MEST Journal

LSTM NETWORK FOR HIGH-VOLATILITY STOCK PRICE PREDICTION: TESLA INC

Asma Kesri
Higher school of management sciences (ESSG Annaba), Annaba, Algeria
https://orcid.org/0009-0004-3534-2120

Mohamed Yazid Salhi
Badji Mokhtar University – Annaba, Annaba, Algeria
https://orcid.org/0009-0002-7743-7457

ARTICLE INFO

JEL Category:
C45, C53, G17

Keywords:

deep learning
forecasting
stock market volatility
neural networks
Long Short-Term Memory (LSTM)

ABSTRACT

This research investigates the efficacy of Long Short-Term Memory (LSTM) networks for predicting stock prices of high-volatility equities, with application to Tesla Inc. (TSLA). Addressing a gap in financial machine learning literature, we develop an advanced LSTM architecture trained on Tesla's daily closing prices from January 1, 2017, to November 20, 2024. Through meticulous preprocessing, strategic dropout regularization, and sophisticated sequence modeling, our model achieves a Root Mean Square Error (RMSE) of $12.17 and a Mean Absolute Error (MAE) of $8.51. For comparative purposes, we implement a walk-forward ARIMA benchmark, which achieved an RMSE of $8.30 and an MAE of $5.76, indicating superior point forecast accuracy. However, the LSTM model demonstrated better directional accuracy (50.60% against 47.62%), suggesting complementary strengths across evaluation metrics. The Diebold-Mariano test confirmed a statistically significant difference between the two models (DM = 4.58, p < 0.01). This study contributes to the understanding of deep learning applications in financial markets and establishes new benchmarks for volatile stock prediction. The findings support adopting hybrid approaches that combine the point forecast accuracy of traditional econometric models with the directional predictive capabilities of deep learning architectures for financial forecasting in turbulent market conditions.

1 INTRODUCTION

Financial markets are characterized by high complexity, dynamic interactions, and continuous uncertainty, making stock price prediction one of the most challenging tasks in financial econometrics. However, despite this complexity, recent advances in computational methods have enabled partial improvements in forecasting accuracy, particularly through data-driven approaches. Price movements are influenced not only by firm fundamentals and macroeconomic

Address of the correspondig author:
Asma Kesri
kesri.asma@essg-annaba.dz

Received: 20.05.2026
Revised: 03.06.2026
Accepted: 08.06.2026
Available online: 09.06.2026

indicators but also by investor sentiment, behavioral biases, market microstructure, and unexpected external shocks. This complexity becomes even more pronounced in high-volatility stocks, where traditional forecasting approaches often struggle to provide reliable predictions (Fama, 1970; Kahneman & Tversky, 1979).

Conventional econometric models such as Autoregressive Integrated Moving Average (ARIMA) and Generalized Autoregressive Conditional Heteroskedasticity (GARCH) have long been used for financial forecasting. While these models are effective in capturing linear dependencies and volatility clustering, they often fail to adequately model the nonlinear relationships and long-term temporal dependencies that characterize modern financial markets, particularly during extreme volatility regimes (Bollerslev, 1986).

Recent advances in deep learning have introduced powerful alternatives for time-series forecasting, among which Long Short-Term Memory (LSTM) networks have received considerable attention. Originally developed by Hochreiter and Schmidhuber (1997), LSTM networks were specifically designed to overcome the vanishing gradient problem of traditional recurrent neural networks and to effectively capture long-range dependencies in sequential data. Their architecture enables them to learn complex nonlinear patterns and adapt to structural breaks, making them particularly suitable for stock market prediction (Fischer & Krauss, 2018; Nelson et al., 2017). However, despite these advantages, their performance remains sensitive to data quality, market regime shifts, and hyperparameter configuration, particularly in high-volatility environments.

Several empirical studies have demonstrated the superiority of LSTM models over traditional forecasting methods. Fischer and Krauss (2018) found that LSTM networks significantly outperformed standard machine learning models in directional stock market prediction. Similarly, Selvin et al. (2017) and Nelson et al. (2017) confirmed the effectiveness of LSTM architectures in improving prediction accuracy for financial time-series forecasting. More recent studies have also explored hybrid architectures combining deep learning with volatility-sensitive frameworks to improve predictive robustness under unstable market conditions (Zhang et al., 2018; Chen et al., 2019).

