Embedding AI Governance In Financial Institutions: A "Governance-First" Strategy

Dr. Rajan Nagarajan
Dept of Artificial Intelligence and Machine Learning
Madras University

Abstract- *Financial firms must embed AI governance and risk management up front, treating AI not as an isolated project but as integral to enterprise risk and compliance frameworks. Recent regulatory signals (Treasury/FSSCC in the U.S., APRA in Australia, EBA in the EU, MAS in Singapore) emphasize integrating AI oversight into existing structures. Industry studies likewise report that banks with centralized, well-defined AI governance outperform peers. This paper surveys regulatory guidance (e.g., NIST AI RMF, EU AI Act mapping, PRA/FCA feedback), industry frameworks (FSSCC's FS AI RMF, FINOS AI Governance Framework), and case studies (JPMorgan's platform approach, industry consortia) to distill three governance-first approaches: (1) Centralized AI Governance Office/CoE, (2) Federated/Embedded Governance, (3) Platform-Centric (Pipeline) Governance. For each, we outline the structure, processes, roles, pros/cons, investment, and maturity, and map suitability to enterprise personas (e.g., global banks vs. fintechs, CIO vs. CRO/CDO). We then propose a reusable "governance-first" framework and step-by-step implementation roadmap (including templates/checklists, KPIs/KRIs, monitoring/audit plans, third-party controls, and change management). Illustrative case studies (e.g., JPMorgan Chase, FINOS consortium, a hypothetical regional bank) highlight practical trade-offs. We conclude with recommended approaches for different organizational types. All statements are supported by recent regulatory and industry sources.*

Keywords: AI Governance, AI Risk Management, AI Frameworks, AI Implementation, Governance Recommendation, and Roadmap

I. INTRODUCTION

AI is reshaping finance (fraud detection, underwriting, robo-advice, trading, etc.), but introduces unique risks (opacity, bias, cyber-threats). Regulators globally now demand that **AI risk be governed as part of enterprise risk**. For example, the U.S. Treasury and industry published six interrelated AI deliverables (including the FS AI Risk Management Framework) signaling that "AI risk should be embedded in existing risk and compliance frameworks, rather than treated as a standalone technology issue"[1]. The EBA finds the new EU AI Act is largely *complementary* to existing banking laws (no major conflicts)[5], implying banks must integrate AI obligations into current governance (e.g, model risk, internal controls) rather than add separate silos. Similarly, UK regulators note that firms favor a *"joined-up" approach* to AI – coordinating across lines of business (data, model risk, compliance) is beneficial[6]. Singapore's MAS (Nov 2025 guidelines) explicitly treats **data governance** (quality, lineage, bias controls) as a *foundational AI risk domain*, requiring board-level oversight[7]. APRA (Australia) warns that AI adoption is outpacing firms' governance: boards lack AI literacy and legacy processes are "fragmented" across domains[2][8]. In practice, banks report governance is the "missing link" in AI scaling: only ~18–50% of executives feel confident in their AI controls[3][9].

Given this context, a **governance-first strategy** means establishing oversight, policies, and controls *before* or alongside any AI deployment. It embeds risk management and compliance check-points in project pipelines (data, model development, deployment) so that AI is managed like any critical enterprise system. This aligns with NIST's AI Risk Management Framework, whose *"Govern"* function is defined as a cross-cutting discipline infused throughout all lifecycle stages[10][11]. Effective governance (per NIST, FINOS, Wharton) requires clear definitions and inventories of AI assets, formal policies/standards, documented processes and controls[12][13]. In short, firms should treat AI governance as a **foundational infrastructure**, not an afterthought.

II. APPROACHES TO GOVERNANCE IMPLEMENTATION

We identify **three archetypal approaches** that financial institutions use to implement governance-first AI risk management. Each involves different structural choices, resource investments, and trade-offs. We describe their organization, key processes and roles, pros/cons, and suitability by organization "persona" (e.g, large global bank vs. regional bank, fintech, insurer; CIO vs. CRO/CDO priorities).

