EVENTUS

The Evolution of AI in Financial Compliance and Trade Surveillance

June 2025

WHITEPAPER

Table of Contents

FOREWORD

Technological Evolution and the Constants of Compliance

Throughout my career in multiple roles, I've witnessed several technological revolutions transform the financial markets. From the emergence of word processing and spreadsheets in the 1980s, to the electronification of markets in the early 2000s, to the rise of social media and cryptocurrencies in the last decade, and now the advent of artificial intelligence — each wave has fundamentally altered how we conduct business, manage risk and ensure compliance.

What strikes me about these technological leaps is not just the differences between them, but the consistent patterns they share. These common threads inform the approach to this paper on AI in compliance and trade surveillance.

First, regulatory frameworks invariably lag behind technological innovation. When electronic trading platforms first emerged, regulators were forced to adapt rules designed for open-outcry trading floors. Similarly, as we navigate the AI revolution, we're applying existing regulatory principles to technologies that weren't contemplated when those rules were written. We must and will adapt over time.

Second, early adoption brings both competitive advantage and heightened risk. Those who embrace new technologies first often reap significant benefits in efficiency and capabilities, but they also serve as test cases for regulatory scrutiny. We saw this with the early adopters of algorithmic trading strategies, who gained market advantages but also faced the first wave of regulatory enforcement actions related to market abuse via algorithms.

Third, the human element remains irreplaceable despite increasing automation. Each technological wave has prompted predictions that human judgment would be rendered obsolete and that machines would “take our jobs,” yet the opposite has proven true. Technology may change the nature of human involvement, but it has consistently elevated the importance of experienced oversight, critical thinking and ethical decision-making.

Fourth, successful implementation has always depended more on organizational culture than on technological sophistication. Organizations that view compliance as fundamental to their strategic objectives, rather than as a cost center, are better positioned to harness new technologies effectively.

Finally, transparency and explainability have consistently emerged as regulatory priorities across technological transitions. Whether explaining an algorithmic trading strategy or an AI-powered risk model, the ability to clearly articulate how a system works and why it makes certain decisions has been essential for regulatory acceptance.

As we stand at the forefront of the AI revolution in financial compliance, these lessons from past technological transitions offer valuable guidance. The tools may change, but the fundamental principles of effective compliance — rigorous risk assessment, clear governance structures, thorough documentation, meaningful testing and continuous improvement — remain constant.

WHAT WE'LL EXPLORE

In this paper, we explore how AI is reshaping compliance and trade surveillance functions, examining current applications, emerging challenges and best practices for implementation. My hope is that by understanding both the transformative potential of AI and the enduring principles of effective compliance, financial institutions can navigate this latest technological revolution successfully.

Joseph Schifano, Global Head of Regulatory Affairs, Eventus

INTRODUCTION

Understanding Artificial Intelligence

The financial services industry stands at a pivotal moment in its adoption of artificial intelligence (AI). What began as experimental initiatives among forward-thinking firms has become mainstream implementation, with AI evolving from a competitive advantage into a strategic imperative. According to CB Insights, and as evidence of a broader trend, AI companies captured an unprecedented 20% of all venture deals globally in Q1 2025 — double the share since OpenAI's launch of ChatGPT in 2022 — with more than half of the $121 billion in quarterly funding flowing to AI-focused enterprises.

For financial compliance and trade surveillance teams, this technological shift arrives amid mounting challenges: increasingly fragmented markets, the proliferation of complex trading strategies and ever-expanding regulatory oversight. As noted in a recent survey by Acuiti, 94% of market surveillance specialists report their roles becoming more complex as markets digitize and data volumes explode. These professionals often spend more than 30 hours per week investigating alerts — many of which prove to be false positives — placing an unsustainable burden on compliance resources.

94%

of market surveillance specialists report their roles becoming more complex as markets digitize and data volumes explode.

Against this backdrop, AI offers compelling solutions to enhance surveillance capabilities while controlling costs. However, the path to effective AI implementation is neither straightforward nor without risk. This paper explores how financial institutions are navigating the evolution of AI in compliance and trade surveillance, examining current applications, regulatory considerations, implementation challenges and best practices for successful adoption.

EXAMINING THE FUTURE

As we examine this landscape, we'll also look at how forward-thinking companies are implementing broader AI strategies across their organizations, regulatory attitudes toward AI use cases and governance and how these general industry trends can inform the specific application of AI in trade surveillance.

Related Reading

GLOSSARY OF KEY AI TERMS

Agentic AI – A system capable of taking independent actions to achieve specified goals, making decisions and executing tasks with minimal human supervision or intervention. These systems can perceive their environment, plan sequences of actions, adapt to changing circumstances and operate with a greater degree of autonomy than traditional AI applications.