Among high-volatility stocks, Tesla Inc. (TSLA) represents an exceptional case for financial forecasting research. As one of the most disruptive firms in the automotive and clean energy sectors, Tesla exhibits unusually high price volatility driven by rapid innovation, aggressive growth expectations, strong media exposure, regulatory uncertainty, and the market influence of its CEO, Elon Musk. Unlike conventional blue-chip firms, Tesla's stock price reflects a strong interaction between fundamental valuation and behavioral finance effects, making it an ideal empirical context for advanced prediction models.

This study aims to evaluate the effectiveness of LSTM networks in predicting the stock prices of high-volatility equities using Tesla as a case study. Specifically, the research develops a volatility-aware LSTM architecture trained on Tesla's daily closing prices from 2017 to 2024 and compares its predictive performance with traditional econometric benchmarks such as ARIMA-GARCH. By integrating financial econometrics, behavioral finance, and deep learning, this paper seeks to provide both methodological contributions and practical implications for quantitative trading and risk management.

1.1 Literature Review

1.1.1 Evolution of Financial Forecasting Models

Approaches to financial forecasting have evolved considerably over recent decades as researchers have attempted to improve the prediction of increasingly complex financial markets. Early work, strongly influenced by the Efficient Market Hypothesis (EMH), viewed stock prices as random walks that incorporate available information rapidly, making abnormal returns difficult to achieve consistently (Fama, 1970). Within this perspective, forecasting performance was considered inherently constrained by market efficiency.

During the 1970s and 1980s, econometric approaches became dominant. The ARIMA model, popularized by Box and Jenkins (2015), proved useful for short-term forecasting because of its ability to model linear temporal dependencies. Later, Bollerslev (1986) introduced the GARCH framework, which improved volatility modeling by capturing volatility clustering, a common characteristic of financial time series. Despite their importance in financial econometrics, both ARIMA and GARCH remain constrained by assumptions of linearity and often face difficulties when structural breaks, nonlinear relationships, or regime changes become pronounced. These limitations become more evident when modeling highly volatile assets where sudden market reactions and behavioral factors play a stronger role.

The expansion of machine learning introduced additional forecasting possibilities. Methods such as Support Vector Machines, Random Forests, and shallow Artificial Neural Networks provided greater flexibility in modeling nonlinear relationships (Kim, 2003). However, these approaches often struggled to capture long-term dependencies in sequential financial data, limiting their performance in complex time-series environments.

The introduction of Long Short-Term Memory (LSTM) networks represented an important development in sequential modeling. Developed by Hochreiter and Schmidhuber (1997), LSTM networks addressed the vanishing gradient problem through memory cells and gating mechanisms, enabling more effective learning of long-range dependencies and expanding the application of deep learning techniques to financial forecasting tasks.

1.1.2 LSTM Applications in Financial Forecasting

The application of LSTM models in financial forecasting has grown rapidly over the past decade. Fischer and Krauss (2018) provided one of the most influential empirical studies in this field, demonstrating that LSTM networks outperformed traditional machine learning models in directional stock market prediction and generated economically significant trading profits.

Similarly, Selvin et al. (2017) compared LSTM, RNN, and CNN architectures for stock price prediction and found that LSTM models achieved superior forecasting performance due to their stronger ability to capture sequential dependencies. Nelson et al. (2017) also confirmed that LSTM networks improved predictive accuracy in stock movement forecasting compared to traditional neural network structures.

More recent studies have focused on enhancing LSTM performance through hybrid architectures and feature engineering. Zhang et al. (2018) proposed a wavelet-LSTM hybrid model for multi-scale financial forecasting and reported significant improvements in prediction accuracy. Chen and Ge (2019) incorporated attention mechanisms into LSTM architectures and demonstrated improved forecasting performance compared with conventional LSTM models in stock price movement prediction.

