



International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 04, April 2026)

The Impact of Artificial Intelligence and Machine Learning on Financial Forecasting and Asset Pricing: A Comprehensive Review of Methodological Advances, Empirical Evidence, Systemic Risks, and Regulatory Challenges

Dr. Pragya Dheer

Associate Professor, Department of Management, University of Kota, Kota, India

Abstract— The integration of artificial intelligence (AI) and machine learning (ML) into financial markets represents one of the most consequential methodological shifts in the history of quantitative finance. This paper provides a comprehensive, multi-disciplinary review of the mechanisms through which AI and ML are reshaping financial forecasting and asset pricing. Drawing on a synthesis of recent empirical literature, we examine the architectural innovations driving performance gains — including recurrent neural networks, transformer models, graph neural networks, and ensemble methods — alongside the growing role of alternative data sources such as satellite imagery, natural language processing of news and social media, and geolocation signals. We further analyse the theoretical implications for market efficiency, the Adaptive Market Hypothesis, and established asset pricing frameworks including the Capital Asset Pricing Model and factor models. Our review documents substantial evidence of superior forecasting accuracy in non-linear and high-dimensional environments, yet we identify critical limitations: model opacity, data quality dependencies, look-ahead bias in backtesting, and signal erosion upon deployment at scale. Most significantly, we evaluate the emerging systemic risks associated with widespread AI adoption, including model herding, crisis acceleration, emergent algorithmic coordination, and third-party concentration risk. We conclude with an assessment of the evolving regulatory landscape and argue that the disciplinary gap between algorithmic capability and governance infrastructure poses the foremost challenge for financial stability in the near term.

Keywords— machine learning; artificial intelligence; financial forecasting; asset pricing; systemic risk; explainability; regulatory oversight; deep learning; natural language processing.

I. INTRODUCTION

The quantitative revolution in finance, broadly dated from the mid-twentieth century, yielded a succession of foundational frameworks — Markowitz's mean-variance optimisation (1952), the Capital Asset Pricing Model (Sharpe, 1964; Lintner, 1965; Mossin, 1966), the Black-Scholes option pricing model (1973), and the Efficient Market Hypothesis (Fama, 1970) — each premised on the assumption that human analysts, armed with statistical tools and sound economic theory, could extract reliable signals from financial data. For several decades, this paradigm, operationalised through ordinary least squares regression, autoregressive integrated moving average (ARIMA) models, and the generalised autoregressive conditional heteroskedasticity (GARCH) family, constituted the methodological frontier of financial forecasting.

The emergence of machine learning as a practical engineering and scientific discipline, catalysed by advances in computational hardware, the proliferation of electronic financial data, and breakthroughs in algorithmic design, has fundamentally disrupted this equilibrium. Academic publication in AI applied to finance has expanded exponentially: from a modest average of approximately five articles per year between 1993 and 2017, the field now generates more than one hundred peer-reviewed contributions annually (Bahoo et al., 2024). The financial services industry allocated roughly \$35 billion toward AI projects in 2023 alone, with the global AI-in-finance market projected to reach \$190 billion by 2030 at a compound annual growth rate of approximately 30 percent (NVIDIA, 2024).



This paper is organised as follows. Point 2 surveys the principal machine learning architectures deployed in financial contexts. Point 3 analyses their application to financial forecasting tasks. Point 4 examines the implications for asset pricing theory and practice. Point 5 considers the role of alternative and unstructured data. Point 6 addresses the systemic risks and market stability concerns that widespread adoption engenders. Point 7 reviews the evolving regulatory landscape. Point 8 identifies open research questions, and Point 9 concludes.

II. MACHINE LEARNING ARCHITECTURES IN FINANCE: A TECHNICAL SURVEY

II.1 Classical and Ensemble Methods

The initial wave of ML applications in finance drew on support vector machines (SVMs) and artificial neural networks (ANNs), which were found to consistently outperform traditional statistical models in point-prediction accuracy for stock returns (Gandhmal & Kumar, 2019; Kumbure et al., 2022). Gradient-boosted decision trees — particularly XGBoost (Chen & Guestrin, 2016) and its successor LightGBM — subsequently emerged as workhorses for structured tabular financial data, offering competitive predictive accuracy alongside relatively tractable hyperparameter tuning and built-in feature importance measures. Their resistance to overfitting on moderately-sized datasets and their computational efficiency rendered them especially suitable for credit-risk modelling, fraud detection, and short-horizon equity return prediction.

