Machine Learning Approaches for Predicting Stock Market Volatility in Emerging Markets

Sarah Jenkins
Department of Economics, Iowa State University
sjenkins@iastate.edu

Marcus Vance
Department of Computer Science, University of Delaware
mvance@udel.edu

Gerald Langston
Department of Financial Engineering, Stevens Institute of Technology
g.langston@stevens.edu

Abstract

Predicting stock market volatility in emerging markets represents a critical challenge for global financial stability, portfolio management, and systemic risk mitigation. Traditional econometric models often fail to capture the highly nonlinear, non-stationary, and regime-shifting characteristics inherent in these developing financial ecosystems. This paper provides a comprehensive, system-level investigation into the deployment of machine learning architectures for volatility forecasting in emerging economies. We explore the architectural trade-offs, infrastructure requirements, and structural vulnerabilities associated with transitioning from classical statistical frameworks to advanced computational paradigms, including deep recurrent networks, ensemble methods, and hybrid neural-econometric systems. Beyond algorithmic performance, this study emphasizes the socio-technical dimensions of algorithmic deployment, analyzing the critical infrastructure needed to handle asynchronous data streams, fragmented regulatory reporting, and low-liquidity regimes. We examine the structural trade-offs between model interpretability and predictive capacity, highlighting how opaque deep learning systems can inadvertently amplify systemic risks during periods of market stress. Furthermore, the paper addresses governance challenges, data fairness, and the policy implications of widespread algorithmic trading in environments characterized by weak institutional safeguards and high susceptibility to capital flight. By evaluating these systems through an interdisciplinary lens that unites computer science, financial economics, and public policy, we outline a robust, sustainable framework for integrating machine learning into the regulatory and operational fabrics of emerging financial markets.

Keywords:

Financial Volatility, Machine Learning, Emerging Markets, Socio-Technical Infrastructure, Systemic Risk, Algorithmic Governance.

1. Introduction

The predictability of financial market volatility remains a cornerstone of modern financial economics, asset pricing, risk management, and regulatory oversight. Volatility serves as a primary metric for uncertainty, dictating capital allocation strategies, portfolio optimization frameworks, and the pricing of complex derivative instruments. While developed financial markets operate under conditions of high liquidity, robust institutional frameworks, and comprehensive informational efficiency, emerging markets present a fundamentally distinct operational environment. Emerging markets are characterized by heightened sensitivity to global macroeconomic shocks, structural vulnerabilities, fragmented regulatory frameworks, and frequent regime shifts. Consequently, understanding and forecasting volatility within these ecosystems is both exceptionally vital and distinctively challenging for global investors and domestic policy makers alike.

Historically, the academic community and financial practitioners have relied on classical econometric models to capture the dynamics of financial volatility. Frameworks such as autoregressive conditional heteroskedasticity and generalized autoregressive conditional heteroskedasticity have long dominated the literature. These models operate on specific structural assumptions regarding the underlying data-generating process, frequently presupposing linearity, stationarity, or particular parametric distributions for asset returns. While these classical approaches offer elegant mathematical tractability and high interpretability, they increasingly struggle to accommodate the structural anomalies present in modern emerging economies. The rapid digitization of financial systems, coupled with the introduction of high-frequency trading and alternative data streams, has introduced degrees of nonlinearity and non-stationarity that violate the core assumptions of classical financial econometrics.

In response to these limitations, machine learning approaches have emerged as a powerful paradigm for financial forecasting. Machine learning frameworks excel at identifying intricate, high-dimensional patterns within vast datasets without requiring rigid, a priori assumptions about the statistical distribution of the underlying data. Technologies ranging from support vector machines and ensemble learning methods to deep neural networks have demonstrated remarkable proficiency in mapping complex historical dependencies to future volatility states. However, the adoption of machine learning in emerging markets cannot be viewed solely as a technical or algorithmic upgrade. Instead, it constitutes a profound shift in the socio-technical infrastructure of financial systems, introducing novel dependencies, operational risks, and governance challenges that demand rigorous academic scrutiny.

