AI in trade surveillance

Turning innovation into impact

Authors

Ben Parker,
CEO and Founder

Ross Pearson,
Head of AI

Alex Parker,
CTO and Founder

Market abuse in 2024: A turning point for surveillance

2024 was a significant year in the fight against market abuse. Enforcement surged to $1.8 billion, the second-highest annual total on record, spanning 163 separate cases. Regulators repeatedly cited weaknesses in firms' trade surveillance systems, with related enforcement value rising by more than 800% year-on-year.[1]

Alongside the scale of penalties, the complexity of market abuse intensified. From sophisticated **cross-market manipulation** schemes to the SEC's first-ever **shadow trading charge**, regulators and firms are confronting new forms of misconduct that are harder to detect, prove, and prevent.

> *It requires a lot of computational brainpower to sift through data and identify relationships that aren't obvious - such as connections between firms that don't share the same product line or industry but are still somehow related in their trading behaviour.*
>
> **Ben Parker, CEO of eflow Global**

This unfolded alongside another defining trend of 2024, the rapid rise of artificial intelligence, which has only escalated in 2025. 70% of financial services executives now believe that AI will directly drive revenue growth in the years ahead.[2]

These parallel developments have prompted the question: how is AI impacting market abuse and trade surveillance, both as a tool for detection and a weapon for bad actors?

The dual role of AI: Tool and threat

Artificial intelligence has become one of the most dynamic forces shaping global finance. Its ability to process data, detect complex patterns, and operate at scale has transformed how firms trade, analyse risk, and monitor behaviour. Generative AI, in particular, has introduced a new level of adaptability, capable of learning, reasoning, and communicating in ways that traditional systems could not - especially where unstructured text data is concerned.

This versatility makes AI a defensive and offensive force in financial markets. It can strengthen surveillance and compliance capabilities, but it can also be exploited, intentionally or unintentionally, to perpetrate market manipulation or amplify systemic risk.

AI as a threat

The risk of AI directly committing market abuse is increasing. Algorithmic and high-frequency trading systems already execute millions of orders per second and withdraw them just as quickly. As AI capabilities are embedded within these systems and surrounding processes, subtle emergent behaviours could translate into large, market-moving effects.

> *I suspect that algorithms are already manipulating markets, whether 'intentional' or not, by simply inflating the order book. Most orders in the order book are just algorithms.*
>
> **Anonymous Quantitative Trader at Proprietary Trading Firm[3]**

While regulators do not yet view AI-driven market abuse as a widespread problem, they are watching closely. The FCA and Bank of England's 2024 joint report on AI adoption found that over 11% of UK financial institutions already use AI for trading, with another 9% planning adoption by 2027.[4] That level of penetration means supervisory frameworks must increasingly account for the emergent behaviour of models operating in complex market environments.

AI as a surveillance tool

AI has the potential to transform how misconduct is detected. Surveillance teams are working hard to ingest and analyse more data: millions of daily orders, trades, and messages, each containing faint signals that could indicate manipulation or collusion. Traditional systems struggle to keep pace, especially when abuse spans instruments, venues, or jurisdictions.

AI can change this dynamic. Machine learning models can learn behavioural fingerprints from historical data, recognise subtle correlations across seemingly unrelated entities, and continuously adapt to new market regimes. Large-language models (LLMs) can infer intent and sentiment in trader communications, helping analysts prioritise genuine threats over noise.

These technologies carry potential to focus, augment or even enhance human expertise. The outcome is faster triage, richer insight, and more proactive surveillance — a necessary evolution as firms respond to both technological and regulatory pressures.

The use of AI in trade surveillance

Regulators' focus on surveillance has translated into a clear financial impact. In 2024, fines linked specifically to trade surveillance failures reached $677 million, with the number of cases doubling from 16 to 33 year-on-year. Deficiencies in detection and monitoring have become a central theme in enforcement, and regulators expect firms to modernise their frameworks in response to elevated risk.

This section of the report explores how the use of AI can improve trade surveillance systems, as firms respond to this regulatory crackdown.