II.2 Deep Learning: Recurrent and Sequential Architectures

The class of recurrent neural network (RNN) variants — most prominently Long Short-Term Memory (LSTM) networks (Hochreiter & Schmidhuber, 1997) and Gated Recurrent Units (GRU) — has demonstrated particular suitability for financial time-series tasks due to their capacity to retain information across variable-length sequences and capture temporal dependencies that fixed-lag autoregressive models cannot represent. Li and Bastos (2020) and Sezer et al. (2020) document the dominance of LSTM implementations across the deep learning literature, a finding consistent with the comprehensive review of 187 Scopus-indexed studies conducted by Caporin et al. (2024). Hybrid architectures combining convolutional layers for local pattern extraction with recurrent layers for temporal integration have shown particularly strong performance on multi-horizon forecasting tasks (Caporin et al., 2024).

II.3 Transformer Models and Attention Mechanisms

The introduction of the transformer architecture (Vaswani et al., 2017) and its associated attention mechanism represented a step-change in sequence modelling capacity. In finance, transformer-based models have been applied to earnings call transcript analysis, macroeconomic indicator forecasting, and high-frequency order book modelling. Their ability to process long-range dependencies without the vanishing-gradient limitations of standard RNNs has proven especially valuable in contexts where market conditions are shaped by narratives and policy signals that may be separated by substantial temporal intervals from their price effects (Sonkavde et al., 2023).

II.4 Graph Neural Networks and Relational Learning

Financial markets are fundamentally relational systems: asset prices are determined not merely by an instrument's own characteristics but by its embeddedness in webs of ownership, supply-chain dependency, cross-shareholding, and macroeconomic exposure. Graph neural networks (GNNs) operationalise this insight by representing assets as nodes in a graph whose edges encode economic or statistical relationships, permitting the model to learn how shocks propagate through the system. Empirical applications have demonstrated superior performance during market dislocations, when standard correlation-based assumptions break down but structural network relationships persist (Sonkavde et al., 2023; Frontiers AI, 2025). This represents a meaningful advance beyond the static factor-based covariance structures that underpin classical portfolio construction.

II.5 Reinforcement Learning and Execution Optimisation

Reinforcement learning (RL) — in which an agent learns a policy by interacting with an environment and maximising a cumulative reward signal — has found traction in trade execution, portfolio rebalancing, and market-making problems. Unlike supervised methods, which require labelled historical data, RL agents can adapt to changing market microstructure conditions. Reported performance improvements of 17–23 percent reduction in implementation shortfall relative to traditional Volume-Weighted Average Price (VWAP) algorithms have been documented in institutional deployments, though rigorous out-of-sample validation remains sparse in the published literature.

Table I:
TRADITIONAL VS. MACHINE LEARNING APPROACHES TO FINANCIAL MODELLING

Dimension	Traditional Methods	ML / AI Methods
Primary models	OLS, ARIMA, GARCH, Fama–French factors	LSTM, GRU, XGBoost, Transformers, GNNs
Data dimensionality	Low–medium; structured only	Very high; structured + unstructured
Non-linearity	Limited (log-linearisation required)	Natively captured
Alternative data	Rarely incorporated	NLP, satellite, geolocation, social media
Interpretability	High (closed-form coefficients)	Low–medium; XAI methods improving
Forecasting accuracy	Baseline; well-understood failure modes	Superior in non-linear regimes; overfitting risk
Regulatory acceptance	High	Evolving; model-risk guidance limited

Note: GNN = Graph Neural Network; XAI = Explainable Artificial Intelligence; NLP = Natural Language Processing.

III. APPLICATIONS TO FINANCIAL FORECASTING

III.1 Equity Return and Volatility Prediction

The preponderance of the empirical literature documents meaningful gains in forecasting accuracy when ML methods are applied to equity return prediction relative to linear factor benchmarks. A synthesis of 187 studies spanning 2020 to 2024 confirms the dominance of RNN-based architectures across stock, index, foreign exchange, commodity, bond, cryptocurrency, and volatility forecasting tasks (Caporin et al., 2024).

