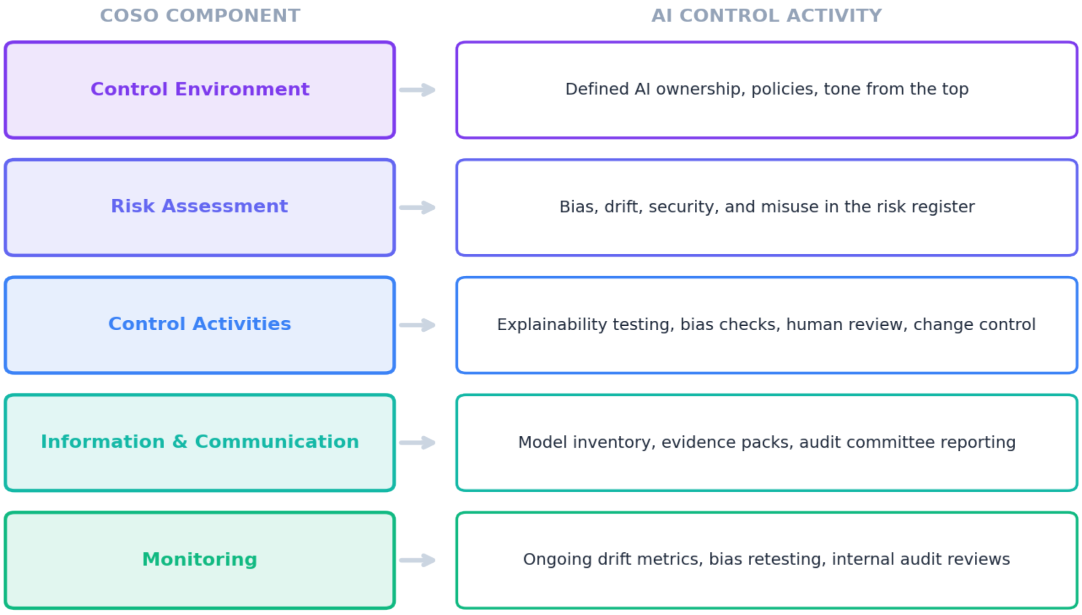

COSO COMPONENT	AI CONTROL ACTIVITY
Control Environment	Defined AI ownership, policies, tone from the top
Risk Assessment	Bias, drift, security, and misuse in the risk register
Control Activities	Explainability testing, bias checks, human review, change control
Information & Communication	Model inventory, evidence packs, audit committee reporting
Monitoring	Ongoing drift metrics, bias retesting, internal audit reviews

Exhibit 5: COSO framework mapped to AI control activities. Source: Adapted from The AI-Ready CFO, Chapter 8; COSO ERM 2017.

By aligning AI oversight to COSO and ICFR, finance teams can prove that every algorithmic control meets the same standards of design, implementation, and operating effectiveness as any manual or automated control in the reporting process. The register also operationalizes NIST's four-function cycle: Govern (policy and ownership), Map (identify use case and context), Measure (assess model performance and bias), and Manage (implement mitigations and track them).

Stop Criteria: Knowing When to Pause Reliance

Even well-governed models can degrade. CFOs should define explicit stop triggers that suspend reliance on AI until re-validation is complete:

- Material drift or unexplained change in feature importance
- Failed bias or stability tests without timely remediation
- Missing or incomplete review and approval evidence
- Security or access breaches that could compromise data integrity

When stop criteria are triggered, the protocol is clear: revert to the fallback or challenger process, document the event, and notify the disclosure and audit committees. Stopping reliance on an AI output is not a failure. It is evidence that the control environment functions as designed.

Implementation Roadmap

Building an AI Risk Register does not require a multi-year program. It requires disciplined execution of five steps, starting with the tools already in use today.

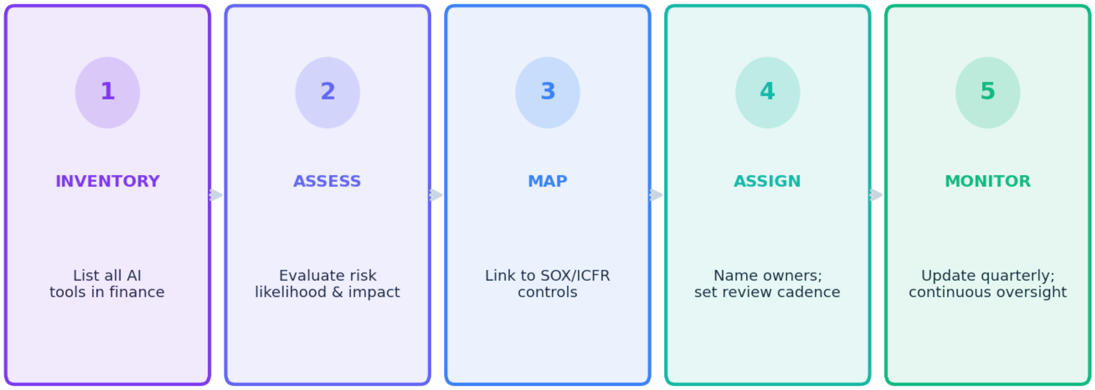
1
INVENTORY
List all AI tools in finance
2
ASSESS
Evaluate risk likelihood & impact
3
MAP
Link to SOX/ICFR controls
4
ASSIGN
Name owners; set review cadence
5
MONITOR
Update quarterly; continuous oversight

Exhibit 6: AI Risk Register lifecycle. Five steps from inventory to continuous monitoring.

1. Inventory. List every AI tool that touches finance: forecasting, reconciliations, disclosures, approvals, reporting commentary. Include shadow AI, the tools teams are using without formal approval. The inventory is the foundation of everything that follows.

2. Assess. Evaluate each tool for likelihood and impact using existing COSO or ERM scoring scales. A generative tool drafting external disclosures carries different risk than an AI that categorizes expenses. Score accordingly.

3. Map. Link mitigations to SOX/ICFR controls already tested by internal audit. If an AI tool produces variance commentary, the existing control for variance review now includes an AI verification step. Update the control narrative.

4. Assign. Name control owners who review evidence quarterly. Exceptions flow to the audit committee. Ownership must sit in finance, not IT. The CFO is accountable for the integrity of financial information, regardless of who built the tool.

5. Monitor. Update the register with new deployments, model changes, or regulatory developments each quarter, on the same cadence as other SOX documentation. Continuous monitoring replaces point-in-time testing.

The same controls that keep reporting credible can also compress the close, improve forecast reliability, and reduce audit noise.

Governance is not the price of AI adoption. It is the mechanism that makes scaling possible. When the register is in place, the CFO can move fast and answer the auditor's questions in a single file.

RoboCFO.ai

a G3 Consulting company

Build Your AI Governance Framework

G3 Consulting helps CFOs build AI Risk Registers, extend SOX controls to AI systems, and implement the governance frameworks that make scaling AI in finance possible without sacrificing auditability.

- AI Implementation Strategy
- Workflow Automation
- Data & Analytics Modernization
- AI Training & Enablement

robocfo.ai/contact

glenn@robocfo.ai

This whitepaper draws on the governance framework from *The AI-Ready CFO: A Strategic Guide to Evaluating and Implementing AI, Automation, and Analytics* by Glenn Hopper (Wiley Finance, September 2026). The AI Risk Register, COSO mapping, and stop criteria frameworks are covered in detail in Chapters 8 and 9 of the book.

Pre-order at amazon.com or barnesandnoble.com. ISBN: 978-1-394-41588-5.