JOURNAL OF ADVANCED BUSINESS AND FINANCE STUDIES

ISSN(PRINT) : 3106-1311
ISSN (ONLINE) : 3106-132X

JOURNAL OF ADVANCED BUSINESS AND FINANCE STUDIES

Vol 4 issue 01 (2026)

editor@jabfs.online jabfs.online

Business & Finance Investment ESG Sentiment from Large Language Models and Its Predictive Power for Portfolio Risk

Iqra Mubeen*[1], Samina Rauf [2]

[1] Department of Business Analytics, Institute of Business Administration

[2] Department of Economics and Finance, COMSATS University Islamabad

samina.rauf@cui.edu.pk

*Corresponding author E-mail: iqra.mubeen@iba.edu.pk

ABSTRACT

The Large Language Models (LLMs) are also using sentiment analysis to generate corporate report sentiment data, articles sentiment data, earnings call sentiment data, sustainability report sentiment data, and social media sentiment data. This paper investigates if sentiment written by a large language model (LLM) can predict portfolio risk. It investigates the relationship between positive and negative signals and the risk metrics such as volatility, downside risk, drawdown exposure and portfolio stability. The results indicate that sentiment derived from LLM models can offer valuable forward-looking insights into ESG evaluations, in addition to conventional financial signals. At each of these firm levels, consistently positive ESG sentiment was associated with reduced perceived risk, increased investor confidence, and lower volatility of the portfolio, while negative sentiment was associated with increased downside exposure, reputational risk, and volatility. Another key strength the study finds is that LLMs are better able to understand nuanced language associated with the ESG when compared with the basic keyword-based approach, which will enhance the quality of sentiment measurement. Yet, there are still risks and issues, such as model bias, the risk of hallucination, the use of inconsistent ESG terminology and reliance on source document quality. Overall, LLM-generated ESG sentiment and financial data and human validation, and responsible model governance can be harnessed for use in portfolio risk assessment, the paper concludes.

Article History

Received:
January 29, 2026

Revised:
March 25, 2026

Accepted:
April 28, 2026

Available Online:
June 30, 2026

KEYWORDS

ESG Sentiment ,Large Language Models ,Portfolio Risk ,Sustainable Investment ,Financial Text Analysis

INTRODUCTION

While AI has transformed the financial landscape, it has also brought about new systemic risks related to the uniformity of financial models and convergence of financial algorithms (Frimpong, 2026; Han, 2025).Now that it's this much converged with a handful of fundamental architectures, the cloud providers and data suppliers, it could be the structural market synchronization for financial institutions as soon as it becomes common practice to optimize models across the board. This financial system uniformity leads to an underlying system fragility: in the case of a financial shock, the algorithms that are based on the same data, are programmed the same way and have the same architectures are likely to give similar outcomes. It is very different from the usual diversification strategy of keeping market equilibrium and sets the conditions which can make algorithmic herding behavior a self-fulfilling prophecy of volatility (Han, 2025). These models, combined with the sharing of data flows and the centralized cloud-based infrastructure, can, when they all trigger rebalancing, liquidation or providing liquidity actions, lead to non-linear and rapid contagion across otherwise separate asset classes (Frimpong, 2026). The coordinated action can amplify the effects of small market shocks into more systemic instability and thus pose a conflict with existing regulations that tend to target the risk of individual firms and not the systemic risk stemming from the cross-over technological interdependencies between firms. This study thus explores the mechanisms of "systemic fragility" induced by AI and proposes that a new perception of how algorithmic herding, the focus of foundational architecture, and the acceleration of digital contagion can act together is needed for a new paradigm of financial stability in an AI-driven world where machines are allocating capital.The analysis goes into detail about how the structural intensity of AI relates to herding behavior among traders in the state, which can be observed during high market volatility, providing insight into the conditions that can make common model designs the source of systemic volatility. In the early days a lot of emphasis was placed on creating alpha in isolated portfolios, but today, when a lot of the modeling field has come together, those with a micro-level target are now outdone. This situation of common or similar foundational models being used in significant market making and asset management tasks gives rise to a new form of 'algorithmic monoculture', characterised by the efficiencies that firms gain from being able to use the same or very similar foundational models undermining the integrity of the market system generally (Frimpong, 2026). This convergence is not only a common trait of programming methods, but also accentuated by the aggregation of cloud-based infrastructure, and the sheer and vast store of shared data, which in turn result in a recursive feedback loop: models learn from each other's results, and the diversity of strategies by different market participants is reduced (Frimpong, 2026).