AI-based predictive analytics has demonstrated a 37 percent improvement in forecasting accuracy alongside a 42 percent reduction in decision latency in financial applications compared to traditional methods, with multimodal systems that integrate market data with alternative sources exhibiting 28–34 percent superior performance in anticipating market movements, particularly during volatile periods.

Nevertheless, reported performance metrics must be interpreted with care. The sensitivity of widely used measures such as Mean Absolute Percentage Error (MAPE) to volatility regimes remains underexplored, and most published results reflect in-sample or rolling-window backtests rather than live deployment, where the act of exploitation may itself erode the anomaly being traded (Khattak et al., 2023). A persistent methodological concern is look-ahead bias: inadvertent contamination of training data with information pioneering works in the 1990s used neural networks to predict quarterly earnings (Callen et al., 1996) and corporate bankruptcy (Jo & Han, 1996), establishing a tradition that has grown dramatically in sophistication. Contemporary models process regulatory filings, management discussion and analysis sections, and audio sentiment features from earnings calls, integrating these unstructured signals with traditional fundamental variables. BlackRock's Aladdin platform reportedly analyses more than 5,000 earnings call transcripts per quarter and over 6,000 broker reports daily — a scale of processing categorically beyond human unavailable at the prediction date, which inflates apparent accuracy figures.

III.2 Macroeconomic and Corporate Earnings Forecasting

Beyond asset returns, ML methods have been applied to macroeconomic indicator forecasting, earnings per share prediction, financial distress classification, and bankruptcy prediction. The latter application was among the earliest: analyst capacity (NVIDIA, 2024).

IV. IMPLICATIONS FOR ASSET PRICING THEORY AND PRACTICE

IV.1 Challenging the Factor Model Paradigm

The dominant paradigm in empirical asset pricing since the early 1990s has been the factor model, originating with the Fama and French (1993) three-factor framework and subsequently extended through momentum (Carhart, 1997), investment and profitability factors (Fama & French, 2015), and a proliferating zoo of additional characteristics documented in the published literature.



A central contribution of ML to asset pricing has been to subject this literature to rigorous, data-driven evaluation. Research applying ML models to over 166 previously identified asset pricing characteristics finds that ensemble and deep learning methods provide superior generalised predictions across multiple return measures, with momentum and trading-based features emerging as particularly robust signals (Gu, Kelly & Xiu, 2020).

These results have several theoretical implications. First, the apparent success of ML methods in predicting returns from a large characteristic set is consistent with the Adaptive Market Hypothesis (Lo & Zhang, 2024; Lo, 2004), which posits that markets are not statically efficient but dynamically evolve as participants adapt their strategies. Second, the ability of ML to identify non-linear interactions among characteristics suggests that the standard linear factor model may be a first-order approximation of a more complex pricing kernel, whose full structure is better captured by flexible function approximators than by pre-specified parametric forms.

IV.2 Alternative Data and Information Asymmetry

Classical theories of market efficiency (Fama, 1970) and asymmetric information (Akerlof, 1978; Stiglitz, 2000) were articulated in environments where the primary information inputs to pricing were financial statements, earnings announcements, and observable macroeconomic data. The emergence of alternative data — satellite imagery of retail parking lots and cargo vessels, geolocation data tracking consumer footfall, credit card transaction aggregates, and social media sentiment — has qualitatively altered the information landscape. Shiller (2020) argues that the type of information that investors process is critical to understanding market efficiency; Big Data substantially expands both the volume and variety of processable information (Hasan et al., 2020).

Research by Ahern and Peress (2023) demonstrates that financial and social media play a critical role in transmitting information between market participants, while Sun et al. (2024) show that social media platforms materially influence investor decisions through rapid information dissemination. Natural Language Processing enables algorithmic sentiment extraction from these sources at millisecond latency, creating an information advantage for well-resourced institutions that compounds existing concerns about market fairness and equal access.

IV.3 Discovery of New Pricing Anomalies

Despite the efficiency concerns noted above, the integration of ML into asset pricing research has demonstrably produced commercially viable investment strategies through the discovery of previously unknown return predictors and non-linear interactions among known ones. The recursive application of ML to identify signals, construct portfolios, and evaluate out-of-sample performance has become a standard component of quantitative investment research. This process, however, raises a fundamental epistemic challenge: the more aggressively researchers mine the data for patterns, the greater the risk that identified anomalies represent statistical artefacts rather than genuine compensation for priced risk.