Classification models: The first generation of AI in surveillance

Early implementations of AI in trade surveillance centred on supervised classification models. These systems were trained on historical examples of confirmed market abuse to identify similar behaviours in new datasets.

They function by analysing large volumes of trading data to detect the statistical patterns, or "fingerprints," that characterise specific forms of manipulation. This allows the system to extract behavioural rules directly from data rather than relying solely on manually defined detection parameters.

The edge

Rules-based surveillance systems rely on human-defined thresholds: a trader cancels a certain number of orders within a given time frame, or a price moves a set percentage before being reversed. These systems are explicit, transparent, and easy to audit, but they can also be rigid.

Classification models offered a potential improvement by learning from historical data. In theory, they:

- Detect subtle, multi-dimensional patterns too complex for static rule logic.
- Adapt dynamically as markets evolve, reducing manual recalibration.
- Weight contextual factors, such as time of day, volatility, and trader behaviour, to differentiate genuine risk from normal activity.

The reality

Despite years of experimentation, few firms have operationalised classification-based surveillance at scale or achieved measurable gains over rules-based systems.

Key challenges include:

- Data scarcity: Large volumes of accurately labelled examples are needed, yet confirmed abuse cases are rare, fragmented, and often confidential.
- Limited transferability: Models trained on one firm's data rarely perform well elsewhere due to differences in trading behaviour and client profiles.
- Validation overhead: Markets evolve constantly; models require frequent retraining and revalidation to prevent drift, adding operational and governance complexity.

Exploring the unknown unknowns

Where classification models search for known forms of market abuse, anomaly detection targets the unknowns — behaviours that deviate from historical norms but have not yet been labelled as suspicious. These techniques are designed to extend the reach of rules-based systems by identifying emerging risks and new abuse typologies that fall outside of existing rules or calibration thresholds.

Unlike supervised learning, which depends on labelled datasets, anomaly detection uses statistical and probabilistic reasoning to detect patterns that are improbable given the data distribution.

Bayesian networks are an example of this approach. These probabilistic graphical models infer relationships between variables and detect low-probability states.

Compared to classification systems, these approaches are more dynamic and context-aware:

- They don't require labelled data. Instead of learning from past abuse cases, they infer conditional probabilities from ongoing market activity.
- They model interdependencies. Bayesian graphs capture how changes in one behaviour (e.g. order cancellations) may affect another (e.g. price movement), allowing for the detection of multi-dimensional anomalies.
- They adapt naturally to new regimes. As trading behaviour evolves, conditional probabilities update, reducing the need for full retraining.

In essence, classification systems are designed to recognise the expected forms of misconduct, while Bayesian models can detect what has never been seen before. This capability makes them well-suited for modern surveillance environments where both trading activity and abusive tactics continuously evolve.

eflow at the FCA Market Abuse Surveillance TechSprint

As part of the FCA's 2024 Market Abuse Surveillance TechSprint (MASTS), eflow showcased how machine learning can make market abuse detection more accurate and adaptive.

eflow demonstrated an AI-driven feedback loop that learns from historical surveillance outcomes – distinguishing high and low-value alerts and dynamically tuning detection parameters to remain fit for changing trading behaviours.

Using the FCA's real-world market dataset, eflow's model identified patterns where both traditional and ML-based methods agreed on high-risk behaviour, while also surfacing new risks that static, parameter-based systems missed.

Watch eflow's FCA presentation here

Prioritisation and triage: The most immediate ROI

The most immediate and measurable impact of these new technologies lies in alert triage and prioritisation; using AI to distinguish between cases deserving of analyst attention and those which can be safely deprioritised.

LLMs have accelerated this capability by automating key stages of the analyst workflow:

- Context assembly: Pulling together relevant trade data, order books, communications, and historical decisions to build a complete picture around each alert.
- Initial assessment: Summarising the evidence and highlighting patterns that suggest whether behaviour is consistent with prior false positives or merits escalation.
- Alert ranking: Scoring and sorting alerts by probable risk level based on contextual and behavioural indicators.
- Feedback learning: Incorporating analyst outcomes to refine future prioritisation, reducing manual workload over time.

