AI as a Systemic Risk Amplifier in High-Frequency Trading: Presenting a Conceptual Regulatory Framework

Santhosh Kumar Shrinivas[1*], Pavithra Shetty[1], Juliet Sophia[1], Mamatha Gnanakumar[1]

[1]*Gulf College, Oman*

Abstract. Artificial Intelligence (AI) has changed how the financial market operates, particularly in High-Frequency Trading (HFT) and autonomous execution. While AI enhances liquidity, speed, and price discovery, it also causes new systemic vulnerabilities. This paper investigates how AI operates as a risk amplifier in financial markets and suggests a novel conceptual framework that integrates transmission risks with governance principles. The paper connects AI-driven mechanisms with systemic outcomes. Drawing on literature, institutional reports, and historical incidents such as the Flash Crash on 6th May 2010, the study underscores important transmission risks such as feedback loops, model opacity, flash crash, algorithmic herding, and regulatory gaps. To mitigate these risks and to reap the benefits of technology, the study proposes a framework that prioritises explainability, accountability, and transparency, model diversity, layered monitoring, stress testing, flexible macroprudential supervision, and international harmonisation. By aligning with previous studies and regulatory warnings, this research contributes to academic discourse and provides practical insights for policymakers aiming to strike a balance between AI-driven efficiency and systemic stability.

Keywords: Artificial intelligence (AI), High frequency trading (HFT), Systemic risks, Transmission risk channels, Regulatory framework

*Corresponding Author:
Santhosh Kumar Shrinivas
santhosh.kumar@gulfcollege.edu.om

1. Introduction

What is the mantra for success in highly volatile financial markes? Part of the answer would be adopting Artificial Intelligence (AI) for analysis and decision-making. AI plays a key role in the success of every business in the modern world. Along with the advantageous outcomes such as increased productivity, better decision-making, and tailored customer experience, it also carries demerits such as job displacements, moral quandaries, and security risks (Shrinivas & Shetty, 2024). AI software can perform almost all the actions that a human being can do, starting from basic mechanical tasks to the most complicated operations, such as investment advice.

AI allows algorithms to develop strategies, missing the benefit of human intuition by processing a large, diverse set of data autonomously (Vincent, 2021). As a result, multiple AI companies deploying similar AI systems may create algorithmic herding behavior, amplifying volatility through synchronized trading behaviour (Serrano, 2020). This connection is supported by strong empirical evidence. A structured review of European data finds that intraday volatility increase of 05 – 0.8 standard deviation (4-6 percentage points anualised) correlates with a standard deviation increase in High Frequency Trading (HFT), especially from purely HFT firms (Serrano, 2020). Agent-based simulations further demonstrate how algorithmic interaction during periods of market stress replicates the mechanics of flash crashes, with market fluctuations matching the real-world scenario like May 6, 2010 (Gao et al., 2024). Feedback loops and liquidity evaporation across the market may be created due to AI strategies, simultaneous triggers of de-risking, or circuit break actions, which reduce liquidity abruptly (International Monetary Fund, 2024). This kind of mechanism will intensify the systemic vulnerability from a localised algorithmic failure to an interdependent market collapse. The regulatory bodies are sounding alarms about these risks. The Bank of England's Financial Policy Committee (FPC) warns that the widespread use of uniform AI models may heighten the correlation among firms and intensify shocks, undermining resilience (Bank of England, 2025). AI-driven trading may culminate in weak shared datasets, increasing systemic risks while posing oversight challenges (Svetlova, 2022). The analysis done by Daníelsson et al. (2022) underscores the misalignment between localised efficiency gains and macroprudential stability. They argue that the pursuit of optimization of AI could lead to results that are less beneficial to society and may lead to possible concealed systemic crises.

Despite the considerable focus on the topic, research still lacks an integrated conceptual framework that systematically connects AI implementation in high-frequency trading or autonomous trading with systematic risk outcomes through defined transmission channels. Most existing research tends to concentrate on isolated incidents or narrow empirical indicators, often neglecting risk governance design, the opacity of AI structures, or policy recommendations guided by systems theory. This paper addresses the gap by presenting a conceptual framework that highlights key risk conduits, such as algorithmic herding, feedback loops, the propagation of flash crashes, model overfitting, and inadequate human oversight. It also illustrates how these conduits lead to systemic consequences, including market volatility, liquidity crises, cross-market contagion, and regulatory blind spots. The study aims to investigate the relationship between AI model opacity and human/regulatory visibility, underscoring the limitations of governance and outlining principles for a risk-based AI governance framework, including layered monitoring, stress testing, encouraging AI model diversity, and requiring transparency. In addition to providing useful information to regulators such as central banks, security authorities, and macroprudential bodies, the aim is to advance the scholarly discourse on how AI-powered automated trading systems may increase systemic risk.