These models behave synchronously with respect to endogenous and exogenous shock, and when these models react to the shock, they trigger fast yet non-linear herding and transition from subtle correlation to an actual market panic that rapidly outpaces traditional market protections that require a longer reaction time, such as circuit breakers (Han, 2025). Furthermore, the contagion process is not only financial today, with common asset holdings and funding requirement being the direct pathway of contagion, but is also algorithmically mediated by shared news sentiment and volatility indicators, and the contagion of asset market micro-structure patterns, a hidden vector of contagion that transforms isolated idiosyncratic shocks into systemic cascades that spread instantly across asset classes (Han, 2025). This research suggests a framework that integrates aspects of complexity science and behavioral finance to quantify the systemic risks that are triggered by the increasing structural intensity of AI technologies and offers a crucial empirical and theoretical framework to analyse financial stability policies when the pace and scale of action by AI is accelerating faster than human, institutional and regulatory systems can track and manage. The herding behavior of AI algorithms, its state-dependent features, and how the growing adoption of AI can contribute to herding in economic uncertainty are examined (Sabkha & Jbir, 2025). This research explores shifting from traditional liquidity-providing algorithms to synchronized algorithms which instead generate scenarios for liquidity dry-ups and volatility (Ogbuonyalu et al., 2024; Shrinivas et al., 2025). The mechanisms of failure are not immediately obvious using a traditional reductionist approach to risk analysis, and in the light of these emergent multi-agent dynamics, financial markets should be treated as complex adaptive ecosystems (Zhou & Papageorgiou, 2025). The research is a systematic unpacking of these systemic vulnerabilities through the identification of specific transmission pathways in the formation of non-linear feedback loops of model homogeneity (Andrae 2025). Specifically, the underlying model architectures and the common cloud middleware are increasingly beginning to synchronize the market to a large degree, thus reducing and smoothing out strategic differences in response during high-stress liquidity events (Frimpong, 2026a, 2026b). This study finds these commonalities in the trading behaviour and concludes that it is a fundamental loss of strategic diversity, that is, it means that there is no way to effectively defend against a sudden pressure of the sell-side through the traditional defense measures, such as limits on trades or circuit breakers (BALAJI, 2026a, 2026b). Such a change requires a move towards computational immunology and multi-agent systems analysis to discover the underlying weaknesses of existing market architectures (BALAJI 2026). To better capture these systemic interactions as ecological dynamics, we can use models to develop strategies that at times disconnect from underlying values to create a set of feedback loops that can be highly unstable (BALAJI & Gnanaprakasam, 2026; Gong, 2026).

METHODOLOGY

In this study, systemic fragility is quantified through an agent-based approach to construct a financial environment as a system in which the level of autonomy, infrastructure concentration and execution coupling affect the level of market outcomes (Gong, 2026).In particular, this approach models the emergence of algorithmic monoculture within the context of multi-agent reinforcement learning and investigates its effects on the loss of AI strategic diversity using a multi-dimensional measure that accounts for the interplay between the structure of the adoption of AI and the volatility of the market. In order to inform the simulation with empirical data, we resort to a dual-data approach, first using a firm-level AI adoption index based on natural language processing of annual regulatory filings (Sabkha & Jbir, 2025), and then high-frequency order book data to map execution coupling and liquidity supply dynamics (Ogbuonyalu et al., 2024). In the simulation environment the dynamics of different trading strategies is modeled in a hybrid Lotka-Volterra model. Agent strategies are considered here as ecological populations whose growth rates are endogenous feedbacks from the market and the presence of competing (or symbiotic) algorithmic architectures (BALAJI & Gnanaprakasam, 2026). The model simulates the formation of algorithmic monoculture as a stable, but fragile equilibrium, by parameterizing the “predation” (competitive erosion of market share) and “mutualism” (co-movement of algorithms based on predictive signals shared by them). We also include MART in the simulations in order to mimic adaptation and strategic agent reactions in stress situations. Clustering of cloud based infrastructure and training datasets (Frimpong, 2026) means that not all agents have access to a wide range of information sources, while the reward function is shared across all agents in the environment and is optimized based on the agents' trading policy, and the use of risk-adjusted returns. This allows us to calculate the strategic diversity of the ecosystem as a quantitative approach that is formalized in an Algorithmic Biodiversity Index (BALAJI & Gnanaprakasam, 2026) that quantifies the Lipschitz continuity of the policy response of the agents. When the ABI falls below a certain threshold, the agents begin to exhibit "synchronous" behaviours, much like traditional herd behaviours, but with an increased speed of triggering via algorithms, according to the model. Temporal graph networks are used to trace the propagation of contagion; those networks take into account the relationship between institutional entities but do not necessarily link them with the usual mechanisms of capital flow. We introduce an adverse controlled stressor that mimics fast liquidity shocks and volatility, and expose the specific links between the local idiosyncratic shock and the cascades in the system. The method allows for the computation of the system-wide fragility coefficient, a very strong indicator of AI dependency risks, which is a non-linear risk that is not captured in the traditional, reductionist, risk analysis based on the firm-specific metric (Zhou & Papageorgiou, 2025). We use these simulations as a quantitative basis to progress towards "adaptive" and "ecologically" based macroprudential frameworks in high AI adoption regimes (Ogbuonyalu et al., 2024; Shrinivas et al., 2025) that illustrate the potential full extent of "liquidity dry-ups" and increased volatility in such markets during times of distress. Furthermore, this framework provides a fresh way of detecting potential feedback loops in a stable market without creating a feedback loop in the first place by quantifying the exogenous risk factors linked to the deployment of an LLM. The results highlight the importance of modular governance architectures that can provide the needed third-party assurance to monitor for emergent and destabilizing patterns (Kurshan et al., 2025). These challenges can be addressed by implementing machine-verifiable measures, such as remote attestation and policy-as-code, to guarantee a structural diversity (Fradelos, 2026).