This research paper provides a systemic, long-form exploration of machine learning approaches for predicting stock market volatility in emerging markets. Rather than focusing exclusively on localized predictive performance metrics, this study situates machine learning models within the broader socio-technical and regulatory ecosystems of developing economies. We investigate the structural trade-offs inherent in model selection, the infrastructure requirements necessary for sustainable deployment, and the broader implications for financial stability and public policy. Through this comprehensive analysis, we aim to bridge the gap between computational innovation and macroeconomic governance, offering a framework for the resilient integration of artificial intelligence into emerging financial systems.

2. Theoretical Foundations of Volatility in Emerging Economies

To evaluate the efficacy of machine learning approaches, it is first necessary to establish the unique theoretical and structural characteristics of volatility within emerging market frameworks. In financial economics, volatility is typically conceptualized as a latent variable representing the intensity of asset price fluctuations over a specific time horizon. In highly developed markets, volatility dynamics are largely driven by continuous information inflows, corporate earnings announcements, and systematic macroeconomic indicators. The market mechanism processes this information rapidly, leading to efficient price discovery.

In emerging economies, the drivers of volatility are considerably more complex and less orderly. These markets are prone to severe structural breaks caused by sudden geopolitical shifts, sovereign debt reconfigurations, currency devaluations, and abrupt regulatory interventions. The phenomenon of volatility clustering, where large shocks are followed by further large shocks, is exceptionally pronounced in emerging markets due to information asymmetry and herd behavior among domestic retail investors. Furthermore, because these markets frequently rely heavily on foreign institutional capital, they are highly susceptible to sudden capital flight triggered by monetary policy changes in developed nations. This vulnerability introduces external, non-domestic variables into the volatility equation, necessitating models that can integrate global macroeconomic signals with domestic market data.

Another defining characteristic of emerging market volatility is the presence of asymmetric effects, often referred to as the leverage effect. This describes the tendency for negative asset returns to generate a greater increase in future volatility than positive returns of an equal magnitude. In emerging economies, this asymmetry is compounded by thinner trading volumes and lower market liquidity. When negative shocks occur, liquidity can evaporate rapidly, leading to extreme bid-ask spreads and discontinuous price jumps. Classical econometric models often attempt to address this through specialized variations, but these models remain limited by their rigid parametric formulations. Machine learning frameworks must therefore be designed to intrinsically capture these asymmetric, threshold-driven relationships without suffering from overfitting in data-scarce environments.

Finally, the institutional environment of emerging markets introduces significant structural noise into asset price series. Regulatory mandates may include sudden short-selling bans, price limits, or capital controls that abruptly alter market microstructure. These interventions create artificial boundaries and non-linearities in historical volatility data. A robust machine learning approach must not only model the organic economic interactions within the market but also adapt to these exogenous institutional constraints. Consequently, understanding the theoretical foundations of emerging market volatility requires an analytical shift away from purely statistical distributions toward a holistic appreciation of the socio-technical, regulatory, and macroeconomic forces at play.

3. Algorithmic Architecture and Model Selection Trade-Offs

Deploying machine learning for volatility prediction involves selecting from a wide array of algorithmic architectures, each presenting distinct trade-offs between computational complexity, predictive accuracy, and interpretability. The architectural spectrum ranges from shallow, tabular machine learning methods to deeply layered neural architectures and hybrid frameworks that blend econometrics with computational intelligence. Navigating this spectrum requires a careful assessment of the specific characteristics of emerging market data, which are often characterized by shorter historical time series, high noise-to-signal ratios, and frequent structural breaks.

Shallow machine learning models, such as random forests, gradient boosted trees, and support vector regression, offer substantial advantages in terms of robust training and resistance to overfitting. Gradient boosting frameworks excel at processing heterogeneous feature spaces that combine technical indicators, macroeconomic variables, and sentiment metrics. These models operate by sequentially correcting the errors of weak decision trees, making them highly effective at capturing non-linear relationships and interactions between features. Furthermore, tree-based ensembles provide inherent feature importance metrics, offering a degree of transparency that is highly valued by risk managers and regulatory bodies. However, these architectures struggle to inherently capture the sequential dependencies and long-term memory structures present in high-frequency financial time series.