2. Theoretical Framework

The intersection of AI, HFT, and systemic risk has drawn considerable scholarly attention. A growing section of research assesses how autonomous trading systems amplify financial fragility through algorithmic uniformity, feedback loops, and model opacity. A systematic review of literature highlights the rapid increase of AI methods, including deep learning and reinforcement learning in financial trading, notably in HFT (Dakalbab et al., 2024). It also reveals how trading methods are fully automated, leveraging AI's ability to process large sets of data. Although AI improves liquidity and prediction accuracy, it also raises concerns over an excessive reliance on historical patterns as well as the model's vulnerability to overfitting and failure. Researchers emphasise algorithmic herding, which happens when several organisations deploy similar AI models based on overlapping data, resulting in coordinated trade that magnifies market volatility (Bank of England, 2025b; Ogbuonyalu et al., 2024). This integration reduces market diversity and may result in flash crashes or liquidity spirals. AI-powered black-box trading algorithms become enraged and all end up selling the same thing at the same time, causing a crash in the market (International Monetary Fund, 2024). It's a known fact that high-frequency trading itself has long been associated with a significant amount of intraday volatility and liquidity instability, especially during market stress. Now, with the introduction of autonomous trading and execution of trade at the same time and the same pattern using AI models has amplified these circumstances. On May 6, 2010, the U.S. financial market experienced an intraday systemic event, a

Flash Crash, due to automated execution of a large selling program in E-mini S&P 500 stock index futures (Kirilenko et al., 2017).

How an AI system can be manipulated or destabilised is illustrated by studies on adversarial vulnerabilities. Research from Nehemya et al. (2021) shows how the erroneous model behaviour can be provoked by an adversarial perturbation to input data streams, opening new channels for systemic risk. Similar to this, a study from Goldblum et al. (2022) highlights the vulnerability of these models in a fragile environment and the prospects of cascading trading errors in their discussion of realistic assaults against machine learning systems in HFT. Regulatory bodies have been giving warnings of AI-related systemic risk. Security and Exchange Commission (SEC) (2023) comments that AI trading shares a weakness of common data sets, which could lead to the convergence of AI systems that result in risky strategies and magnify systemic fragility. Likewise, the FPC of the Bank of England points out the vulnerabilities, such as collusion and correlated model errors, to illustrate the risk that autonomous AI actors might unintentionally trigger market crises and profit from them (Bank of England, 2025a). In their systematic literature review of AI governance, Batool et al. (2025) stress the importance of aligning policy mechanisms, accountability, and transparency across model life cycles. However, they note that there are few frameworks specifically addressing the risk in trading contexts. By contrasting international regulatory approaches to AI in finance, such as the European Union Artificial Intelligence (EU AI) Act, U.S. SEC policy, and Financial Conduct Authority (FCA) principles, Mirishli (2024a) deepens this gap and highlights the necessity of flexible risk-based governance capable of connecting innovation with systemic safety.

While the current literature discusses the discrete aspects like execution efficiency, liquidity volatility, adversarial vulnerabilities, and regulatory intent, it lacks a comprehensive conceptual model that connects AI-driven HFT systems to systemic risk outcomes via a well-defined transmission mechanism. The majority of studies are descriptive or empirical and do not incorporate conceptual risk architecture or governance analysis. Theoretical concerns related to the macro-level effects of model homogenisation and optimisation misalignment are expressed in scholarly work done by Daníelsson et al. (2022) and Ozili (2024), but they stop short of operationalising these mechanisms for governance analysis. The systematic review by Lakhchini et al. (2022) and Ahmed et al. (2022) emphasize the technical application and the success of predictive modelling, but emergent fragility and risk aggregation are rarely discussed. Thus, the need for a conceptual risk transmission framework that bridges the fields of AI model behaviour, trading dynamics, systemic outcomes, and regulatory responses has been identified by this theoretical framework, which closes an obvious gap in the discourse that is relevant to both academia and policy.

3. Methodology

This study adopts a qualitative, conceptual methodology based on the literature review from peer-reviewed journals, regulatory reports, and institutional publications. For the literature review, the papers were selected from the reputed databases such as Scopus, Web of Science, and Science Direct with keywords including "AI in finance", "high frequency trading", "systemic risks", "risk amplifier", flash crash", "algorithmic herding", 'financial instability", and "autonomous trading". The papers were considered only if they addressed AI-enabled or algorithmic trading mechanisms, systemic vulnerabilities, transmission risks, governance, or regulatory responses. Only English language sources were selected. Studies focusing only on model architecture without systemic implications were excluded.