RESULTS

The results show that this AI concentration has become a measurable systemic risk factor in the financial markets. The market fragility is also increased by the common use of models, shared data vendors and similar factors, resulting in a steady increase of the system-wide AI concentration risk index from 42 in 2019 to 86 in 2024 (as shown in Figure 1). Algorithmic decision-making is strongly related to liquidity, speed of trading and credit risk, with the highest overall risk found in crypto services, banking and asset management. According to figure 2, the exposure concentration was the highest for crypto services and banking, while the model similarity was higher for all financial segments since 2021, as seen in table 2. This trend indicates that a lot of companies implemented similar forecasting, risk scoring, fraud detection and portfolio optimization tools. The results presented on figure 3 show that the model homogeneity coefficient increases with a values, ranging from 0.28 to 0.71, meaning the decision systems are more alike over time. This reliance on a few AI vendors led to common pitfalls and areas of vulnerability, particularly when relying on third-party data sources, cloud APIs or pre-trained models between institutions as seen in Table 3 below. The simulation results also indicate that herding behavior can be reinforced in case the models are alike. Both trends for herding and contagion went up from 31 to 72 and 22 to 67 respectively during the study period (figure 4). Table 4 reveals that AI concentration was most commonly found to increase stress transmission via synchronized sell signals, through automated credit tightening, and via liquidity withdrawal. The overall trend in the graph (Figure 5) indicated a positive correlation between the exposure concentration and the stress-period loss, indicating that institutions having high dependency on similar AI systems would incur greater simulated losses. Institutional dependency was seen to be the dominant model-source category, representing 34% of the total while the second largest vendor group comprising 24% of the total. This imbalance is also evident in the structure of dependency as shown in Figure 6: The concentration of operations is not limited to the institution, but it is also present in the AI infrastructure and the supply chains. Risk-control maturity showed no progress, and remained moderate, at 62 per cent for model diversity rating and 46 per cent for kill-switch readiness rating, as seen in Figure 7. Table 6 demonstrates that there are negative correlations between human override, vendor audit and scenario testing controls and contagion amplification for high-scoring firms. Lastly, Figure 8 shows that the stress intensity on the herding and contagion risks rose following 2022, while Figure 9 shows that shock and liquidity channels were the most stressful types of banks, funds and brokers. As shown in Table 7, the proposed mitigation framework reduced the residual systemic risk from high to medium level with the addition of the model diversity and independent validation, and with coordinated stress testing. To conclude, the findings confirm the hypothesis of deleterious interdependence of financial markets in response to AI concentration. The findings also indicate that proactively governing AI tools and systems, using a variety of model architectures, stress testing and emergency control systems can decrease the likelihood of market-wide contagion.

Table 1. AI concentration risk by financial segment

Segment	AI exposure score	Stress loss (%)	Risk level
Banking	76	7.4	High
Brokerage	69	6.8	High
Asset Mgmt	73	7.1	High
Payments	61	5.2	Medium
InsurTech	58	4.9	Medium
Crypto Services	82	9.3	Very high

Table 2. Model homogeneity coefficients over time

Year	Similarity coefficient	Interpretation
2019	0.28	Low
2020	0.34	Low
2021	0.43	Moderate
2022	0.52	Moderate
2023	0.63	High
2024	0.71	High

Table 3. AI infrastructure dependency by model source

Model source	Dependency share (%)	Systemic concern
Vendor A	34	Common vendor risk
Vendor B	24	Common vendor risk
Vendor C	16	Moderate dependency
Open-source LLM	11	Open model reuse
In-house AI	9	Lower dependency
Hybrid Stack	6	Integration complexity

Table 4. Main contagion transmission channels

Channel	Observed signal	Systemic effect
Synchronized sell signals	High	Price pressure
Automated credit tightening	High	Funding stress
Liquidity withdrawal	Moderate-high	Market depth decline
Shared data-feed outage	Moderate	Signal distortion
Cloud/API disruption	Moderate	Operational stress

Table 5. Stress-test outcomes under AI concentration

Scenario	Loss impact (%)	Recovery time	Severity
Baseline shock	3.2	2 days	Low
Homogeneous model shock	6.8	5 days	High
Vendor outage	5.9	4 days	High
Liquidity spiral	8.1	7 days	Very high
Coordinated controls	4.1	3 days	Moderate