AI Assistant (or AI Agent) – A software system designed to interact with users, understand their requests and perform specific tasks or provide information based on its programming and available data. It typically focuses on completing specific workflows with varying degrees of autonomy.

Artificial Intelligence (AI) – A broad field of computer science focused on building systems that can perform tasks normally requiring human intelligence — such as pattern recognition, natural language processing and decision-making.

Deterministic AI – AI systems that produce predictable, rule-based outputs. These systems are favored in compliance settings due to their transparency and auditability.

Explainability – The ability to clearly articulate how an AI system arrives at its outputs or decisions — a key requirement in regulated industries like finance.

Generative AI (GenAI) – AI models (such as large language models) capable of producing human-like text, code or other outputs based on prompts — used with caution in compliance due to concerns over hallucination and reliability.

Human-in-the-Loop – A model of AI implementation that keeps human oversight central to decision-making, especially for risk assessment and regulatory interpretation.

Large Language Model (LLM) – An advanced AI model trained on massive datasets of human language, enabling capabilities such as summarization, sentiment analysis and text generation.

Machine Learning (ML) – A subset of AI that enables systems to learn patterns from historical data and improve performance over time without being explicitly programmed for every scenario.

Model Risk Management – The process of validating, testing and overseeing AI models to ensure they perform as intended and do not introduce systemic or regulatory risk.

Natural Language Processing (NLP) – A branch of AI that enables machines to understand, interpret and generate human language. Used in compliance for analyzing unstructured communications.

Probabilistic AI – AI systems that make decisions based on statistical likelihood rather than predefined rules. While flexible, these models can be harder to explain and audit.

Rule-Based System – A traditional surveillance model that triggers alerts based on predefined thresholds or conditions. Often combined with AI for more effective detection.

Unsupervised Learning – A type of machine learning where the system identifies patterns or anomalies in data without labeled training examples — useful for detecting novel forms of misconduct.

The Evolving AI Landscape in Financial Compliance

Evolution vs. Revolution

In a recent speech, Federal Reserve Vice Chair for Supervision Michael Barr outlined two potential scenarios for AI's impact on financial services:

1. ***Under the "Incremental Progress" scenario, AI primarily enhances existing processes, delivering steady efficiency improvements without fundamentally altering industry structures. This evolutionary approach sees AI augmenting human capabilities in compliance and risk management while potentially magnifying existing vulnerabilities like algorithmic bias or market herding.***

2. ***The alternative "Transformative Change" scenario envisions AI triggering a fundamental reshaping of financial markets, defined by hyper-personalized services, real-time intermediation and radical shifts in market structure and regulatory frameworks.***

Today, most financial institutions appear to be following the incremental path, carefully integrating AI into existing frameworks rather than replacing them wholesale. A notable example is Moody's, whose Research Assistant AI tool became the most rapidly adopted solution in the company's history. According to Cristina Pieretti, General Manager of Digital Insights at Moody's Analytics, this success stems from a measured approach to AI that emphasizes practical applications over technological hype: "It's always better to be the disruptor than to be the disrupted one."

Moody's methodology for AI implementation started with small, practical applications that demonstrated clear value before scaling. This approach, more characteristic of a technology startup than a century-old firm, enabled Research Assistant to move from concept to market in record time while maintaining the regulatory rigor required in financial services.

AI Technologies in Compliance and Trade Surveillance

From Rule-Based Systems to Machine Learning

Trade surveillance technology has historically relied on rule-based systems: platforms that flag activities exceeding predetermined thresholds or parameters. While powerful in their simplicity and directness, these systems face significant limitations in today's complex trading environment. They struggle to adapt to new market conditions, generate excessive false positives during volatile periods and often miss sophisticated manipulation tactics that span multiple products or markets.

As one industry report noted, during the unprecedented volume increase experienced during the COVID-19 pandemic, surveillance alert messages doubled in some cases, overwhelming compliance teams with false positives. This operational challenge has driven many firms to seek more adaptable solutions that can enhance, rather than replace, their rule-based systems.

This operational challenge has driven many firms to seek more adaptable solutions that can enhance, rather than replace, their rule-based systems.

Machine Learning Approaches in Surveillance

The first wave of AI adoption in surveillance involved supervised machine learning (ML) models that could recognize patterns in historical data and improve alert accuracy. These systems helped reduce false positives and prioritize alerts, but they required extensive labeled datasets and still operated within relatively rigid frameworks. Most importantly, these early ML implementations complemented, rather than replaced, rule-based systems, providing an additional layer of analysis while maintaining the deterministic nature of the underlying surveillance mechanisms.

Supervised learning models remain prevalent today for detecting known patterns of market abuse. These models are trained on labeled examples of activities like spoofing or front-running, enabling them to identify similar patterns in new data. The advantage of this approach is its high accuracy for known violation types, but it requires extensive training data and cannot detect novel forms of misconduct.