To identify systemic risk channels, evidence from recent studies was integrated with historical incidents such as the 2007 Quant Meltdown and 2010 Flash Crash. By combining these observations, the study proposes a conceptual framework that connects systemic vulnerabilities and AI-driven HFT mechanisms. The study also provides governance practices to minimise the risk.

## 4. Proposed Risk-Transmission Channels in AI-Augmented Trading

Systemic risk can be amplified through a new transmission mechanism brought about by the growing use of AI in financial markets, especially in HFT and autonomous trading systems. The main conceptual pathways through which AI can worsen weakness in market stability and structure are identified and

explained in this part (see Figure 1). The channels can interact to produce intricate feedback loops that exacerbate systemic volatility; they are not always independent. Similar data sets and goals, such as maximising profits and minimising volatility, are usually used to train the AI trading systems, especially those based on machine learning and deep learning. Even if it is not an intended outcome, algorithmic convergence can result from the deployment of AI agents by multiple market participants that are designed on overlapping strategies, producing herding behaviour (Ozili, 2024). In times of market stress, this type of autonomous herding can cause systemic risk. For example, several AI systems may execute large volumes of trade in the same direction if they simultaneously detect a market shift or anomaly, which could worsen price swings and possibly lead to flash crashes (Gayduk & Nadtochiy, 2017). As opposed to conventional herding, AI-driven convergence can happen almost swiftly, providing market players and regulators little opportunity to respond.

The feedback loops are created due to relying on real-time data to update AI models and make trading decisions. When AI algorithms influence market price through autonomous trading, it becomes input for future decisions. The market can become uncertain due to recursive feedback mechanisms, particularly when volatility is high. Agent-based simulations illustrate how recursive learning makes disproportionately huge market movements by amplifying reactions to the small shocks (Gao et al., 2024). In an inactive market, where even a small AI-driven transaction may have a significant impact on prices, this is especially problematic. Moreover, contagion across the market could result from these feedback effects spreading to different asset classes and geographical regions (Serrano, 2020). AI-powered high-frequency trading has previously been linked to a number of market anomalies, including flash crashes. The 2010 Flash Crash showed how algorithmic trading may quickly drain liquidity and lead to market volatility. As AI has become more common in the present market, the speed and execution risks have also grown (Leal et al., 2015). AI trading systems may unintentionally create instability in the market if their risk models fail to adequately account for uncommon or outlier events. Moreover, many AI trading systems operate as "Black Boxes", meaning it is difficult to decipher the decision logic used by them, which in turn results in unpredictable behaviour when the market is under stress (International Monetary Fund, 2024; Khan et al., 2025). Model homogeneity is another risk transmission channel. Financial organisations use similar AI models or frameworks for operations such as risk assessment and trade execution due to industry norms and regulatory compliance. This results in less diverse decision-making and raises the chances of systemic failure during times of stress. When dealing with external factors such as cyberattacks or geopolitical events, monoculture AI models may result in synchronised trading decisions and systemic shocks. The market might show a non-linear collapse dynamic rather than a normal steady decline as a result (Daníelsson et al., 2022).

Sometimes the financial market is truly unpredictable. AI models used in these financial markets to overoptimise patterns in historical data, which may not hold true in future market conditions. This may result in short-term gains but makes models fragile to black swan events (rare, unpredictable, and high impact) or regime shifts (Vancsura et al., 2025). It is very important to note that AI models outside of their training environments can become vulnerable. Whenever there is an unprecedented change in the market, the overoptimised system might take unwarranted risks or not react swiftly, leading to systemic disruptions (Zhang et al., 2018). The intensity with which AI models are being developed surpasses the regulatory oversight. The algorithms develop way faster than the regulations. The authorities might not have enough technical literacy or real-time data to detect systemic risk arising from AI autonomous trading. This diminishes the effectiveness of macroprudential policies by creating information asymmetry between regulators and market participants (International Monetary Fund, 2024). Furthermore, the dynamic and adaptive nature of AI systems is not yet accommodated by most regulatory frameworks. As the financial market is interconnected globally, AI-driven risk in one jurisdiction can quickly spread to others, creating transnational systemic risk. The absence of harmonised global standards further complicates this matter (Mirishli, 2024b).

The above-discussed risk transmission channels show that AI is just not another technological upgrade in the financial market, but a paradigm-shifting force that has the capacity to cause systemic risks through novel mechanisms, rather than merely being another tool built by human beings to ease the tasks relating to financial markets. Understanding these transmission pathways is necessary to design an efficient and effective governance framework, regulatory reactions, and risk mitigation techniques.