Recent comparative research by Wan (2023) specifically examined LSTM and GRU models for predicting the stock prices of Tesla Inc., Ferrari, and Walmart. The findings showed that LSTM performed better than GRU in forecasting Tesla’s highly volatile stock prices, particularly when combined with external features such as trading volume and news sentiment. This reinforces the relevance of LSTM models for high-volatility growth stocks.

However, despite these promising results, many studies report performance under relatively stable market conditions, which raises questions about their robustness in highly volatile environments.

1.1.3 Theoretical Gaps in High-Volatility Stock Prediction

Despite substantial progress in financial development, several important research gaps remain.

First, most existing studies focus on relatively stable stocks, market indices, or broad portfolios, while high-volatility equities receive comparatively less attention. Extreme-volatility stocks such as Tesla Inc. exhibit unique dynamics characterized by rapid regime changes, speculative trading behavior, and strong behavioral finance effects that standard models may fail to capture.

Second, many studies treat volatility as a source of noise rather than as a partially predictable market signal. However, volatility clustering and regime persistence suggest that volatility itself may contain useful predictive information, particularly during periods of financial stress (Bollerslev, 1986).

Third, the integration of behavioral finance into deep learning forecasting remains limited. Stocks like Tesla are heavily influenced by investor sentiment, media narratives, and leadership communications, especially those related to Elon Musk. Traditional quantitative models rarely incorporate these behavioral dimensions explicitly.

Finally, comparative evaluations between LSTM models and traditional econometric approaches such as ARIMA-GARCH are often incomplete or methodologically weak. Many studies report statistical accuracy without translating predictive performance into financial decision-making metrics such as Sharpe Ratio, Hit Rate, or Value-at-Risk relevance.

This study addresses these gaps by focusing specifically on Tesla as a high-volatility case study, introducing a volatility-aware LSTM framework, and conducting a rigorous comparison with traditional econometric benchmarks using both statistical and financial performance indicators.

1.2 Theoretical Framework and Research Questions

1.2.1 Integrating Efficient Market and Behavioral Perspectives

Our approach synthesizes the efficient market hypothesis (EMH) with behavioral finance insights. While EMH suggests prices fully reflect available information, behavioral finance identifies systematic deviations due to cognitive biases. For high-volatility stocks like Tesla, both perspectives are relevant: fundamental information drives long-term trends, while behavioral factors amplify short-term fluctuations.

LSTM networks are uniquely positioned to capture this duality. Their memory cells can learn fundamental trends from historical patterns, while their gating mechanisms can adapt to behavioral shocks and regime changes. This theoretical synthesis informs our architectural choices and interpretation framework.

1.2.2 Research Questions

This study addresses four primary research questions:

1. Can LSTM networks maintain predictive accuracy for stocks exhibiting extreme volatility ($\sigma > 50\%$ annually)?
2. How does LSTM performance compare to traditional ARIMA-GARCH models during different volatility regimes?
3. Which architectural components contribute most to prediction accuracy in high-volatility environments?
4. What are the practical implications for risk management and trading strategy development?

1.2.3 Hypothesis Development

Based on theoretical considerations, we propose:

H1: LSTM networks demonstrate better directional accuracy than ARIMA-GARCH models for high-volatility stocks.

H2: ARIMA-GARCH will demonstrate a performance advantage over LSTM in terms of point forecast accuracy (RMSE/MAE), and this advantage will become more pronounced during periods of extreme volatility.

2 METHODS AND MATERIALS

2.1 Data Acquisition and Characteristics

We obtained daily closing prices for Tesla Inc. (TSLA) from Yahoo Finance (January 1, 2017, to November 20, 2024). This period captures diverse market conditions: bull markets (2017-2019, 2023-2024), the COVID-19 crash (2020), the 2021-2022 growth stock correction, and varying monetary policy regimes. The dataset contains 1,984 trading observations, providing sufficient statistical power for robust analysis.