This doesn't replace human expertise; rather, it creates a tiered workflow in which AI handles operationally heavy, repetitive tasks, allowing analysts to focus on higher-value judgement and escalation decisions.

As these models work within existing governance and validation frameworks, they offer the most immediate, low-risk opportunity for firms to realise operational benefits from AI. The result is faster time-to-insight, lower investigation costs, and a clearer path from alert to action; a practical application of AI that helps firms act on real risks faster.

LLMs for eComms surveillance

Historically, firms have relied on lexicon-based systems, built on static lists of trigger words and phrases. These systems are not built to interpret context, intent, or sentiment, and can easily miss suspicious conversations that use coded language, slang, emojis or other non-English expressions.

LLMs, on the other hand, can interpret the semantic meaning of communications, detecting behavioural nuance even when no explicit red flags are present. Firms are now layering these models on top of existing keyword filters, using sentiment analysis, toxicity scoring, and anomaly detection to prioritise alerts and surface genuinely suspicious behaviour for faster review.

More generally, firms struggled with consistently handling inconsistent, unstructured data. LLMs excel with unstructured data, and can even derive insight into certain structured fields. This helps with both identifying new risks and prioritising existing cases (e.g. those based on trade signals). If firms can combine insight from eComms data with their trade surveillance data, they can maximise the context available within each alert.

AI co-pilots: Talk to your data

Thinking beyond workflow automation, AI provides better accessibility of insight. Firms are sitting on extensive knowledge bases (trade data, communications, reference sources, or past decisions) from which they can draw valuable inference. LLMs allow them to better make use of this approach.

Retrieval-augmented generation (RAG) and LLM-based co-pilots combine natural language querying with secure access to structured and unstructured surveillance data, allowing analysts to talk to their data directly.

The result is a democratisation of knowledge. Whereas certain information was once several structured queries away (e.g. SQL), now non-technical individuals can access insights that would have previously required specialist support.

The challenges that lie ahead

The potential of AI in trade surveillance is enormous. When deployed effectively, it can reduce operational costs through workflow automation, strengthen compliance and avoid enforcement penalties by improving alert coverage and configuration, and ultimately support revenue by enabling firms to expand with confidence. In this sense, AI is not just a technological advantage, it is a strategic one.

Yet firms' readiness to realise these benefits varies widely. The landscape of technological maturity across the industry is uneven, and early optimism can easily give way to costly missteps. The next phase of AI adoption in surveillance will demand not only innovation, but rigour: in data management, in talent and governance, and in the ability to balance automation with accountability.

Data maturity

AI is only as strong as the data it learns and infers from. Traditional machine learning relies on highly structured, well-labelled datasets, something most financial institutions still lack.

While LLMs are more tolerant of unstructured and heterogeneous data, they don't eliminate the need for data discipline. Firms must first decompose surveillance problems into their atomic data elements: define what each variable represents, how it behaves under different market conditions, and how it should be standardised.

Achieving this level of data maturity requires significant investment in data engineering, taxonomy design and governance before AI can be effectively deployed. Without it, models risk amplifying noise.

Skillset and resourcing

There is a global shortage of professionals who combine AI engineering expertise with deep financial domain knowledge. While low-code and "vibe coding" platforms can accelerate prototyping, scaling to enterprise-grade systems demands teams that understand both the mathematical foundations of models and the regulatory context in which they operate.

Effective implementation depends on cross-disciplinary collaboration between data scientists, compliance officers and risk managers to ensure that model outputs align with surveillance objectives and supervisory expectations.

Cost and operational overheads

Even for firms that have the talent, the economics of AI deployment remain challenging. Training, fine-tuning, validating, and maintaining models require substantial investment in compute, data infrastructure, and governance tooling.

Even with pre-trained models, inference, retraining, testing and performance monitoring create recurring costs that must be justified through measurable efficiency gains or accuracy improvements.