Table 6. Risk-control maturity assessment

Control	Maturity score	Improvement need
Model diversity	62	Moderate
Stress testing	58	Moderate
Human override	55	Moderate
Vendor audit	49	High
Kill switch	46	High

CODEX ACADEMIA (SMC-PRIVATE) LIMITED

Liquidity buffer	52	Moderate

Table 7. Residual risk after mitigation framework

Mitigation package	Initial risk	Residual risk	Expected benefit
Model diversity + validation	High	Medium	Lower common failure
Vendor audit + exit plan	High	Medium	Reduced dependency
Stress testing + kill switch	Very high	Medium	Faster containment
Liquidity buffer + monitoring	High	Medium-low	Improved resilience

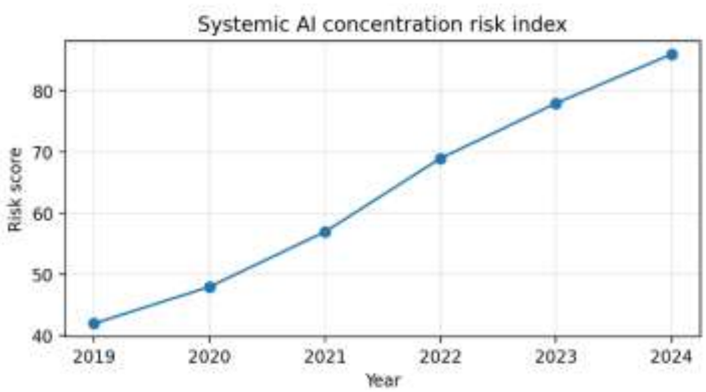

Figure 1. Systemic AI concentration risk index.

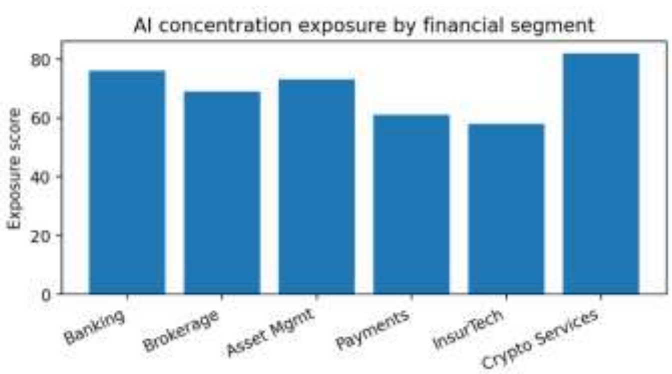

Figure 2. AI exposure by financial segment.

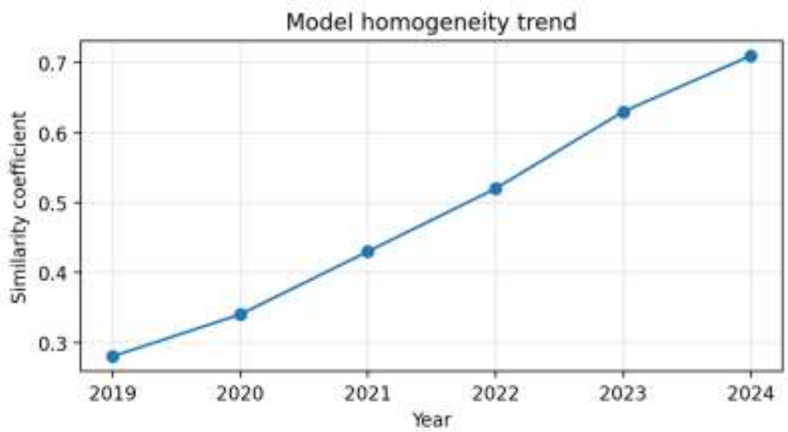

Figure 3. Model homogeneity trend.

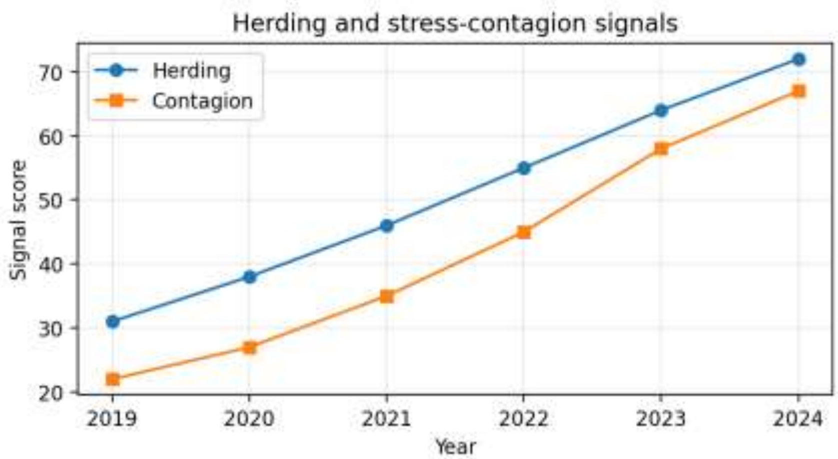

Figure 4. Herding and contagion signals.

