Account and Financial Management Journal e-ISSN: 2456-3374
Volume 11 Issue 02 February 2026, Page No.- 4270-4284
DOI: 10.47191/afmj/v11i2.11, Impact Factor: 8.167

Code, Compliance and Crashes: A Risk-Based Framework for AI Trading Systems

Dr. G. Prasanna Kumar[1], Dr. Manoj Kumar[2], P.V.S Swamy[3]
[1]Assistant Professor, Department of Management Studies, Pragati Engineering College, Surampalem, Andhra Pradesh, India. Pin - 533437, ORCID: https://orcid.org/0000-0002-4753-7580
[2]Assistant Professor, Department of Business Management and Commerce, Desh Bhagat University Mandi Gobindgarh, Punjab
[3]Assistant Professor, Department of Management Studies, Pragati Engineering College, Surampalem, Andhra Pradesh, India Pin - 533437

ABSTRACT: The use of AI in securities markets has transformed the design of securities markets by introducing adaptive machine-learning and reinforcement-learning algorithms in order generation, routing and execution. This chapter focuses on the re-definition of market, operational and conduct risks by such systems and suggests a risk-based system of internal governance of institutions that use AI trading strategies. The analysis involves a combination of the analysis of the doctrines of regulatory regimes of key financial centres and syntheses of academic literature, case facts of flash crashes and near-misses, and conceptual maps of risk channels throughout the data-model-execution lifecycle. The results indicate a closed risk universe where market integrity issues, model and data risk, technology failures, cyber threats and data-protection concerns are intertwined. The current algorithmic trading regulations, which are mainly based on deterministic code, partially cover these dynamics. The chapter advances a governance framework based on human-in-the-loop supervision, proportional clarification, explicit duty all the way along the three lines of defence, and graded oversight that is aligned with the risk of strategy. Tools such as concrete tools, improved testing regimes, dynamic pre-trade limits, kill-switches, ongoing monitoring with explainability analytics, and detailed model inventories and audit trails are all concrete tools. The discussion admits shortcomings of a conceptual and legal-analysis method: there is yet to be empirically calibrated proposed controls, particularly in realistic multi-agent AI market settings. However, the framework has definite implications on regulators and companies. The regulators are urged to improve standards of algorithmic and AI trading, increase system wide stress testing and facilitate cross-border coordination. The companies are urged to incorporate AI-specific risk policies, ethics codes and psychosocial support to supervisors facing opaque and high-speed systems. Besides, ethical guidelines of fairness, non-manipulation, transparency and inclusiveness are also applied to design decisions concerning data selection, validation practises, client communication and internal ethics management. The difference is in the fact that microstructural risk analysis, comparative regulation, internal governance design and human-factor considerations are combined within one, coherent structure of AI trading. The chapter therefore adds to the current discussions on reliable financial AI by describing how trading technology innovation can be balanced with systemic stability and investor protection.

KEYWORDS: Algorithmic trading, Artificial intelligence, Market integrity, Systemic risk, Regulatory governance, Model risk
JEL Codes: G14; G18; G28; K22; O33

1. INTRODUCTION: AI TRADING AND THE NEW RISK LANDSCAPE

1.1 Background and Motivation: From Code to Crashes

Within the past 30 years, securities markets have ceased to be floor-based trading and are now dissolved into fragmented electronic ecosystems, in which orders are created, transmitted and cancelled at machine speed. A significant portion of equity and derivatives volumes is currently mediated by algorithmic and high-frequency traders, which is changing the way liquidity and price discovery is practised (Gomber and Haferkorn, 2013; Virgilio, 2019). Simultaneously, machine-learning techniques have driven trading beyond fixed-rule automation into adaptive, high-dimensional, real-time adaptive, or AI trading systems (Bahoo et al., 2024; Dakalbab et al., 2024).

This change will offer narrower spreads and more effective execution, but it will also increase the apprehension regarding control, fairness and stability of the system (Kirilenko and Lo, 2013; Jarrow and Protter, 2012). AI models are opaque and brittle: any error in configuration, unexpected feedback loop or change of market regime can create an outsized flow of orders before human risk doctors can respond (Johnson et al., 2013). The May 6, 2010 “Flash Crash” remains a paradigmatic illustration of how interacting automated strategies can trigger extreme, short-lived dislocations in key indices (Kirilenko et al., 2017; Easley et al., 2011). Subsequent mini-flash events and anomalous order-book dynamics suggest that such instability is a recurring feature of highly automated market microstructure rather than a unique accident (Virgilio, 2019).

AI techniques now diffuse beyond specialist proprietary traders to banks, brokers, hedge funds and fintech platforms (Bahoo et al., 2024; Alcázar-Blanco et al., 2024). Deep learning and reinforcement learning are deployed not only for signal generation and execution, but also to automate risk limits and intraday margin management (Dakalbab et al., 2024). This raises difficult questions about responsibility, explainability and auditability when trading systems are partly self-learning and continuously updated in production, and about the continued adequacy of traditional three-lines-of-defence models.

Regulators have responded with algorithmic trading regimes that emphasise pre-trade controls, kill switches, testing and governance expectations for automated systems. Yet these frameworks were largely designed for deterministic algorithms rather than adaptive, data-driven models, and their effectiveness in containing extreme events remains uncertain (Kirilenko and Lo, 2013; Virgilio, 2019). The emerging risk landscape is therefore defined by a tension between speed and complexity on the one hand, and the regulatory need for control, robustness and fairness on the other. This chapter takes that tension as its starting point and asks what a risk-based internal governance framework for AI trading systems should look like.

1.2 Objectives, Scope and Research Questions

Against this backdrop, the chapter has three linked objectives. First, it analyses how AI-enabled trading reshapes the market, liquidity, operational and model risk profile of firms active in equity and exchange-traded derivatives (Jarrow and Protter, 2012; Johnson et al., 2013). Second, it develops a risk-based internal governance framework for AI trading systems that integrates model lifecycle management, technical safeguards and human oversight. Third, it derives policy implications for supervisors and market operators seeking to align firm-level governance with concerns about systemic fragility (Kirilenko et al., 2017; Virgilio, 2019).

The scope is bounded to institutional settings—banks, broker-dealers, proprietary trading firms, hedge funds and fintech brokers—active in listed equity and derivatives markets. Over-the-counter markets, retail robo-advisory and crypto-asset trading are considered only where they illuminate governance issues for regulated securities markets (Bahoo et al., 2024). Within this perimeter, the chapter addresses three guiding questions:

- How do AI-enabled trading architectures alter the propagation of market, liquidity, operational and model risks relative to earlier algorithmic trading?
- Which governance arrangements across risk, compliance, technology and front-office functions are needed to control these risks proportionately?
- How should regulatory and supervisory frameworks evolve to recognise the distinctive properties of AI trading systems while preserving innovation and competition?

1.3 Methodology and Structure of the Chapter

Methodologically, the chapter combines doctrinal and policy analysis with a selective review of literature on algorithmic and AI trading, market microstructure and financial regulation (Kirilenko and Lo, 2013; Virgilio, 2019; Bahoo et al., 2024). It further provides a comparative evaluation of the regulatory regimes that guide automated trading in the United States and in the European Union in the perspective of risk-based supervision. The presentation, throughout, is based upon concrete episodes (the Flash Crash and the string of instability that followed), but also on vignettes based on enforcement actions and supervisory statements in different jurisdictions (Kirilenko et al., 2017; Virgilio, 2019).

On this basis, Section 2 maps the evolution of electronic and AI trading architectures, tracing pipelines for data, signals, execution and risk control. Section 3 develops a typology of risk channels specific to AI trading systems, linking micro-level model failures to market-wide stress scenarios. Section 4 proposes a risk-based internal governance framework comprising model governance, technical safeguards and human-in-the-loop oversight. Section 5 draws the regulatory and policy implications, and Section 6 concludes with a research agenda on AI trading, governance and systemic stability for the reader.