Another key benefit of supervised learning models is their ability to provide risk-scoring capabilities that enhance rule-based systems without unnecessarily narrowing their detection parameters. Rather than generating binary outcomes (alert/no alert), these models can assign confidence scores to potential violations, helping compliance teams better prioritize their investigations. This risk-based approach enables analysts to see "near misses" that might not trigger a traditional alert but could represent emerging patterns of misconduct.

Over time, these risk scores create a valuable historical record that compliance teams can analyze to identify behavioral trends. For instance, a trader consistently scoring in the moderate risk range across multiple scenarios might warrant closer examination, even if they never trigger a high-priority alert. This longitudinal view helps identify subtle manipulation tactics that might otherwise go undetected.

As human analysts provide feedback on these risk scores, they effectively train the models to improve over time, creating a kind of reinforcement learning. This approach means the system learns continuously from expert judgment, adapting its risk assessments based on confirmed true and false positives. Importantly, this feedback loop can be tailored to specific contexts, enabling models to develop specialized expertise for different companies, asset classes, trading strategies or business types. For example, a surveillance model monitoring options trading might develop different risk patterns than one focused on fixed income – each can be refined through analyst feedback specific to those markets. This adaptive capability ensures that surveillance becomes increasingly effective and efficient as it accumulates more contextual training data from human experts.

✓ *This adaptive capability ensures that surveillance becomes increasingly effective and efficient as it accumulates more contextual training data from human experts.*

Current State: NLP, LLMs and Generative AI

Today's general financial compliance landscape goes far beyond ML, incorporating more sophisticated AI capabilities across a wider range of use cases. Natural language processing (NLP), large language models (LLMs) and various forms of ML can analyze vast datasets across not only different markets and asset classes, but across both structured and unstructured data, enabling ever more efficient and effective monitoring.

E-communications surveillance represents one of the most promising areas for AI applications in compliance. NLP and ML tools can analyze email, chat and voice communications to identify potential misconduct — a task that would be impossible to perform manually at scale. LLMs have been trained to analyze such communications in many languages, which is a major advancement in capability. Several e-communication vendors leverage AI to monitor for indicators of manipulative behavior, insider trading or other compliance violations across various channels.

Anti-money laundering (AML) represents another significant application of AI in compliance. ML models can analyze transactions to identify potentially harmful patterns that might be missed by traditional rule-based systems. These systems can detect complex relationships and anomalies across large datasets, significantly reducing false positives while improving detection rates for genuinely suspicious activities.

In addition, regulatory change management has emerged as a key use case for AI in compliance. Financial institutions face an ever-growing volume of regulatory shifts across multiple jurisdictions, creating significant challenges for compliance teams. AI solutions can automatically monitor, analyze and categorize regulatory updates from various sources, helping firms identify relevant changes, assess their impact and make necessary adjustments to compliance programs. According to industry experts, automating regulatory change management has become something of a "holy grail" in the use of AI for compliance.

As these diverse AI applications continue to advance, one common theme is the need for explainability. Financial institutions must ensure they can justify the findings of their AI-based surveillance systems to regulators and internal stakeholders in clear, understandable terms. This emphasis on explainability has been a key reason for the cautious integration of generative AI (GenAI) solutions in core compliance functions.

✓ *This emphasis on explainability has been a key reason for the cautious integration of generative AI (GenAI) solutions in core compliance functions.*

Promising Applications of AI in Trade Surveillance

While many advanced AI applications in trade surveillance remain in development or testing stages, financial institutions are already leveraging the technology for several key functions:

Report Generation and Analysis

Many compliance teams use LLMs to draft standardized reports, analyze investigation results and compile metrics for supervisory review. This reduces the manual effort required for documentation while ensuring consistency.

Pattern Recognition Enhancement

ML algorithms complement traditional rule-based systems by identifying subtle patterns that might otherwise go undetected, particularly in analyzing communications or trade data for potential misconduct.

Process Automation

AI streamlines routine compliance workflows, such as data collection, initial alert triage and preliminary investigation, enabling analysts to focus on more complex cases.

The Critical Distinction: Probabilistic vs. Deterministic AI

Perhaps the most important difference in applications of AI for compliance is between probabilistic and deterministic models.

Probabilistic (or statistical) AI models score behaviors and make predictions based on statistical likelihoods. While powerful for certain applications, these models can produce variable outputs and may operate as "black boxes" that resist easy explanation — a serious vulnerability for regulatory compliance.

Deterministic AI, by contrast, provides predictable, structured and repeatable outputs. While still capable of learning from patterns and adapting to new risks, deterministic models ensure that every flagged trade, alert and risk score is supported by logic that compliance teams can review and regulators can audit.