To address temporal dependencies, deep recurrent neural network architectures, particularly long short-term memory networks and gated recurrent units, have become widely adopted. Long short-term memory networks utilize specialized gating mechanisms to regulate the flow of information over extended temporal horizons, allowing the model to retain historical volatility contexts across weeks or months while discarding short-term noise. This capability is uniquely valuable in emerging markets, where the memory of a major macroeconomic shock can influence investor behavior for prolonged periods. Despite their predictive prowess, deep recurrent networks demand massive datasets for optimal parameter tuning, rendering them highly susceptible to overfitting when applied to smaller emerging markets with limited historical data. Moreover, their computational training footprint is substantial, requiring dedicated infrastructure.

More recently, hybrid architectures that combine classical econometric models with deep learning frameworks have emerged as a sophisticated compromise. In these systems, a classical generalized autoregressive conditional heteroskedasticity model is utilized to capture the linear, baseline volatility dynamics, while a neural network is stacked on top to model the residual, highly non-linear components. This approach effectively anchors the machine learning model within established financial theory, ensuring that the system generates stable, economically viable predictions even when encountering unprecedented market conditions. The trade-off shifts from purely maximizing algorithmic accuracy to optimizing system resilience, ensuring that the model benefits from the flexible pattern recognition of deep learning while retaining the structural stability of classical econometrics.

4. Socio-Technical Infrastructure and Data Governance

The successful implementation of machine learning approaches for volatility prediction depends heavily on the underlying technical infrastructure and the socio-technical governance systems that manage data lifecycles. In emerging markets, data infrastructure is frequently fragmented, characterized by disparate reporting standards, varying levels of transparency, and technological disparities among market participants. Building a reliable predictive system requires an infrastructure capable of ingesting, cleaning, and synchronizing highly heterogeneous data streams under stringent time constraints.

Financial data in emerging markets often arrives via asynchronous channels, encompassing high-frequency order book data from domestic exchanges, lower-frequency macroeconomic disclosures from central banks, and unstructured text from local news outlets and social media. The infrastructure must implement robust data pipelines capable of handling missing values, adjusting for frequent corporate actions, and mitigating the impact of erroneous data entries resulting from manual reporting processes. This requires advanced data preprocessing layers that can dynamically impute missing information without introducing artificial patterns that could distort the machine learning model's perception of market volatility.

Furthermore, data governance frameworks within emerging economies face unique regulatory and operational constraints. Unlike developed markets governed by comprehensive frameworks, emerging markets may exhibit regulatory gaps regarding data privacy, algorithmic accountability, and institutional access. The governance infrastructure must ensure compliance with fluctuating local legal standards while simultaneously maintaining data integrity. Issues of data sovereignty also arise, particularly when domestic financial data is processed using cloud infrastructure located in foreign jurisdictions. This creates a critical tension between the computational efficiency of global cloud providers and the regulatory mandates of domestic oversight bodies.

To achieve sustainability, the socio-technical infrastructure must incorporate continuous monitoring and auditing mechanisms. Financial markets are dynamic systems where the underlying data-generating process is constantly evolving—a phenomenon known as concept drift. In emerging markets, concept drift can occur abruptly due to policy shifts or sudden capital reallocations. The infrastructure must therefore feature automated anomaly detection systems that flag when the incoming data distribution deviates significantly from the training baseline, signaling the need for immediate model recalibration. Without these architectural safeguards, the predictive system risks generating highly inaccurate forecasts during periods of structural transition, potentially exacerbating the very market instability it was designed to monitor.

5. Robustness, Generalization, and Systemic Vulnerabilities

The deployment of predictive machine learning models in financial systems introduces complex challenges related to robustness and generalization, alongside the potential for creating systemic vulnerabilities. A model that exhibits exceptional predictive accuracy during stable economic periods may fail catastrophically when subjected to the extreme tail-risk events that periodically characterize emerging economies. Ensuring the structural resilience of these models is paramount, as failure can lead to severe capital losses, mispriced risk, and broader economic destabilization.

Generalization is a particularly difficult objective in emerging markets due to the scarcity of historical data encompassing full macroeconomic cycles. A machine learning model trained on a five-year period of sustained economic growth will likely fail to generalize when confronted with a sudden currency crisis or a global pandemic. This limitation necessitates the use of advanced validation techniques, such as combinatorial purified cross-validation and non-stationary time-series splitting, to prevent data leakage and evaluate the model's performance across distinct economic regimes. Furthermore, models must be subjected to rigorous stress-testing frameworks where synthetic, extreme-loss scenarios are injected into the feature space to evaluate system behavior under duress.

