

# Generative AI and the Transformation of Risk Management and Compliance

## RISK MANAGEMENT AND COMPLIANCE ARE AT THE FOREFRONT OF ENTERPRISE AI ADOPTION

Risk management and compliance functions represent some of the most compelling and immediately viable applications of generative artificial intelligence (AI) in financial services.<sup>1</sup> In contrast to client-facing advisory tools or trading systems, these functions operate almost entirely within the firm, remain subject to structured oversight, and do not directly handle client assets or execute external transactions. As a result, they combine high potential value with comparatively limited downside risk. In many respects, risk and compliance are already information-intensive, judgment-driven, and documentation-heavy disciplines—precisely the domains in which generative AI performs best when deployed as a decision support tool.

**These functions are prime candidates for generative AI not simply because they are lower in risk but also because of how this work is organized.** The core outputs of second-line functions are not trades or product decisions but interpretations, assessments, narratives, and escalations. These are developed based on regulatory texts, supervisory guidance, internal policies, investigation notes, alerts, audit findings, and examination correspondence, all of which are overwhelmingly unstructured. Much of the value created by risk and compliance teams lies in their ability to synthesize this information across silos, apply judgment, and produce defensible explanations for senior management and regulators. Generative AI is uniquely suited to this task because it is designed to reason across large volumes of unstructured material and translate fragmented inputs into coherent, context-aware outputs.<sup>2</sup>

1 Rahul Agarwal, Andreas Kremer, Ida Kristensen, and Angela Luget, "How Generative AI Can Help Banks Manage Risk and Compliance," McKinsey & Company, March 1, 2024, <https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/how-generative-ai-can-help-banks-manage-risk-and-compliance>.

In April 2026, the Cambridge Centre for Alternative Finance at the University of Cambridge published "The 2026 Global AI in Financial Services Report: Adoption, impact, and risks." The Report was led by Bryan Zhang and Kieran Garvey and done in collaboration with a consortium of global institutions including the Bank for International Settlements (BIS), the World Economic Forum (WEF), and the Arab Monetary Fund (AMF). It captures the perspectives of 628 financial institutions. The study provides a robust foundation for forward-looking dialogue on AI in financial services.

2 Hariharan Pappil Kothandapani, "AI-Driven Regulatory Compliance: Transforming Financial Oversight Through Large Language Models and Automation," *Figshare* 3 (2025): 12–24, <https://doi.org/10.6084/m9.figshare.28251608>.

James Da Costa and Angela Strange, "Everything, Everywhere is Compliance," *A16Z News*, May 26, 2026, <https://www.a16z.news/p/everything-everywhere-is-compliance>.

**Equally important, risk management and compliance are already built around human review and accountability.** Second-line professionals are trained to evaluate outputs produced by systems—transaction monitoring engines, surveillance tools, risk models—and to decide when escalation is warranted. Generative AI fits naturally into this operating model as an augmentation layer rather than as a sole decision maker. It can assemble information, propose analyses, and draft narratives, but it does not displace the human obligation to review, advise, challenge, and approve. This embedded human-in-the-loop structure significantly constrains downside risk and makes governance more tractable than in autonomous or client-facing use cases.

**Risk and compliance functions also face persistent scalability pressures.** Regulatory expectations continue to expand in volume, granularity, and speed, while institutions confront structural limits on headcount growth and increasing competition for experienced compliance talent.<sup>3</sup> Generative AI offers a way to manage this complexity and scale without linear increases in staffing by reducing the manual effort required to gather information, reconcile inconsistencies, and produce documentation.<sup>4</sup> In this sense, these functions are not merely early adopters of AI; they are among the areas where AI adoption is increasingly necessary to maintain baseline effectiveness.

Generative AI and agentic AI change the risk and compliance function not by replacing existing detection systems but by **transforming how information flows through the organization and how judgments are formed and documented.** Traditional compliance technology excels at identifying signals—alerts,

thresholds, breaches, and exceptions—but leaves humans to perform the cognitively demanding work of interpretation, correlation, and explanation. This division of labor has long constrained second-line effectiveness, particularly as data volumes and regulatory expectations have grown.

**Generative AI addresses this limitation by operating at the layer where risk and compliance historically struggle: synthesis.** It can ingest outputs from multiple systems, connect related facts across time and business lines, and produce structured summaries that reflect both context and nuance. This capability fundamentally alters how risk and compliance professionals spend their time, shifting effort away from manual compilation and toward evaluation, challenge, and escalation. The result is not simply faster decision making, but better-informed and more consistent decision making<sup>5</sup> (see figures 1 and 2).

**Agentic AI extends this shift by coordinating complex, multistep workflows that previously required extensive manual orchestration.** In a compliance context, this may include gathering relevant data from multiple sources, assembling an investigation file, drafting an initial narrative, identifying missing information, and preparing materials for review—all before a human makes a substantive determination.<sup>6</sup> Crucially, agentic systems do not eliminate checkpoints; they reorganize work so that human review occurs at the point of highest value rather than at every intermediate step. This reconfiguration of workflow materially increases the capacity of second-line teams without altering accountability.