2. EVOLUTION OF ALGORITHMIC AND AI TRADING SYSTEMS

2.1 From Rule-Based Automation to Machine Learning Models

The evolution of electronic trading began with relatively simple program trading strategies implemented by large institutional investors in the 1980s. These early systems automated index arbitrage and portfolio insurance by encoding fixed decision rules that could submit baskets of orders when prices deviated from pre-defined thresholds. Their logic was transparent—if-then conditions based on observable price relationships—and human supervisors could typically anticipate their behaviour, even if they could not match their speed (Aldridge, 2013).

During the 1990s and early 2000s, demutualisation of exchanges and the spread of electronic limit order books created fertile ground for more sophisticated algorithmic trading. Order handling functions that had previously relied on floor brokers were progressively delegated to automated "execution algorithms," which sliced large institutional orders into smaller child orders to minimise market impact and implementation shortfall. These algorithms were still mostly rule based and were based on historical volume profiles and naive statistical models of intraday volatility.

In the mid-2000s a major change happened with the emergence of high-frequency trading (HFT) which was facilitated by co-location, direct market access and low-latency specialised infrastructure. In HFT, ultra-fast strategies continuously update and cancel quotes in response to order-book changes, exploiting fleeting price discrepancies across venues (Gomber and Haferkorn, 2013; Virgilio, 2019). Latency competition became a central dimension of market structure, with firms investing heavily in hardware acceleration, microwave links and optimised code paths. Rule-based logic still dominated, but the complexity of strategies and their tight coupling to market microstructure made their aggregate behaviour difficult to predict.

The diffusion of machine-learning techniques has further transformed trading. Initially, ML models were used to enhance signal generation—for example, classifying short-term return patterns, predicting order-book imbalance or inferring hidden liquidity from order-flow dynamics. Over time, deep learning architectures, including convolutional and recurrent neural networks, have been applied to high-dimensional market and alternative data, enabling non-linear feature extraction and pattern recognition beyond traditional factor models (Dakalbab et al., 2024; Anuar et al., 2025). Reinforcement learning (RL) and deep RL approaches go further by modelling the trader as an agent learning to maximise a reward function, such as risk-adjusted returns or execution quality, through interaction with a simulated or real market environment (Zhang et al., 2021; Tamma and Bégin, 2023).

The key qualitative change from earlier generations of automation lies in the iterative and adaptive nature of AI models. Whereas rule-based algorithms embody fixed decision logic designed ex ante, ML and RL models are trained on historical or streaming data and updated through retraining or online learning. Their performance is inherently path-dependent: model parameters evolve with new information, and interaction effects between multiple adaptive agents can generate complex emergent behaviour. This adaptivity offers potential performance gains but also introduces new forms of model risk and makes ex ante validation more challenging than for static rules.

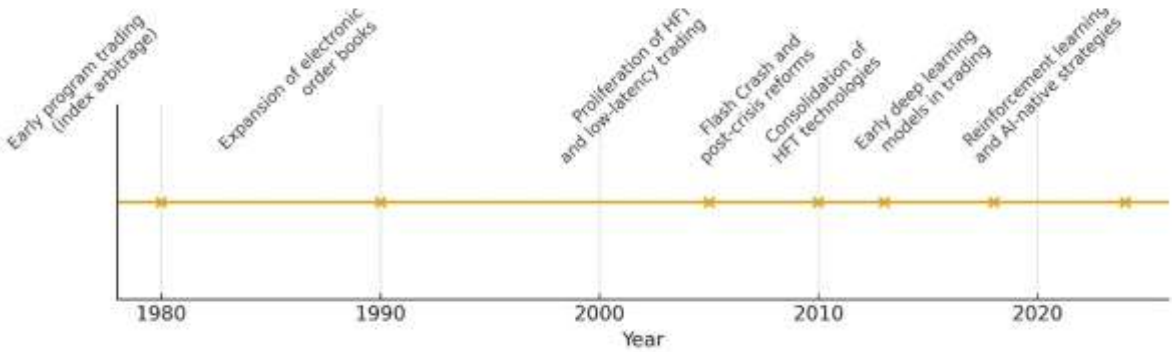

Figure 1. Timeline of Algorithmic and AI Trading Milestones.
Source: Author's elaboration based on Gomber and Haferkorn (2013), Virgilio (2019), Dakalbab et al. (2024) and Anuar et al. (2025).

2.2 Trading Architecture, Market Microstructure and Organisational Dynamics

A contemporary AI trading system typically consists of several tightly coupled layers. At the base is the data pipeline, which ingests, cleans and time-aligns multiple data sources: limit order book updates, trades, reference and corporate action data, and increasingly alternative data such as news feeds and sentiment indices (Dakalbab et al., 2024). Feature engineering then transforms raw inputs into predictors—order-flow imbalance measures, volatility estimates, technical indicators or text-derived sentiment scores.

Selection, calibration and validation of forecasting or decision models are in the model training layer. It can include gradient-boosted trees and random forests, to deep neural networks and RL policies, and frequently used in an ensemble architecture to make it more robust. Back-test, stress-testing of models on past data is sometimes done with agent-based simulations of market microstructure to capture effects of interaction (Virgilio, 2019). Upon approval, models are marketed to the deployment layer where they produce real time trading signals within risk limits, throttles and execution restrictions.

The execution layer interprets signals into specific order, and sends them through venues and liquidity pools based on the dynamic cost and risk measures. The aspects of venue selection, choice of order type and pacing are controlled by smart order routers and execution algorithms to control the market impact and information leakage (Aldridge, 2013). Surrounding this basic architecture are monitoring dashboards, kill switches and logging, which are meant to give real-time access to model outputs, exposures and order-flow anomalies.

These technical layers map onto specific organisational roles. Trading desks and quantitative teams design strategies, specify model objectives and evaluate performance. IT and data engineering teams are responsible for infrastructure resilience, latency management and data quality. Compliance, risk management and internal audit functions provide second- and third-line oversight, focusing on model governance, limit setting, control design and post-incident review (Anuar et al., 2025). In principle, this division of labour should ensure that AI trading systems are both innovative and properly controlled.

In practice, however, organisational silos can obscure end-to-end risk. Quants may optimise model performance against historical backtests without fully considering how their strategies will interact with other internal or external algorithms in stressed markets. Risk managers may focus on portfolio-level metrics without seeing the detailed behaviour of individual models at microsecond timescales. Compliance teams may lack the technical visibility needed to interpret anomalous patterns in order flow, especially when models are opaque deep learning or RL agents (Dakalbab et al., 2024; Zhang et al., 2021). These frictions create blind spots where technical, model and conduct risks can accumulate unnoticed until crystallised by a market shock.

Table 1. Typology of AI Trading Strategies and Techniques.

Strategy Type	Core Techniques (ML / Rules)	Typical Data Inputs	Key Risk Drivers
Market making / liquidity provision	Rule-based quoting; predictive models for short-term volatility and order-book imbalance	Limit order book (LOB) data, trade prints, venue fee/rebate schedules	Adverse selection, inventory risk, sudden loss of liquidity, latency shocks
Statistical arbitrage	Time-series and cross-sectional models; machine learning for spread/mean reversion signals	Historical prices and volumes, factor data, cross-asset relationships	Model misspecification, regime shifts, crowded trades, funding and leverage risk
Execution algorithms	Rule-based slices (VWAP, TWAP, POV); adaptive scheduling with ML forecasts of volume and volatility	Intraday volume curves, real-time LOB and trade data, client order constraints	Implementation shortfall, market impact, malfunctioning throttles, interaction with venue microstructure
Sentiment-driven strategies	Natural language processing (NLP), deep learning classifiers and regressors	News feeds, social media streams, earnings calls, macro announcements	Noisy or manipulated text data, sentiment regime shifts, event misclassification
Reinforcement learning (RL) agents	Deep Q-learning, policy gradients, deep deterministic policy gradient and related DRL methods	Multivariate price series, technical indicators, volatility and liquidity measures	Exploration–exploitation trade-off, overfitting to backtests, unstable policies in changing environments

Source: Author's compilation based on Dakalbab et al. (2024), Virgilio (2019) and Anuar et al. (2025).