Approach 1: Centralized AI Governance Office / Center of Excellence (CoE). A dedicated central team (often under the CIO/CRO or Chief Data Officer) owns enterprise-wide AI policy, standards, review,w and oversight. This team develops the governance framework, codes of conduct, and approval processes, and works with all business units. Typically,y it includes AI risk officers, data scientists, compliance experts,s and legal advisors. Project teams funnel through the CoE for policy guidance, model validation, bias testing, etc. For example, JPMorgan has established an Explainable AI CoE to research accountability, and IBM notes that AICoEs demand *"defining clear oversight responsibilities, implementing proper controls and developing policies that balance innovation with risk mitigation"*[3].

Pros: Uniform standards and consistency; strong executive visibility and accountability; reuse of best practices. *Cons:* Can be resource-heavy and slow if too bureaucratic; risk of bottlenecks or detaching from business needs; requires cultural change (breaking silos, cross-team collaboration[3]). *Investment:* High (people and processes); best when AI is strategic across the firm. *Maturity:* Most mature large banks and insurers operate such teams (though only ~31% of firms report having any AI CoE yet[14]). *Best for:* Large global banks, multinational insurers,s or finance firms with complex products. **Persona fit:** CIO/CDO – prioritize enterprise-wide governance. CRO/Model Risk – want central review of all high-impact models. CFO (in finance-led orgs) – supports cross-functional CoE.

Approach 2: Federated / Embedded Governance. Governance responsibility is distributed among the business lines or units. Each unit has its own risk & compliance function (e.g, an AI Risk Committee within Retail Banking, another in Credit), but guided by enterprise-wide policy. A central steering group still exists (perhaps as a policy board or committee of senior executives) to set standards and resolve cross-cutting issues. Front-line units embed data stewards, model validators, or ethics advisors in their teams. Under this model, governance is "baked into" existing domains (e.g, fraud, lending, trading) rather than a separate function.

Pros: Faster decision-making and domain knowledge (teams tailor controls to their context); reduces single point of failure. *Cons:* Risk of inconsistent practices across units; potential duplication of effort; some lines may under-invest in controls. *Investment:* Moderate (focus on training and tools per unit, less on centralized structure). *Maturity:* Common in mid-size banks, fintechs, or firms that trust local autonomy. *Best for:* Regional banks, digital banks,s or fintechs where agility is key, and budgets are lean. **Persona fit:** Business-unit leaders (Head of Retail, CDAO) – want control in their unit. CIO (in a federated model) – supports infrastructure but lets units govern AI use. Model Risk/CRO – maintains an oversight "light touch" role.

Approach 3: Platform-Centric (Pipeline/Reusable) Governance. Here, the firm builds or acquires a common AI/ML platform with integrated controls. Governance is enforced via the platform: standardized data pipelines, versioned model repositories, audit logs, and approval workflows are built in. For example, JPMorgan's CIO Lori Beer emphasizes "building governance into platforms from the start," e.g., integrating security intelligence into software dev processes[15]. The platform team (often under IT or Data) collaborates with risk to embed checks – e.,g. automated bias scans, log routines, and documentation templates – so that any AI model built on the platform automatically complies with policy.

Pros: Scalability and consistency via technology; avoids re-inventing controls per project; easier reporting and metrics across all use-cases. *Cons:* Significant upfront cost to develop or procure; requires talented data engineers; less flexible for one-off projects. *Investment:* High (platform dev, tool licenses, training). *Maturity:* Advanced; suitable once the organization has many AI projects. *Best for:* Largest institutions (global banks, major insurers) or tech-savvy firms that can leverage platforms. **Persona fit:** CIO/CDO – favors enterprise platforms; CTO/CIO might lead implementation. Model Risk/CRO – supports automated controls. Data Scientists – benefit from reusable pipelines.