V. GENERATIVE AI AND LARGE LANGUAGE MODELS IN FINANCE

The emergence of Large Language Models (LLMs) — most prominently the GPT family, Claude, and domain-specific variants trained on financial corpora — introduces a qualitatively distinct category of AI capability. The financial LLM literature organises around three strands: AI as an analytical tool that automates tasks including information extraction, regulatory filing summarisation, and financial forecasting; AI as an external shock that affects market efficiency and investor behaviour by lowering the cost of information processing; and AI as an economic agent capable of simulating financial reasoning and decision-making (Cao, 2023).

Practical deployments already in operation include earnings call analysis, automated compliance documentation, regulatory capital computation, robo-advisory portfolio construction, and customer-facing chatbot interfaces. Bi, Deng, and Xiao (2024) examine ChatGPT's potential for financial forecasting, finding meaningful capacity for zero-shot qualitative assessment of earnings release sentiment and corporate event classification, while documenting limitations in quantitative precision and susceptibility to hallucination. Liu (2024) surveys the broader landscape of financial AI architectures and identifies the integration of retrieval-augmented generation with structured financial databases as a promising direction for improving factual grounding.



VI. SYSTEMIC RISKS AND FINANCIAL STABILITY IMPLICATIONS

VI.1 Model Herding and Homogeneity

Perhaps the most consequential systemic concern arising from the mass adoption of ML in financial markets is the risk of model herding: the tendency for independently developed models, trained on similar datasets using similar architectures, to generate correlated predictions and, consequently, correlated trading behaviour. This dynamic is not merely theoretical. Rule-based algorithmic trading was implicated in cascade dynamics during the 1987 US equity market crash; with ML models exhibiting greater sophistication but also greater statistical similarity across institutions, the potential for synchronised responses to common shocks is structurally elevated (IOSCO, 2021). When a shock occurs, AI systems rapidly and simultaneously parse the same streams of market, macroeconomic, and news data, potentially executing correlated responses before human supervisors register an abnormal market move (IMF, 2025).

VI.2 Emergent Algorithmic Coordination

A more subtle and theoretically troubling risk concerns emergent coordination among AI agents. Simulation research has demonstrated that AI-driven trading agents can achieve near-cartel-like outcomes in terms of joint profitability without explicit collusion programming — through spontaneous behavioural convergence arising from repeated interactions in the same market environment. This phenomenon, sometimes termed emergent communication, poses a challenge to competition law and market regulation frameworks designed to identify and prosecute intentional collusive behaviour, since no such intent need be present (Calvano et al., 2020).

VI.3 Third-Party Concentration Risk

The cloud-computing infrastructure underpinning most institutional ML deployment is highly concentrated among a small number of providers. A May 2025 study by the United States Government Accountability Office (GAO) warned that financial instability may arise from this structural dependency: a failure at a major AI infrastructure provider could simultaneously impair the operational capacity of a large fraction of the financial sector, creating a novel and underappreciated channel of systemic risk (GAO, 2025). This concern is compounded by the reliance of many institutions on common third-party model providers and data vendors, which further homogenises the risk landscape.

VI.4 The Black Box Problem and Model Risk

Regulatory frameworks for model risk management, typified by the Federal Reserve's SR 11-7 guidance in the United States, require that models be validated, documented, and understood by qualified personnel. Deep neural networks present a fundamental challenge to these requirements: their internal representations are high-dimensional, non-linear, and not reducible to interpretable economic coefficients without specialist explainability methods (Mohsin & Nasim, 2025). The field of Explainable AI (XAI) — encompassing SHAP values, LIME, and attention visualisation — is actively developing tools to address this opacity, but a meaningful gap remains between algorithmic capability and regulatory-standard interpretability. Financial data quality compounds the problem: training data that is noisy, incomplete, or historically biased can produce models that appear accurate on validation samples while embedding structural vulnerabilities (Bi et al., 2024).