Figure 1
Conceptual Framework: AI as Risk Amplifier in High-Frequency Trading
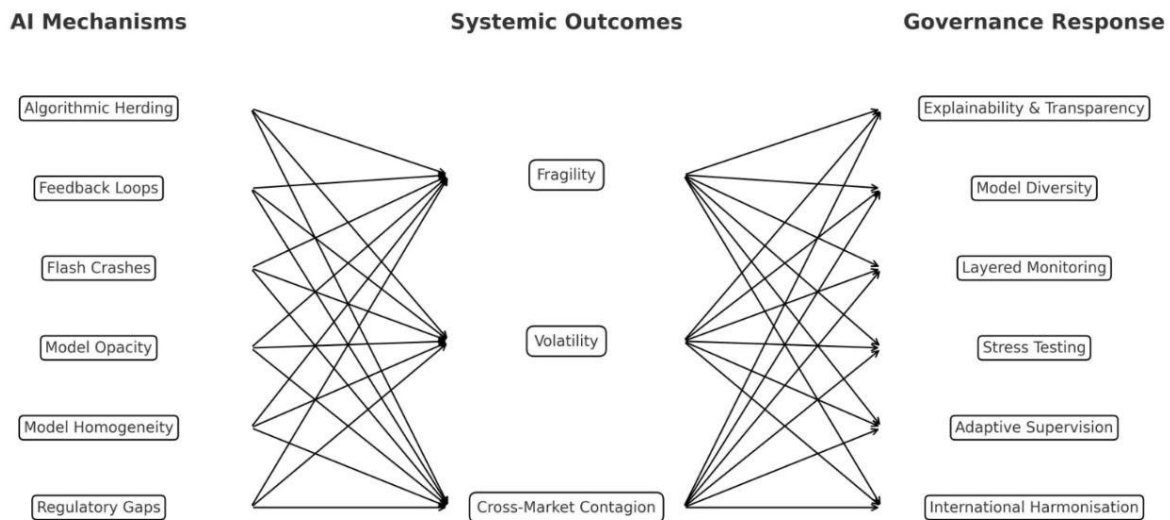
5. Supporting Empirical and Institutional Evidence

Several transmission channels through which AI augmented trading can cause systematic vulnerabilities have been identified in this conceptual framework. Both academic research and institutional white papers provide evidence supporting these mechanisms. According to historical data, automated trading can cause rapid, self-reinforcing price moves when liquidity vanishes. The joint Commodity Futures Trading Commission (CFTC)-Securities and Exchange Commission (SEC) forensic report on the May 6, 2010 "Flash Crash" states that an archetypal feedback loop was created when automated order execution, aggressive selling, and withdrawal of liquidity providers integrated to cause an immediate price dislocation followed by an equally swift rebound (Commissions, 2010). Using audit trail data from CME's E-mini S&P 500 futures during the Flash Crash, Kirilenko et al. (2017) demonstrated that HFT initially provided liquidity before swiftly switching to consuming it, which increased the intraday price dislocation once selling pressure spiked. Based on European Securities and Markets Authority (ESMA)'s risk surveillance, liquidity is still brittle during periods of market stress, which raises the possibility that algorithmic behaviour will accelerate price changes on European Union markets (Authority, 2024). The Bank for International Settlements (BIS) recorded how venue microstructure and execution algorithms can magnify order flow imbalances during market stress by channelling shocks via order book dynamics, such as abrupt changes in market-making behaviour and queue priority (Bank for International Settlements, 2020). A "fragile market liquidity" backdrop where sharp corrections are still possible is frequently highlighted in ESMA's Trends, Risks and Vulnerabilities (TRV) report (Authority, 2024). While automated trading improves liquidity and informational efficiency, it also raises short-horizon volatility, which is a significant channel for intraday amplification when shocks occur (Boehmer et al., 2021). These findings collectively show that liquidity risks are not historical anomalies but remain a present-day concern in AI-augmented markets.

Algorithmic herding and convergence are the most common issues discussed by researchers and institutions. Similar quantitative models can amplify losses when market conditions change, as evidenced by scholarly work on strategy crowding. According to Khandani and Lo (2011), funds employing comparable factor models unwound concurrently during the August 2007 "Quant Meltdown," resulting in a chain reaction of losses. Authorities warn that the widespread use of comparable AI models may lead to correlated behaviour that intensifies shocks. Focusing on the shared data, features, or model providers can make it more likely for the firms to "move together," increasing market fragility and opening the door for correlated errors or even unintentional collusion among agents (Bank of England, 2025b). Supervisory concerns about the rapid scaling of AI deployments can

synchronise decision rules across venues and intermediaries, possibly reducing diversity of responses under market stress, and are also documented in the International Organization of Securities Commissions (IOSCO)'s 2025 Capital-Market Report (International Organization of Securities Commissions, 2025). In this context, the main risk is not just individual model failure but the systemic impact of synchronized behavior across firms.

