SOPHIZO

INDUSTRY PLAYBOOK · FINANCIAL SERVICES & BANKING

You can audit a balance sheet. Can you audit an algorithm?

Agentic AI for fraud, regulatory reporting, and the back office that built modern finance.

PUBLISHED 2026 · **SOPHIZO RESEARCH**

Fractional AI & RevOps Leadership
sophizo.net

Compliance moves in days. Markets move in microseconds.

Banks, insurers, and fintechs sit on the cleanest transactional data in any industry. And the most fragmented control environment for using it. The opportunity is not another model. It is an operating model where agents act in real time inside a control framework your board can defend.

Why this matters now

Fraud teams investigate yesterday's losses. Compliance teams spend two weeks producing a report regulators read in twenty minutes. Trading desks chase signals that have already decayed. An AI Officer reframes the problem from which model do we deploy to how do we instrument autonomy without losing the plot in front of a regulator.

> *Your model risk inventory has 137 entries. How many are agents, and how many know it?*

Model risk is not a new idea in banking. AI is forcing every existing framework to extend.

Beyond the six global regimes, financial services & banking carries the overlays below. Each one has its own enforcement model and its own evidence expectation.

SR 11-7 Federal Reserve SR 11-7 Model Risk Management

UNITED STATES, FEDERAL RESERVE + OCC

CRITICAL

Applies to. Any model used in credit decisions, capital calculations, anti-money-laundering, fraud detection, or material business decisions at a regulated bank.

Key obligation. Three pillars: model development standards, validation independent of development, governance and policies. Applies to ML and generative AI as models.

Evidence. Model inventory, validation reports, conceptual soundness documentation, ongoing monitoring outputs.

NYDFS AI NYDFS Part 500 + AI Guidance (Oct 2024)

UNITED STATES, NEW YORK HIGH

Applies to. Banks, insurers, and other entities licensed by the NY Department of Financial Services.

Key obligation. Treat AI as both a tool and a threat. Address AI-enabled social engineering, deepfake-based fraud, and model vulnerabilities in the cybersecurity program.

Evidence. Risk assessment covering AI threats, training records, third-party AI vendor due diligence, incident reporting.

CFPB AI Credit CFPB ECOA + Reg B AI Adverse Action Guidance

UNITED STATES, CFPB HIGH

Applies to. Any AI or ML model used in consumer credit decisions, pricing, or adverse actions.

Key obligation. Specific, accurate adverse action reasons even for complex models. Cannot use "black box" as a reason.

Evidence. Adverse action reason logic documentation, model explainability reports, fair lending monitoring.

Basel + ECB AI Basel Committee Principles + ECB Guide on AI

GLOBAL, PRUDENTIAL

ELEVATED

Applies to. Globally systemic banks and ECB-supervised institutions using AI in risk, finance, or treasury.

Key obligation. Operational risk treatment of AI, third-party model risk, board-level oversight expectations.

Evidence. Board AI risk reporting, third-party model assessments, operational risk register entries.

Four capability domains. One operating layer.

01 Real-Time Fraud Investigation

- Multi-agent investigation orchestration
- Case-prep with evidence chains pre-assembled
- Cross-channel signal correlation
- SAR drafting with analyst review

02 Regulatory Reporting and Compliance

- Reporting cycle: 15 days to under 1 hour
- EU AI Act, GLBA, MAR, AML alignment
- Explainability layer on every model output
- Continuous control monitoring

> *If the OCC asked tomorrow which decisions a model made unsupervised, could you answer in 48 hours?*

Capability domains, continued.

03 Trading and Risk Operations

- Risk-tiered execution autonomy
- Pre-trade compliance checks at the agent layer
- Scenario simulation with replay logging
- Exception-handling agents for the ops desk

04 Claims, Lending and Back-Office

- Claims triage and adjudication assistance
- Loan-decision packet preparation
- Reconciliation with auto-resolved breaks
- Vendor-invoice processing

05 · OUTCOMES AND ROI REALITY

What production deployments look like at scale.

192%
US AVERAGE AI ROI IN FINANCIAL SERVICES

7 to 12 mo
MEDIAN PAYBACK PERIOD

66%
OF FIRMS REPORTING PRODUCTIVITY GAINS IN PRODUCTION

Production-stage benchmarks compiled from industry research, Federal Reserve supervisory reporting, and Deloitte 2024 State of AI in Financial Services. Your spread depends on data lineage maturity, model governance posture, and the regulator you answer to.