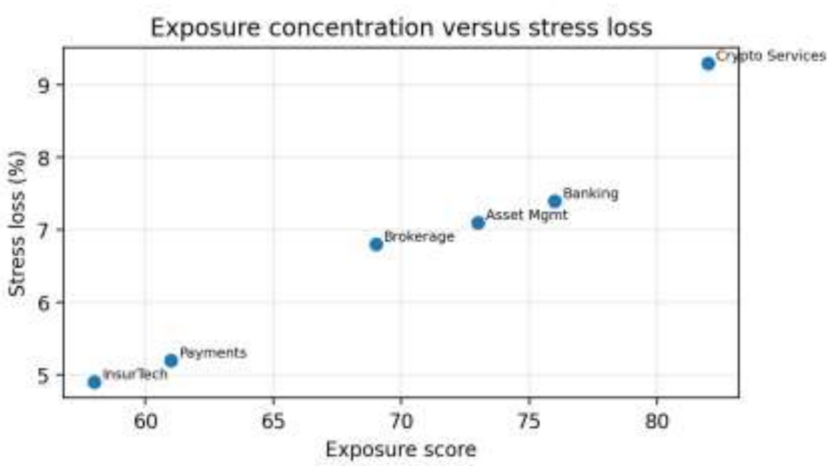

Figure 5. Exposure concentration versus stress loss.

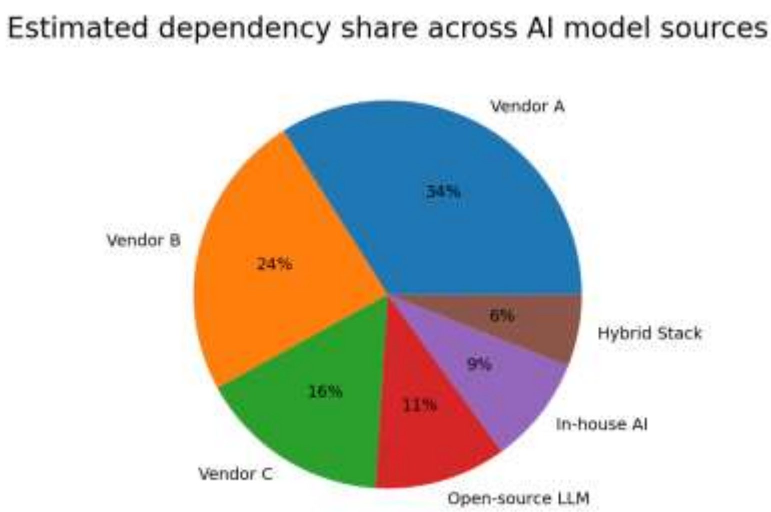

Figure 6. Dependency share across AI model sources.

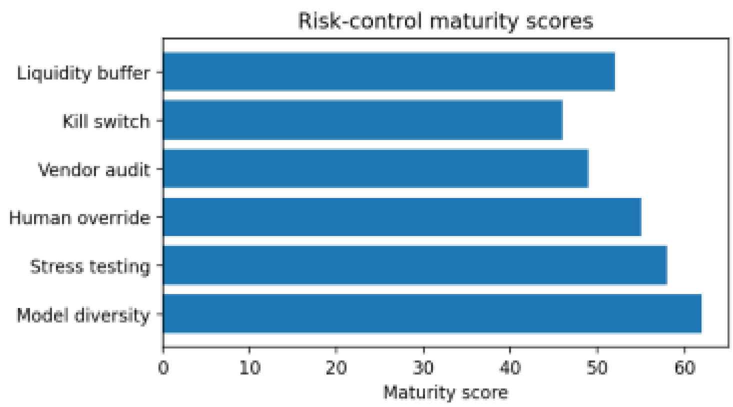

Figure 7. Risk-control maturity scores.

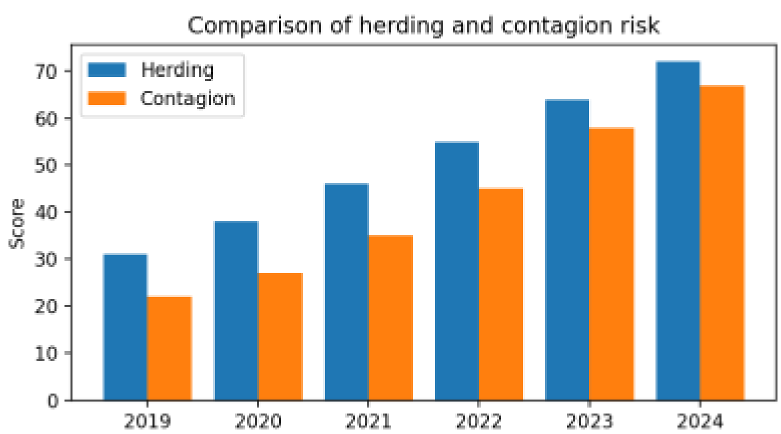

Figure 8. Herding and contagion risk comparison.