Explainability and accountability

Accountability in financial regulation cannot be delegated to a machine. It resides squarely with the firm, and, more specifically, with individuals. Under regimes such as the UK's Senior Managers and Certification Regime (SM&CR), this principle is explicit. The FCA is also clear that the rules emphasising accountability for senior managers are "relevant to the safe use of AI."[5]

The opacity of AI decision-making is one of the biggest barriers to adoption. Surveillance systems that generate alerts or conclusions without transparent reasoning make it difficult for firms or regulators to understand how those outcomes were reached or to verify their compliance.

Firms are therefore reluctant to give AI significant autonomy because the inability to explain or justify its outputs translates directly into personal and institutional risk.

Best practices for AI adoption

The question is no longer whether AI can improve surveillance programmes, but how to implement AI to achieve a true return on investment. Based on experience implementing surveillance solutions across global financial institutions, eflow recommends a set of pragmatic best practices that balance innovation with governance, and value with control.

Leverage human-first AI

AI should be implemented to enhance human decision-making. Surveillance outcomes carry significant regulatory, financial and reputational implications. Accountability must remain with the firm and its people.

- Keep humans in the loop at key decision points. AI can pre-classify or rank alerts, but escalation and closure decisions should remain analyst-owned.

- Build feedback loops where analyst actions (dismissing, escalating, or re-categorising alerts) are captured and used to retrain or fine-tune models.

- Build traceability into AI-driven insights that show analysts exactly what data the model draws on to reach its conclusion. This expedites and enhances the human review process, moving away from blind acceptance.

Prioritise data governance

Invest in building solid data foundations that ensure consistency and transparency.

- Consolidate trade, order and communications data into a unified, well-structured repository. Fragmented or duplicated datasets can generate inconsistent results.

- Define clear taxonomies and metadata standards so that every field — from order type to trader ID — has a format aligned to trade reporting standards.

- Implement quality controls to track how data moves, transforms and is used by models. Automated validation checks should flag missing or uncategorised data before it reaches the model layer.

Choose security over convenience

Surveillance data is sensitive. Privacy and data protection must take precedence over speed or ease of use.

- Keep surveillance data within secure environments and apply least-privilege principles
- Ensure sensitive records are not transferred to open or unvetted Shadow AI models
- Ensure only authorised users and systems can access AI outputs or underlying data

Partner with technology experts

Deploying AI for trade surveillance requires expertise across multiple domains. Few firms can maintain these capabilities in-house at the level needed for regulatory-grade performance. Strategic partnerships can close the gap while preserving internal oversight.

- Define your core competency. Self-reflection at the organisational level is important. Firms should assess their appetite and capacity to operate as a technology-driven organisation. Understanding this helps determine which functions should be built internally and which are better outsourced to partners with the scale and specialisation to deliver securely and efficiently.

- Adopt a hybrid approach by combining in-house subject matter expertise with specialist technology partners who bring proven models, scalable infrastructure, data governance frameworks and ongoing validation support.

Conclusion

AI is set to transform trade surveillance, but it is not a silver bullet. Firms should focus on practical, high-impact applications that deliver measurable returns today, rather than chasing fully autonomous systems tomorrow. Human-first AI — where automation strengthens rather than replaces oversight — is becoming the regulatory baseline, as supervisors themselves experiment with AI and raise expectations around governance, explainability, and accountability. Strong data foundations, model validation, and clear ownership remain essential to building trust in AI-driven surveillance.

Partner with experience. Work with vendors who understand both surveillance and supervision — who know where AI fits and where it doesn't — and who prioritise data governance with no compromise on security.

About eflow

Since 2004, eflow has had a clear mission: to help financial institutions meet their regulatory obligations in the most robust and efficient way possible.

To achieve this, we first had to identify why so many firms either struggled to demonstrate their compliance or spent far too much time, effort and money in doing so. We found that for many institutions, their regulatory processes were broken. An over-reliance on spreadsheets and siloed data. Slow, legacy reporting systems that were no longer fit for purpose. Or, an unscalable point of failure in the form of one person 'who has always looked after compliance'.