2.3 Case Snapshots: Crashes, Anomalies and Learning Events

The interaction of complex algorithms with market microstructure has periodically produced dramatic price dislocations. The 6 May 2010 Flash Crash remains the most widely studied example: within minutes, major US equity indices plunged by about 9 per cent before rebounding almost as quickly. The event was later attributed to a massive sell algorithm engaging with high-frequency liquidity providers whose quick order cancellations made the prices and liquidity head into a downward spiral (Easley et al., 2011; Kirilenko et al., 2017). Though this episode came before massive use of deep learning and RL, it demonstrated how automated strategies can work together to create extreme, non-linear dynamics which are beyond human intervention capability.

Subsequent mini flash crashes in single securities and ETFs have followed the same pattern: sudden bursts of price or a crash, usually due to interaction between fragmented liquidity, aggressive types of orders and strategies sensitive to latency (Virgilio, 2019). Such incidents normally indicate shortcoming in circuit breakers, volatility auctions or kill-switch mechanisms, and resulting small-scale modifications to exchange-level protections. In governance terms, they emphasise the value of modelling not only the individual strategy performance, but also feedback loops on the system level across venues and participants.

There are also internal near-miss situations, which offer an extra frame of education. As an illustration, there have been incidences of misconfigured algorithms or model drift that led to massively increasing order submissions or position acquisition, only to be halted when pre-trade risk-related controls or kill switches were activated (Aldridge, 2013; Anuar et al., 2025). The post-incident reviews often point to flaws in the change-management procedures, lack of stress testing in unfavourable liquidity situations, or lack of communication between the quantitative developers and control functions.

All these instances have valuable lessons to be learned in the following chapters. To begin with, they demonstrate that speed and automation increase errors and feedback, thus that local mis-specifications can spread to market-wide stress before human supervisors can intervene. Second, they demonstrate that technical safeguards are necessary but not sufficient: circuit breakers, throttles and kill switches must be embedded within a broader risk-based governance framework that includes robust model validation, clear accountability and cross-functional communication (Dakalbab et al., 2024; Zhang et al., 2021). Third, they underline the need for regulators and firms alike to adopt a systems perspective, recognising that the stability of AI-enabled markets depends on interactions among diverse agents, venues and rules rather than on any individual algorithm. These insights motivate the risk mapping in Section 3 and the governance architecture developed in Sections 4 and 5.

3. MAPPING THE RISK UNIVERSE IN AI TRADING

3.1 Market Integrity and Systemic Risk

AI-enabled trading operates in an environment already vulnerable to sophisticated market-abuse strategies such as spoofing, layering, quote stuffing and momentum ignition. In spoofing and layering, traders place large visible orders they do not intend to execute, creating a false impression of supply or demand before cancelling those orders and trading on the induced price move (Lee, 2013; Chullamonthon and Tangamchit, 2023). High-frequency quote stuffing floods venues with rapidly placed and cancelled orders, degrading data-feed quality and impairing price discovery (Dalko, 2019). These manipulations jeopardise market integrity by manipulating prices out of the fundamentals, and undermining trust in the order book fairness.

There are systemic implications when the manipulative or poorly controlled strategies are used on a large scale in highly automated markets. The Flash Crash of 6 May 2010 demonstrated that a single large automated selling programme could interrelate with high-frequency liquidity provision to cause a rapid liquidity breakdown and extreme and short-lived price distortions in equity indices and exchange-traded products (Easley et al., 2011; Kirilenko et al., 2017). The occurrence of similar, smaller-scale, flash crashes in single securities underscores the importance of correlated algorithmic reactions to order-flow shocks to increase intraday volatility and spread stress across venues (Virgilio, 2019). During these episodes, order-flow toxicity increases dramatically, and human bosses find it difficult to decipher and interfere within time frames within which prices are moving (Easley et al., 2011).

Sophisticated A.I. will only intensify these weaknesses. Deep learning and reinforcement learning systems are able to deduce lucrative trends out of sophisticated information without clear, human comprehensible regulations (Bhuiyan et al., 2025; Sadeghi et al., 2025). In cases where such systems are reward-optimising, they might find strategies that use market microstructure, in a manner that would be would be considered to be, akin to, but not necessarily equal to, manipulation, without having been explicitly told to do so. According to simulation-based research, the process of a multi-agent RL system learning tacit collusive or destabilising behaviour by accommodating the behaviour of the other agent, makes it difficult to regulate behaviour by regulators who need to distinguish abusive behaviour with aggressive but legal trading (Khodabandehlou and Golpayegani, 2022; Alfajeer et al., 2025).

Supervisory-wise, the critical issue is that AI can create new types of market abuse as well as disguise the already existing ones. Orchestrated spoofing or momentum ignition are hidden behind the noisy order-flow patterns and thus the traditional surveillance algorithms are not as effective with complex models (Khodabandehlou and Golpayegani, 2022). Meanwhile, a false positive can rise when the detection mechanisms are not able to take into consideration the richer dynamics that AI trading brings. This calls for an upgraded "risk universe" that treats market integrity and systemic stability risks as tightly coupled in AI-intensive markets.

3.2 Model, Data and Technology Risk in AI Trading Pipelines

Within AI trading pipelines, **model risk** arises from misspecification, overfitting and instability when market regimes change. Deep learning architectures can fit complex non-linear patterns but may generalise poorly beyond the training window, particularly when volatility, liquidity or correlation structures shift abruptly (Bhuiyan et al., 2025). Model drift gradual degradation of predictive performance as data distributions evolve can be hard to detect when monitoring focuses solely on headline P&L rather than more granular performance diagnostics. Opaque decision logic further complicates ex ante validation and ex post attribution of losses.

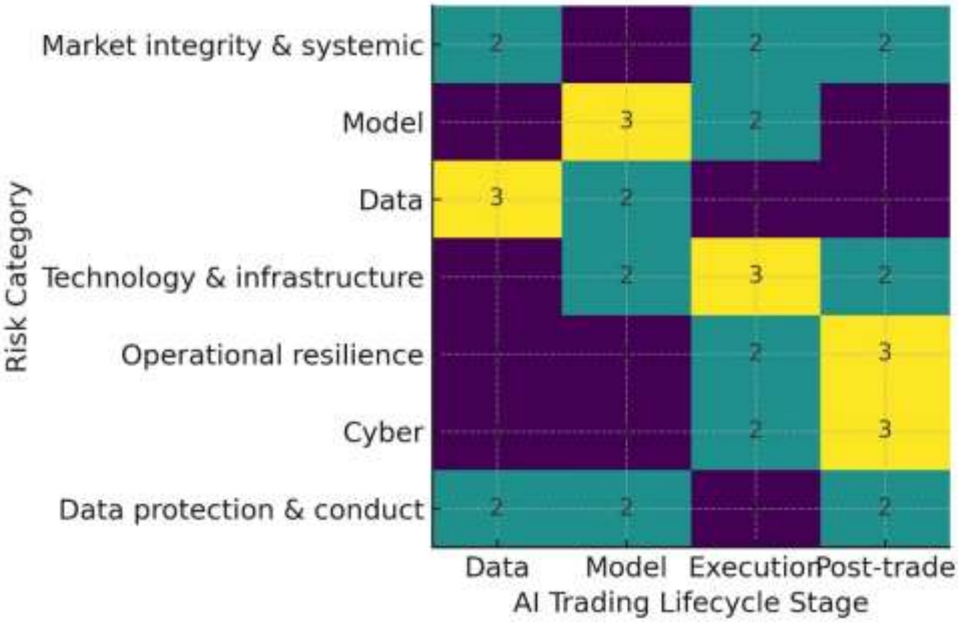

Figure 2. Risk Map for AI Trading Systems along the Trading Lifecycle (data → model → execution → post-trade).
Source: Author's schematic synthesis based on Bhuiyan et al. (2025) and related literature.