For market participants, the need to balance innovation and explainability often favors deterministic AI, particularly for core surveillance functions. This approach enables firms to demonstrate to regulators that their surveillance controls remain transparent and accountable even as they leverage advanced pattern recognition and data analysis capabilities.

✓ *This approach enables firms to demonstrate to regulators that their surveillance controls remain transparent and accountable even as they leverage advanced pattern recognition and data analysis capabilities.*

Related Reading

FIRSTHAND PERSPECTIVE: VALIDUS AND FRANK AI

Against this backdrop of evolving AI applications, Eventus has developed Frank AI – a new interactive tool to address specific challenges faced by compliance teams. Frank represents a significant advancement in how analysts interact with surveillance data, enabling natural language queries into the Validus platform without requiring SQL or coding expertise.

Unlike general-purpose GenAI, Frank operates in a structured data environment, enabling analysts to directly query Validus-specific data tables. This approach unlocks several key advantages:

- ✓ **Deterministic Results**
  Rather than providing probabilistic answers based on pattern recognition, Frank queries real-time structured data and returns precise, fact-based results.
- ✓ **Security and Control**
  Frank operates in a manner that ensures sensitive compliance data never leaves a client's secure environment — a critical consideration for financial institutions.
- ✓ **Guided Insights for Analysts**
  Frank helps analysts find the right data points to resolve alerts, making it a powerful tool for compliance and risk analysis – no "guessing" responses.
- ✓ **Freeform Querying**
  Users can make requests like "Show me all high-risk transactions flagged in the last seven days" without writing SQL queries, enabling complex data retrieval for non-technical users.
- ✓ **Report Building and Decision Support**
  Frank assists in generating reports with structured, actionable insights rather than open-ended text interpretation.

Above all, Frank represents a thoughtful approach to AI adoption that prioritizes accuracy, security and practical utility for financial compliance professionals. By focusing on natural language interaction with structured data rather than generative outputs, Frank delivers the benefits of AI while avoiding the risk of hallucination or misleading responses that can plague general-purpose AI systems.

Future iterations of Frank may incorporate additional capabilities, including more advanced pattern recognition and enhanced visualization tools. However, the emphasis on deterministic outputs and secure, controlled data access will remain central to its design philosophy.

Regulatory Perspectives and Concerns on AI in Financial Markets

Regulatory attitudes toward AI in surveillance represent a complex balance between encouraging innovation and controlling risks. Financial regulators worldwide are increasingly focused on AI applications within their jurisdictions and developing frameworks to address both the opportunities and challenges.

U.S. Regulators' Scrutiny of AI-Driven Trade Surveillance

U.S. agencies have moved from observation to action on how AI is described and used in financial markets.

SEC – EARLY ENFORCEMENT AND PROPOSED RULES

On March 18, 2024, the SEC announced its first AI-focused enforcement, settling with two advisers that marketed, but were not using, AI models – a practice the agency labeled "AI washing." Chair Gary Gensler warned that firms must not deceive investors about AI capabilities: "Investment advisers or broker dealers should not mislead the public by saying they are using an AI model when they're not, nor say that they're using an AI model in a particular way, but not do so."

Well before that case, the SEC had begun to create new guardrails. Its July 2023 Predictive Data Analytics (PDA) proposal would mandate that broker-dealers and advisers eliminate or neutralize conflicts arising from any algorithm that interacts with investors, signaling that mere disclosure is no longer enough. The Division of Examinations' FY 2024 priorities echo that stance, flagging "artificial intelligence and trading algorithms" as an exam focus alongside crypto and cybersecurity. Given new Chair Paul Atkins' deregulatory stance, these proposals may be reconsidered. His leadership may steer the SEC toward a more innovation-friendly regulatory environment, potentially impacting future AI-related policies.

FINRA – JUST ANOTHER SUPERVISORY TOOL

Regulatory Notice 24-09 (June 2024) reminded member firms that FINRA's rulebook is "technology-neutral." If a market participant uses GenAI or other models to review emails or monitor trading, Rule 3110 still demands a "reasonably designed supervisory system," including pre-deployment vetting, ongoing testing and documented model-risk controls — whether the code is homegrown or purchased from a vendor. FINRA's 2025 Oversight Report reinforces that message, listing AI among the year's "emerging risk areas" and previewing more in-depth examinations of governance and vendor management.

CFTC – FROM GENERAL QUESTIONS TO CONCRETE EXPECTATIONS

A January 2024 Request for Comment asked derivatives markets how they use AI in trading, risk and compliance, highlighting fairness, explainability and market-manipulation risks. Eleven months later, a staff advisory told exchanges, FCMs and other registrants that all existing Commodity Exchange Act duties continue to apply "irrespective of their use of AI." The advisory urges human oversight, robust testing and clear policies before and during any AI deployment.