Here at eflow, we took a different approach. eflow technology is built on PATH, our robust and standardised digital ecosystem that integrates seamlessly with each of our specialist regtech modules. This unique technological model offers firms the speed, convenience and efficiency of an off-the-shelf software solution, combined with a level of customisation that is typically only associated with a bespoke platform. This means that as new regulatory challenges arise, as they inevitably will, you can rest assured that eflow's regulatory tools will already be one step ahead.

Explore our regulatory technology solutions at www.eflowglobal.com.

Trade Surveillance

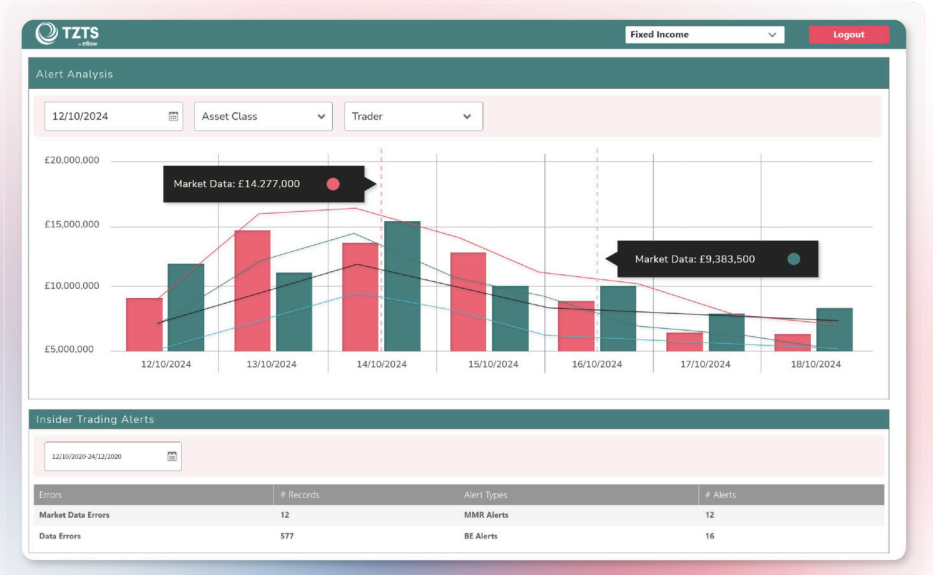