Model risk is tightly linked to data risk. AI trading systems increasingly rely on alternative data—news, social media and ESG signals—whose statistical properties differ from traditional market data (Harun and Naafey Sardar, 2021; Çetin et al., 2025). Biases in training data, such as over-representation of calm periods or particular sectors, can lead to systematically misleading forecasts. Latency or outages in data feeds may cause models to operate on stale information, while data poisoning or subtle corruption of input streams could misdirect trades at scale. The research on financial textual sentiment indicates that the incorrectly measured sentiment may have a significant impact on trading decisions and short-term price movements, which is why strong data validation is essential (Harun and Naafey Sardar, 2021; Yang et al., 2025).

Technology risk relates to the stability and dependability of the infrastructure to which AI trading is located. Hardware failures, network congestion, bugs in software and poorly configured gateways are especially vulnerable to latency-sensitive strategies. Empirical studies of algorithmic and high-frequency trading report how access and message-processing delays of a few microseconds can influence the execution and the dynamics of price books (Virgilio, 2019; Rafi et al., 2025). Once several ingredients are interacting in real time - data ingestion, model inference engines, smart order routers - across fragmented venues, even minor integration flaws can lead to run-away order submissions or unwanted self-trading feedback loops. These can lead to both direct losses and exchange-level protection, and spillovers on the other market participants.

3.3 Operational Resilience, Cyber and Data-Protection Issues.

In addition to front-line trading, AI increases the operational resilience issues. The regulatory regimes of key financial hubs now oblige companies to single out material business services, impose impact thresholds and exhibit the capacity to withstand the thresholds in extreme yet realistic disruption conditions (Klumpes and Mann, 2025). In the case of AI trading, it means that critical services, such as data pipelines, model inference, risk controls, and so forth, are tested during outages at cloud providers, market-data vendors or internal systems. Since numerous institutions are growing more dependent on fewer technology and data suppliers, concentration risk will become a local failure into cross-firm shocks to price-making (ESMA, 2025).

The AI risk universe has especially a salient dimension of cyber risk. Financial market infrastructures and trading platforms have been the leading targets of advanced cyber-attacks, such as ransomware, credential theft and data exfiltration (CPMI–IOSCO, 2016; ESMA, 2025). The AI components present new attack surfaces: the adversaries can seek to steal or reverse-engineer proprietary models, poison training data, or manipulate real-time inputs to cause mispricing. The international principles of cyber resilience in financial market infrastructures focus on the importance of advanced detection, segmentation and recovery features, as cyber attacks can easily spread within payment and settlement chains (CPMI–IOSCO, 2016). In the case of AI trading companies, it would imply that cyber controls are part of the model design and implementation, rather than an IT problem.

The use of AI trading also creates complicated data-protection and behaviour concerns. When models are based on granular customer data, such as transaction history and behavioural profile, they can enter an area controlled by privacy laws, e.g. the GDPR. Studies on financial data analytics highlight tensions between the predictive value of fine-grained data and legal requirements around purpose limitation, data minimisation and consent (Baker and Della Valle, 2023). Unclear governance over data lineage and access rights can lead to inadvertent misuse of personal data in trading or risk models, exposing firms to enforcement action, remediation costs and reputational damage. In parallel, opaque AI-driven decisions may raise fairness and discrimination concerns if certain client groups systematically face worse execution outcomes.

To synthesise these threads, Table 2 summarises the main risk categories, typical indicators or triggers, illustrative incidents and potential impacts at investor, firm and market levels. It is based on the conceptual and empirical literature reviewed above and represents the author’s synthesis rather than a direct transcription of any single source.

Table 2. Key Risk Categories, Indicators and Illustrative Incidents in AI Trading.

Risk Category	Typical Indicators/Triggers	Example Incidents	Potential Impact
Market integrity & systemic risk	Cancellation spikes, short-lived price jumps, vanishing depth	Index/ETF mini-flash crash after algo interaction	Distorted prices, volatility shock, loss of confidence
Model risk	Live–backtest drift, unstable parameters, repeated limit hits	AI model failure after regime change	Investor losses, desk P&L hit, forced deleveraging
Data risk	Missing/lagged feeds, noisy alternative data	Sentiment error driving one-sided positioning	Mispriced risk, capital misallocation, model distrust
Technology & infrastructure risk	Latency jumps, outages, high reject ratios	Exchange or broker gateway failure	Inability to trade or hedge, disorderly markets
Operational resilience & third-party risk	Vendor concentration, unresolved incidents	Cloud or data-vendor outage	Service disruption, remediation costs, cross-firm contagion
Cyber risk	Strange logins, odd traffic, malware alerts	Ransomware or credential theft	Rogue trades, direct loss, liquidity strain
Data protection & conduct risk	Weak consent, opaque lineage, poor deletion	Enforcement for unlawful profiling	Privacy harms, fines, tighter use of data

Source: Author’s synthesis based on Bhuiyan et al. (2025); Khodabandehlou and Golpayegani (2022); Alfajeer et al. (2025); Klumpes and Mann (2025); ESMA (2025).

Using the same synthesis data, Figure 3 provides a simple visual summary of the **relative systemic salience** of the different risk categories, on a qualitative scale from 1 (low) to 3 (high). Market integrity and systemic risk, operational resilience and cyber risk score highest, reflecting their potential to trigger widespread disruption, while model, data, technology and data-protection risks are rated as medium but still material.

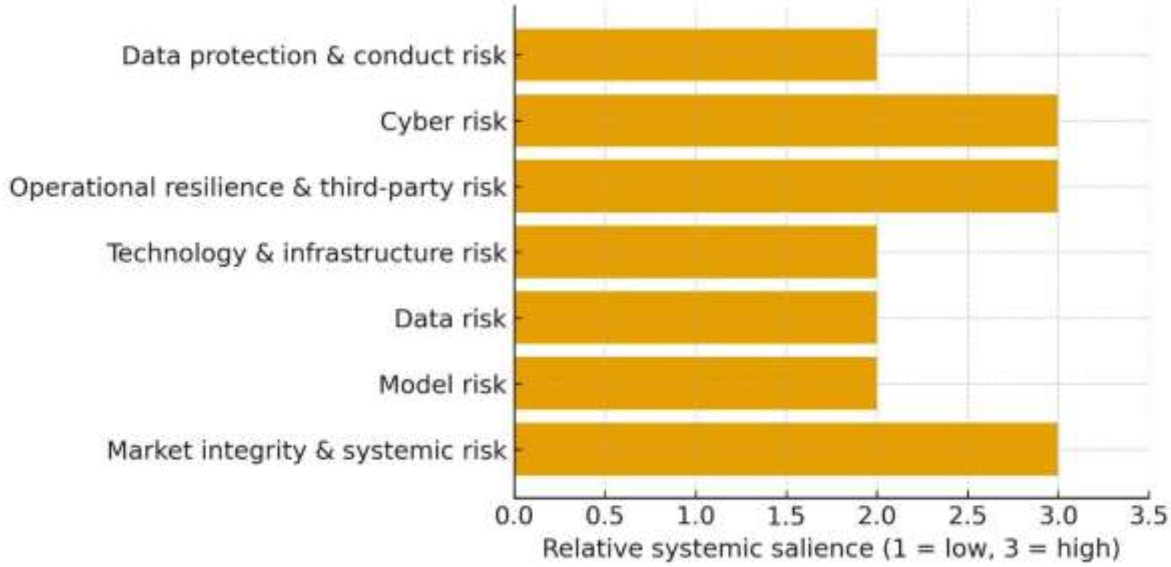

Figure 3. Relative Systemic Salience of Key AI Trading Risks (Author's Synthesis).
Source: Author's synthesis based on literature underlying Table 2.