CFTC Commissioner Kristin Johnson has advocated for a comprehensive approach to AI regulation that balances innovation with appropriate safeguards. As she noted in a speech, "The first [priority] is protecting the integrity of the trading markets so that they fairly serve the interests of participants and the larger public. The second is welcoming and encouraging the development and application of the newest technologies with responsible guardrails." Johnson has also called for the creation of an interagency task force composed of market and prudential regulators to coordinate AI oversight across financial markets.

DOJ – PROSECUTORS ADD AI GOVERNANCE TO THE CHECKLIST

The Criminal Division's September 2024 update to its Evaluation of Corporate Compliance Programs instructs prosecutors to ask whether companies performed risk assessments on AI tools, maintained "human oversight" controls and continuously monitored models to ensure they work "as intended." Failure to govern AI systems can weigh against companies in charging and penalty decisions.

The Trump Administration's AI and Crypto Strategy

The landscape of AI regulation in financial markets is likely to further evolve with the Trump administration's appointment of David Sacks, a tech entrepreneur and venture investor known for his work at PayPal and Yammer, as the "White House A.I. & Crypto Czar." Sacks is expected to guide key administration policies, with a focus on making the U.S. "the clear global leader in both areas."

According to venture investor Steve Jang, "Sacks will likely have a light touch on regulation, but not without some guardrails." His approach may emphasize regulating how AI is used in critical applications, as opposed to regulating the development of the models themselves. This could create a more innovation-friendly regulatory environment while maintaining appropriate oversight for high-risk applications like financial surveillance.

The Trump administration's January 23, 2025 executive order on AI directs certain White House advisers to develop an action plan to "sustain and enhance America's global AI dominance." This suggests a potential shift toward more pro-innovation policies that may impact how financial regulators approach AI oversight.

Global Regulatory Perspectives on AI Surveillance

UNITED KINGDOM – BALANCING CONTINUITY AND CHANGE

An April 2024 FCA Artificial Intelligence Update stresses that firms using AI "remain responsible for ensuring compliance with our rules" and that existing regimes — Senior Managers & Certification, Operational Resilience, Consumer Duty — already give the FCA enforcement tools. The watchdog promises "close scrutiny" of governance, data integrity and third-party outsourcing as firms adopt AI.

Accordingly, the FCA has announced plans to invest more in AI for its own regulatory functions, including "market surveillance purposes" and potentially further use cases involving "Natural Language Processing to aid triage decisions, assessing AI to generate synthetic data or using LLMs to analyse and summarise text." This underscores the regulator's commitment to understanding and leveraging AI while ensuring appropriate oversight. The fact that leading regulators themselves are deploying AI in surveillance signals a broader legitimization of these tools across financial services.

To support innovation, the FCA has established several initiatives, including a Regulatory Sandbox that allows firms to test innovative propositions in a controlled environment, as well as an AI Sandbox to help firms develop and test AI models. In October 2024, the regulator hosted a TechSprint where trade surveillance specialists accessed FCA trading datasets to develop and test AI-powered surveillance solutions. Further, the Bank of England has established an Artificial Intelligence Consortium to provide a platform for public-private engagement on the development, deployment and use of AI in UK financial services.

EUROPEAN UNION – RAMPING UP OVERSIGHT

The Artificial Intelligence Act (Reg. 2024/1689), finalized in June 2024, designates each Member State's financial services supervisor as the market surveillance authority for high-risk AI systems used by banks, brokers and trading venues. Those authorities have the right to inspect documentation, datasets and even source code when compliance is in doubt — a preview of more intense scrutiny once the regulation starts applying in 2026-27.

Further, in May 2024, in its first statement on the topic, the European Securities and Markets Authority (ESMA) stressed the importance of abiding by established regulations in adopting AI: “Importantly, firms' decisions remain the responsibility of management bodies, irrespective of whether those decisions are taken by people or AI based tools.”

INTERNATIONAL COORDINATION – FOCUS ON ALIGNMENT

IOSCO’s March 2025 consultation report finds that firms worldwide are expanding AI use in trade surveillance, AML and market abuse detection and urges regulators to align expectations on governance, testing and transparency.

Key Themes and Legal Concerns Emerging

1 **Technology-neutral enforcement**

U.S. regulators emphasize that using an AI label does not change statutory duties. Failures of an AI surveillance system are still misstatements or supervision lapses.

2 **Governance, testing and human oversight**

Agencies will likely expect written model inventories, pre-deployment validation, ongoing back-testing and “kill switches” so humans can override errant code. Examination priorities and DOJ guidance both say inadequate oversight can itself be a compliance weakness.

3 **Third-party and model-risk management**

Outsourcing surveillance to a vendor does not outsource liability; FINRA and the CFTC explicitly require due diligence, transparency and continuous monitoring of external AI models.

4 **Transparency and explainability**

The CFTC’s RFC, the EU AI Act and IOSCO all highlight the need to understand and explain why an AI system flags — or misses — certain trading patterns. Lack of explainability can itself draw regulatory scrutiny.

