AI Governance for Financial Institutions: A Strategic Framework for Evaluation, Adoption, and Implementation

Authored By: Caleb J. Mabe, Anna Rogers, Grant Goodwin, Benjamin Patton, Thomas Hill & Bill Stuart

Executive Summary

Artificial Intelligence (AI) is transforming the financial services landscape, offering unprecedented opportunities to enhance customer experience, streamline operations, and drive innovation. However, with these opportunities come significant responsibilities around governance, risk management, and regulatory compliance.

This whitepaper provides financial institutions with a framework for understanding, evaluating, and implementing AI technologies responsibly. Drawing on nCino's extensive experience supporting over 2,700[1] financial institutions globally, this guide addresses the critical questions every financial leader must answer: What exactly is AI? How do we evaluate AI by use-case? What are the actual risks versus perceived risks? How can we successfully implement AI to benefit our institution and customers?

Table of Contents

Understanding AI in Financial Services

The Current State of AI Adoption

Financial institutions are at varying stages of AI adoption, from early exploration to full-scale implementation. The industry has long relied on statistical models and automated processes for functions like credit scoring, fraud detection, and risk assessment. Today's AI technologies represent an evolution of these capabilities, offering enhanced predictive power, pattern recognition, and decision-making support.

Why AI Matters for Financial Institutions

Operational Efficiency: AI automates routine tasks, reduces manual processing, and accelerates decision-making processes across lending, account opening, and customer service operations.

Enhanced Customer Experience: Intelligent systems enable personalized service delivery, faster response times, and more accurate product recommendations.

Risk Management: Advanced analytics improve fraud detection, credit risk assessment, and regulatory compliance monitoring.

Competitive Advantage: Early adopters of responsible AI gain market advantages through improved efficiency, better customer insights, and innovative service offerings.

The nCino Perspective

nCino has observed that successful AI implementation requires more than just technology deployment. It demands a strategic approach that balances innovation with governance, ensuring that AI initiatives align with institutional values, regulatory requirements, and customer expectations.

Defining AI: The Critical First Step in Risk Management

As your organization embraces technology to drive efficiency and optimize processes, a fundamental question emerges: How do you distinguish true AI from other automated systems? This distinction matters because it determines which technologies require enhanced risk assessments and governance oversight.

The Definition Dilemma

Without a clear definition of AI, organizations face two risks:

Too broad: Traditional statistical models like linear regression get swept into unnecessary AI governance processes, creating administrative burden without corresponding benefit.

Too narrow: Complex AI systems slip through governance frameworks, exposing the organization to unmanaged risks.

Drawing the Line

The challenge lies in establishing boundaries that capture the technologies requiring special attention while avoiding unnecessary complexity. Consider the spectrum:

Basic automation (rules-based systems)	Traditional analytics (statistical models)	Machine learning (pattern recognition)	Generative AI (content creation)	Agentic AI (autonomous decision-making)

nCino's Approach

We adopt the Organization for Economic Co-operation and Development's (OECD) definition as our foundation:

> "An AI system is a machine-based system that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments."[2]

This definition captures the essence of AI–systems that infer and adapt–while providing practical boundaries for governance frameworks.

Action Step: Before implementing any AI governance framework, clearly define what constitutes AI for your institution. This definition will guide every subsequent decision about risk assessment, compliance, and oversight.