4. REGULATORY AND LEGAL RESPONSES IN LEADING FINANCIAL CENTRES

4.1 Core Regulatory Themes: Conduct, Transparency and Control

Across major financial centres, the regulation of algorithmic and AI-enabled trading builds on long-standing objectives of fair and orderly markets, investor protection and systemic stability (Crisanto et al., 2024; Avgouleas, 2016). Frameworks aimed at high-frequency and algorithmic trading treat AI as an intensification of existing challenges rather than a wholly new category of activity: regulators worry about disorderly trading, information asymmetries and the possibility that complex systems reduce effective human control (Moloney, 2018).

Conduct rules emphasise that firms must not allow their trading systems to create false or misleading signals or otherwise compromise market integrity. Best-execution and client-protection obligations require investment firms to obtain the most advantageous terms reasonably available while managing conflicts of interest and information leakage (Busch and Rittler, 2014). Record-keeping and post-trade transparency obligations support surveillance and enforcement by ensuring that orders and decisions can be reconstructed, even when initiated or shaped by AI models (Moloney, 2018; Avgouleas, 2016).

A second theme is governance and control obligations specific to automated trading. MiFID II's algorithmic trading regime, US Regulation SCI and comparable standards impose requirements for resilient systems, capacity planning and trading thresholds designed to prevent erroneous orders and disorderly markets (Moloney, 2018; Amtenbrink and Kleinen, 2019). They typically mandate documented development and testing processes, change-management controls, and clear allocation of responsibility for algorithms within the firm (Crisanto et al., 2024).

These provisions increasingly incorporate ideas such as "human in the loop" and ex ante risk controls. Human in the loop does not necessarily imply manual approval of each order, but it requires that responsible personnel can understand, challenge and, if necessary, override AI-driven strategies (Crisanto et al., 2024). Pre-trade controls include credit and position limits, price collars, maximum order sizes and throttles, complemented by kill-switches and venue-level circuit breakers as last-resort safeguards. Real-time monitoring, including automated alerts for unusual patterns in order-to-trade ratios or self-trading, is expected to be coupled with human review, particularly where AI models may encode opaque decision logic.

4.2 Comparative Perspectives: US, EU, UK and Selected Asian Markets

Although jurisdictions share these core objectives, they operationalise them in different ways. In the United States, the Securities and Exchange Commission (SEC) and the Commodity Futures Trading Commission (CFTC) regulate algorithmic trading primarily through general securities and derivatives law, supplemented by technology-focused rules such as Regulation SCI and exchange rulebooks. Scholarship emphasises how this framework relies on a mix of disclosure, exchange-level control and ex post enforcement to manage high-frequency and algorithmic trading risks (Lee, 2016; Chung and Lee, 2016).

In the European Union, MiFID II/MiFIR and the Market Abuse Regulation (MAR) provide a more explicit, granular regime for algorithmic and high-frequency trading. Article 17 MiFID II and related technical standards require firms engaging in algorithmic trading to maintain resilient systems, perform conformance testing, keep an inventory of algorithms and notify regulators (Moloney, 2018; Busch and Rittler, 2014; Amtenbrink and Kleinen, 2019). ESMA guidance further clarifies expectations on kill-switches, stress testing and real-time monitoring. The EU strategy can be described as largely prevention-based, aimed at minimising the likelihood and magnitude of unruly markets by using elaborate organisational regulations (Moloney, 2018).

Since Brexit the United Kingdom has onshored MiFID II ideas, but has indicated that it may recalibrate to achieve more intelligent regulation. Empirical research on the changing nature of the high-frequency traders as market makers demonstrates the impacts of UK and EU design options, including penalties on high order-to-trade ratios, on market-making incentives and algorithmic business models (Boehmer et al., 2023). They are also employed by the UK Financial Conduct Authority (FCA) in the form of thematic reviews and letters to the CEO in which they draw attention to the deficiencies in the regulation of algorithmic trading, stating that senior managers should be held more responsible and that technically informed compliance teams are required.

The major Asian financial centres pursue similar objectives using their instruments. The Monetary Authority of Singapore (MAS) has provided concepts of risk management of electronic trading and the application of AI and other data analytics in financial services in Singapore and emphasises the responsibility and oversight of boards (Crisanto et al., 2024). In its Code of Conduct, the Hong Kong Securities and Futures Commission (SFC) provides a list of requirements of electronic and algorithmic trading that covers due diligence of trading systems, testing, and real-time monitoring. The emerging markets like India are expanding algorithmic trading systems to retail investors and it has been proposed that the retail algorithm should be approved by the exchange and have a unique identifier of each retail algorithm so that it can be traced and audited.

In these jurisdictions, convergence is around:

- **Licensing and authorisation**, where firms providing algorithmic trading services must be appropriately authorised and, in some cases, separately notify regulators of their algorithmic activities;
- **Testing and stress testing**, including pre-deployment testing in production-like environments, conformance tests against venue rules and stress scenarios for abnormal market conditions;
- **Governance and accountability standards**, such as designated responsible officers or senior managers for algorithmic trading and requirements to maintain algorithm inventories; and
- **Reporting and audit-trail expectations**, including detailed order-message records and, increasingly, requirements to document algorithm logic, parameters and change history.

Table 3. Selected Jurisdictional Approaches to Algorithmic/AI Trading.

Jurisdiction	Key Legal Instruments	Governance Requirements	Testing/Control Mechanisms	Enforcement Trends
US (SEC/CFTC)	Exchange Act, Reg SCI, Reg ATS, abuse and best-execution rules	Senior manager accountability; policies for design and change	Pre/post-trade limits, kill switches, resilience checks	Cases on abuse, systems failures, disclosure weaknesses
EU (ESMA)	MiFID II/MiFIR, MAR, algo trading RTS/ITS	Board oversight; algorithm inventory; clear responsibility	Stress tests, circuit breakers, pre-trade thresholds	Focus on disorderly markets, HFT, reporting quality
UK (FCA)	UK MiFID framework, FCA Handbook, MAR	Senior Management Functions; governance for development/deploy	Scenario tests, kill switches, surveillance of conduct	Thematic reviews and public letters on governance gaps
Singapore (MAS)	Securities and Futures Act; MAS e-trading and AI guidelines	Fit-and-proper management; allocation of oversight roles	Pre-deployment tests, throttles, resilience and recovery	Supervisory guidance and inspections on outsourcing, risk
Hong Kong (SFC)	Securities and Futures Ordinance; SFC codes, circulars	Responsible officers; documented controls for trading systems	Production-like testing, real-time monitoring, kill switches	Enforcement around e-trading controls and market abuse

Source: Author's synthesis based on Moloney (2018); Amtenbrink and Kleinen (2019); Lee (2016); Boehmer et al. (2023); Crisanto et al. (2024).

4.3 Interaction with Cross-Cutting AI, Data Protection and Workplace Laws

Algorithmic trading rules do not operate in isolation. They increasingly interact with AI-specific regulatory initiatives and broader data-governance frameworks. In the EU, the emerging AI Act will classify certain financial-sector AI systems as "high-risk," subjecting them to requirements on risk management, data governance, human oversight and transparency (Hacker et al., 2023; Crisanto et al., 2024). Scholars emphasise that many risks raised by AI in trading opacity, bias, lack of contestability can be addressed by strengthening existing governance and conduct rules rather than creating entirely separate regimes (Avgouleas, 2016; Alexander, 2020).