Best Execution

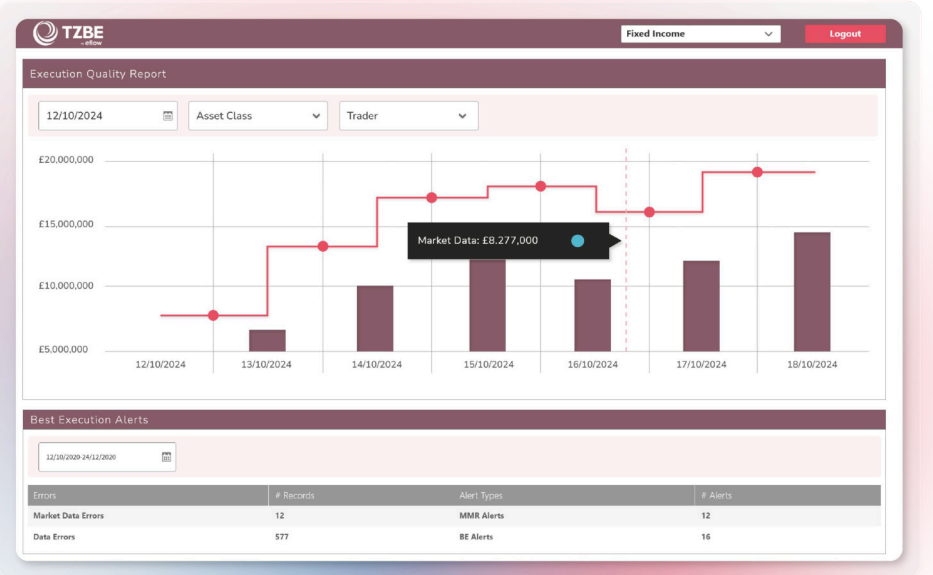

eComms Surveillance

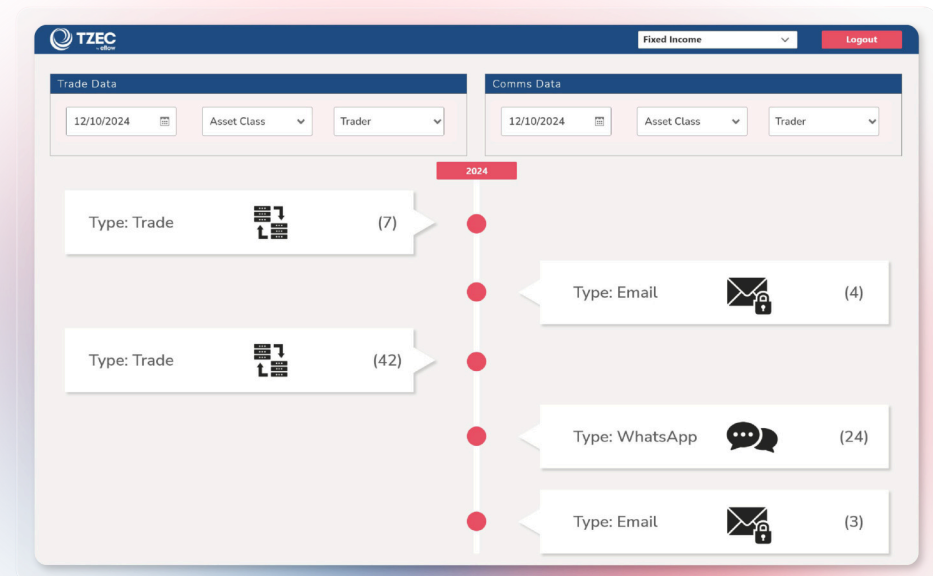

Transaction Reporting

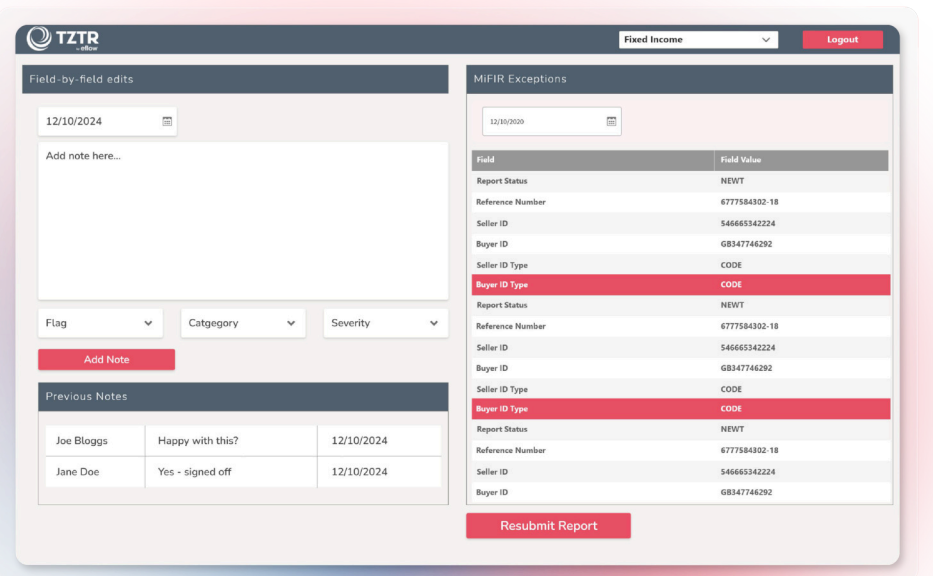

References

1. *Global trends in market abuse and trade surveillance 2025 - Report published by eflow, March 2025.*

2. *Artificial Intelligence in Financial Services - Report published by World Economic Forum, January 2025.*

3. *Global trends in market abuse and trade surveillance 2025 - Report published by eflow, March 2025.*

4. *Artificial intelligence in UK financial services - Report published by the FCA and Bank of England, November 2024.*

5. *AI and the FCA: our approach - Article published by the FCA, September 2025.*