VII. REGULATORY AND GOVERNANCE RESPONSES

Regulatory responses to AI adoption in financial markets remain fragmented, jurisdiction-dependent, and, in most cases, lagging significantly behind the pace of technological change. The International Organisation of Securities Commissions (IOSCO) has documented that some jurisdictions are adapting existing regulatory frameworks to AI activities — applying conduct-of-business, model risk, and operational resilience rules developed for earlier generations of algorithmic tools — while others are developing bespoke frameworks to address AI-specific challenges (IOSCO, 2021). The European Union's AI Act, which entered force in 2024, introduces a risk-based classification system that designates AI systems used in credit scoring, insurance risk assessment, and certain financial advisory functions as high-risk, requiring conformity assessments and human oversight obligations.

In the United States, the Securities and Exchange Commission launched an AI Task Force in August 2025 charged with enhancing oversight of AI-related fraud, cybersecurity, and governance. The SEC's AI roundtable convened in May 2025 focused specifically on the tension between innovation, market integrity, and investor protection. The Federal Reserve and the Office of the Comptroller of the Currency have updated model risk management guidance to require more rigorous validation of ML models used in credit underwriting and risk management.

Across jurisdictions, a common regulatory theme is the insistence on human accountability: that an identifiable human decision-maker must remain responsible for AI-generated outputs in consequential financial contexts, regardless of the opacity of the underlying algorithm.

VIII. OPEN RESEARCH QUESTIONS AND FUTURE DIRECTIONS

Several significant gaps in current understanding merit priority in future research. First, the relationship between ML adoption and market efficiency remains underspecified. If ML uniformly accelerates the incorporation of information into prices, it should reduce predictability and narrow arbitrage opportunities; yet the simultaneous expansion of alternative data sources continuously creates new informational asymmetries. Empirically distinguishing these effects requires careful identification strategies.

Second, the deployment gap between academic model development and live trading strategy implementation represents a systematic limitation of the published literature. Most reported forecasting results are based on simulated backtests; rigorous evaluation of live performance — accounting for transaction costs, market impact, and signal decay — is substantially rarer. Liu et al. (2023) identify this as a critical research priority.

Third, federated learning and quantum machine learning represent frontier methodological directions. Federated learning enables model training across decentralised financial data without centralising sensitive client information, addressing both privacy and concentration concerns. Quantum machine learning, while largely pre-commercial, promises computational advantages for portfolio optimisation and derivative pricing in exponentially large state spaces (Chang et al., 2024; Najem et al., 2024).

Fourth, the governance of AI in finance requires interdisciplinary theoretical development that does not currently exist in adequate form. Economic theory of mechanism design, computer science research on AI safety and alignment, and financial regulation scholarship must be more deeply integrated to produce regulatory frameworks capable of governing the risks identified in Section 6.

IX. CONCLUSION

Artificial intelligence and machine learning have irrevocably altered the information architecture and methodological toolkit of financial markets.

The empirical evidence reviewed in this paper is documents substantial and genuine advances in forecasting accuracy, the identification of previously undiscovered return predictors, and the operational efficiency of trade execution. These gains are concentrated in domains characterised by high dimensionality, non-linearity, and the availability of alternative, unstructured data sources — conditions under which classical parametric methods are fundamentally limited.

At the same time, the rapid and uneven adoption of these technologies introduces systemic vulnerabilities that existing regulatory frameworks are ill-equipped to address. Model herding, crisis acceleration, emergent algorithmic coordination, and third-party concentration risk represent structurally novel channels through which AI could amplify, rather than dampen, financial instability. The opacity of deep learning models undermines established model-risk governance frameworks, and the pace of algorithmic development continues to outstrip both regulatory capacity and academic understanding.

The central conclusion of this review is not that AI adoption in finance is net-harmful — the efficiency and resilience benefits are real and substantive — but that the discipline-wide underinvestment in governance infrastructure relative to model development constitutes the foremost challenge for financial stability in the near term. Closing this gap requires coordinated action across regulators, industry participants, and academic researchers, as well as sustained intellectual investment in the interdisciplinary domain of financial AI governance.

REFERENCES

- [1] Ahern, K. R., & Peress, J. (2023). Information network dynamics in financial markets. *Journal of Financial Economics*, 148(3), 654–682.
- [2] Akerlof, G. A. (1978). The market for lemons: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3), 488–500.
- [3] 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(2), 1–25.
- [4] Bi, S., Deng, T., & Xiao, J. (2024). The role of AI in financial forecasting: ChatGPT's potential and challenges. *arXiv preprint arXiv:2411.13562*.
- [5] Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), 637–654.
- [6] Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review*, 110(10), 3267–3297.
- [7] Caporin, M., et al. (2024). Deep learning for financial forecasting: A review of recent trends. *Expert Systems with Applications*. doi:10.1016/j.eswa.2025.08822.