Data-protection law, particularly the GDPR, also shapes AI trading. Articles on algorithmic decision-making stress that automated profiling using personal data must satisfy requirements of lawfulness, transparency and the possibility of meaningful human review (Brkan, 2019). This raises questions where trading systems incorporate granular client-order histories or behavioural data, or where firms reuse data collected for other purposes to train models. Recent work highlights how the GDPR, Digital Services Act and AI Act together form a multi-pillar framework for algorithmic security, privacy and accountability in the EU, with implications for financial institutions deploying AI at scale (van Dijk et al., 2025; Hacker et al., 2023).

Finally, AI trading intersects with labour and workplace law. Algorithmic systems are increasingly used to monitor staff performance and code quality, allocate tasks and even support disciplinary decisions. Literature on algorithmic management warns that such uses may undermine employee autonomy and privacy if not carefully governed (De Stefano and Taes, 2023). In financial firms, the combination of trading surveillance and AI-enabled employee monitoring heightens psychosocial risks and may require consultation mechanisms, impact assessments and rights to explanation or contestation. As centres experiment with "Senior Managers" regimes and individual accountability, firms must harmonise obligations under financial regulation with duties arising from employment and data-protection law, ensuring that responsibility for AI trading does not translate into unreasonable personal exposure for individual staff.

Taken together, these cross-cutting regimes mean that boards and senior managers cannot treat AI trading as a narrow technical issue. Instead, they must navigate an integrated landscape in which market, prudential, data-protection and labour laws jointly determine what "responsible" AI trading looks like. The risk-based internal governance framework developed in the next section is designed precisely to help institutions manage this intersection.

5. DESIGNING A RISK-BASED INTERNAL GOVERNANCE FRAMEWORK

5.1 Guiding Principles: Human-in-the-Loop, Explainability and Accountability

A risk-based governance framework for AI trading must convert regulatory aims into decisions about how models are built, deployed and retired. Reviews of machine-learning use in banking risk management emphasise that governance, not algorithms alone, determines whether AI strengthens resilience (Heß and Damásio, 2025; Leo et al., 2019). The first principle is robust human oversight. "Human-in-the-loop" means that accountable staff can understand model purpose, see diagnostics and intervene quickly through throttles, parameter changes or kill-switches when behaviour looks abnormal (Kahya et al., 2025).

The second principle is explainability proportionate to risk. Systematic reviews on financial XAI show that complex models are acceptable only when firms provide decision-relevant explanations to risk managers, auditors and supervisors (Yeo et al., 2025; Tülü and Tüzün, 2024; Pimentel and Pisoni, 2025). For low-impact tools, high-level performance metrics may suffice. For strategies with material market or client impact, explanations must show why signals were generated and how behaviour changes under stress, enabling meaningful challenge and regulatory dialogue.

The third principle is clear accountability and sign-off across the model lifecycle. Evidence from operational-risk governance shows that unclear ownership of controls is a recurring feature of major incidents (Mabwe et al., 2017). For each AI trading model, firms should specify an accountable owner, an independent validation authority and escalation paths to senior management and the board, reflecting the trend to treat model risk as a distinct risk type (Heß and Damásio, 2025).

Finally, the framework must be proportionate to each strategy's risk profile. Work on ML in banking risk management stresses that governance intensity should scale with leverage, complexity and market-disruption potential (Heß and Damásio, 2025; Leo et al., 2019). Mapping strategies to the risk categories in Section 3—market integrity, model and data risk, technology and operational resilience, cyber and data-protection risk—allows firms to prioritise deeper explainability, tighter limits and more frequent review where externalities are greatest.

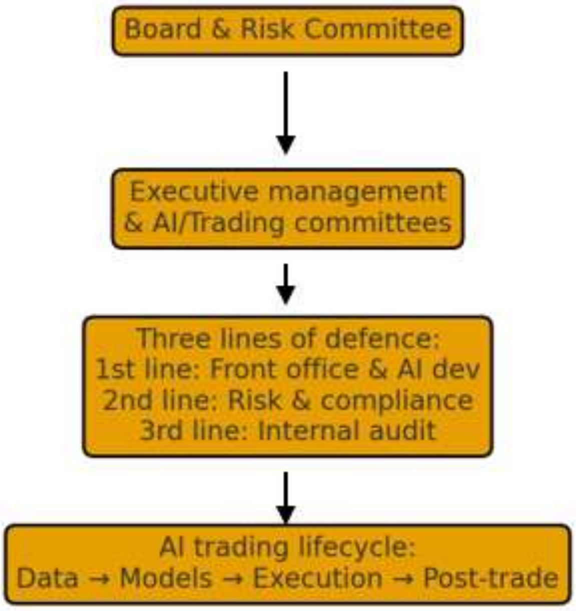

Figure 3. Proposed Risk-Based Governance Architecture for AI Trading Systems.
Source: Author's schematic synthesis based on Heß and Damásio (2025); Yeo et al. (2025); Mabwe et al. (2017).

5.2 Role Clarity Across Front Office, Risk, Compliance and Technology

To operationalise these principles, firms need sharp role definition across the "three lines of defence". Empirical work on UK financial institutions finds that ambiguity over what belongs to each line undermines operational-risk governance (Mabwe et al., 2017). In an AI trading context, the first line comprises traders, quantitative researchers and AI engineers. They own strategy design, data selection, model development and day-to-day performance management, including documenting assumptions, implementing limits and ensuring that code changes follow controlled lifecycles with testing.

The second line—risk management and compliance provide independent challenge and policy setting. Reviews of AI and ML in risk management stress that second-line teams must have enough technical literacy to interrogate models, not simply tick checklists (Heß and Damásio, 2025; Kahya et al., 2025). Market and operational-risk functions should define standards for validation, stress testing and data governance, approve higher-impact strategies and monitor risk concentrations across models. Compliance should focus on conduct and market-abuse risk, review evidence for abusive patterns and ensure that algorithm inventories, parameter changes and overrides are documented to a standard that meets supervisory expectations.

The third line, internal audit, offers independent assurance on both design and operation of the framework. For AI trading, audit programmes should periodically sample high-risk models and trace them from business case through validation, deployment, monitoring and retirement, checking that escalation and remediation occur when thresholds are breached (Mabwe et al., 2017).

Practical clarity is crucial around three cross-cutting processes: model approval, change management and incident response. Material new models should require joint sign-off from business sponsors, model-risk or validation teams and, where relevant, compliance. Change-management standards should distinguish material changes—new data sets, architectures or objective functions from minor parameter tuning, with the former triggering re-validation. Incident-response playbooks should specify who can pause or disable a model, how quickly, how clients and venues are informed, and how lessons are embedded in subsequent policies and training.

5.3 Tools, Processes and Controls: From Kill-Switches to Audit Trails

The framework must rest on a coherent set of tools and processes. At the sharp end are operational controls. Pre-trade limits on price, quantity and exposure, together with message-rate throttles, are standard features of algorithmic trading control frameworks and remain central in AI settings (Kahya et al., 2025). Venue-level circuit breakers and firm-level kill-switches provide backstops against runaway behaviour, but they only work if calibrated, tested and accessible to responsible staff under stress conditions.

Testing and monitoring form the second pillar. Literature on ML in banking risk management highlights the need for robust back-testing across regimes, forward-testing in sandbox or paper-trading environments, and scenario analysis under extreme but plausible conditions (Heß and Damásio, 2025; Leo et al., 2019). For AI trading, test suites should cover the risk channels mapped in Section 3: data outages or corrupt feeds, latency spikes, cyber incidents, and stressed liquidity and volatility. Continuous monitoring should track data quality, model-confidence diagnostics, limit utilisation, error and reject rates, and activations of throttles or kill-switches. Work on XAI in finance shows that pairing such monitoring with explanation tools helps detect drift and emerging biases before losses escalate (Yeo et al., 2025; Pimentel and Pisoni, 2025).