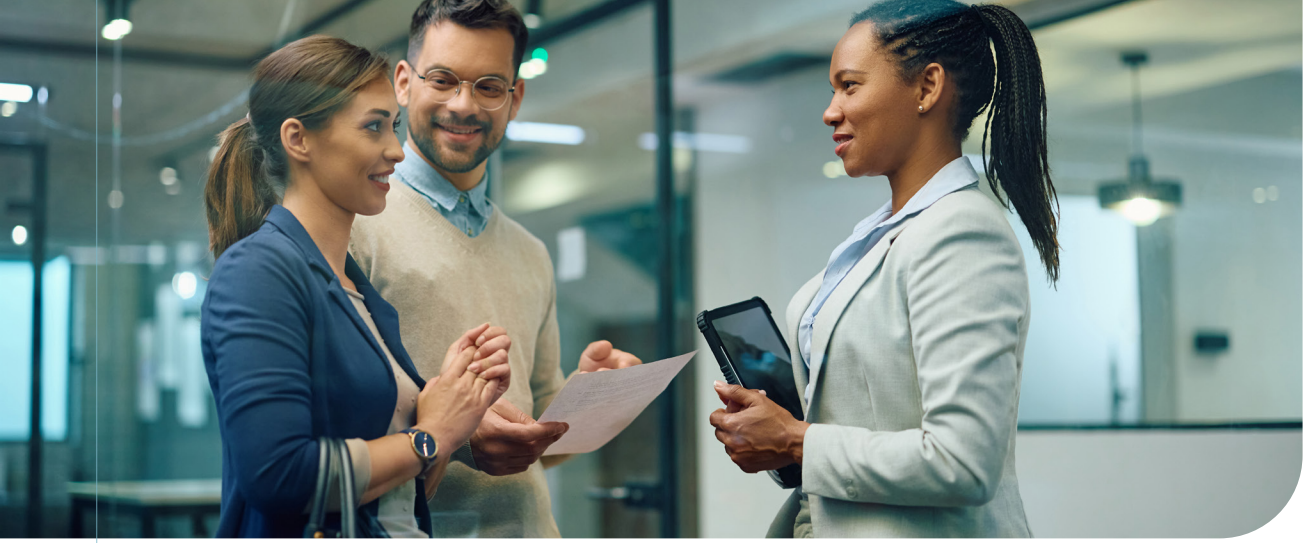

Risk-Based AI Governance: Matching Scrutiny to Impact

Not all AI applications carry equal risk or require identical oversight. A one-size-fits-all approach to AI governance creates unnecessary friction, slows innovation, and frustrates teams trying to deliver value. The key is implementing risk-proportionate governance that aligns scrutiny with potential impact.

Incorporating into Existing Privacy Frameworks

Most organizations have current privacy or data governance programs. Rather than starting from scratch, these frameworks can be built upon or replicated for AI Governance. The Financial Services – Information Sharing and Analysis Center (FS-ISAC) encourages reviewing existing privacy controls and standards and applying or expanding as needed for AI systems[3].

The Cost of Over-Governance

When organizations apply the same intensive risk assessment to every AI tool–treating email sorting algorithms with the same rigor as automated lending decisions–they create several problems:

Innovation paralysis: Teams avoid exploring beneficial AI applications due to bureaucratic overhead.

Resource drain: Limited compliance resources get consumed by low-risk evaluations.

Competitive disadvantage: Slower adoption of efficiency-enhancing tools while competitors move ahead.

Employee frustration: Technical teams lose faith in governance processes that seem disconnected from reality.

A Tiered Approach to AI Risk

Consider this risk hierarchy for financial institutions:

Tier 1 - Low Risk (Streamlined Review)

- Internal process automation (email sorting, document filing)
- Data visualization and reporting tools
- Employee productivity assistants
- Back-office optimization

Governance approach: Basic documentation, standard security review

Tier 2 - Medium Risk (Standard Review)

- Customer service chatbots with human oversight
- Marketing personalization engines
- Fraud detection systems with human review
- Internal risk monitoring tools

Governance approach: Full risk assessment, regular monitoring, clear escalation paths

Tier 3 - High Risk (Comprehensive Review)

- Automated credit decisions
- Customer-facing financial advice
- Regulatory compliance determinations
- Agentic AI with decision-making authority

Governance approach: Board-level approval, continuous monitoring, extensive documentation, regular audits

The Business Impact

Organizations that implement risk-proportionate governance see:

- Faster deployment of low-risk AI tools.
- Better resource allocation for high-risk applications.
- Higher employee engagement with AI initiatives.
- Maintained regulatory compliance without unnecessary burden.

Actual vs. Perceived Risks in AI Implementation

Common Misperceptions About AI Risks

4.1 The "Black Box" Myth

Perception: All AI systems are black boxes that can't be understood or explained.

Traditional AI Reality: Established AI techniques offer robust explainability tools that help us understand how models make decisions.

These include:

- Model-agnostic explanation methods that work across different algorithms
- Feature importance analysis that identifies which inputs most influence outcomes
- Surrogate models that approximate complex systems with simpler, interpretable ones
- Rule extraction techniques that translate model behavior into human-readable logic.