5 **Misleading claims (“AI washing”)**

The SEC’s March 2024 settlement shows that overstating AI capabilities can violate antifraud provisions, just as greenwashing cases did for ESG. Accurate, evidence-based disclosures are now a compliance must.

6 **Market integrity and bias concerns**

Global regulators warn that poorly trained models may miss novel manipulation or generate biased alerts, undermining fair markets. Guidance stresses the need for data quality controls and regular bias testing.

7 **Enforcement trajectory**

Early actions are widely seen as warning shots. As regulators and enforcement agencies investigate and examine AI-driven trading and surveillance tools, relevant enforcement matters are expected.

The Human Element:
Keeping People in the Process

Despite significant advancement in AI technologies, human judgment remains essential to effective surveillance. AI should augment, not replace, skilled compliance professionals, who bring contextual understanding and ethical judgment that machines cannot replicate.

The Human-in-the-Loop Approach

The concept of "human in the loop" has become central to responsible AI implementation in surveillance. This approach ensures that significant decisions — particularly those with regulatory or legal implications — incorporate human judgment alongside algorithmic analysis. One example is the reinforcement learning approach to AI-generated risk scores outlined earlier in this paper. Regulators increasingly expect this level of human oversight, viewing it as essential to prevent overreliance on potentially flawed or biased models.

In practice, this means designing workflows where AI identifies potential issues, gathers relevant information and presents it to human analysts for review and decision-making. This approach leverages the complementary strengths of both forms of intelligence: machines excel at processing vast amounts of data and identifying patterns, while humans bring contextual understanding, ethical judgment and accountability.

Essential Skills for Compliance Professionals

As surveillance technology evolves, compliance professionals must develop new skills to work effectively with AI systems:

TECHNICAL LITERACY

While compliance staff don't need to become data scientists, they should have a basic understanding of how AI technologies operate, including their capabilities and limitations.

CRITICAL EVALUATION

The ability to assess AI outputs critically, recognizing potential biases or errors, will be crucial. This includes understanding when to trust and when to question algorithmic findings.

INTERDISCIPLINARY COMMUNICATION

Compliance professionals will increasingly need to collaborate with data scientists and technology specialists, requiring effective communication across disciplinary boundaries.

ETHICAL JUDGMENT

As AI raises new questions around privacy, fairness, and accountability, compliance staff must develop strong ethical reasoning skills to navigate these complex issues.

ADAPTABILITY

Perhaps most importantly, compliance professionals must embrace continuous learning to keep pace with rapidly evolving technologies and regulatory expectations.

Risks and Ethical Considerations

While AI offers significant benefits for surveillance, it also introduces new risks that firms must manage carefully:

1

Algorithmic Bias

Models trained on historical data may perpetuate or amplify existing biases, potentially leading to discriminatory outcomes or uneven enforcement.

2

Overreliance on Systems

Excessive confidence in AI can create a false sense of security, particularly if models have unforeseen weaknesses or blind spots.

3

Systemic Risk

Widespread adoption of similar models among market participants could lead to industry-wide vulnerabilities or synchronized responses to market events.

4

Privacy and Data Protection

Advanced surveillance capabilities raise questions around appropriate boundaries and the risk of excessive monitoring.

+

One particularly sensitive area involves predictive analytics for trader profiling.

While technically feasible, using AI to forecast potential misconduct by specific individuals based on their past behavior raises significant ethical and legal concerns. Such applications could create an unpleasant corporate culture, potentially expose firms to litigation and replicate biased findings through opaque models. Concerns about predictive analytics mirror debates arising in other sectors, such as the use of AI for employee monitoring or creditworthiness assessments, where opacity and bias have triggered regulatory scrutiny.

Organizations implementing AI in surveillance must develop ethical frameworks governing its use, with clear boundaries, oversight mechanisms and processes for addressing edge cases. These frameworks should evolve as technology advances and regulatory expectations change, ensuring surveillance practices remain both effective and responsible.

A Practical Roadmap for AI Adoption in Surveillance

Financial institutions considering AI for surveillance should follow a structured approach:

ASSESSMENT AND STRATEGY

- Evaluate current surveillance capabilities and pain points
- Identify specific use cases where AI could deliver meaningful improvements
- Determine build vs. buy decisions based on organizational capabilities and resources

PILOT IMPLEMENTATION

- Start with limited-scope applications where success can be clearly measured
- Establish baseline metrics for comparison with pre-AI performance
- Collect feedback from users and stakeholders

SCALING AND INTEGRATION

- Expand successful applications to broader surveillance activities
- Integrate AI capabilities with existing systems and workflows
- Develop governance frameworks for ongoing oversight

CONTINUOUS IMPROVEMENT

- Regularly assess model performance against changing market conditions
- Refine models based on analyst feedback and emerging threats
- Maintain documentation of model changes and validations

CROSS-FUNCTIONAL GOVERNANCE

- Establish an internal governance group spanning business, technology, compliance, legal, risk and security
- Ensure AI initiatives align with firmwide risk frameworks and regulatory obligations
- Incorporate input from diverse stakeholders to meet both operational and oversight goals

A 2024 study from Microsoft and LinkedIn found that:

79%

of leaders believe AI adoption is critical to competitiveness...