Automated trading increases short-horizon volatility while improving average liquidity. This is the mechanism through which shocks in one venue quickly shifted to others via statistical-arb and latency-arb links (Boehmer et al., 2021). Disturbances don't have to stay local when risk controls and strategies are coordinated through similar AI pipelines. The political and macro shocks quickly spread across the instruments and venues when liquidity is limited and systems are closely connected (Authority, 2024). AI-enabled capital market structure raises the possibility of interconnectedness. When firms source models or model components from a concentrated set of providers, operational or modeling errors can scale across asset classes (International Monetary Fund, 2024). Additionally, the Bank of England also warns that relying too much on a limited number of cloud platforms or third-party model providers can lead to common dependence that can create issues that become system-wide disruptions (Bank of England, 2025a). This suggests that AI is accelerating cross-market contagion by tightening both technical and behavioral linkages among trading systems.

When AI automated trading executes an order, there is a problem of explainability. Advanced AI models are dynamic and adaptive, learning from new input data production, which makes ex-post reconstruction of decisions more complicated and increases governance demand (Bank of England, 2025b). When models change after they are deployed, AI and machine learning in finance create unique challenges related to accountability, transparency, and validation, suggesting supervisory blind spots in the absence of improved controls (Board, 2017). A study from Fritz-Morgenthal et al. (2022) highlights the black box's explanatory gaps. Supervisors demand Explainable AI (XAI) setups like SHAP clustering because they cannot understand or they don't have the ability to question how AI systems make up decisions, but even these setups require supervisors to learn new skills and assume additional modeling risk. Also, it's very important to note that XAI techniques help, but they don't completely eliminate the risk of model opacity (Vancsura et al., 2025). The disparity between regulatory toolkits and the quick development of AI is a common theme. Model concentration risk, agent autonomy, and cross-firm dependencies are among areas where the perimeter must adjust, and the FPC specifically frames AI as a financial stability issue (Bank of England, 2025a). The use of AI in finance creates systematic risks that are impossible for microprudential tools to identify. They expose a regulatory misalignment between local efficiency and global resilience and demonstrate how AI optimization can jeopardise systemic stability (Daníelsson et al., 2022). According to IOSCO's reports, intermediaries and asset managers should be held responsible for proportionate, risk-based governance of AI and machine learning, with supervisors concentrating on testing, oversight, accountability, and keeping records to allow for post-hoc review of automated decisions (International Organization of Securities Commissions, 2021). According to the Financial Stability Board's (FSB) (2017) foundational work, in order to capture system-wide externalities, the financial stability perspective (macroprudential) must be used along with micro conduct and model-risk controls. To keep up with rapidly changing markets, regulation needs to become more adaptive, driven by data and AI-enabled (O'Halloran & Nowaczyk, 2019). Truby et al. (2020) argue that there is insufficient reactive regulation. The precautionary principle must be adopted by regulators toward AI in the financial market. The research also points out that the laws around innovation usually emerge after the crisis, but in markets like high-frequency trading, delayed intervention may lead to risks that are uncontrollable. Goldblum et al. (2022) demonstrate experimentally that robustness in HFT environments can be undermined by even a small adversarial perturbation, which can destabilise deep learning trading models. This is supported by Nehemya et al. (2021), who show that order-book data can be manipulated in the opposite direction to distorted outcomes. Financial AI and machine learning systems are seriously impacted by data poisoning attacks. The discussion about systemic fragility in the presence of hostile data intervention is further supported by the fact that even a small, well-planned tweak of input labels can compromise deep learning architectures in particular (Gallagher et al., 2022). Together, these studies highlight a widening gap

between technological capability and regulatory oversight, underscoring the urgency for adaptive supervision.

This evidence from academic publications and institutional reports focuses on one finding. AI trading improves efficiency in normal market conditions, but during market stress, it amplifies systemic fragility. This signifies the need for strong macroprudential supervision and flexible governance frameworks. Systemic risks such as herding, feedback loops, opacity, and adversarial vulnerabilities are further worsened by regulatory gaps.