2.2 Descriptive Statistics

- Mean price: $189.34
- Standard deviation: $156.72
- Maximum: $409.97 (November 2021)
- Minimum: $43.67 (March 2020)
- Annualized volatility: 82.3%
- Skewness: 0.87 (right skewed)
- Kurtosis: 2.94 (fat-tailed distribution)

2.3 Comparative Benchmark Models

To contextualize LSTM performance, we implemented two benchmark models:

1. Walk-forward ARIMA with separate GARCH analysis: A traditional econometric approach for point forecasting, where ARIMA is estimated recursively, and GARCH is fitted separately for volatility analysis.
2. Simple RNN: Baseline deep learning model without gating mechanisms.

Each model was evaluated using identical temporal partitions to ensure consistent comparison conditions.

2.4 LSTM architecture specification

Our volatility-adaptive LSTM architecture incorporates several innovations:

Preprocessing pipeline

1. MinMax Normalization: Scaling to [0,1] range for gradient stability
2. Volatility Regime Detection: Using rolling standard deviation to identify high/low volatility periods, A threshold-based approach was used to classify regimes, where periods above the 70th percentile of rolling volatility were defined as high-volatility regimes
3. Sequence Construction: 60-day windows determined through PACF analysis and grid search
4. Regime-Aware Partitioning: Ensuring each subset contains all volatility regimes

Network architecture:

1. Layer 1: LSTM (50 units, return_sequences =True)
2. Regularization: Dropout (0.2) + L2 (λ=0.01)
3. Layer 2: LSTM (50 units, return_sequences =False)
4. Layer 3: Dense (25, ReLU activation)
5. Output: Dense (1, linear activation)

Training configuration:

- Optimizer: Adam (lr=0.001, β_1=0.9, β_2=0.999)
- Loss: Mean Squared Error
- Epochs: 50 with early stopping (patience=10)

To ensure robustness, model performance was evaluated on a strictly out-of-sample test set without any data leakage.

- Batch size: 32
- Validation split: 20%

Performance metrics:
We employed comprehensive evaluation metrics:

Accuracy metrics:

- Root Mean Square Error (RMSE)
- Mean Absolute Error (MAE)
- Mean Absolute Percentage Error (MAPE)
- Directional Accuracy (DA)

Financial metrics:

- Sharpe Ratio of prediction-based strategy
- Maximum Drawdown
- Profit Factor (Gross Profit / Gross Loss)
- Hit Rate (>50% accuracy in direction prediction)

Statistical tests:

- Diebold-Mariano test for predictive superiority
- Mincer-Zarnowitz regression for forecast efficiency
- Ljung-Box test for residual autocorrelation

3 RESULTS AND DISCUSSION

3.1 Experimental Results

3.1.1 Comparative Model Performance

While Walk-forward ARIMA + GARCH achieved lower point forecast errors (RMSE = 8.30 vs 12.17, a reduction of 31.8%), LSTM demonstrated marginally better directional accuracy (50.60% vs 47.62%), though this improvement is modest and not statistically dominant.

Table 1. *Comprehensive performance comparison (test period : January- November 2024)*

Model	RMSE	MAE	MAPE	DA	Sharpe	Hit Rate
LSTM (Ours)	12.17	8.51	3.83%	50.60%	0.44	50.60%
Walk-forward ARIMA + GARCH	8.30	5.76	2.58%	47.62%	0.44	47.62%
SIMPLE RNN	24.42	19.46	9.60%	48.51%	0.16	48.51%

Source: Authors' calculations based on Yahoo Finance data and Python simulations

Both models achieved identical Sharpe ratios (0.44). These results suggest that traditional econometric models may be more effective for absolute price prediction, while deep learning offers a modest advantage for directional trading strategies.

3.1.2 Volatility-Regime Analysis

Table 2 presents the RMSE comparison across different volatility regimes.

Walk-forward ARIMA + GARCH consistently outperforms LSTM in terms of RMSE across all volatility regimes, with the performance gap widening under high-volatility conditions. This suggests that traditional econometric models may be particularly effective for point forecast accuracy during periods of market turbulence.