- 
- 3 Tom Rollauer, "Demand for Compliance and Risk Skills Leaves a Talent Shortage," Deloitte and *The Wall Street Journal*, August 13, 2014, [https://deloitte.wsj.com/riskandcompliance/demand-for-compliance-and-risk-skills-leaves-a-talent-shortage-1407902503?gaa\\_at=eafs&gaa\\_n=AWETsqf5gIrmO1y05HWXjLAlzlB9f54XdFkdU4vdjX0eO1Dy\\_NJTfIfsaEYqRT-zXU%3D&gaa\\_ts=6957d25b&gaa\\_sig=D2JFBSaocpj\\_-WJm29UMASVl4bNMGmEFw--GjQkD42Rn5yBagJ0PQLivT9lFon8RMn\\_h6DyBCsE6HFZJThvVzA%3D%3D](https://deloitte.wsj.com/riskandcompliance/demand-for-compliance-and-risk-skills-leaves-a-talent-shortage-1407902503?gaa_at=eafs&gaa_n=AWETsqf5gIrmO1y05HWXjLAlzlB9f54XdFkdU4vdjX0eO1Dy_NJTfIfsaEYqRT-zXU%3D&gaa_ts=6957d25b&gaa_sig=D2JFBSaocpj_-WJm29UMASVl4bNMGmEFw--GjQkD42Rn5yBagJ0PQLivT9lFon8RMn_h6DyBCsE6HFZJThvVzA%3D%3D).
- 4 Hariharan Pappil Kothandapani, "AI-Driven Regulatory Compliance: Transforming Financial Oversight Through Large Language Models and Automation," *Figshare* 3 (2025): 12–24, <https://doi.org/10.6084/m9.figshare.28251608>.
- 5 Natalie Dytrych, "How AI Helps Solve the 4 Biggest Challenges in Regulatory Compliance," *Optro*, December 4, 2025, <https://auditboard.com/blog/how-ai-helps-solve-the-4-biggest-challenges-in-regulatory-compliance>.
- 6 Bernd Federmann, "How AI Is Poised to Reshape Compliance Functions," KPMG Law, October 30, 2025, <https://kpmg.com/xx/en/our-insights/ai-and-technology/how-ai-is-poised-to-reshape-compliance-functions.html>.

**Figure 1. Traditional versus generative AI compliance workflow**

**TRADITIONAL (BEFORE)**



**GENERATIVE AI (AFTER)**



Generative AI enables continuous real-time analysis and automates manual aggregation and synthesis tasks. Human review becomes more efficient and of higher value, allowing experts to focus on oversight, judgment, and final decision making rather than manual processing.

**Agentic AI** refers to generative AI systems designed not just to produce content but also to take actions toward goals—such as planning tasks, using tools, making decisions, or interacting with software and data sources with some level of autonomy. Created by combining LLMs with other tools (APIs, code), agentic systems are proactive

where generative AI is reactive, and are more capable of complex reasoning and managing multilevel workflows. The reduction of human oversight in agentic AI use has important implications for accountability and security, particularly in how it is integrated into a real company’s workflows.

**Figure 2. Generative AI use cases in finance: Value versus risk spectrum**

| Use Case                      | Value | Risk     |
|-------------------------------|-------|----------|
| Anti-money laundering / Fraud | High  | Moderate |
| Cybersecurity                 | High  | Moderate |
| Compliance                    | High  | Low      |

In many of these applications, generative AI systems can operate continuously (24/7), significantly improving monitoring capacity and response speed. Human expertise remains essential for review, oversight, and final decision making, while AI reduces repetitive manual work and streamlines analysis.

## STRATEGIC APPLICATIONS

Within this framework, several use cases stand out as particularly well suited to generative AI deployment.

One of the most mature and valuable applications lies in **cybersecurity risk management**, especially given the increasing sophistication of cybercrime—see the ransomware-as-a-service business model as an example.<sup>7</sup> Security operations generate vast quantities of alerts and telemetry, much of which must ultimately be interpreted, escalated, and documented for management and regulators. Generative AI can summarize incidents,<sup>8</sup> correlate signals across systems,<sup>9</sup> and draft incident reports, escalation memos, and regulator-facing narratives. For second-line functions, this enhances oversight of first-line security operations, improves the quality and consistency of incident documentation, and strengthens readiness for supervisory scrutiny. As institutions increasingly rely on AI systems themselves, cybersecurity risk management must also encompass the security of models, data pipelines, and third-party dependencies—an area where generative tools can assist in mapping controls, identifying gaps, and maintaining defensible governance records.<sup>10</sup>

**Anti-money laundering (AML), sanctions, and Financial Action Task Force (FATF) Travel Rule compliance** represent another high-impact domain. The United Nations Office on Drugs and Crime estimates that between 2 and 5 percent of global GDP is lost to money laundering every year.<sup>11</sup> Existing compliance programs are heavily text-based, reliant on complex judgment, and subject to intense supervisory scrutiny. Generative AI can support second-line oversight across enterprise risk assessments, transaction

monitoring governance, investigation quality assurance, and regulatory reporting. In AML and sanctions contexts, these tools can summarize alerts, standardize investigative narratives, and support the drafting and review of suspicious activity reports while preserving human decision-making authority.<sup>12</sup> For Travel Rule compliance, particularly in digital asset markets, generative AI could potentially assist with jurisdictional applicability analysis, exception handling documentation, audit preparation, and regulator engagement. The goal would be greater consistency, clearer audit trails, and improved scalability as transaction volumes and regulatory expectations increase.

More broadly, supervisory expectations in AML have increasingly shifted toward whether monitoring systems can operate effectively at the scale and complexity of modern financial activity. The consequences of failing to invest in adequate compliance infrastructure were illustrated starkly in October 2024, when TD Bank entered a \$3.1 billion resolution with the US Department of Justice (DOJ), FinCEN, the Office of the Comptroller of the Currency (OCC), and the Federal Reserve—the largest penalty ever imposed under the Bank Secrecy Act. According to the DOJ, TD Bank’s transaction monitoring systems failed to cover approximately 92 percent of transaction volume—roughly \$18.3 trillion in activity—in part because leadership maintained a “flat cost paradigm” despite known control gaps and rising transaction volumes. Three money laundering networks moved more than \$670 million through these monitoring blind spots. The case illustrates both the regulatory cost of underinvestment in compliance technology and the growing expectation

7 Syed Jawwad, “The Rise of Ransomware-as-a-Service: How Cybercrime Has Become a Business,” IBM, accessed May 11, 2026, <https://www.ibm.com/think/insights/the-rise-of-raas>.