III. MAPPING APPROACHES TO ORGANIZATIONAL PERSONAS

Different organizations (and even roles) will find one approach more suitable than others:

- **Large Global Bank / Multinational Insurer:** Typically gravitate to *centralized CoE + platform* hybrid. They have resources to fund a CoE and build common platforms. With complex products and strict regulators, they need uniform controls[3]. CIO/CTOs push platform solutions; CRO/CDOs insist on central oversight.
- **Regional Bank / Credit Union:** Often uses a *federated approach*. Smaller budgets make heavy CoEs or platforms impractical. Embedding risk officers in each department (guided by core policies) lets them adopt AI use-cases fast while meeting basic controls. CIOs still coordinate data/classification.
- **Fintech / Digital Bank:** Likely *platform-focused* (if tech-enabled) or *embedded*. A lean fintech may build

on a cloud AI platform with automated governance or rely on vendor controls. They prioritize agility: for example, a fintech CDO might prefer using a platform-as-a-service with built-in bias testing. They will still align with supervisor expectations but may not mirror big-bank bureaucracy.

- **Insurers:** With heavy actuarial/math roles, insurers may adopt either a *central CoE* (especially large firms) or *federated* in underwriting vs. claims. They face similar data challenges, so many are now establishing enterprise AI committees under the CRO or CIO.
- **Specific Roles:** A **CIO/CTO** favors solutions (platforms, automated tools) to streamline deployment across units. A **Chief Data Officer** focuses on data governance and architectures supporting any approach. A **Head of Model Risk** or **AI Risk Officer** needs visibility into all AI models – they often drive the establishment of inventory controls (which any approach must enable). A **Chief Compliance Officer** or **Legal Counsel** will insist on policies/checklists and audit trails (again common to all approaches).

In practice, many firms use a hybrid: e.,g. a central CoE defines policy and supports platform build, while individual units carry out local compliance tasks. The key is to clearly delineate roles: e.g, **Board/Senior Management** owns strategy and resources; **AI Steering Committee** (cross-functional) approves use-cases; **AI Risk/Compliance Group** performs audits; and **Project Teams** (data scientists) build models following the established playbook.

IV. GOVERNANCE-FIRST IMPLEMENTATION FRAMEWORK

Based on industry best practices and regulator guidance, we propose a reusable **framework** for embedding governance at each stage of the AI project lifecycle. Firms can adapt this as a "governance-first pipeline". Key steps include:

1. **Executive Commitment & Policy:** Board and C-suite endorse an AI governance policy (covering ethics, risk appetite, data use, etc.). Define clear oversight roles (e.g., AI Steering Committee, AI Risk Officer)[16][17]. Align AI strategy with risk appetite and strategy; consider regulations (e.g., EU AI Act, SR 11-7, etc.).
2. **Inventory & Classification:** Create an inventory of all AI/ML use-cases. For each, record data sources, model purpose, output audience, and impact level. Classify use-cases by risk (high-impact, moderate, low) based on criteria like customer outcome sensitivity or complexity[7][11]. Use tools or spreadsheets to register models, datasets, and vendors.
3. **Risk Assessment:** Perform risk/gap analysis for each high/medium use-case. Evaluate potential harms (bias, privacy, resilience, governance) and control gaps. For example, assess data bias risks, model opacity, and cyber vulnerabilities[13][2]. Score/rank risks using a methodology (some firms use a *risk scorecard*, as McKinsey suggests[18]). Determine which models need rigorous review (e.,g. by Model Risk Management) vs. those needing basic checks.
4. **Governance Structure & Controls:** Set up or refine the governance body (central vs. federated as chosen) and processes for oversight. Develop or update policies: data governance (quality, lineage, labeling standards)[7]; model development protocols; documentation templates; ethics guidelines. Establish review gates: e.g., all high-risk models require independent validation. Embed continuous monitoring policies (data drift, performance). Incorporate third-party risk management: require vendor attestations, due diligence questionnaires, and contractual AI-risk clauses for outsourced models.
5. **Platform & Tools:** Invest in technology to operationalize governance. This may include: model registries (for inventory), audit logs, bias/fairness testing tools, and code review pipelines. Ensure repositories enforce version control and lineage. For data, build data catalogs with metadata and classification. Deploy monitoring dashboards (for performance metrics, data drift, usage logs). The Governance team and IT should ensure these tools align with policy (e.g., automated alerts when a model crosses risk thresholds).
6. **Training & Change Management:** Train staff (data scientists, devs, business users) on governance policies and the "why" of controls[3]. Roll out the new processes and tools gradually, with pilot projects and feedback loops. Communicate the benefits (risk mitigation, regulator confidence). Define KPIs/KRIs (e.,g. % of models validated, incidents resolved, time to deployment, number of models meeting explainability targets) and report them to stakeholders.
7. **Monitoring, Audit & Feedback:** Continuously monitor AI use-cases in production (performance, drift, compliance with policy)[19]. Conduct periodic audits of AI systems and governance processes. Update the framework as regulations and technology evolve (NIST calls for a "live" approach with