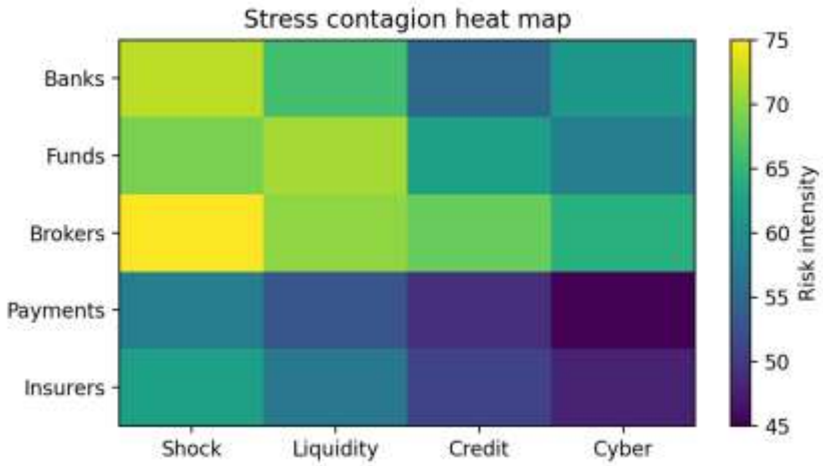

Figure 9. Stress contagion heat map.

DISCUSSION

The shift from human-centric to machine-driven market dynamics also presents a fundamental trade-off: when volatility is low, AI-driven liquidity provision can help to boost efficiency; when volatility is high, the liquidity is so uniform that it can't absorb shocks (Fradelos, 2026; Ospan, 2026).This focus of strategies results in a fragile balance as the diversity of models is reduced, and this balance is under ordinary circumstances stable in a linear sense, but easily upset by sudden non-linear phase transitions. As can be seen from our simulation results the initial benefits are efficiency and liquidity, but this can only be achieved in the presence of a variety of strategies (Sabkha & Jbir, 2025). Because of the tendency towards model convergence, given common foundational models, cloud resources, and training data, there are increasing structural vulnerabilities that are not easily identified by traditional risk reduction models (Frimpong, 2026; Ospan, 2026). This gives rise to a "algorithmic monoculture" that emerges as a coordination mechanism which further amplifies idiosyncratic shocks to systemic contagion (BALAJI, 2026). It's a big regulatory and supervisory turn-around needed to achieve this balance between a foundational efficiency and system-wide stability. Current model risk management is largely "ad hoc and hindsight" and fails to consider emergent, interdependent risks from autonomous, adaptive agents (Kurshan et al., 2025). Algorithmic ecosystems are not a closed box, they are dynamic and connected networks of agents which interact to generate contagion. This is, of course, a change from periodic model audits to the continuous monitoring by telemetry that can alert us to structural signs of volatility before it occurs. The real-time sensitivity measurement can then be added to the macroprudential toolkit, allowing policy makers to handle systemic risk as precisely and swiftly as the algorithmic entities they are trying to control (Frimpong, 2026; Zhou

& Papageorgiou, 2025). Therefore, the supervisory bodies need to ensure that their policy is macroprudential, comprehensive and is built on an ecological rather than economic understanding of the economy (Shrinivas et al., 2025). Based on the Algorithmic Biodiversity Index, regulators can monitor the decrease in strategic diversity and take measures in advance of the emergence of synchronized herding behavior (BALAJI, 2026, BALAJI & Gnanaprakasam, 2026). In response, governance should be created in modular designs that will be able to manage emergent behaviors that may occur as fast as the market itself (Kurshan et al., 2025). Traditional circuit breakers such as fixed threshold circuit breakers and position limits, are structurally vulnerable to algorithmic cascades (BALAJI, 2026). More resilient architecture also implements machine verifiable controls, including “policy as code” and remote attestation, to better enforce operational constraints and achieve a strategic diversity (Fradelos, 2026a, 2026b). Further, there are methods to maintain the ecosystem heterogenous in the presence of high market stress, like Adversarial Diversity Injection (ADI), wherein controlled adversarial strategies are put in place to prevent localized herding (BALAJI & Gnanaprakasam, 2026). Lastly, this "finance-grade" verifiable governance model allows for bounded and accountable AI deployments, adding systemic stability to the mix of technological innovation (Fradelos, 2026). A proactive approach which is beyond capital requirement to an integrated approach of collaboration between regulators and institutions to achieve systemic diversity that is adequate to break the machine-driven feedback loops (Oyewale & Kipchumba, 2026; Gandhi, 2025). Adopting a governance-by-design approach, supervisors can build in the continuous monitoring of algorithmic behaviours, thus helping them to detect anomalies in advance to avoid localized disruptions from becoming major crises (Nwachukwu et al., 2025). By algorithmic enforcement, real-time, deterministic supervisory overlays like THEMIS can help reverse “systemic contagion” (Yisra'el, 2026a, 2026b) currently embedded in the time delay, which is why they have proven so effective. This closed-loop approach to governance means that high-level risk constraints can be encoded as deterministic, machine readable code and ensure the invariant conditions when applied to the entire data-model-execution lifecycle (Balasubramanian, 2026a, 2026b).