The AI Officer Mandate.

Three responsibilities a Fractional AI Officer owns from day one in financial services & banking.

01
Model-risk management aligned to SR 11-7 and your internal governance committee.

02
Bias and fairness audits built into deployment pipelines. Not retrofitted after a regulator asks.

03
Risk-tiered autonomy: agents make small decisions, escalate the consequential ones, log everything.

06 · A 90-DAY PATH

How a Sophizo engagement starts in Financial Services & Banking.

DAYS 1 TO 30

Diagnose

MAP THE OPERATING REALITY

- Model inventory review with second-line risk team
- Governance committee charter alignment (SR 11-7, EU AI Act, NIST AI RMF)
- Fraud and AML signal-loss baseline, dollarized
- Joint readiness session with CRO, CIO, and head of compliance

DAYS 31 TO 60

Architect

DESIGN THE AUTONOMY BOUNDARY

- Agent permissions and escalation policy
- Evidence file and audit trail design
- First production pilot scoped with rollback plan
- Cross-functional governance committee charter

DAYS 61 TO 90

Operate

SHIP AND INSTRUMENT

- First agent in production with HITL controls
- Operator coaching and policy refinement
- P&L instrumentation by use case
- Quarterly review cadence established

What we will not do.

We do not take over your model risk inventory. That is your second line of defense and it should stay there. We do not write your SR 11-7 documentation for you, and we do not do regulatory exam coaching. We pass on engagements where the CRO sees agentic AI as an IT initiative, because that framing alone tells us governance will be retrofitted at exam time, which is too late.

07 · WHAT THE BOARD NEEDS TO HEAR

Five things the board needs to hear about AI in banking.

Five cited insights for the next risk-committee meeting. Each one is sourced. Each one is what an experienced AI Officer would put in front of the board if they walked in tomorrow.

01 · MODEL RISK

SR 11-7 already governs your agents. You may not have called them models yet.

The Federal Reserve's SR 11-7 has been the model-risk standard since 2011. A retrieval-augmented agent that influences a credit decision, a fraud disposition, or a regulatory filing is a model under SR 11-7. Full stop. The fastest path to defensible agent governance is to put the agent through the same inventory, validation, and ongoing-monitoring discipline your model risk function already runs.

Source. Federal Reserve SR 11-7 (2011 , OCC 2011-12.

02 · THIRD-PARTY RISK

OCC Bulletin 2023-17 makes vendor-agent risk a bank operational risk event.

The June 2023 interagency third-party risk guidance applies to every AI vendor a bank uses. Concentration risk on a single foundation model provider, prompt-routing dependency, and sub-processor visibility now belong in the operational risk register. Boards that cannot answer the concentration question on exam will lose that exchange.

Source. OCC Bulletin 2023-17 (Joint, FRB and FDIC), June 2023.

03 · STATE PRESSURE

NYDFS Part 500 and the NAIC AI bulletin set the template other states copy.

New York's Second Amendment to 23 NYCRR 500 (November 2023) names AI vendor risk explicitly. The NAIC December 2023 model bulletin extends the same logic to insurers. Multi-factor authentication, vendor due diligence, and governance now extend to generative AI by reference. The states are not waiting for federal action.

Source. NYDFS 23 NYCRR 500, Second Amendment, November 202 , NAIC AI Model Bulletin, December 2023.

Two more, then the framework.

04 · ADVERSE ACTION

CFPB has already said the algorithm is not a defense.

CFPB Circular 2022-03 and its 2024 follow-ups make clear that creditors using complex models still owe specific, accurate adverse-action notices under ECOA. "The model decided" is not a notice. The implication for credit, lending, and pricing agents is that explainability is a first-class deployment requirement, not a post-launch project.

Source. CFPB Circular 2022-03, May 202 ,CFPB Spring 2024 supervisory highlights.

05 · AML LEVERAGE

Financial crime compliance is where agents pay back fastest.