Generative AI Reality: Large language models and other generative systems present unique interpretability challenges that current explainability methods struggle to address. The complexity of these models and their emergent behaviors make traditional explanation techniques less effective. Since we can't fully explain generative AI behavior, we must focus on comprehensive risk management. This means:

- Red team early and often: test edge cases, jailbreaks, and failure modes before deployment.
- Build guardrails at multiple layers: input filtering, output monitoring, and human oversight loops.
- Log everything: conversations, decisions, failures. You can't govern what you can't see.
- Start with lower-risk use cases: learn how your specific LLM behaves before scaling to critical applications.
- Have humans in the loop for high-stakes decisions, especially early on.

Best Practice: Implement layered explainability approaches appropriate to the use case and regulatory requirements.

4.2 The Bias Amplification Fear

Perception: AI inevitably amplifies bias and discrimination.

Reality: AI can actually help identify and reduce bias when properly implemented:

- Bias detection tools can identify disparate impact
- Diverse training data can improve fairness
- Ongoing monitoring can catch bias drift
- Human oversight provides necessary checks

Best Practice: Adopt partners who routinely test for these types of issues.

Actual Risks Requiring Attention:

4.1 Model Performance Degradation

Risk: AI models may lose accuracy over time due to data drift or changing conditions.

Mitigation Strategies:

- Create comprehensive logging and observability infrastructure to capture system behavior
- Establish performance alert thresholds with escalation procedures
- Implement A/B testing frameworks for gradual rollout of system updates
- Maintain rollback capabilities to previous system versions when performance degrades
- Conduct regular performance audits and stakeholder reviews
- Document performance baselines and acceptable degradation thresholds
- Establish feedback loops from end-users and downstream systems to detect performance issues early
- Implement automated testing suites that run continuously to catch regressions
- Create performance dashboards for real-time visibility into system health

4.2 Data Quality and Governance

Risk: Poor data quality can lead to unreliable AI outputs and biased decisions.

Mitigation Strategies:

- Establish comprehensive data governance programs
- Implement data quality monitoring tools
- Create clear data lineage documentation
- Regularly audit data sources and processes

4.3 Regulatory Compliance Gaps

Risk: AI implementations may not meet evolving regulatory requirements.

Mitigation Strategies:

- Stay current with regulatory guidance and requirements
- Implement comprehensive documentation practices
- Establish clear governance and oversight structures
- Engage with regulators proactively

4.4 Cybersecurity Vulnerabilities

Risk: AI systems may introduce new attack vectors or security vulnerabilities.

Mitigation Strategies:

- Mitigation strategies should be applied based on scope of use; developer, deployer, or end-user.
- On a defined frequency (monthly is ideal, based on AI innovation), update risk and opportunity register with new risks, opportunities, sources, impact, likelihood, and mitigation.
- Keep an inventory of AI models and tools used along with evaluation methods, why they are used, and how they are used (i.e. which solution).
- Create data diagrams showing how data flows to and from AI models and tools, including data type and data lineage.
- Define acceptable data types used for each model and/or tool.
- Using data flow diagrams, apply zero trust concepts to create necessary security controls and detections to mitigate risks (defined in AI risk register).
- For relevant risks (from AI risk register), perform Incident Response tabletop exercises.
- Perform penetration tests on AI infrastructure to test controls.
- Train staff on AI security best practices.

Building Your AI Governance Framework: A Practical Blueprint

Here are some suggestions based upon nCino's governance model on how your institution can establish an effective AI governance structure that balances innovation with responsible oversight.