International Journal of Recent Development in Engineering and Technology
Website: www.ijrdet.com (ISSN 2347-6435 (Online) Volume 15, Issue 04, April 2026)

- [8] Carhart, M. M. (1997). On persistence in mutual fund performance. *Journal of Finance*, 52(1), 57–82.
- [9] Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794.
- [10] Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2), 383–417.
- [11] Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1), 3–56.
- [12] Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1), 1–22.
- [13] Gandhmal, D. P., & Kumar, K. (2019). Systematic analysis and review of stock market prediction techniques. *Computer Science Review*, 34, 100190.
- [14] Government Accountability Office. (2025, May). AI in financial services: Risks to stability from third-party concentration. GAO-25-106851.
- [15] Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *Review of Financial Studies*, 33(5), 2223–2273.
- [16] Hasan, M. M., et al. (2020). Big data analytics and financial markets: A systematic review. *Journal of Risk and Financial Management*, 13(10), 225.
- [17] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
- [18] International Organisation of Securities Commissions. (2021). The use of artificial intelligence and machine learning by market intermediaries and asset managers. IOSCO Report.
- [19] Jensen, M. C., & Meckling, W. H. (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4), 305–360.
- [20] Khattak, M. A. K., et al. (2023). Machine learning methods in financial forecasting: A review of recent advances. *International Journal of Financial Studies*, 11(2), 68.
- [21] Kumbure, M. M., et al. (2022). Machine learning techniques and data for stock market forecasting: A literature review. *Expert Systems with Applications*, 197, 116659.
- [22] Kureljusic, M., & Karger, E. (2024). Forecasting in financial accounting with artificial intelligence: A systematic literature review and future research agenda. *Journal of Applied Accounting Research*, 25(1), 81–104.
- [23] Li, B., & Bastos, G. (2020). Stock market forecasting using deep learning: A review. *Expert Systems with Applications*, 152, 113404.
- [24] Lintner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics*, 47(1), 13–37.
- [25] Liu, J. (2024). A survey of financial AI: Architectures, advances and open challenges. *arXiv preprint arXiv:2411.12747*.
- [26] Liu, Y., Wang, H., & Zhang, X. (2023). Enhancing financial forecasting with machine learning: A comprehensive review. *Journal of Financial Analytics*, 15(2), 45–62.
- [27] Lo, A. W. (2004). The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management*, 30(5), 15–29.
- [28] Lo, A. W., & Zhang, R. (2024). Adaptive markets and the dynamics of financial innovation. *Annual Review of Financial Economics*, 16, 1–28.
- [29] Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1), 77–91.
- [30] Mohsin, M. T., & Nasim, N. B. (2025). Explaining the unexplainable: A systematic review of explainable AI in finance. *arXiv preprint arXiv:2503.05966*.
- [31] Mossin, J. (1966). Equilibrium in a capital asset market. *Econometrica*, 34(4), 768–783.
- [32] NVIDIA Financial Services. (2024). State of AI in financial services. NVIDIA Corporation white paper.
- [33] Sezer, O. B., Gudelek, M. U., & Ozbayoglu, A. M. (2020). Financial time series forecasting with deep learning: A systematic literature review. *Applied Soft Computing*, 90, 106181.
- [34] Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19(3), 425–442.
- [35] Shiller, R. J. (2020). *Narrative economics: How stories go viral and drive major economic events*. Princeton University Press.
- [36] Sonkavde, G., et al. (2023). Forecasting stock market prices using machine learning and deep learning models: A systematic review, performance analysis and discussion of implications. *International Journal of Financial Studies*, 11(3), 94.
- [37] Stiglitz, J. E. (2000). The contributions of the economics of information to twentieth century economics. *Quarterly Journal of Economics*, 115(4), 1441–1478.
- [38] Sun, Z., et al. (2024). Social media and investor decisions: Evidence from financial market platforms. *Journal of Financial and Quantitative Analysis*, 59(1), 212–238.
- [39] Vaswani, A., et al. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.