HOWEVER,

60%

of leaders fear that their company lacks a vision and plan to implement it.

✓ **Effective change management is the key to this process. Compliance professionals may initially view AI with skepticism of its abilities or concern about job displacement. Clear communication of how AI will augment rather than replace human expertise can help build organizational buy-in and ensure successful adoption.**

Related Reading

DATA QUALITY AND GOVERNANCE

Recent enforcement actions highlight the regulatory focus on comprehensive surveillance coverage, regardless of the technology used. In 2024, U.S. regulators levied record fines on JPMorgan Securities — collectively over half a billion dollars — for failing to capture billions of order messages over a 10-year period.

This case serves as a stark reminder that surveillance systems are only as good as the data they ingest. As firms implement more sophisticated AI-based surveillance technologies, they must ensure comprehensive data capture and retention to avoid similar enforcement actions.

The foundation of effective AI-powered surveillance is high-quality data. Poor data quality leads to false positives, missed risks and potential regulatory sanctions. As the industry adage goes: "Surveillance platforms are only as good as the data they ingest."

Financial institutions face significant challenges in this area:

✕

Data Volume and Variety

Surveillance systems must process diverse data types from multiple sources, including order management systems, execution platforms, reference data feeds and market data feeds.

✕

Inconsistent Formats

Different systems use varied data structures and formats, requiring normalization before analysis.

✕

Organizational Silos

Relevant data often resides in isolated systems across different internal business units, making it difficult to perform comprehensive surveillance at scale.

Best practices for addressing these challenges include:

✓

Standardized Data Taxonomies

Establishing consistent naming conventions and data structures across the organization.

✓

Automated Quality Monitoring

Implementing tools to flag gaps or inconsistencies in surveillance data.

✓

Validation Checks

Regularly testing data completeness and accuracy against known benchmarks.

✓

Data Volume and Variety

Creating clear ownership structures and accountability for data quality across the surveillance lifecycle.

Future Outlook:
Industry Transformation and Agentic AI

The future of AI in financial compliance and trade surveillance promises both evolutionary improvements and potentially revolutionary changes. As both Moody's and Shopify have demonstrated through their strategic AI implementations, companies that effectively integrate these technologies stand to gain significant competitive advantages. Market participants can apply lessons from these successful implementations to surveillance, broader compliance needs and beyond.

Case Studies in AI Transformation

MOODY'S

The Moody's example highlighted earlier in this paper illustrates how a traditional financial services firm can successfully implement AI technologies. Rather than waiting to be disrupted, Moody's has proactively embraced AI as a transformative opportunity. The company has combined deterministic AI approaches with robust and trusted datasets to create industry-leading solutions that deliver precise, timely insights.

shopify

Similarly, Shopify has aggressively integrated AI into its commerce platform through tools like Shopify Magic, which provides natural language capabilities for store owners. Users can tap into GenAI for product descriptions, email campaigns, FAQ suggestions, image editing and more. Beyond that, Shopify CEO Tobi Lütke recently took a bold step by requiring teams to "demonstrate why they cannot get what they want done using AI" before requesting additional headcount, indicating the company's strong commitment to AI-first operations. This approach emphasizes the potential for AI to enhance productivity and efficiency across all business functions.

AI's Expanding Role in Financial Services

Across the financial sector, leading firms are publicly signaling that AI is not just a tool for efficiency, but a core driver of strategic transformation. JPMorgan CEO Jamie Dimon has stated AI could be "as transformational" as the printing press or electricity, emphasizing its potential across fraud detection, trading optimization and even software development. With over 2,000 AI technology experts embedded in the bank, JPMorgan is actively integrating AI across business lines while also investing in governance and infrastructure to support large-scale deployment. Similarly, BlackRock CEO Larry Fink has called AI a foundational shift for the global economy and a critical enabler of productivity in asset management. Both leaders highlight AI not as a marginal experiment, but a capability that will define the next era of financial services.

Other firms are taking similarly assertive positions. Morgan Stanley has deployed GenAI directly into its wealth management operations, enabling financial advisors to retrieve research and documentation via ChatGPT-powered tools. CEO James Gorman has framed this as enhancing — not replacing — human expertise. UBS, in the context of its merger with Credit Suisse, has pointed to AI as a key technology for scaling oversight and compliance, especially in an increasingly complex regulatory environment. Even digital-native players like Revolut are leaning heavily into AI-first operations, with CEO Nik Storonsky suggesting the company is building toward a future where autonomous AI systems handle the bulk of customer interaction.