periodic reviews[20]). Involve internal/external auditors as needed. Use findings to refine policies and training (continuous improvement loop).

These steps can be organized in an implementation **roadmap**. **Mermaid Flowchart: AI Governance Lifecycle** – The lifecycle of AI governance involves continuous loops of oversight and improvement. The flowchart below summarizes the key stages of a governance-first process:

This cycle emphasizes **governance ("A"-"B") at the top**, feeding into systematic risk analysis and control design (green). Monitoring (blue) feeds back to revise risk assessments and policies as needed.

V. ILLUSTRATIVE CASE STUDIES

Case: JPMorgan Chase (Large Global Bank). *Approach:* Centralized CoE + Platform. CIO Lori Beer describes how JPMorgan embeds governance into its software pipelines from the outset[15]. The bank's centralized AI Governance office works with IT to build shared AI/ML platforms. GenAI models and external tools must pass through the risk assessment process. For example, JPMorgan built an in-house cyberthreat model and integrated security reviews into its development process[15]. They treat generative AI systems as requiring *"governance and risk assessments… they are must-dos"*[15]. Outcome: By making governance a platform feature, JPMorgan can deploy AI at scale (e.g, their Connect Coach app has embedded AI improving advisor productivity) while satisfying regulators' expectations. *Key takeaway:* Enterprises leading on AI tie governance to technology infrastructure and get executive buy-in (e.g, JPMorgan measures success not by short-term ROI but by strategic resilience)[21].

Case: Industry Consortium (FINOS / Banking Forum). *Approach:* Collaborative hybrid. Several large financial institutions (e.g, Bank XYZ, Insurer ABC) participated in FINOS's AI Governance Framework project and the FSSCC AI Executive Oversight Group. They contributed to open catalogs of AI risks and controls (the FINOS Risk Catalog). Such consortia serve as "virtual governance centers" where members share best practices. For example, FINOS emphasizes the unique challenges of GenAI (hallucinations, unpredictability) and recommends *adaptive governance* models to mitigate them[22]. Member firms adopt these shared controls (e.g, bias detection methods, data documentation templates) into their own pipelines. Outcome: Firms' consortia can accelerate development of their frameworks by reusing vetted templates and learning from peers. *Key takeaway:* Even without a formal CoE, firms can leverage cross-industry governance work (FINOS, FSSCC's FS AI RMF) to jump-start policies; this federated collaboration reduces duplicative effort.

Case: Mid-Sized Regional Bank (Hypothetical). *Approach:* Federated/Embedded. Consider a regional bank with ~$20B in assets. They have a few dedicated AI experts. Instead, the bank's risk and compliance office issues an **AI Governance Handbook** for business units. Each business unit forms an "AI Working Group" (fraud unit, lending unit, etc.) that inventories its projects and reports monthly to the Risk Committee. The bank invested modestly in a cloud AI platform (provided by a vendor) that has built-in logging and bias checks, but most day-to-day governance is done via checklist approvals before deployment. The Board reviews an annual AI report prepared by the Chief Risk Officer, covering inventory, incidents, and audit findings. Outcome: The bank manages to comply with regulatory expectations (monitored by its supervisor) without a heavy CoE. However, it faces challenges in ensuring consistency: some units push proprietary models too quickly, causing a few near-misses. This highlights the *cons* of a federated approach: strong policy communication and occasional audits are needed to avoid gaps. *Key takeaway:* Smaller institutions can start with lightweight embedded governance (checklists, unit-level oversight) and still strengthen compliance; they should scale up formal structures only as cases grow.