Finally, documentation and audit trails provide the spine of accountability. Model-risk literature emphasises centralised model inventories that record purpose, owners, validation status and known limitations (Heß and Damásio, 2025). For AI trading, this should be complemented by version-controlled code repositories, configuration histories, test results and logs of manual overrides or exceptional interventions. These artefacts allow internal audit and supervisors to reconstruct behaviour around stress events and to judge whether the three-lines structure and control environment functioned as intended.

Table 4. Mapping Risk Categories to Governance Controls and Metrics.

Risk Type	Key Controls	Metrics / KPIs	Responsible Function(s)
Market integrity & systemic	Pre-/post-trade limits, circuit breakers, kill-switches	Order-to-trade ratios, cancelled-order spikes	Front office, risk, compliance, venue
Model	Model inventory, independent validation, use restrictions	Back test vs live drift, validation exceptions	Quants, model risk, risk committee
Data	Data quality checks, vendor SLAs, redundancy	Missing/late data rates, feed error counts	Data engineering, vendors, risk
Technology & infrastructure	Capacity planning, failover, change controls	Latency bands, outage minutes, failed deployments	IT operations, architecture, first line
Operational resilience & third-party	Impact-tolerance setting, BCP/DR tests, vendor oversight	Recovery times, vendor incidents, test pass rates	Business owners, resilience function, risk
Cyber	Network segmentation, detection, response playbooks	Incident frequency, dwell time, patch coverage	Cybersecurity, IT, risk, CISO
Data protection & conduct	Data minimisation, access controls, DPIAs	Breach counts, access-request cases, retention breaches	Legal, privacy, compliance, data owners

Source: Author's synthesis based on Heß and Damásio (2025); Yeo et al. (2025); Mabwe et al. (2017).

6. HUMAN SUPERVISORS, ETHICS AND POLICY DIRECTIONS

6.1 Human Supervisors and Psychosocial Hazards in AI-Intensive Trading Floors

AI-intensive trading floors change what it means to supervise markets and desks. Human supervisors can no longer rely on intuition about familiar strategies; they must understand systems that learn, adapt and interact across venues. Three broad competency clusters are emerging. First, data literacy: supervisors need to interpret dashboards and metrics on data quality, model confidence, limit utilisation and control activations. Second, model interpretability skills: they should be comfortable with concepts such as training data, features, overfitting, drift and scenario analysis, and know what questions to ask when explanations appear weak. Third, systems thinking: supervisors need to understand the fact that strategies, venues, vendors and controls constitute an ecosystem where minor changes can create non-linear results.

These requirements pose psychosocial threats. The constant real-time surveillance, as well as the fear of neglecting the appearance of the first symptoms, may cause chronic stress and hyper-vigilance fatigue. Supervisors can feel that they are responsible to opaque algorithms that they did not write, and that they are anxious of being personally liable. Robots and predictive analytics have the potential to instil the fear of displacement: employees fear that machines might take their jobs as they will be held responsible in case of errors. Another risk is cognitive overload. Dashboards usually have numerous pointers and notifications and little assistance with triage; in swift markets this may drive supervisors to paralysis or routinised clicking that weakens control.

To combat these risks, there is need to invest in wellbeing, training and support. Training programmes must incorporate both technical skills and exercise based on scenarios, which practise incident response. Front office/ risk and technology rotations may enhance mutual understanding, as well as alleviate the supervisory role isolation. The feeling of solitude can be eliminated with clear escalation procedures and access to the support of specialists, including model-risk experts and cyber teams. Wellbeing metrics such as realistic staffing, screen time and psychological support, understand that effective supervision is not a control box; it is a human performance activity.

6.2 Ethical Guidelines for Trustworthy AI Trading

Reliable AI trading needs more than minimum legal compliance of ethical guidelines. The important dimensions are fairness and non-discrimination, transparency and explainability, harmful manipulation avoidance, and inclusiveness. Fairness means that AI-based execution plans must not be designed in a systematic way to disfavour specific groups of clients, by directing smaller clients to less well-off venues or exploitative exploitation of behavioural information. Transparency and explainability are associated with internal stakeholders who should be aware of how models work and customers who should be provided with valuable information on how automation will be used in their orders.

The prevention of dangerous manipulation is the core of the markets in which AI has the power to identify and use minor patterns. The screening of strategies must be conducted so as to avoid strategies which depend on getting other participants to make foreseeable errors, causing unwarranted volatility spikes or creating synthetic liquidity premises. The element of inclusiveness refers to the fact that the definition of what harm and benefit are in specific situations should take into account a variety of views, including the different disciplines and levels of seniority.

Ethics should also be incorporated in design. Selection of data must avoid those that are acquired unlawfully, those sources that are too intrusive or prone to instil discriminatory trends. Disparate impact tests across various segments of clients should be part of model validation practises and not only aggregate performance. Communication with clients must not confuse human guidance with machine implementation and must not exaggerate predictability in unpredictable markets. These promises will be formalised by many firms by codes of ethics, professional standards regarding quants and engineers and internal AI ethics committees who will review high impact projects and offer an escalation route when employees have concerns.

6.3 Policy Recommendations, Implementation Roadmap and Future Research

This analysis leads to policy directions. To regulators, there are harmonised standards of algorithm and AI regulation, cross-border supervisory forums, and the application of sandboxes or controlled live experiments to experiment with safeguards. The stress-testing regimes may not be limited to the balance sheets and may include the simulation of the interaction of AI strategies when stressed. To companies, it has been suggested to have a board-approved AI risk policy, form interdisciplinary forums of governance, invest in supervisor education, and review ethics as a product-development gate.

Quick wins should be separated by an implementation roadmap of structural reforms. Short term, firms can map existing AI use, formalise model inventories and tighten documentation of overrides. Medium-term steps include enhancing monitoring infrastructure, deploying explainability tools, and revising evaluation systems so that staff are rewarded for prudent challenge, not just P&L. Longer-term reforms may involve modernising legacy technology, strengthening standards on incident reporting and data sharing, and clarifying liability regimes for AI-related failures.

Several research gaps remain. The behaviour of interacting AI agents in realistic markets is not well understood, especially when models adapt to one another's strategies. Cross-border supervisory coordination faces questions about data access, mandates and crisis management. Measurement of AI-induced systemic risk will require indicators that capture concentration in algorithms, dependencies on shared data and infrastructure, and common model failures under stress. Dialogue between practitioners, regulators and researchers will be essential to ensure that AI trading supports innovation and financial stability.

Conflict of interest

The author declares no conflict of interest and confirms that no personal, financial, professional or institutional relationships influenced the conception, analysis or interpretation of findings.

Funding

This research received no specific grant from funding agencies, commercial entities or not for profit organisations, and all research expenses were supported by the author.

REFERENCES

1. Alcázar-Blanco, A., et al. (2024). Incorporating artificial intelligence into finance: A bibliometric approach. *Journal of Risk and Financial Management, 17*(12), 556. https://doi.org/10.3390/jrfm17120556
2. Bahoo, S., Cucculelli, M., Goga, X., & Mondolo, J. (2024). Artificial intelligence in finance: A comprehensive review through bibliometric and content analysis. *SN Business & Economics, 4*(1), 23. https://doi.org/10.1007/s43546-023-00618-x
3. Dakalbab, F., Abu Talib, M., Nasir, Q., & Saroufil, T. (2024). Artificial intelligence techniques in financial trading: A systematic literature review. *Journal of King Saud University – Computer and Information Sciences, 36*(3), 102015. https://doi.org/10.1016/j.jksuci.2024.102015