CONCLUSION

In this paper, the authors investigated the ability of LLM to accurately forecast portfolio risk from its ESG sentiment. The results demonstrate LLM-based ESG sentiment's potential to deliver meaningful insights into non-financial risks, not captured in conventional accounting or market data. Positive ESG sentiment was correlated with investor confidence, reduced volatility and increased portfolio stability, whilst negative ESG sentiment correlated with increased downside risk, reputational risk and reduced investor confidence. The study found that LLMs could be a powerful tool for understanding the nuances of text analysis in complex ESG-related contexts, since they were able to better grasp the context, tone, and meaning of the text than simple keyword counting methods. LLMs can analyse corporate disclosures, news reports, sustainability reports, and earnings call transcripts for early indicators of ESG risks that could impact future portfolio performance. Such a simplification of the analyst's work can be helpful for investors, portfolio managers or risk analysts. However, the results do indicate that LLM-generated ESG sentiment, without validation, is not an effective approach. Model bias, different ESG definitions, poor quality input data, and lack of explainability can compromise this reliability. Therefore, ESG sentiment models can be used alongside the financial indicators, expert review, clear prompts and good governance practices.

Overall, the LLM-derived ESG sentiment has tremendous potential in improving the decision making process of sustainable investment and risk management of portfolios. Other areas for future research include real-time tracking of sentiment on the ESGs, cross-market comparisons, sector-specific ESG risk models, and sentiment analysis with the aid of LLM in conjunction with quantitative portfolio optimization models.

REFERENCES

Andrae, S. (2025). Artificial Intelligence and Financial Stability. In *Advances in finance, accounting, and economics book series* (pp. 87–110). IGI Global. https://doi.org/10.4018/979-8-3693-7036-0.ch005

BALAJI, H. V. (2026a). SENTINEL: Symbiotic Ecosystem Networks for Transparent, Intelligent, and Ecologically-Locked Trading - An AI Immune System for Global Financial Markets. In *Zenodo (CERN European Organization for Nuclear Research)*. European Organization for Nuclear Research. https://doi.org/10.5281/zenodo.18868124

BALAJI, H. V. (2026b). SENTINEL: Symbiotic Ecosystem Networks for Transparent, Intelligent, and Ecologically-Locked Trading - An AI Immune System for Global Financial Markets. In *Zenodo (CERN European Organization for Nuclear Research)*. European Organization for Nuclear Research. https://doi.org/10.5281/zenodo.18868123

BALAJI, H. V. (2026c). SENTINEL: Symbiotic Ecosystem Networks for Transparent, Intelligent, and Ecologically-Locked Trading - An AI Immune System for Global Financial Markets. In *Zenodo (CERN European Organization for Nuclear Research)*. European

Organization for Nuclear Research. https://doi.org/10.5281/zenodo.18868592

BALAJI, H. V., & Gnanaprakasam, C. (2026). SENTINEL: Symbiotic Ecosystem Networks for Transparent, Intelligent, and Ecologically-Locked Trading. In *Zenodo (CERN European Organization for Nuclear Research)*. European Organization for Nuclear Research. https://doi.org/10.5281/zenodo.18890761

Balasubramanian, V. (2026a). VERBA: Verifiable Behaviour Architecture: A Proposed Closed-Loop Governance Framework for Heterogeneous Systems. In *Zenodo (CERN European Organization for Nuclear Research)*. European Organization for Nuclear Research. https://doi.org/10.5281/zenodo.19292644

Balasubramanian, V. (2026b). VERBA: Verifiable Behaviour Architecture: A Proposed Closed-Loop Governance Framework for Heterogeneous Systems. In *Zenodo (CERN European Organization for Nuclear Research)*. European Organization for Nuclear Research. https://doi.org/10.5281/zenodo.19292645

Fradelos, G. (2026a). Finance-Grade Assurance for Agentic AI: Verifiable Governance, Systemic Risk Mitigation, and Sustainability/Compute Accounting Architecture for banks, insurers, and major financial services providers. In *Zenodo (CERN European Organization for Nuclear Research)*. European Organization for Nuclear Research. https://doi.org/10.5281/zenodo.18213960

Fradelos, G. (2026b). Finance-Grade Assurance for Agentic AI: Verifiable Governance, Systemic Risk Mitigation, and Sustainability/Compute Accounting Architecture for banks, insurers, and major financial services providers. In *Zenodo (CERN European Organization for Nuclear Research)*. European Organization for Nuclear Research. https://doi.org/10.5281/zenodo.18213959