8 Nachaat Mohamed, “Artificial Intelligence and Machine Learning in Cybersecurity: A Deep Dive into State-of-the-Art Techniques and Future Paradigms,” *Knowledge and Information Systems* 67 (2025): 6969–7055, <https://doi.org/10.1007/s10115-025-02429-y>.

9 Ibid.

10 Clarke Rodgers, Moumita Saha, Dimple Ahluwalia, Kevin Skapinetz, and Gerald Parham, “Securing Generative AI,” IBM, accessed May 15, 2026, <https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/securing-generative-ai>.

11 “Money Laundering,” United Nations Office on Drugs and Crime, accessed May 22, 2026, <https://www.unodc.org/unodc/en/money-laundering/overview.html>.

12 Carol Beaumier and Constantine Boyadjiev, “From Hindsight to Insight to Foresight: FAQ Guide on the Use of AI for Financial Crime Compliance,” *Protiviti*, accessed May 26, 2026, <https://www.protiviti.com/us-en/research-guide/faq-use-of-ai-for-financial-crime-compliance>.

that AML monitoring systems be capable of operating at the scale of contemporary financial markets.<sup>13</sup>

At the same time, emerging empirical research suggests that well-governed AI deployments may materially improve not only the efficiency of AML operations, but also the quality, consistency, and defensibility of investigative outputs. A 2025 study by Gu and colleagues<sup>14</sup> examined a human-machine collaborative AML system deployed at a major financial institution, in which generative AI was integrated into alert triage and investigative narrative generation across a daily workflow of approximately 53 million transactions. The institution reported a 29 to 36 percent improvement in the early detection of suspicious activity, a 33 percent reduction in alert queue backlog, and a 17 percent gain in per capita analyst productivity. Equally notable, the system produced “auditable investigation skeletons” that preserved an explicit chain of reasoning and evidence, directly addressing supervisory expectations for defensibility and transparency. These findings suggest that, in mature and well-governed implementations, generative AI does not merely accelerate existing workflows, but may enhance the analytical rigor, consistency, and scalability of financial crime surveillance itself.

**Anti-fraud** oversight can similarly benefit from generative AI’s ability to synthesize behavioral and network-based information, especially as criminals increasingly deploy AI in fraud attacks.<sup>15</sup> The Federal Trade Commission reported that Americans lost \$12.5 billion to fraud in 2024, a 25 percent increase from the previous year.<sup>16</sup> As fraud schemes become more

synthetic, socially engineered, and cross-channel, traditional rule-based approaches generate increasing volumes of alerts that require human interpretation. Generative AI can enrich alerts with contextual summaries, identify recurring patterns across cases, and draft investigation narratives and remediation recommendations. Having piloted an AI-driven AML system in 2021, HSBC reported in 2024 that their crime detection rates have improved by two to four times, and false positives have been reduced by 60 percent.<sup>17</sup> In 2021, JPMorgan Chase integrated large language models into their payment validation process, reporting improved review times and lower levels of fraud as a result.<sup>18</sup> Used appropriately, these tools allow compliance teams to focus their expertise on judgment, escalation, and control enhancement rather than manual data assembly, improving both effectiveness and efficiency.

**Regulatory compliance and supervisory readiness constitute a fourth core use case.** Financial institutions face a constant flow of new rules, interpretive guidance, and enforcement requirements across jurisdictions. Generative AI can support regulatory horizon scanning, extract obligations from dense regulatory text, and map requirements to products, controls, and policies.<sup>19</sup> Generative AI can extend these functions by assisting in drafting and maintaining policies, procedures, and disclosures, as well as in preparing regulatory filings, exam responses, and evidence binders. By turning unstructured regulatory information into actionable compliance outputs, generative AI helps institutions respond more quickly and consistently to

13 “TD Bank Pleads Guilty to Bank Secrecy Act and Money Laundering Conspiracy Violations in \$1.8B Resolution,” US Department of Justice, news release, October 10, 2024, <https://www.justice.gov/archives/opa/pr/td-bank-pleads-guilty-bank-secrecy-act-and-money-laundering-conspiracy-violations-18b>.

14 Xiaoxiong Gu, Jingwen Yang, Xia Tian, and Min Liu. “Research on the Construction of a Human-Machine Collaborative Anti-Money Laundering System and Its Efficiency and Accuracy Enhancement in Suspicious Transaction Identification,” In *Proceedings of the 2025 5th International Conference on Big Data, Artificial Intelligence and Risk Management*, (2025): 530–34, <https://doi.org/10.1145/3800227.3800306>.

15 “10 Global Compliance Concerns for 2026,” Thomson Reuters Institute, December 10, 2025, <https://www.thomsonreuters.com/en/reports/10-global-compliance-concerns-for-2026>.

16 “New FTC Data Show a Big Jump in Reported Losses to Fraud to \$12.5 Billion in 2024,” Federal Trade Commission, news release, July 31, 2025, <https://www.ftc.gov/news-events/news/press-releases/2025/03/new-ftc-data-show-big-jump-reported-losses-fraud-125-billion-2024>.

17 Jennifer Calvery, “Harnessing the Power of AI to Fight Financial Crime,” HSBC, June 10, 2024, <https://www.hsbc.com/news-and-views/views/hsbc-views/harnessing-the-power-of-ai-to-fight-financial-crime>.

18 “How AI Will Make Payments More Efficient and Reduce Fraud,” J.P.Morgan: Payments Security & Trust, November 20, 2023, <https://www.jpmorgan.com/insights/payments/security-trust/ai-payments-efficiency-fraud-reduction>.