4. Easley, D., López de Prado, M., & O'Hara, M. (2011). The microstructure of the "Flash Crash": Flow toxicity, liquidity crashes, and the probability of informed trading. *Journal of Portfolio Management, 37*(2), 118–128. https://doi.org/10.3905/jpm.2011.37.2.118
5. Gomber, P., & Haferkorn, M. (2013). High-frequency-trading: Technologies and their implications for electronic securities trading. *Business & Information Systems Engineering, 5*(2), 97–99. https://doi.org/10.1007/s12599-013-0255-7
6. Jarrow, R., & Protter, P. (2012). A dysfunctional role of high frequency trading in electronic markets. *International Journal of Theoretical and Applied Finance, 15*(3), 1250022. https://doi.org/10.1142/S021902491250022
7. Johnson, N., Zhao, G., Hunsader, E., Meng, J., Ravindar, A., Carran, S., & Tivnan, B. (2013). Abrupt rise of new machine ecology beyond human response time. *Scientific Reports, 3*, 2627. https://doi.org/10.1038/srep02627
8. Kirilenko, A., & Lo, A. W. (2013). Moore's Law versus Murphy's Law: Algorithmic trading and its discontents. *Journal of Economic Perspectives, 27*(2), 51–72. https://doi.org/10.1257/jep.27.2.51
9. Kirilenko, A. A., Kyle, A. S., Samadi, M., & Tuzun, T. (2017). The Flash Crash: High-frequency trading in an electronic market. *Journal of Finance, 72*(3), 967–998. https://doi.org/10.1111/jofi.12498
10. Virgilio, G. P. M. (2019). High-frequency trading: A literature review. *Financial Markets and Portfolio Management, 33*(2), 183–208. https://doi.org/10.1007/s11408-019-00331-6
11. Aldridge, I. (2013). *High-frequency trading: A practical guide to algorithmic strategies and trading systems* (2nd ed.). Wiley.
12. Anuar, A. A., Ismail, N. A., & Ahmad, S. (2025). Artificial intelligence techniques in investment trading: A systematic literature review. *Asia-Pacific Journal of Information Technology and Multimedia, 14*(1), 1–30. https://doi.org/10.17576/apjitm-2025-1401-02
13. Zhang, Z., Zohren, S., & Roberts, S. (2021). Deep reinforcement learning for trading. *Expert Systems with Applications, 177*, 114802. https://doi.org/10.1016/j.eswa.2021.114802
14. Alfajeer, A., Altaweel, A., Bouridane, A., Mouheb, D., & Aslam, S. (2025). On detecting stock price manipulation attacks: A comprehensive systematic literature review. *Multimedia Tools and Applications*. Advance online publication. https://doi.org/10.1007/s11042-025-20641-4
15. Bhuiyan, M. S. M., Mir, M. N. H., Rafi, M. A., Ishraq, A., Rodrigues, G. N., Mridha, M. F., & Shin, J. (2025). Deep learning for algorithmic trading: A systematic review of predictive models and optimization strategies. *Array, 26*, 100390. https://doi.org/10.1016/j.array.2025.100390
16. Chullamonthon, P., & Tangamchit, P. (2023). Ensemble of supervised and unsupervised deep neural networks for stock price manipulation detection. *Expert Systems with Applications, 228*, 119698. https://doi.org/10.1016/j.eswa.2023.119698
17. Dalko, V. (2019). High-frequency trading: Order-based innovation or manipulation? *Journal of Banking Regulation, 21*(4), 289–304. https://doi.org/10.1057/s41261-019-00115-y
18. Harun, M., & Naafey Sardar, M. (2021). Social media sentiment and the stock market. *Journal of Financial Services Marketing, 26*(4), 185–197. https://doi.org/10.1057/s41264-021-00112-3
19. Khodabandehlou, S., & Hashemi Golpayegani, S. A. (2022). Market manipulation detection: A systematic literature review. *Expert Systems with Applications, 210*, 118330. https://doi.org/10.1016/j.eswa.2022.118330
20. Klumpes, P. J. M., & Mann, S. (2025). Operational resilience in the UK financial sector: Practical guidance. *British Actuarial Journal*. Advance online publication. https://doi.org/10.1017/S1357321725100202
21. Alexander, K. (2020). Regulating AI in financial markets: Risk-based approaches. *Law and Financial Markets Review, 14*(3), 167–180. https://doi.org/10.1080/17521440.2020.1811123
22. Amtenbrink, F., & Kleinen, J. (2019). The European legal framework for algorithmic and high-frequency trading: MiFID II and MAR. *European Journal of Risk Regulation, 10*(3), 501–525. https://doi.org/10.1017/err.2019.27
23. Avgouleas, E. (2016). Systemic risk, complex systems and macroprudential regulation. *European Journal of Risk Regulation, 7*(2), 227–246. https://doi.org/10.1017/S1867299X00005839
24. Boehmer, E., Li, D., & Saar, G. (2023). High-frequency traders' evolving role as market makers. *Journal of Financial Markets, 63*, 100739. https://doi.org/10.1016/j.finmar.2023.100739
25. Brkan, M. (2019). Do algorithms rule the world? Algorithmic decision-making and data protection in the framework of the GDPR and beyond. *International Journal of Law and Information Technology, 27*(2), 91–121. https://doi.org/10.1093/ijlit/eay017
26. Busch, D., & Rittler, T. (2014). Best execution and high-frequency trading. *Capital Markets Law*

Journal, 9(3), 203–232. https://doi.org/10.1093/cmlj/kmu017

27. Crisanto, J. C., Leuterio, C. B., Prenio, J., & Yong, J. (2024). Regulating AI in the financial sector: Recent developments and main challenges. *Journal of Banking Regulation, 25*(4), 321–340. https://doi.org/10.1057/s41261-024-00200-3
28. De Stefano, V., & Taes, S. (2023). Algorithmic management and workers' rights. *International Labour Review, 162*(3), 345–370. https://doi.org/10.1111/ilr.12389
29. Hacker, P., Cordes, J., & Rochon, J. (2023). Regulating gatekeeper artificial intelligence and data: Transparency, access and fairness under the DMA, the GDPR and beyond. *European Journal of Risk Regulation, 15*(1), 1–24. https://doi.org/10.1017/err.2023.1
30. Lee, C. (2016). High-frequency trading: Review of the literature and regulatory initiatives around the world. *Journal of Applied Finance & Banking, 6*(2), 1–30. https://doi.org/10.2139/ssrn.2725521
31. Moloney, N. (2018). EU securities and financial markets regulation (3rd ed.). *Journal of Financial Regulation, 4*(2), 312–318. https://doi.org/10.1093/jfr/fjy010
32. Heß, V. L., & Damásio, B. (2025). Machine learning in banking risk management: Mapping a decade of evolution. *International Journal of Information Management Data Insights, 5*(1), 100324. https://doi.org/10.1016/j.jjimei.2025.100324
33. Leo, M., Sharma, S., & Maddulety, K. (2019). Machine learning in banking risk management: A literature review. *Risks, 7*(1), 29. https://doi.org/10.3390/risks7010029
34. Mabwe, K., Ring, P. J., & Webb, R. (2017). Operational risk and the three lines of defence in UK financial institutions: Is three really the magic number? *Journal of Operational Risk, 12*(1), 53–69. https://doi.org/10.21314/JOP.2017.187
35. Yeo, W. J., et al. (2025). A comprehensive review on financial explainable AI. *Artificial Intelligence Review*. Advance online publication. https://doi.org/10.1007/s10462-024-11077-7
36. Tülü, Ç., & Tüzün, E. (2024). Explainable artificial intelligence (XAI) in finance: A systematic literature review. *Artificial Intelligence Review*. Advance online publication. https://doi.org/10.1007/s10462-024-10854-8
37. Pimentel, N., & Pisoni, G. (2025). Rethinking explainable AI in financial services. *AI & Society*. Advance online publication. https://doi.org/10.1007/s00146-025-02315-9
38. Kahya, A. E., et al. (2025). Artificial intelligence and machine learning in risk management—Current challenges and ethical considerations in the financial sector. In *Artificial intelligence and cybersecurity in finance* (Lecture Notes in Electrical Engineering and Computer Technology). Advance online publication. https://doi.org/10.1007/978-3-031-89718-4_20