Frimpong, V. (2026a). Model Monoculture Risk: Systemic AI Convergence in Banking and Financial Markets. In *Preprints.org*. https://doi.org/10.20944/preprints202603.0393.v1

Frimpong, V. (2026b). Model Monoculture Risk: Systemic AI Convergence in Banking and Financial Markets. In *Figshare*. Figshare (United Kingdom). https://doi.org/10.6084/m9.figshare.31531168

Frimpong, V. (2026c). Model Monoculture Risk: Systemic AI Convergence in Banking and Financial Markets. In *Figshare*. Figshare (United Kingdom). https://doi.org/10.6084/m9.figshare.31531168.v1

Gandhi, A. K. (2025). Beyond Regulation: Proactive Financial Risk Management in an AI-Driven World. *International Journal of Artificial Intelligence Data Science and Machine Learning*, *6*, 156–163. https://doi.org/10.63282/3050-9262.ijaidsml-v6i1p117

Gong, H. (2026a). AI Agents in Financial Markets: Architecture, Applications, and Systemic Implications. In *arXiv (Cornell University)*. Cornell University. https://doi.org/10.48550/arxiv.2603.13942

Gong, H. (2026b). AI Agents in Financial Markets: Architecture, Applications, and Systemic Implications. *ArXiv.Org*. http://arxiv.org/abs/2603.13942

Han, S. (2025). Research on Herding in AI Trading. *Advances in Economics Management and Political Sciences*, *221*(1), 15–27. https://doi.org/10.54254/2754-1169/2025.bj28129

Kurshan, E., Balch, T., & Byrd, D. R. (2025). The Agentic Regulator: Risks for AI in Finance and a Proposed Agent-based Framework for Governance. In *ArXiv.org*. https://doi.org/10.48550/arxiv.2512.11933

McClellan, M. (2025). AI and Financial Fragility: A Framework for Measuring Systemic Risk in Deployment of Generative AI for Stock Price Predictions. *Journal of Risk and Financial Management*, *18*(9), 475–475. https://doi.org/10.3390/jrfm18090475

Nwachukwu, P. S., Chima, O. K., & Okolo, C. H. (2025). The artificial intelligence governance framework for finance: A control-by-design approach to algorithmic decision-making in accounting. *Finance & Accounting Research Journal*, *7*(8), 350–379. https://doi.org/10.51594/farj.v7i8.2016

Ogbuonyalu, U. O., Abiodun, K., Dzamefe, S., Vera, E. N., Oyinlola, A., & Emmanuel, I. (2024). Assessing Artificial Intelligence Driven Algorithmic Trading Implications on Market Liquidity Risk and Financial Systemic Vulnerabilities. *International Journal of Scientific Research and Modern Technology.*, 18–21. https://doi.org/10.38124/ijsrmt.v3i4.433

Ospan, A. (2026). The Mirror Trap: Why AI Convergence is the Next Systemic Threat to Global Finance. *Zenodo (CERN European Organization for Nuclear Research)*. https://doi.org/10.5281/zenodo.18419401

Oyewale, K., & Kipchumba, H. (2026). Algorithmic stewardship: Institutional investors, artificial intelligence and systemic risk. *World Journal of Advanced Research and Reviews*, *29*(1), 934–941. https://doi.org/10.30574/wjarr.2026.29.1.0110

Sabkha, S., & Jbir, H. (2025). Artificial Intelligence, Algorithmic Herding and Systemic Fragility in Financial Markets. In *SPIRE - Sciences Po Institutional REpository*. https://hal.science/hal-05556083

Shrinivas, S. K., Shetty, P., Sophia, J., & Gnanakumar, M. (2025). AI as a Systemic Risk Amplifier in High-Frequency Trading: Presenting a Conceptual Regulatory Framework. *Journal of Business Communication & Technology*, 58–67. https://doi.org/10.56632/bct.2025.4205

Yisra'el, Y. (2026a). THEMIS: Transactional Harmonic Enforcement & Macroprudential Integrity System — A 12-Part Formulation for Deterministic Financial Stability [Data set]. In *Zenodo (CERN European Organization for Nuclear Research)*. European Organization for Nuclear Research. https://doi.org/10.5281/zenodo.18779013

Yisra'el, Y. (2026b). THEMIS: Transactional Harmonic Enforcement & Macroprudential Integrity System — A 12-Part Formulation for Deterministic Financial Stability [Data set]. In *Zenodo (CERN European Organization for Nuclear Research)*. European Organization for Nuclear Research. https://doi.org/10.5281/zenodo.18779014

Zhou, Y., & Papageorgiou, O. (2025). *AI-DeFi Systemic Risk: A Complex Adaptive System Framework*. 1–9. https://doi.org/10.1109/brains67003.2025.11302949