19 Hariharan Pappil Kothandapani, “AI-Driven Regulatory Compliance: Transforming Financial Oversight Through Large Language Models and Automation,” *Figshare* 3 (2025): 12–24, <https://doi.org/10.6084/m9.figshare.28251608>.

regulatory change while maintaining a clear record of rationale and accountability.<sup>20</sup>

The efficiency gains in this domain can be substantial. A 2026 industry account documented two complementary compliance deployments at Marqeta, a large fintech. In monitoring and testing, the firm replaced statistical sampling with AI-enabled 100 percent population testing, allowing existing staff to manage significant volume increases while returning two previously allocated headcount to the business. Separately, the team's annual policy review—historically a multi-week exercise of parsing statutes, regulations, and supervisory guidance—was compressed by approximately tenfold, freeing senior compliance capacity for judgment-intensive work.<sup>21</sup>

Across these domains, the benefits of generative AI are largely shared. These tools offer meaningful scalability, allowing risk and compliance functions to manage growing complexity without proportional increases in headcount. They promote consistency by standardizing narratives, analyses, and documentation across teams and time. They improve speed, enabling faster response to incidents, regulatory developments, and supervisory inquiries. And they enhance defensibility by creating clearer audit trails and more coherent explanations of decisions and controls. Perhaps most importantly, they allow skilled professionals to focus on judgment and oversight rather than manual synthesis of information.

At the same time, the effectiveness of generative AI in risk management and compliance depends critically on the expertise of the professionals who use it, the quality of the data, and the ability to easily move the data around the organization. These tools do not reduce

the need for highly skilled compliance, cybersecurity, AML, or anti-fraud practitioners; rather, they heighten it. Generative and agentic AI systems amplify human judgment, but they do not substitute for it.<sup>22</sup>

To use these models responsibly, institutions must rely on sophisticated professionals who have a deep understanding of the underlying regulatory frameworks, industry practices, threat typologies, and control environments—and who also possess deep fluency in AI capabilities and limitations to evaluate outputs critically. These systems cannot be safely or effectively operated by untrained staff or novices.<sup>23</sup> They require a next phase of professional skill in the industry: the ability to frame the right questions, challenge model assumptions, interpret outputs in context, and recognize when escalation is warranted. In this sense, generative AI significantly raises the cognitive and analytical demands of second-line functions, shifting the role from process execution toward higher-order reasoning and oversight rather than simple prompt entry or mechanical review.<sup>24</sup>

Effective use of generative AI in risk management and compliance also depends on strong data governance and integrated data infrastructure. Firms must be able to ensure the quality and integrity of data and enable seamless data flow across first- and second-line functions and across business lines such as commercial banking, asset management, and wealth management. Where data and technology remain siloed across divisions—or where firms lack unified, authoritative enterprise data that can be accessed consistently across the organization—the ability to deploy generative AI effectively at scale will be significantly constrained.

---

20 See, for example, CUBE Global Regulatory Compliance Solutions, <https://cube.global/solutions>.

21 S. Cameron, "From Skeptic to Evangelist: A First Mover's Playbook for AI in Compliance," *ComplyAdvantage*, May 15, 2026, <https://complyadvantage.com/insights/playbook-for-ai-in-compliance/>.

22 Hariharan Pappil Kothandapani, "AI-Driven Regulatory Compliance: Transforming Financial Oversight Through Large Language Models and Automation," *Figshare* 3 (2025): 12–24, <https://doi.org/10.6084/m9.figshare.28251608>.

23 "10 Global Compliance Concerns for 2026," Thomson Reuters Institute, December 10, 2025, <https://www.thomsonreuters.com/en/reports/10-global-compliance-concerns-for-2026>.

24 The capacity-building requirements discussed here extend beyond second-line professionals. C-suite executives and boards of directors must develop sufficient fluency in the capabilities and limitations of generative and agentic AI to ensure its appropriate and effective deployment across the enterprise. A related question concerns the expertise of regulators and supervisors themselves, whose ability to evaluate firms' AI-enabled compliance programs—and to deploy AI tools in their own supervisory work—depends on comparable technical and analytical depth. These broader governance and supervisory dimensions lie outside the scope of the present paper but warrant dedicated treatment in subsequent work.

## Data Governance

**AI governance** is inseparable from data governance because the effectiveness, reliability, and integrity of AI systems depend fundamentally on the quality and stewardship of the data on which they rely. As financial institutions increasingly deploy AI across risk management and compliance operations, governance frameworks must extend beyond models themselves to encompass the full data life cycle—including data collection, access, usage, retention, protection, and provenance. Weak data governance can undermine model performance, create operational and legal risk, and expose institutions to privacy, cybersecurity, and regulatory failures at enterprise scale.

**Data management** policies are the first key component of effective data governance. Policies covering data quality, ownership, privacy, and life cycle management should explicitly address AI use cases. For any AI project, the input data should be assessed for quality (accuracy, completeness) and suitability. Companies should define what data are permissible to use—for instance, prohibiting use of data collected without proper consent or containing biased sampling. If AI models use personal data, data protection impact assessments may also be appropriate in certain jurisdictions or under company policy. Strong data governance mitigates the “garbage in, garbage out” problem while reducing legal and operational risk by ensuring that AI systems rely on accurate, authorized, and appropriately governed data.

**Privacy considerations** are increasingly becoming integral to AI governance and system design. As institutions

deploy AI systems that rely on large volumes of customer, employee, and transactional data, issues surrounding data minimization, retention, access, consent, and individual rights become more significant. Organizations may incorporate techniques such as anonymization or pseudonymization in model training, establish standards for retaining training data, and develop processes to address privacy obligations, including deletion requests under regimes such as the European Union’s General Data Protection Regulation (GDPR). More broadly, privacy-by-design principles encourage institutions to embed privacy protections—including controlled access, data minimization, and safeguards for sensitive information—into AI systems from the earliest stages of development.