VI. IMPLEMENTATION ROADMAP AND TOOLS

In practice, firms will tailor the above framework. Key recommendations and resources include:

- **Templates/Checklists:** Use standardized templates for **model documentation** (purpose, data lineage, performance metrics) and **risk-control checklists**. For example, leverage industry toolkits (FSSCC's AI RMF user guide, or MAS risk templates) to ensure no key steps are missed. Create an AI inventory template (metadata, owner, risk rating) to track progress.
- **KPIs/KRIs:** Define metrics to measure governance effectiveness: e.g, number of AI projects cataloged, % of high-risk models reviewed and approved, data quality scores, audit findings, time to deploy models, and number of incidents or biases detected. Baseline these before launch and track monthly. Regulators may require proof of board reporting and incident logs.
- **Monitoring & Audit:** Establish automated monitoring where possible: drift detection for models, business-outcome tracking, and log reviews. Plan periodic audits by internal model risk or a third-party to verify compliance. NIST and OCC guidance suggest integrating continuous monitoring tools (e.,g. MLflow, Delta Lake lineage) into the platform[23].
- **Vendor/Third-Party Controls:** Ensure any third-party AI (or data) provider is subject to due diligence. Maintain an "approved vendors" list. Conduct AI-focused reviews: check the vendor's governance (does it align with our policy?), legal rights (IP, retraining consent), and data practices. MAS and APRA stress oversight of outsourcing in AI contexts[7][2]. Embed contractual clauses requiring audit rights and compliance certification for AI services.
- **Change Management:** Communicate widely: include C-suite on the steering committee, involve business sponsors for each project, and train end-users on new model outputs. Use internal forums (Town Halls, newsletters) to share successes and reinforce controls. Leadership endorsement is critical: as one IBM executive noted, AI cannot be "a one-off project that a single department owns"[24].

VII. RECOMMENDATIONS AND NEXT STEPS

Match approach to organization: Large, diverse banks should lean on centralized CoEs (often with robust platforms) to ensure consistency across markets. Midsize or fast-moving firms may start with federated models, then gradually build central expertise. All firms should embed data governance at the core (as MAS requires[7]).

Iterative implementation: Begin with high-impact pilots using the full governance cycle (inventory → assess → control → monitor). Refine the framework before scaling. Use audit findings to fill gaps.

Use industry resources: Adopt or adapt existing frameworks (NIST AI RMF, FS AI RMF, FINOS controls). Participate in industry groups (FSSCC, BCBS forums) to stay aligned on standards.

Prepare for regulation: Even absent explicit AI-specific rules, expect scrutiny under existing regimes (model risk, consumer protection, information security). For example, U.S. agencies have indicated that enhanced model risk principles (SS1/23) will cover AI[25]. The EU's AI Act (in force Aug 2024) designates key AI uses (credit scoring, trading) as "high-risk" requiring strict governance[26][5]. Plan ahead to meet these requirements by 2026–27.

Document and Report: Keep thorough records of governance activities (inventory logs, approval forms, audit reports). This not only aids internal tracking but also prepares for possible regulator questions or audits.

Gaps and Sources: We based this paper on public regulatory documents and industry analyses. Some primary guidance (e.g, MAS 2025 draft guidelines) is not fully accessible; we rely on summaries (e.g, Alation[7]). Most information comes from official reports (NIST AI RMF[10], EBA mapping[5], APRA media releases[2][27], OCC bulletin[23]) and reputable industry sources (Grant Thornton/Treasury brief[1], McKinsey[4], FINOS[22], BizTech/IBM[3]). Wherever possible, we have cited authoritative or peer-reviewed sources.