6. Discussion

The objective of this paper was to examine how AI in high-frequency and autonomous trading can amplify systemic volatility, and to propose a governance framework that reduces these vulnerabilities. While automated trading systems may improve liquidity, speed, and price discovery under stable market conditions, they also tend to magnify systemic fragility during stress events. Risks such as algorithmic herding, flash crashes, feedback loops, model opacity, and regulatory blind spots emerge through multiple channels. The findings of the study are consistence with the previous academic research conducted by Serrano (2020), who argued that algorithmic herding intensifies intraday fragility; Khandani and Lo (2011) linked cascading losses to strategy convergence; how automated execution of orders resulted the 2010 Flash Crash was illustrated by Kirilenko et al. (2017) and Daníelsson et al. (2022) highlighting the misalignment between local efficiency and global resilience. When taken together, these perceptions offer solid evidence confirming that AI is not just a technological tool but a systemic amplifier, one that can quickly transform the localised trading volatility into global financial disturbances if left unchecked.

Based on the above discussions and the scientific evidence, we propose the following regulatory framework to be adopted by regulators, stock exchanges, financial institutions, authorities, supervisors, and cross-border institutions to mitigate the systemic risk arising from AI-augmented trading. Building on the evidence from Gao et al. (2024), regulators and exchanges under extreme but plausible conditions must conduct agent-based stress simulators of AI trading to anticipate flash crash dynamics and liquidity evaporation. The financial institutions should be motivated or mandated to use diverse AI architectures and data sources to avoid systemic monocultures. This reduces the possibility of correlated failures as warned by Daníelsson et al. (2022). Using modern techniques such as XAI and post-hoc interpretability frameworks, regulatory authorities should enforce basic explainability standards for AI trading. As pointed out by Vancsura et al. (2025), they may not eliminate opacity completely but will strengthen accountability and oversight. AI-assisted regulatory systems, which are capable of real-time surveillance and adaptive response, must be adopted by supervisors, and they must progress beyond static, rule-based monitoring systems as suggested by O'Halloran and Nowaczyk (2019). Both the International Organization of Securities Commissions (2025) and the International Monetary Fund (2024) highlight that fragmented oversight leaves crucial blind spots. So, the regulators must pursue international cooperation and harmonise standards across jurisdictions, since AI-driven risk transmission is a global phenomenon.

Although the proposed framework highlights important governance principles, it becomes more valuable when operationalised through quantifiable metrics and case-based applications. To measure fragility due to AI-trading, regulators could use intraday volatility indices, liquidity coverage ratios (LCRs), or order-to-trade ratios (OTRs). Model concentration scores could be used to gauge the dependence on shared datasets, cloud providers, or model vendors to monitor algorithmic convergence. To evaluate model opacity, explainability audits can be used, and companies must show that interpretable methods like SHapley Additive exPlanations (SHAP) or Local Interpretable Model Agnostic Explanations (LIME) are used to create AI decision pathways. In addition to traditional balance sheet simulations, agent-based market simulators should be included in stress testing. This will allow regulators to foresee feedback loops and liquidity spirals in the event of a crisis. For instance, the framework would provide directions to regulators to trigger real-time circuit-breaker stress tests across AI trading models in the event of a flash crash. Supervisors could keep an eye on whether feedback loops accelerate market contagion and whether liquidity evaporates disproportionately across particular

asset classes. To restore stability in these situations, adaptive supervision may require staggered trade execution or temporarily throttle algorithmic order submission. These operational scenarios show how the proposed framework can be converted into a quantifiable instrument and useful regulatory measures.

Thus, the study contributes to the academic research and regulatory discussion by presenting not only a conceptual framework but also practical pathways for its implementation. While existing studies considered systemic risks as an isolated incident, this study offers an integrated conceptual framework that systematically links transmission risk channels with governing principles. The paper's dual focus is to identify how AI and autonomous trading amplify systemic risks and to outline operational regulatory responses, such as explainability audits, model diversity requirements, and stress-testing tools. It stresses how AI in HFT presents systemic risks during periods of market stress, even though it can be beneficial in stable market conditions. Diverse models, layered monitoring, explainability and transparency, dynamic supervision, and cross-border regulatory cooperation will be necessary for effective governance. Policy makers can only strike a balance between the need for financial stability and efficiency gains from AI by establishing such a framework. Future research should test this conceptual framework empirically using real trading data and simulation models. Comparative research across jurisdictions may show the impact of regulatory variation on systemic risk associated with AI. Resilient governance strategies will be informed by a deeper understanding of explainable AI, stress testing, and cross-market contagion mechanisms.

Disclosure Statement

The authors claim no conflict of interest.