Because AI systems often rely on large volumes of sensitive data and highly valuable models, cybersecurity and **data protection** are increasingly central to AI governance. Many institutions are extending established cybersecurity frameworks, such as ISO 27001 and the National Institute of Standards and Technology Cybersecurity Framework, to AI environments by restricting access to training data and models, maintaining audit trails and version controls, and implementing safeguards against threats such as data poisoning, unauthorized model extraction, and compromised third-party AI components. As organizations increasingly rely on external data sources, APIs, and third-party models, vendor diligence and ongoing security monitoring become important elements of AI risk management.

## AI applications are not without risk, but the nature of that risk is often misunderstood.

In risk management and compliance contexts, the most material vulnerabilities do not arise from autonomous execution or direct financial exposure, but from how models are designed, supervised, and integrated into decision-making processes. Overreliance on AI-generated outputs, misuse and/or insufficient challenge by reviewers, poor data quality, and weak model governance can all degrade effectiveness. In these settings, failure modes often tend to be subtle rather than obvious and explicit: incomplete analysis,

unchallenged assumptions, inconsistent reasoning, or erosion of professional skepticism.

Crucially, these risks reinforce the case for human-centered deployment. Generative AI does not remove responsibility from risk and compliance professionals; it concentrates it. Institutions must therefore ensure that model outputs are treated as analytical inputs, not conclusions, and that escalation thresholds, review standards, and accountability lines are clearly defined. When properly governed, with trained professionals firmly in the loop and clear ownership for

outcomes, generative AI can strengthen the integrity of second-line functions by improving insight and consistency without displacing judgment. In contrast, when

poorly governed, it risks becoming a force multiplier, exacerbating the consequences of existing weaknesses in controls, training, or oversight.<sup>25</sup>

## Agentic Monitoring

Agentic monitoring represents the next stage of AI-enabled compliance. Rather than using AI only to summarize information or draft discrete outputs, agentic monitoring embeds AI systems into the workflow itself, allowing them to identify relevant obligations, test activity against those obligations, route issues for review, and maintain evidence of the steps taken. Companies such as IBM<sup>26</sup> and Norm.ai have already begun building in this direction. Their agentic monitoring tools are designed to provide real-time assurance, visibility, and control across AI systems; connect AI assets to risks, regulatory requirements, company policies, and controls; and automate compliance mapping, evidence collection, and audit-ready reporting.

This model is important because effective agentic monitoring is not simply a technical exercise. The central challenge is not only whether an AI system can monitor activity, but also whether the institution has correctly determined what the system should monitor, which rules apply, how those rules should be interpreted, what evidence is required, and when escalation is warranted. That requires domain-specific knowledge of financial services, deep familiarity with the applicable regulatory framework, strong understanding of the firm's control environment, and appreciation for the characteristics and limitations of the AI system itself.

For that reason, agentic monitoring increases the need for AI governance and risk management expertise. Effective oversight requires AI professionals who understand not only law, regulation, and compliance, but also how AI models function technically, including their

architectures, limitations, failure modes, and monitoring requirements. The role of second-line professionals therefore shifts away from manual review and toward calibration, supervision, escalation, and continuous challenge of AI-enabled systems.

Agentic monitoring also has important implications for how financial institutions organize and staff control functions. As AI systems increasingly perform routine monitoring, documentation, workflow coordination, and information synthesis, the role of second-line professionals will shift toward judgment-intensive activities such as governance, oversight, and strategic decision making. This transition will require institutions to redesign compliance workflows around human supervision of AI-enabled systems rather than manual process execution alone. It will also require the development of future-ready teams with expertise spanning financial regulation, risk management, AI governance, and the technical operation and supervision of AI models.

If anything, the emergence of client-facing AI systems—including AI-driven financial assistants and advisory tools—strengthens the case for agentic monitoring. As these systems begin producing vast volumes of individualized outputs, recommendations, disclosures, and interactions at machine speed, traditional human review models become increasingly impractical. Effective oversight will therefore require monitoring agents capable of reviewing outputs continuously, identifying potential regulatory or policy concerns, escalating anomalies, and maintaining defensible records of supervisory review.

25 Rahul Agarwal, Andreas Kremer, Ida Kristensen, and Angela Luget, "How Generative AI Can Help Banks Manage Risk and Compliance," McKinsey & Company, March 1, 2024, <https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/how-generative-ai-can-help-banks-manage-risk-and-compliance>.

26 IBM's agentic monitoring tool is watsonx.governance, at [https://www.ibm.com/products/watsonx-governance?utm\\_content=SRCWW&p1=Search&p4=372724786258&p5=p&p9=161518572624&gclid=Cj0KCQjw2YDQBhD\\_ARIsAE1qeSdUNBS-nY7kdjh8Y22f\\_16OVj0uUwCcQMUNvxQ1Eng33fOTYTq8LU-caAv90EALw\\_wcB](https://www.ibm.com/products/watsonx-governance?utm_content=SRCWW&p1=Search&p4=372724786258&p5=p&p9=161518572624&gclid=Cj0KCQjw2YDQBhD_ARIsAE1qeSdUNBS-nY7kdjh8Y22f_16OVj0uUwCcQMUNvxQ1Eng33fOTYTq8LU-caAv90EALw_wcB).

## GOVERNANCE

### **The use of generative AI in risk management and compliance is already governed by a dense and well-established body of financial regulation.**

As noted in the appendix, long-standing US federal supervisory frameworks covering model risk management, technology risk, third-party risk, cybersecurity, AML, sanctions, fraud, supervision, and books-and-records apply fully to AI-enabled tools used within control functions. Guidance from federal banking regulators on model risk management, interagency third-party risk expectations, and technology and operational resilience

standards all apply based on function, not technology. When generative AI is used to support alert review, investigation narratives, regulatory analysis, reporting, or supervisory oversight, it is subject to the same governance, validation, documentation, and accountability requirements that apply to any other analytical system used in these roles. The regulatory premise is therefore activity based and technology neutral: Firms remain responsible for the outputs of their risk and compliance functions regardless of whether those outputs are produced by humans, traditional models, or generative AI.