REFERENCES

[1] NIST, *AI Risk Management Framework (AI RMF 1.0)* (Jan 2023): defines core functions,s including *"Govern"* as a cross-cutting discipline[10][11].
[2] FSSCC/Treasury, *Financial Sector AI RMF* (2024): an operationalization of NIST for finance[28].
[3] Bank of England/PRA/FCA, *Artificial Intelligence & Machine Learning* (DP5/22, 2022) and *Feedback Statement* (FS2/23, Oct 2023)[6][29].
[4] EU EBA, *AI Act: Implications for Banking* (Nov 2025)[5].
[5] MAS, *AI Risk Management Guidelines* (2025) – see secondary analyses[7].
[6] APRA, *Letter on AI Risk Management* (Apr 2026)[2][27].

[7] Grant Thornton/Treasury, "AI Governance Shift" (May 2026) – summary of U.S. AI deliverables[1][28].
[8] FINOS AI Governance Framework (Oct 2025)[22].
[9] BizTech/IBM (Dec 2025), *How to Build an AI CoE in Financial Services*[3].
[10] McKinsey (2025), *Improving Governance of Generative AI in Finance*[4][30].

These sources informed the analysis.

[1] Treasury guidance brings urgency to AI governance | Grant Thornton
https://www.grantthornton.com/insights/articles/banking/2026/treasury-guidance-brings-urgency-to-ai-governance
[2] [8] APRA calls for a step-change in AI-related risk management and governance | APRA
https://www.apra.gov.au/news-and-publications/apra-calls-for-a-step-change-ai-related-risk-management-and-governance
[3] [14] [24] How Financial Services Can Build an AI Center of Excellence | BizTech Magazine
https://biztechmagazine.com/article/2025/12/how-financial-services-can-build-ai-center-excellence
[4] [18] [30] How financial institutions can improve their governance of gen AI | McKinsey
https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/how-financial-institutions-can-improve-their-governance-of-gen-ai
[5] [26] AI Act implications for the EU banking sector _updated 20/11/2025
https://www.eba.europa.eu/sites/default/files/2025-11/d8b999ce-a1d9-4964-9606-971bbc2aaf89/AI%20Act%20implications%20for%20the%20EU%20banking%20sector.pdf
[6] [20] [25] [29] FS2/23 – Artificial Intelligence and Machine Learning | Bank of England
https://www.bankofengland.co.uk/prudential-regulation/publication/2023/october/artificial-intelligence-and-machine-learning
[7] MAS Data Governance Compliance: 2026 Guide for Financial Institutions | Alation
https://www.alation.com/blog/mas-data-governance-guidelines-compliance/
[8] Banking insights: 2026 AI Impact Survey | Grant Thornton
[9] https://www.grantthornton.com/insights/survey-reports/banking/2026/banking-insights-2026-ai-impact-survey
[10] [10] [11] AI RMF Core - AIRC
[11] https://airc.nist.gov/airmf-resources/airmf/5-sec-core/
[12] [12] Artificial Intelligence Risk & Governance - Wharton AI & Analytics Initiative
https://ai-analytics.wharton.upenn.edu/industry/finance/artificial-intelligence-risk-governance/
[13] [16] [19] AI Governance in Financial Services
https://www.holisticai.com/blog/ai-governance-in-financial-services
[14] [21] How Goldman Sachs, JPMorgan, and AIG Are Actually Deploying AI
https://www.govinfosecurity.com/how-goldman-sachs-jpmorgan-aig-are-actually-deploying-ai-a-31643
[15] [27] APRA Letter to Industry on Artificial Intelligence (AI) | APRA
https://www.apra.gov.au/apra-letter-to-industry-on-artificial-intelligence-ai
[16] FINOS AI Governance Framework:
https://air-governance-framework.finos.org/
[17] Model Risk Management: Revised Guidance | OCC
https://www.occ.treas.gov/news-issuances/bulletins/2026/bulletin-2026-13.html
[18] fsscc.org
https://fsscc.org/wp-content/uploads/2026/03/Financial-Sector-AI-Deliverable-Reference-and-Application-Guide.pdf