Funding

References

Ahmed, S., Alshater, M. M., El Ammari, A., & Hammami, H. (2022). Artificial intelligence and machine learning in finance: A bibliometric review. *Research in International Business and Finance*, *61*, Article 101646. https://doi.org/10.1016/j.ribaf.2022.101646

Authority, E. S. (2024). *Trends, risks and vulnerabilities (TRV) Report, No. 2, 2024*. ESMA. https://www.esma.europa.eu/sites/default/files/2024-08/ESMA50-524821-3444_TRV_2_2024.pdf

Bank for International Settlements. (2020). *FX execution algorithms and market functioning*. Markets Committee. https://www.bis.org/publ/mktc13.pdf

Bank of England. (2025a). *Financial stability in focus: Artificial intelligence in the financial system*. https://www.bankofengland.co.uk/financial-stability-in-focus/2025/april-2025?utm

Bank of England. (2025b). *AI in financial services: Written evidence submitted by the Bank of England*. House of Commons Treasury Committee. https://committees.parliament.uk/writtenevidence/140106/pdf/

Batool, A., Zowghi, D., & Bano, M. (2025). AI governance: A systematic literature review. *AI and Ethics*, *5*, 3265-3279. https://doi.org/10.1007/s43681-024-00653-w

Board, F. S. (2017). *Artificial intelligence and machine learning in financial services: Market developments and financial stability implications*. Financial Stability Board. https://www.fsb.org/uploads/P011117.pdf

Boehmer, E., Fong, K., & Wu, J. (2021). Algorithmic trading and market quality: International evidence. *Journal of Financial and Quantitative Analysis*, *56*(8), 2659–2688. https://doi.org/10.1017/S0022109020000782

Commissions. (2010). *Findings regarding the market events of May 6, 2010: Report of the staffs of the CFTC and SEC to the joint advisory committee on emerging regulatory issues*. https://www.cftc.gov/sites/default/files/idc/groups/public/@otherif/documents/ifdocs/staff-findings050610.pdf

Dakalbab, F., Talib, M. A., 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), Article 102015. https://doi.org/10.1016/j.jksuci.2024.102015

Daníelsson, J., Macrae, R., & Uthemann, A. (2022). Artificial intelligence and systemic risk. *Journal of Banking & Finance*, 140, Article 106290. https://doi.org/10.1016/j.jbankfin.2021.106290

Fritz-Morgenthal, S., Hein, B., & Papenbrock, J. (2022). Financial risk management and explainable, trustworthy, responsible AI. *Frontiers in Artificial Intelligence*, 5, Article 779799. https://doi.org/10.3389/frai.2022.779799

Gallagher, M., Pitropakis, N., Chrysoulas, C., Papadopoulos, P., Mylonas, A., & Katsikas, S. K. (2022). Investigating machine learning attacks on financial time series models. *Computers & Security, 123*, Article 102933. https://doi.org/10.1016/j.cose.2022.102933

Gao, K., Vytelingum, P., Weston, S., Luk, W., & Guo, C. (2024). High-frequency financial market simulation and flash crash scenarios analysis: An agent-based modelling approach. *The Journal of Artificial Societies and Social Simulations*, 27(2), Article 8. https://doi.org/10.18564/jasss.5403

Gayduk, R., & Nadtochiy, S. (2017). Liquidity effects of trading frequency. *Mathematical Finance, 28*(3), 839-876. https://doi.org/10.1111/mafi.12157

Gensler, G. (2023, July 17). *"Isaac Newton to AI" remarks before the national press club.* U.S. Securities and Exchange Commission. https://www.sec.gov/newsroom/speeches-statements/gensler-isaac-newton-ai-remarks-07-17-2023?stream=top#_ga=2.67425366.639542815.1755666062-573009676.1755666062

Goldblum, M., Schwarzschild, A., Patel, A., & Goldstein, T. (2022). Adversarial attacks on machine learning systems for high-frequency trading. In A. Calinescu & P. Szpruch (Eds.), *Proceedings of the Second ACM International Conference on AI in Finance* (pp. 1-9). Association for Computing Machinery. https://dl.acm.org/doi/10.1145/3490354.3494367

International Monetary Fund. (2024). *Advances in artificial intelligence: Implication for capital market activities.* https://www.elibrary.imf.org/display/book/9798400277573/CH003.xml?utm

International Organization of Securities Commissions. (2021). *The use of artificial intelligence and machine learning by market intermediaries and asset managers: Final report.* https://www.iosco.org/library/pubdocs/pdf/IOSCOPD684.pdf

International Organization of Securities Commissions. (2025). *Artificial intelligence in capital markets: Use cases, risks, and challenges.* https://www.iosco.org/library/pubdocs/pdf/IOSCOPD788.pdf