### **Model Risk Oversight**

Model risk oversight refers to the governance and control processes used to ensure that models operate as intended and produce reliable, appropriately governed outcomes. For more than a decade, US banking organizations largely approached these issues through the framework established by the Federal Reserve's Supervisory Letter 11-7 (SR 11-7), issued jointly with the OCC in 2011 following weaknesses exposed during the global financial crisis. SR 11-7 was intended to strengthen oversight of the growing use of quantitative models in banking by establishing supervisory expectations around validation, ongoing monitoring, and governance. Although originally focused on traditional financial and risk models, the framework increasingly has become an important reference point for broader AI governance as institutions deployed machine learning and generative AI systems across business and control functions.

Against this backdrop, on April 17, 2026, US banking regulators substantially revised the federal model risk-management framework when the OCC, Federal Reserve, and FDIC jointly issued new supervisory guidance superseding SR 11-7, OCC Bulletin 2011-12, the 2021 BSA/AML interagency statement on model risk, and related issuances that had governed bank model risk practices for many years.<sup>27</sup> The new framework adopts a principles-based, materiality-driven

approach. It narrows the definition of "model" to a complex quantitative method grounded in statistical, economic, or financial theories; introduces a \$30 billion asset threshold for primary applicability; eliminates requirements such as the annual model validation cadence, the formal "model owner" role, and explicit board-level approval expectations; and states that "noncompliance with this guidance will not result in supervisory criticism." Notably, the revised guidance explicitly excludes generative AI and agentic AI from its scope on the grounds that these technologies are "novel and rapidly evolving." The agencies have signaled, however, that they will issue a separate request for information addressing model risk management generally and, in particular, banks' use of AI, including generative AI, agentic AI, and AI-based models.

For institutions deploying generative and agentic AI in second-line functions, this development is meaningful but should not be misread. Arguably, the explicit carve-out does not place these tools outside the supervisory perimeter: Although excluded from this specific guidance, generative and agentic AI remain subject to the broader risk management and governance expectations that apply to banking organizations in general.<sup>28</sup> Patomak Global Partners similarly emphasized that, while the removal of SR 11-7's prescriptive specifics provides institutions with the "opportunity... to evaluate their documented

<sup>27</sup> "Model Risk Management: Revised Guidance," Office of the Comptroller of the Currency, OCC Bulletin 2026-13, <https://www.occ.gov/news-issuances/bulletins/2026/bulletin-2026-13.html>.

<sup>28</sup> "Federal Banking Agencies Issue Revised Guidance on Model Risk Management," Sullivan & Cromwell LLP, April 29, 2026, <https://www.sullcrom.com/insights/memo/2026/April/OCC-Fed-FDIC-Issue-Revised-Guidance-Model-Risk-Management>.

requirements and rightsize them based on model risk,” the removal of the annual review cadence, model owner language, and explicit board-level expectations does not relieve firms of the underlying obligation to maintain sound policies and procedures.<sup>29</sup> When issued, the request for information is likely to elicit industry input on how existing model risk management principles should be adapted to generative and agentic AI, and to inform the next phase of supervisory expectations for these tools.

Looking ahead, however, the supervisory challenges posed by generative and agentic AI are likely to extend well beyond the scope of current model risk management guidance. As agentic AI becomes increasingly embedded in first-line activities—including trading, operations, payments, accounting, and other core business functions—financial institutions, regulators, and policymakers will face significant challenges in maintaining effective oversight. Second-line risk management functions will need to evolve in parallel, developing agentic capabilities capable of monitoring, challenging, and governing activities occurring at a speed, scale, and level of complexity that may exceed traditional oversight models. Nevertheless, institutions should treat generative and agentic AI deployments as subject to established safety-and-soundness principles, third-party risk management standards, and operational resilience requirements.

At the same time, AI-specific and cross-sector governance frameworks increasingly clarify how these existing obligations should be applied to generative and agentic systems. Horizontal AI risk management guidance—such as model life cycle governance, transparency expectations, data controls, and human-oversight requirements—functions as an overlay rather than as a replacement for financial regulation. In the European Union and the United Kingdom, supervisory authorities have emphasized that AI systems used in internal control functions must comply with existing organizational, governance, and operational resilience requirements, while also meeting emerging expectations for AI risk management and accountability. Global standard setters have reinforced this approach, focusing on governance, explainability, and oversight rather than prescribing use-case-specific prohibitions. Taken together, this layered regulatory architecture supports the conclusion that the primary challenge is not regulatory gaps, but execution: ensuring that institutions apply familiar supervisory principles—governance, validation, monitoring, escalation, and accountability—with sufficient rigor to a new class of tools. An inventory of applicable financial and AI-specific rules is set out in the appendix.

**Taken together, risk management and compliance illustrate how generative AI can be deployed responsibly, incrementally, and at scale within financial institutions.**

These use cases demonstrate that meaningful value does not require direct customer interaction or autonomous execution. Instead, they show how generative AI can reinforce the very control functions designed to safeguard market integrity and financial stability, offering a pragmatic and credible path for broader adoption across the financial system.

---

<sup>29</sup> “Banking Regulators Replace Comprehensive Model Risk Management Guidance with Risk-Based Approach,” Patomak Global Partners, April 24, 2026, <https://patomak.com/2026/04/24/banking-regulators-replace-comprehensive-model-risk-management-guidance-with-risk-based-approach/>.

The author would like to thank Allyson Ye for her assistance in completing this paper.