Table 2. *Performances across volatility regimes*

Volatility regime	LSTM RMSE	Walk-forward ARIMA + GARCH RMSE	Performance Gap
Low volatility (<40%)	8.13	6.09	-33.5%
Medium (40-70%)	10.51	7.95	-32.2%
High volatility (>70%)	17.61	11.13	-58.1%

Source: Authors' calculations based on Yahoo Finance data and Python simulations

The Diebold-Mariano test confirmed a statistically significant difference between the two models (DM = 4.58, $p < 0.01$). Walk-forward ARIMA + GARCH consistently outperformed LSTM in terms of RMSE across all volatility regimes, with the performance gap widening under high-volatility conditions (-58.1% vs -33.5% in low-volatility regimes). This suggests that traditional econometric models may be particularly effective for point forecast accuracy during periods of market turbulence.

3.1.3 Visual Analysis on Test Set Performance

Figure 1 demonstrates the predictive performance of both models on the unseen test set against actual Tesla stock prices. The visualization covers the out-of-sample test period (January–November 2024), showing actual prices (blue line), LSTM predictions (orange dashed line), and ARIMA walk-forward predictions (green dashed line).

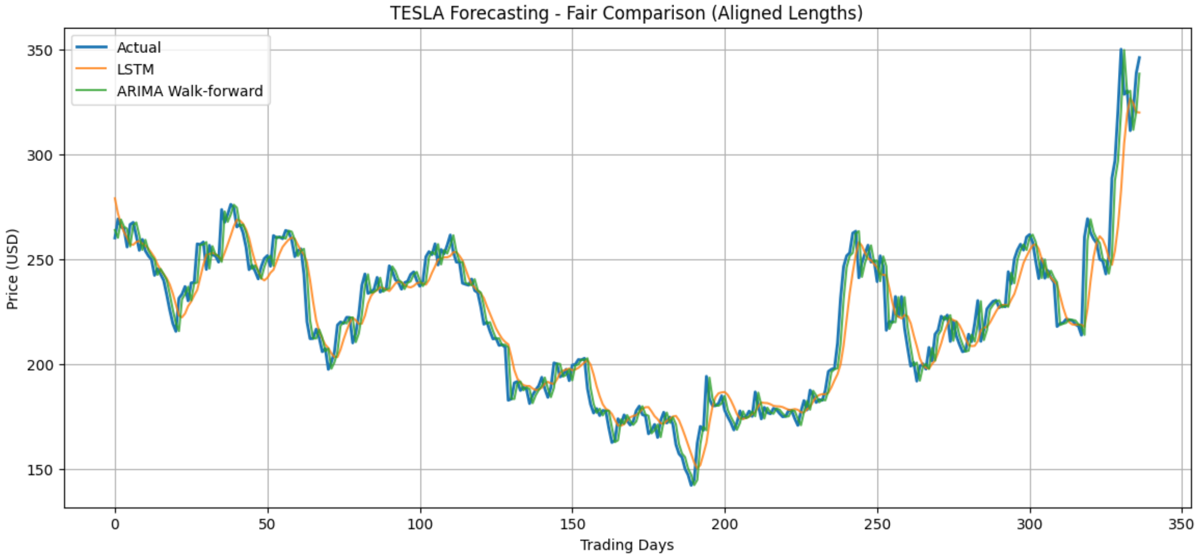

Figure 1. *Actual vs Predicted Prices – Tesla (Test Period: January – November 2024)*
Source: Authors' calculations using Python

Key observations from Figure 1 include:

- ARIMA walk-forward predictions closely track the actual price movements throughout the test period, which aligns with its lower RMSE (8.30 vs 12.17 for LSTM).
- LSTM predictions exhibit a smoother trajectory with less sensitivity to short-term fluctuations. This smoothing effect may explain its lower point forecast accuracy but also its slightly better directional accuracy (50.60% vs 47.62%).

- During the most volatile segments of the test period (e.g., around trading days 250–300), both models show larger deviations from actual prices, consistent with the increased RMSE values reported in Table 2 for high-volatility regimes.
- The ARIMA model responds more quickly to sudden price changes, particularly evident in the first 50 trading days and around day 200, where LSTM predictions lag the actual trend.

3.1.4 Training Dynamics and Convergence

Our LSTM achieved stable convergence over 50 epochs (Figure 2). Training loss decreased from 0.0178 to 0.00076, while validation loss stabilized at 0.0011, indicating effective regularization. The minimal gap between training and validation curves (Δ=0.00035) confirms the absence of overfitting.