Beyond individual model robustness, the widespread, homogenous adoption of machine learning approaches can create profound systemic vulnerabilities within financial ecosystems. When multiple major institutional investors and market makers deploy similar deep learning architectures trained on identical historical datasets, their predictive outputs tend to align. This convergence can induce highly correlated trading behaviors, leading to algorithmic herding. In periods of market stress, if multiple independent machine learning systems simultaneously predict an unprecedented surge in volatility, they may collectively execute automated risk-reduction strategies, such as mass asset liquidations or the widening of market-maker spreads.

This collective algorithmic response can trigger a self-fulfilling prophecy, where the automated actions designed to protect individual portfolios directly cause a catastrophic collapse in market liquidity. In thin emerging markets, where the total volume of capital is limited, such algorithmic feedback loops can be devastating. The system transitions from predicting volatility to actively generating it, creating an endogenous systemic risk that is extraordinarily difficult for human regulators to intercept. Consequently, the analysis of machine learning in financial markets must extend beyond individual algorithmic accuracy to consider the macroeconomic consequences of widespread system deployment.

6. Algorithmic Fairness, Market Manipulation, and Inclusion

The integration of machine learning into emerging financial markets intersects directly with issues of algorithmic fairness, market manipulation, and socio-economic inclusion. Financial markets are not merely abstract computational environments; they are socio-technical institutions that distribute wealth, influence borrowing costs, and shape national economic development. The introduction of highly sophisticated predictive technologies has the potential to alter the playing field, creating new disparities between technologically advanced global actors and local market participants.

Algorithmic fairness in this context relates to how predictive models evaluate and impact different sectors of the economy. Machine learning models trained on historical data may inherently capture and perpetuate institutional biases, such as under-allocating risk capacity to smaller, domestic firms in favor of large, multinational conglomerates. In emerging markets, where small and medium enterprises form the backbone of economic growth, a systemic bias within volatility forecasting models can artificially elevate the borrowing costs or depress the equity valuations of these critical entities. This occurs because the model may interpret the lower trading volumes and higher informational opacity of smaller firms as signals of unmanageable volatility risk, ignoring their underlying economic viability.

Furthermore, the vulnerability of emerging markets to informational asymmetries creates fertile ground for algorithmic market manipulation. Sophisticated actors can exploit the known vulnerabilities of machine learning models—such as their sensitivity to sudden, localized anomalies—by intentionally injecting misleading signals into alternative data streams, social media platforms, or thin order books. These adversarial attacks can deceive predictive systems, triggering automated trading responses that alter asset prices to the manipulator's advantage. In markets lacking robust surveillance infrastructure, identifying and prosecuting such technologically advanced forms of market manipulation is exceptionally difficult.

Finally, the technological divide risks exacerbating issues of financial inclusion. The development and deployment of high-performance machine learning models require substantial capital, computational infrastructure, and specialized human expertise—resources that are predominantly concentrated within international investment banks and global hedge funds. Local retail investors and domestic financial institutions in emerging markets may find themselves at a persistent informational disadvantage, unable to compete with the predictive precision of global algorithmic systems. This dynamic can lead to a colonial extraction of wealth from emerging financial ecosystems, where international capital utilizes superior computational power to extract profits while leaving local economies to bear the structural systemic risks.

7. Policy Implications and Algorithmic Governance

The systemic challenges, technical trade-offs, and socio-economic impacts highlighted throughout this study underscore the critical need for comprehensive regulatory oversight and algorithmic governance frameworks. Emerging market regulators face the daunting task of fostering technological innovation to enhance market efficiency while simultaneously implementing safeguards to protect financial stability, investor integrity, and national sovereignty. Achieving this balance requires moving away from reactive, post-hoc regulations toward proactive, technologically aligned governance structures.

Regulatory bodies must establish clear standards for model interpretability and explainability. Given that completely opaque black-box neural networks can mask underlying vulnerabilities and propagate systemic risks, regulators should mandate that machine learning models utilized by major financial institutions incorporate explainable artificial intelligence methodologies, such as Shapley additive explanations or local interpretable model-agnostic explanations. These frameworks allow risk managers and regulatory auditors to deconstruct individual volatility forecasts into their constituent feature attributions, ensuring that model predictions are grounded in verifiable economic fundamentals rather than statistical anomalies or data artifacts.