# APPENDIX

## I. United States—Financial-Services Regulations

*(These are technology-neutral rules that apply regardless of AI but squarely cover Generative AI use in risk and compliance.)*

### A. Model risk management

- **OCC Bulletin 2026-13; Federal Reserve SR 26-2, Model Risk Management: Revised Guidance.** Adopts a principles-based, materiality-driven approach that narrows the definition of a model to a complex quantitative method grounded in statistical, economic, or financial theory; introduces a \$30 billion asset threshold for primary applicability; eliminates requirements such as an annual model-validation cadence, a formal model-owner role, and explicit board-level approval expectations; and states that “non-compliance with this guidance will not result in supervisory criticism.” The guidance expressly excludes generative and agentic artificial intelligence at this time, with an AI-focused Request for Information to follow.
- **Federal Reserve SR 11-7; OCC Bulletin 2011-12; FDIC FIL-22-2017 (rescinded April 17, 2026, by OCC Bulletin 2026-13 and Federal Reserve SR 26-2), Model Risk Management.** Applies when generative artificial intelligence is used as a model supporting risk-management or compliance decisions, such as alert enrichment, suspicious activity report (SAR) drafting, or regulatory analysis. Covers model development, validation, governance, and ongoing monitoring.

### B. Technology risk management

- **Federal Financial Institutions Examination Council (FFIEC) IT Examination Handbook.** Establishes examiner expectations for IT governance, development, operations, and information security; applies to systems supporting risk, compliance, AML, and reporting, including AI embedded within them.

- **OCC Heightened Standards (12 CFR Part 30, Appendix D).** Establishes a risk governance framework for large national banks; requires a defined risk appetite and independent risk and audit oversight that extend to AI used in control functions.
- **Interagency Paper on Sound Practices to Strengthen Operational Resilience (SR 20-24; OCC 2020-94; FDIC FIL-103-2020).** Establishes resilience expectations for large firms, including mapping critical operations, identifying third-party dependencies, and ensuring recovery capability where AI supports those operations.

### C. Third-party and vendor risk management

- **Interagency Guidance on Third-Party Relationships: Risk Management (2023) (Federal Reserve SR 23-4; OCC Bulletin 2023-17; FDIC FIL-29-2023).** Establishes a full life-cycle framework for vendor risk management, including planning, due diligence, contracting, ongoing monitoring, and termination. Applies to generative AI sourced from vendors or relying on external models, data, or APIs.

### D. Cybersecurity and data protection

- **Federal Reserve, FDIC, and OCC—Computer-Security Incident Notification Rule (2022).** Requires banking organizations to notify their primary regulator within 36 hours of a significant “notification incident,” including incidents involving AI systems that affect operations or data.
- **Interagency Guidelines Establishing Information Security Standards (12 CFR Part 30, Appendix B).** Implements GLBA §501(b) safeguards for customer information; applies to AI systems that access or process such data.
- **National Institute of Standards and Technology (NIST)—Cybersecurity Framework (CSF 2.0).** A voluntary but widely adopted framework used by regulators and institutions to identify, protect, detect, respond to, and recover from cyber risk, including AI-enabled cybersecurity and compliance systems.

### *E. AML, sanctions, and fraud compliance*

- **Financial Crimes Enforcement Network (FinCEN)—Bank Secrecy Act / AML Program Rules.** Requires a risk-based AML program, including internal controls, independent testing, a designated compliance officer, training, and customer due diligence. Generative AI used in monitoring, investigations, SAR drafting, or quality assurance must operate within these requirements.
- **Office of Foreign Assets Control (OFAC)—A Framework for Sanctions Compliance Commitments.** Establishes expectations for sanctions compliance programs, including management commitment, risk assessment, internal controls, testing, and training. Applies where generative AI supports screening, escalation, or documentation.
- **U.S. Department of Justice (DOJ)—Evaluation of Corporate Compliance Programs (ECCP).** Provides criteria used to assess whether compliance programs are well designed and effective, increasingly including governance of AI-enabled compliance systems.

### *F. Broker-dealer and market regulation (where applicable)*

- **Financial Industry Regulatory Authority (FINRA) Rules 3110 and 2210.** Require supervisory systems and oversight of firm communications; apply where generative AI is used for communications review, surveillance, or supervisory oversight.
- **Securities and Exchange Commission (SEC) Exchange Act Rules 17a-3 and 17a-4.** Establish books-and-records requirements for broker-dealers; apply to AI-generated compliance artifacts, reports, and supervisory documentation.

## **II. United States—AI-Specific and Horizontal AI Governance**

- **National Institute of Standards and Technology (NIST)—AI Risk Management Framework (AI RMF 1.0).** A voluntary framework for governing, mapping, measuring, and managing AI risk across the lifecycle; widely treated as a de facto baseline in regulated sectors.
- **Executive Order 14110 on Safe, Secure, and Trustworthy AI (2023).** Revoked on January 20, 2025, and replaced by Executive Order 14179.
- **Executive Order 14179, Removing Barriers to American Leadership in AI (2025).** Establishes the current federal policy orientation toward AI, emphasizing innovation and reduced regulatory friction while retaining voluntary reliance on frameworks such as NIST AI RMF.
- **Federal Trade Commission (FTC)—Section 5 of the FTC Act (Unfair or Deceptive Practices).** Applies to AI systems engaged in unfair or deceptive conduct, including misleading claims about AI capabilities or consumer-facing decisions based on flawed or unexplainable outputs.

## **III. European Union**

### *A. Financial-services regulation*

- **European Securities and Markets Authority (ESMA)—MiFID II Organizational and Conduct Requirements.** Requires investment firms to maintain effective compliance, risk management, and internal control functions; applies to AI systems used within those functions.
- **Digital Operational Resilience Act (DORA) (in force 2025).** Establishes an EU-wide framework for ICT risk management, incident reporting, resilience testing, and oversight of third-party ICT providers, including cloud and AI vendors.