Looking ahead, this trajectory points toward increasingly autonomous and adaptive systems — what some in the industry are beginning to describe as "agentic" AI. Morgan Stanley has referred to 2025 as the "year of agentic AI," signaling a shift from passive, prompt-based assistants to proactive agents capable of goal-driven planning, task execution and multi-system coordination. While most firms remain in the early stages of experimentation, the direction is clear: financial institutions are laying the groundwork for AI systems that not only support human decisions, but can begin to make and act on them — with oversight, but less direct instruction. This evolution will raise new questions around accountability, explainability and trust — but also offers a glimpse into a future of financial services shaped by intelligent, initiative-taking machines.

Focus on Surveillance: Emerging Technologies and Methodologies

1

Unsupervised Learning Advancements

The continued advancement of algorithms that can identify anomalies without labeled training data will improve detection of novel manipulation tactics. Research from the University of Technology Sydney shows promising results with new clustering algorithms inspired by how galaxies merge in space, achieving 97.7% accuracy across diverse datasets.

2

Cross-Platform and Cross-Market Surveillance

More sophisticated integration across traditionally siloed markets will enable better detection of manipulation schemes that span multiple venues or asset classes, aligning with an increasingly important regulatory priority.

3

From AI Assistants to Agentic AI in Compliance

AI is beginning to find its way into compliance workflows, primarily in the form of AI assistants designed to support, rather than replace, human decision-making. These early-stage AI assistants can summarize large volumes of data, suggest responses to alerts or help interpret regulatory texts — but they still require direct prompting and human oversight.

As organizations grow more comfortable with these capabilities, the next phase may involve AI agents akin to those advocated by Morgan Stanley. These agents can autonomously carry out multi-step workflows, adapt to new inputs and interact with other systems in the compliance stack. Ultimately, this progression could lead to an environment where agentic AI — systems that act with a high degree of autonomy to pursue specific compliance objectives — is the industry norm. In such a scenario, agentic AI might continuously monitor trading activity, detect emerging risks in real time and even initiate preliminary investigations without human initiation.

✓ **While the potential efficiencies are enormous, deploying agentic AI in compliance raises serious questions around oversight, accountability and trust in machine-led judgment that will be hotly debated in the years to come.**

The Path Forward

As Michael Barr noted in his speech on AI's future impact, we could potentially see incremental progress or transformative change. While the financial industry tends toward cautious evolution, technological advancement may accelerate this timeline, pushing firms to pursue more rapid adoption than initially planned.

Regardless of the pace of innovation, success will depend not on implementing every new capability, but on applying the right technologies in ways that enhance rather than compromise regulatory compliance. Organizations that maintain this balance — leveraging AI's power while ensuring transparency, accountability and human oversight — will be best positioned to thrive in the evolving landscape.

CONCLUSION

Strategic Considerations for Successful AI Adoption

The integration of AI into financial compliance and trade surveillance represents a significant opportunity to enhance detection capabilities, improve efficiency and manage the growing complexity of market oversight. However, realizing these benefits requires thoughtful implementation that balances technological innovation with regulatory expectations and ethical considerations.

Strategic Principles for Implementation

- ✓ Prioritize explainability and transparency in AI applications, ensuring that surveillance findings can be clearly articulated to regulators and stakeholders.
- ✓ Keep humans in the loop, leveraging AI to augment rather than replace skilled compliance professionals.
- ✓ Develop comprehensive governance frameworks that address both technical performance and ethical considerations.
- ✓ Maintain strong data governance as the foundation of effective AI, recognizing that even the most sophisticated models cannot overcome poor inputs.
- ✓ Take a phased, iterative approach to implementation, starting with well-defined use cases and expanding based on demonstrated success.

By following these principles, financial institutions can harness the potential of AI in surveillance while managing associated risks. In time, the result should be more effective compliance functions that protect firms and markets from manipulation while allowing legitimate trading activities to flourish — all at a sustainable cost in an increasingly complex regulatory environment.

As this technology continues to evolve, ongoing dialogue between market participants, technology providers and regulators will be essential in shaping systems that are not only intelligent but also ethical, adaptable and aligned with the ever-changing nature of the markets. Organizations that view AI not merely as a technology tool, but as a strategic capability that requires thoughtful governance and integration with human expertise, will be best positioned to succeed in this new era of financial compliance.

Eventus is a leading global provider of multi-asset class trade surveillance and market risk solutions.

Our powerful, award-winning Validus platform is easy to deploy, customize and operate and is proven in the most complex, high-volume and real-time environments.

EVENTUS.COM

EVENTUS