Furthermore, central banks and financial authorities in emerging markets should implement centralized algorithmic stress-testing centers. Before deploying automated trading or risk management systems at scale, institutional actors should be required to subject their machine learning models to standardized, simulated market stress scenarios designed by regulators. These scenarios should replicate historical crises, simulate sudden capital flights, and introduce synthetic liquidity freezes. By evaluating how diverse algorithmic architectures interact within a controlled, simulated ecosystem, regulators can identify potential algorithmic herding behaviors and systemic feedback loops before they manifest in real-world markets.

International cooperation also represents a vital pillar of effective governance. Because capital flows and data streams cross national boundaries, isolated domestic regulations are insufficient to address the globalized nature of algorithmic trading. Emerging market regulators must collaborate with international bodies, such as the International Organization of Securities Commissions and the Financial Stability Board, to establish unified standards for algorithmic transparency, data sharing, and cross-border market surveillance. Additionally, policies must be enacted to encourage domestic technological capacity, including investing in local educational infrastructure and computational resources. By building local expertise, emerging economies can reduce their dependence on external technological architectures and ensure that the deployment of machine learning aligns with sustainable national economic development goals.

8. Future Directions and Emerging Paradigms

As the field of financial machine learning evolves rapidly, several emerging technological paradigms are poised to reshape the landscape of volatility prediction in emerging markets. Tracking these forward-looking trends is essential for designing systems that remain resilient, adaptive, and sustainable over long-term horizons amidst continuous socio-technical transformations.

One of the most promising avenues of research lies in the integration of federated learning architectures. Traditional machine learning models require the centralization of data into a single repository, a requirement that frequently conflicts with data privacy regulations, competitive proprietary interests, and national data sovereignty mandates in emerging markets. Federated learning allows multiple financial institutions to collaboratively train a shared volatility forecasting model without ever exchanging their raw, localized transaction data. The model parameters are updated locally at each institution and then aggregated globally via a secure cryptographic protocol. This approach preserves data privacy and institutional confidentiality while enabling the collective model to benefit from a vastly expanded, heterogeneous dataset, significantly enhancing its ability to generalize across diverse market regimes.

Another major paradigm shift involves the incorporation of natural language processing and multimodal large language models to ingest and synthesize alternative data streams. Financial markets in emerging economies are heavily influenced by qualitative narratives, political announcements, and local regulatory updates that are poorly captured by structured numerical time series alone. Advanced language models can continuously parse local news publications, legislative dockets, and social media discourse in native languages, converting unstructured text into dense sentiment vectors. When integrated with numerical neural architectures via multimodal fusion layers, these textual features provide critical early-warning signals of impending policy shifts or socio-political instability, allowing the predictive system to anticipate volatility spikes long before they reflect in historical asset price series.

Finally, the exploration of quantum machine learning represents a long-term, transformative frontier for financial systems research. The optimization of massive neural network architectures and the execution of high-dimensional portfolio simulations represent immense computational challenges that stretch the limits of classical silicon-based hardware. Quantum computing architectures utilize quantum bits to perform complex, parallel computational operations at speeds unimaginable with classical systems. In the context of emerging markets, quantum machine learning algorithms could process high-frequency order book dynamics and global macroeconomic interactions simultaneously in real time, providing instantaneous, highly precise volatility forecasts and risk optimizations. While still in its infancy, preparing financial infrastructure for the eventual transition to quantum-resistant cryptography and quantum-enhanced analytics is a critical prerequisite for long-term systemic sustainability.

9. Conclusion

The deployment of machine learning approaches for predicting stock market volatility in emerging markets represents a profound socio-technical paradigm shift that extends far beyond the boundaries of algorithmic optimization. As demonstrated throughout this study, while advanced computational architectures—including deep recurrent networks, ensemble methods, and hybrid neural-econometric systems—offer unprecedented capacity to model the highly non-linear, non-stationary dynamics of developing economies, their introduction fundamentally alters the operational, structural, and regulatory risk profiles of financial ecosystems.