## B. AI-specific regulation

- **EU Artificial Intelligence Act (EU AI Act).** Establishes a risk-based horizontal AI regulatory framework. Certain compliance and risk management systems may qualify as “high-risk,” triggering obligations relating to governance, transparency, data quality, human oversight, and documentation.
- **General Data Protection Regulation (GDPR), Articles 5, 22, and 25.** Governs personal data processed by generative AI, including principles of lawful processing (Article 5), limits on solely automated decision-making (Article 22), and data protection by design and by default (Article 25).

## IV. United Kingdom

### A. Financial regulation

- **Financial Conduct Authority (FCA)—SYSC (Senior Management Arrangements, Systems and Controls).** Requires firms to maintain adequate systems and controls, including effective risk and compliance functions; applies to AI used in these contexts.
- **Financial Conduct Authority (FCA)—Operational Resilience Framework (PS21/3).** Requires firms to identify important business services, set impact tolerances, and remain within them, including where AI supports those services.

### B. AI and technology governance

- **UK Government—A Pro-Innovation Approach to AI Regulation (2024).** Establishes five cross-sector principles—safety, security, transparency, fairness, accountability, and contestability—implemented through existing regulators rather than a single AI statute.
- **UK GDPR.** Mirrors EU GDPR obligations and applies to AI-driven personal data processing.

## V. Global and Standard-Setting Bodies

- **International Organization of Securities Commissions (IOSCO)—The Use of AI and Machine Learning by Market Intermediaries and Asset Managers (FR06/2021).** Establishes six governance measures covering oversight, development and testing, staff competence, and transparency in AI and machine learning use.
- **Basel Committee on Banking Supervision (BCBS)—Principles for the Sound Management of Operational Risk (2021).** Sets operational risk governance and resilience principles applicable to AI-supported risk and compliance functions.
- **Financial Stability Board (FSB)—Artificial Intelligence and Financial Stability Reports (2017; 2024).** Provides analytical assessments of AI’s systemic implications for financial stability and informs supervisory approaches without imposing binding requirements.

## VI. Professional and Control-Function Standards

- **Institute of Internal Auditors (IIA)—Global Internal Audit Standards (with GTAG technology guidance).** Establishes professional standards for internal audit; applies where AI supports assurance, oversight, and control testing functions.







---

## Artificial Intelligence Working Group

CO-CHAIRS: William C. Dudley and Sarah Hammer

WORKING GROUP: Ghassan K. Abou-Alfa, Oluyomi Adebote, Fanny Solano Agramonte, David Alcala, Paul Andrews, Sandeep Arora, Isadora Arredondo, Jenner Balagot, Beverly Barrett, Daniela Bassan, Richard Berner, Spyridon Bisisidis, Barbara Boehler, David Borcsok, Yisroel Bruner, Camille Bryan, Onorina Buneanu, Josh Burek, F. Christopher Calabia, Christian Callender-Easby, Massimo Carassinu, Jake Carney, Keith B. Carter, Christian Castro, Uldis Cerps, David W. Cervantes, Alexander James Chesney, Aleksandar Colic, Wesley Brown de Colstoun, Himamauli Das, Gurunath Dasari, Antonio U. De Lorenzo, Laura De Rivera, Jag DeSaint Dior, Qahir Dhanani, Alexandra Dimitrijevic, Mona Elliot, Anthony Elson, Jonathan R. Everhart, Omer Fein, Maria Jose Fernandez, Pilar Cristina Frank, Carl-Benedikt Frey, Gerald Gallagher, Nitin Gaur, Samim Ghamami, Daniel L. Gleizer, Jeffrey Greenbaum, Michael B. Greenwald, Jon Hartley, Naureen Hassan, Sarah Hirsch, Ulrike Hoffmann-Burchardi, Olive Hsiao, Jason Hsu, Michael Hsu, Linda Jeng, Gregory C. Johnson, Peter E. Johnson, Oliver Kellman, Ron Korland, Mahesh Kotecha, Abhesh Kumar, Pat LaVecchia, June Lau, Maximilien Pucci-Sisti Maisonrouge, Fergus McCormick, Tetsuharu Minowa, Ricardo Mora, Margarita Musante, Neha Narula, Fabio Natalucci, Patrick Noonan, Andrzej Nowak, Gregory J. Peters, Caroline Pham, Sebastian Robertson, David Rogovic, Gursimran Rooprai, Joan Rosas Xicota, Nilmini Rubin, Sunil Sabharwal, Paul R. Samson, Jason Schenker, Lemma W. Senbet, Tarik Senhaji, Irina Shaorshadze, Deepika Sharma, Alexandra de Hoop Scheffer, Adnan Pasha Siddiqui, Michael F. Silva, Divya Singaravelu, Maggie Sklar, Andrew Slack, Danica Starks, Radek Stech, Keerthika M. Subramanian, Carrie Suen, Rajesh Swaminathan, Zhenqian Tay, Simon Taylor, Raffaella A. Tenconi, Hilary F. Till, Siddharth Tiwari, Helen Turner, Tobin G. Van Ostern, Marsha Vande Berg, Markus Veith, Manuela Veloso, Andreas Veneris, Akshay Vijayendiran, Flavia Villanueva, Shawei Wang, Adam Wang-Levine, David Webster, Antonio Weiss, Benjamin S. Weiss, Rex Wempen, Carolyn A. Wilkins, David P. Willard, Wai Ting Sandy Yeung, Linda Yueh, Bryan Zhang



THE BRETTON WOODS COMMITTEE  
1701 K St NW #950, Washington, DC 20006  
[www.brettonwoods.org](http://www.brettonwoods.org)