Step 1: Establish Your Structure – nCino Uses a Two-Team Model

Executive AI Council (Strategic Level)

- Who to include: CEO or designated C-suite sponsor; heads of Technology, Risk, Legal, Operations; and your Data/Privacy Officer
- Charter: Set AI strategy, approve major investments, resolve escalated issues
- Meeting cadence: Start monthly during setup, then shift to quarterly

AI Operations Team (Tactical Level)

- Who to include: Senior managers from IT, Security, Data Science, Product, and Compliance
- Charter: Implement policies, conduct assessments, monitor performance
- Meeting cadence: Bi-weekly initially, adjust based on AI activity volume

Step 2: Define Your Core Processes

- ✔ AI Impact Assessment Framework: Create a standardized evaluation process for all AI initiatives:
- ✔ Business justification: Why this AI? What value does it deliver?
- ✔ Technical assessment: Can we implement and maintain it effectively?
- ✔ Risk evaluation: What could go wrong? How do we mitigate?
- ✔ Compliance check: Does it meet regulatory and contractual requirements?
- ✔ Resource planning: Which people, what technology, and what budget do we need?

Model Governance Protocol: Establish controls for any AI models you develop or deploy:

- ✔ Pre-deployment: Data quality checks, performance testing, bias assessment
- ✔ Approval gates: Clear signoffs before production use
- ✔ Post-deployment: Continuous monitoring, drift detection, incident response

Step 3: Start Small, Scale Smart

Month 1-2: Form committees, draft charter documents

Month 3-4: Pilot governance process with 1-2 low-risk AI projects

Month 5-6: Refine processes based on lessons learned

Month 7+: Scale to all AI initiatives

Key Success Factors:

- Executive sponsorship is non-negotiable
- Document decisions and rationales
- Regular communication across teams
- Flexibility to adapt as you learn

Remember: The goal isn't perfection on day one—it's establishing a foundation upon which you can build as your AI maturity grows.

Conclusion and Next Steps

As financial institutions navigate the rapidly evolving AI landscape, success depends on taking a thoughtful, strategic approach that balances innovation with responsibility. The opportunities for AI to transform banking operations, enhance customer experiences, and drive competitive advantage are substantial, but they must be pursued with appropriate governance, risk management, and regulatory compliance.

Essential Success Factors

1. Clear AI Definition and Scope

- Establish what constitutes AI for your institution
- Avoid definitions that are too broad or too narrow
- Use industry standards like the OECD definition as a foundation
- Let this definition guide all governance decisions

2. Risk-Proportionate Governance

- Implement tiered risk assessments matching scrutiny to impact
- Avoid one-size-fits-all approaches that stifle innovation
- Focus intensive oversight on high-risk applications
- Streamline processes for low-risk efficiency tools

3. Strong Governance Structure

- Separate strategic oversight from operational management
- Ensure C-suite engagement and sponsorship
- Build cross-functional teams with diverse expertise
- Establish clear escalation paths and decision rights

4. Practical Implementation Approach

- Start with pilot programs on low-risk use cases
- Build on existing governance frameworks
- Document decisions and learn from experience
- Scale gradually based on proven success

5. Focus on Real vs. Perceived Risks

- Address actual risks like model degradation and data quality
- Implement appropriate explainability measures
- Establish continuous monitoring and improvement processes
- Avoid overreacting to misconceptions about AI

Partnership with **nCino**

As your institution embarks on its AI journey, nCino stands ready to be your trusted partner. With our comprehensive cloud banking platform, built-in AI capabilities, and commitment to responsible AI development, we can help you realize the full potential of AI while maintaining the highest standards of governance, compliance, and customer protection.

How nCino Can Help

Platform Implementation: Our platform provides built-in AI capabilities with comprehensive governance and compliance controls.

Best Practice Sharing: Learn from our experience supporting over 2,700 financial institutions in their digital transformation journeys.

Ongoing Support: Our dedicated support and success teams ensure your AI initiatives deliver maximum value with minimal risk.

The Future of AI in Banking

The financial services industry stands at an inflection point. Institutions that thoughtfully embrace AI while maintaining strong governance and ethical standards will thrive in the digital future. Those that hesitate or implement AI without appropriate safeguards risk falling behind or facing significant regulatory and reputational challenges.

The choice is clear: embrace AI responsibly or risk irrelevance in an increasingly AI-driven marketplace. With the right approach, governance framework, and technology partners, your institution can harness the power of AI to better serve customers, improve operations, and drive sustainable growth.

1. https://www.ncino.com/en-US/our-customers
2. https://oecd.ai/en/ai-principles
3. https://www.fsisac.com/hubfs/Knowledge/AI/FSISAC_ResponsibleAI-Principles.pdf

ncino.com