Successfully navigating this transition requires a holistic approach that bridges the gap between computational data science, financial economics, and public policy. The structural trade-offs between model complexity and interpretability must be managed to prevent the propagation of black-box vulnerabilities. Robust data infrastructure and stringent governance frameworks are essential to maintain data integrity and protect national sovereignty amid fragmented reporting environments. Furthermore, system designers and regulators must actively mitigate the systemic risks of algorithmic herding and market manipulation, ensuring that technological advancement does not come at the expense of financial stability, algorithmic fairness, and socio-economic inclusion.

Ultimately, machine learning should not be viewed merely as a tool for extracting short-term trading profits, but as a foundational infrastructure for enhancing macro-financial resilience. By implementing proactive regulatory frameworks, cultivating local technological capacity, and embracing emerging paradigms such as federated learning and multimodal data fusion, emerging economies can harness the power of artificial intelligence to build more efficient, transparent, and inclusive financial institutions. The future of global financial stability depends on our collective ability to design and govern these intelligent systems with equal measures of technological innovation and socio-technical responsibility.

References

1. Akaike, H. (1974). A new look at the statistical model identification. IEEE Transactions on Automatic Control, 19(6), 716-723.

2. Bandi, F. M., & Russell, J. R. (2006). Separating microstructure noise from volatility. Journal of Financial Economics, 79(3), 655-692.

3. Barndorff-Nielsen, O. E., & Shephard, N. (2002). Econometric analysis of realized volatility on high-frequency financial data. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 64(2), 253-280.

4. Bengio, Y., Ducharme, R., Vincent, P., & Jauvin, C. (2003). A neural probabilistic language model. Journal of Machine Learning Research, 3(Feb), 1137-1155.

5. Bera, A. K., & Higgins, M. L. (1993). ARCH models: Properties, estimation and testing. Journal of Economic Surveys, 7(4), 305-366.

6. Black, F. (1976). Studies of stock price changes. Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economical Statistics Section, 177-181.

7. Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, 31(3), 307-327.

8. Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5-32.

9. Brooks, C. (2019). Introductory econometrics for finance. Cambridge University Press.

10. Campbell, J. Y., Lo, A. W., & MacKinlay, A. C. (1997). The econometrics of financial markets. Princeton University Press.

11. 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.

12. Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. Quantitative Finance, 1(2), 223-236.

13. Dietterich, T. G. (2000). Ensemble methods in machine learning. Multiple Classifier Systems, 1-15.

14. Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. Econometrica, 50(4), 987-1007.

15. Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. The Journal of Finance, 25(2), 383-417.

16. Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. Annals of Statistics, 1189-1232.

17. Glosten, L. R., Jagannathan, R., & Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. The Journal of Finance, 48(5), 1779-1801.

18. Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.

19. Hansen, P. R., & Lunde, A. (2005). A forecast comparison of volatility models: Does GARCH (1, 1) choose the winner?. Journal of Applied Econometrics, 20(2), 251-273.

20. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735-1780.

21. Kariuki, D., & Mwangi, J. (2024). Algorithmic trading and capital flight in emerging markets. Journal of Financial Regulation and Compliance, 32(1), 45-62.

22. Kim, K. J. (2003). Financial time series forecasting using support vector machines. Neurocomputing, 55(1-2), 307-319.

23. Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. Advances in Neural Information Processing Systems, 30, 4765-4774.

24. Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. Econometrica, 59(2), 347-370.

25. Poon, S. H., & Granger, C. W. (2003). Forecasting volatility in financial markets: A review. Journal of Economic Literature, 41(2), 478-539.

26. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1135-1144.

27. Scholkopf, B., & Smola, A. J. (2002). Learning with kernels: Support vector machines, regularization, optimization, and beyond. MIT Press.

28. Silva, A., & Santos, M. (2025). Multimodal learning with textual sentiment for emerging equity markets. International Review of Financial Analysis, 94, 103-118.

29. Taylor, S. J. (1986). Modelling financial time series. John Wiley & Sons.

30. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998-6008.

31. Zhang, G., Patuwo, B. E., & Hu, M. Y. (1998). Forecasting with artificial neural networks: The state of the art. International Journal of Forecasting, 14(1), 35-62.