Khan, F. S., Mazhar, S. S., Mazhar, K., AlSaleh, D. A., & Mazhar, A. (2025). Model-agnostic explainable artificial intelligence methods in finance: A systematic review, recent developments, limitations, challenges and future directions. *Artificial Intelligence Review*, 58, Article 232. https://doi.org/10.1007/s10462-025-11215-9

Khandani, A. E., & Lo, A. W. (2011). What happened to the quants in August 2007? Evidence from factors and transactions data. *Journal of Financial Markets*, 14(1), 1-46. https://doi.org/10.1016/j.finmar.2010.07.005

Kirilenko, A., Kyle, A. S., Samadi, M., & Tuzun, T. (2017). The flash crash: High-frequency trading in an electronic market. *The Journal of Finance*, 72(3), 967-998. https://www.jstor.org/stable/26652722?utm

Lakhchini, W., Wahabi, R., & El Kabbouri, M. (2022). Artificial intelligence & machine learning in finance: A literature review. *Inernational Journal of Accounting, Finance, Auditing, Management and Economics*, 3(6-1), 437-455. https://doi.org/10.5281/zenodo.7454231

Leal, S. J., Napoletano, M., Roventini, A., & Fagiolo, G. (2015). Rock around the clock: An agent-based model of low- and high-frequency trading. *Journal of Evolutionary Econimics*, 26, 49-76. https://link.springer.com/article/10.1007/s00191-015-0418-4

Mirishli, S. (2024a). Ethical implications of AI in data collection: Balancing innovation with privacy. *Ancient Land International Online Scientific Journal*, 6(8), 40-55. https://doi.org/10.36719/2706-6185/38/40-55

Mirishli, S. (2024b). Regulating AI in financial services: Legal frameworks and compliance challenges. *Qanun*, 8(358), 29–40. https://arxiv.org/pdf/2503.14541

Nehemya, E., Mathov, Y., Shabtai, A., & Elovici, Y. (2021). Taking over the stock market: Adversarial perturbations against algorithmic traders. In Y. Dong, N. Kourtellis, B. Hammer, & J. A. Lozano

(Eds.), *Machine learning and knowledge discovery in databases. Applied data science track* (pp. 221-236). Springer. https://link.springer.com/chapter/10.1007/978-3-030-86514-6_14

Ogbuonyalu, Abiodun, K., Dzamefe, S., Vera, E. N., Oyinlola, A., & Emmanuel, I. (2024). Assessing artificial intelligence driven algorithmic trading on market liquidity risk and financial systemic vulnerabilities. *International Journal of Scientific Research and Modern Technology*, *3*(4) 18-21. https://doi.org/10.38124/ijsrmt.v3i4.433

O'Halloran, S., & Nowaczyk, N. (2019). An artificial intelligence approach to regulating systemic risk. *Frontiers in Artificial Intelligence*, *2*, Article 7. https://doi.org/10.3389/frai.2019.00007

Ozili, P. K. (2024, September 05). *Artificial intelligece (AI), financial stability and financial crisis*. SSRN. https://doi.org/10.2139/ssrn.4970195

Serrano, A. S. (2020). High-freqency trading and systemic risk: A structured review of findings and policies. *Review of Economics*, *71*(3), 169-195.

Shrinivas, S. K., & Shetty, P. (2024). Impact of artificial intelligence on business, business operations and strategy. *International Journal of Creative Research Thoughts*, *12*(8), 152-160.

Svetlova, E. (2022). AI ethics and systemic risks in finance. *AI and Ethics*, *2*, 713-725. https://link.springer.com/article/10.1007/s43681-021-00129-1?utm

Truby, J., Brown, R., & Dahdal, A. (2020). Banking on AI: mandating a proactive approach to AI regulation in the financial sector. *Law and Financial Markets Review*, *14*(2), 110-120. https://doi.org/10.1080/17521440.2020.1760454

Vancsura, L., Tatay, T., & Bareith, T. (2025). Navigating AI-driven financial forecasting: A systematic review of current status and critical research gaps. *Forecasting*, *7*(3), Article 36. https://doi.org/10.3390/forecast7030036

Vincent, V. U. (2021). Integrating intuition and artificial intelligence in organizational decision-making. *Business Horizons*, *64*(4), 425-438. https://doi.org/10.1016/j.bushor.2021.02.008

Zhang, C., Vinyals, O., Munos, R., & Bengio, S. (2018). *A study on overfitting in deep reinforcement learning*. arXiv preprint. https://arxiv.org/abs/1804.06893v2