Banks spend roughly 12 to 15 percent of revenue on financial crime compliance, against alert false-positive rates that frequently exceed 95 percent on legacy systems. Agentic triage of AML alerts has cut analyst time per case by 40 to 60 percent in early production deployments. The Wolfsberg Group's 2024 technology statement signals industry alignment.

Source. BCG 2024 Global Risk Repor ,Wolfsberg Group 2024 Statement on Technology in Financial Crime Compliance.

Two-Pillar Agent Governance.

Every banking agent answers to two pillars. Model risk on one side. Operational and third-party risk on the other. Build the artifact set for both before launch and the exam team has nothing to negotiate.

PILLAR A1

Model inventory

SR 11-7 inventory entry with tier, owner, validator, and decision boundary.

PILLAR A2

Validation file

Bias and fairness, performance, robustness, explainability. Reviewed by second line.

PILLAR A3

Ongoing monitoring

Drift thresholds, revalidation cadence, and clear retire criteria.

PILLAR B1

Operational risk register

Mapped to your existing loss-event taxonomy and risk-appetite statement.

PILLAR B2

Third-party file

OCC 2023-17 vendor due diligence, sub-processors, concentration analysis.

PILLAR B3

Incident playbook

Tabletop-tested against your CRO's tolerance. Reviewed annually.

From John Utley.

> *Banks already know how to govern a model. The mistake is treating agents as IT projects so they sidestep the model risk machine you spent fifteen years building. Put the agent through SR 11-7. The exam team will not negotiate that on your behalf, and the OCC reads the same press releases you do.*

John Utley

FOUNDER, SOPHIZO · SEATTLE, WA

John Utley founded Sophizo to give growth-stage companies the AI and revenue architecture work historically reserved for the Fortune 500. He writes and advises on agentic AI governance, predictive forecasting, and operating-model design for boards and operators across financial services & banking and adjacent sectors.

Test your operating picture against these.

1

Your model risk inventory has 137 entries. How many are agents, and how many know it?

2

If the OCC asked tomorrow which decisions a model made unsupervised, could you answer in 48 hours?

3

You budget two FTEs to a quarterly report regulators read in 20 minutes. Why?

Frequently asked questions.

How do you handle explainability for regulated decisions?

Every agent decision is logged with the inputs, the model version, the policy that allowed the action, and the human who could have intercepted it. We generate explanation artifacts in the form your model risk team and regulators already expect.

Can you work alongside our existing model risk function?

Yes. We partner with model risk, not around it. Most engagements start with a joint review of your existing inventory and governance committee charter so we layer agentic controls on top of, not parallel to, your established framework.

What about EU AI Act and SR 11-7 alignment?

Both. We treat agentic systems as high-risk by default and apply the disclosure, monitoring, and human-oversight requirements of the EU AI Act, mapped to your existing SR 11-7 governance.

If this maps to your operating reality, we should talk.

The Diagnostic Sprint is two weeks. Board-ready output. Tailored to financial services & banking.

ENGAGE
sophizo.net/checkout/diagnostic-sprint

INDUSTRY PAGE
sophizo.net/industries/financial-services-banking

EMAIL
john@sophi

12 · SOURCES

Primary research behind this brief.

Every claim, statistic, and citation in this playbook traces back to one of the primary sources below. Pressure-test any of them with your team. We have done the same.

01. Board of Governors of the Federal Reserve System.

SR 11-7: Guidance on Model Risk Management, April 2011.

02. Office of the Comptroller of the Currency.

Bulletin 2023-17, Third-Party Relationships: Interagency Guidance, June 2023.

03. New York Department of Financial Services.

23 NYCRR 500 Cybersecurity Regulation, Second Amendment, November 2023.

04. National Association of Insurance Commissioners.

Model Bulletin on the Use of Artificial Intelligence Systems by Insurers, December 2023.

05. Consumer Financial Protection Bureau.

Circular 2022-03, Adverse Action Notification Requirements, May 2022.

06. BCG.

2024 Global Risk Report.

07. Wolfsberg Group.

2024 Statement on the Application of Technology in Financial Crime Compliance.

08. Deloitte.

2024 State of AI in Financial Services.

Editorial note. This brief is a field reference compiled by Sophizo Research. It is not legal, accounting, or clinical advice. Cite the primary regulator guidance for binding interpretation. Where statistics are quoted, the most recent published figure as of early 2026 is used.