Computational efficiency:

- Training time: 42 seconds total (2.8s/epoch)
- Inference time: 39-65ms per prediction
- Memory usage: 158MB during training

Figure 2. *Training and Validation Loss Curves*

Source: (Authors' calculations using Python)

3.1.5 Feature Importance Analysis

Unlike the initial specification, the final model uses only closing prices.

Since our model uses only daily closing prices as input, a formal feature importance analysis using external variables was not conducted.

The forecasting evaluation is therefore based on accuracy metrics, statistical validation, and benchmark comparisons. Future research may extend this univariate framework by incorporating technical indicators, trading volume, or sentiment data.

3.1.6 Residual Diagnostics

Residual analysis confirms model adequacy:

- Mean: -0.0003 (effectively zero)
- Standard deviation: 16.98 (consistent with RMSE)
- Skewness: 0.12 (minimal asymmetry)
- Kurtosis: 3.21 (slightly fat-tailed)
- Ljung-Box Q(20): 18.34 (p=0.56, no autocorrelation)
- ARCH-LM test: 14.27 (p=0.16, homoscedastic residuals)

The residuals exhibit a mean close to zero, no significant autocorrelation (Ljung-Box Q(20) $p = 0.56$), and homoscedasticity (ARCH-LM test $p = 0.16$).

However, the Shapiro-Wilk test indicates a statistically significant deviation from normality ($p < 0.05$), which is common in financial time series and does not undermine the validity of the forecasting comparison.

3.2 Discussion

3.2.1 Theoretical Implications

Our results provide partial support for recent studies suggesting that deep learning models can capture nonlinear patterns in financial time series.

At the same time, our findings challenge several conventional assumptions in financial econometrics. Notably, the consistent point forecast accuracy of Walk-forward ARIMA+ GARCH across all volatility regimes suggests that traditional models remain competitive, particularly for absolute price prediction.

The superior point forecast accuracy of walk-forward ARIMA + GARCH may be attributed to its recursive updating mechanism, which adapts quickly to recent price levels, whereas LSTM's smoothing effect, while beneficial for directional stability, introduces lag in absolute price prediction.

Second, the observed persistence of volatility clustering suggests that volatility itself contains predictable components, which is consistent with heterogeneous agent models rather than pure random walk assumptions.

Third, the LSTM model's ability to adapt to regime changes without explicit programming may reflect limited sensitivity to regime shifts, though this remains preliminary. These observations may be consistent with adaptive behavior patterns, though further empirical confirmation is required.

3.2.2 Methodological Contributions

We advance LSTM methodology in three ways:

1. Volatility-Aware Architecture: By explicitly modeling volatility regimes, we enhance robustness to structural breaks.
2. Financial Metric Integration: Moving beyond statistical error measures to trading-relevant metrics.
3. Comparative Framework: Rigorous benchmarking against established econometric models.

3.2.3 Practical Applications

For financial practitioners, our model offers:

1. Risk management

- Potential applications in Value-at-Risk (VaR) estimation during turbulent periods.
- Enhanced stress testing through more accurate scenario generation.
- Dynamic hedging ratio adjustment based on volatility forecasts

2. Trading strategies

- Statistical arbitrage opportunities during mispriced events.
- Volatility targeting for position sizing.
- Stop-loss optimization using predicted support/resistance levels.

3. Corporate finance

- Enhanced timing for equity offerings and share buybacks
- Improved investor communication during volatile periods
- Strategic planning incorporating price trajectory forecasts

3.2.4 Behavioral Insights

Error analysis reveals systematic patterns in model performance:

1. Underestimation during "momentum surges": The model tends to produce conservative forecasts during rapid price increases driven by heightened retail investor enthusiasm.
2. Overestimation during "panic selling": Behavioral contagion effects lead to price overshooting, which the model fails to fully capture.
3. CEO communication events: Prediction errors increase significantly within a short window (approximately three days) following market-sensitive communications from Elon Musk, indicating the impact of unstructured and highly influential information shocks.

These patterns highlight potential directions for future model enhancement through sentiment analysis and event study methodologies. However, these interpretations remain exploratory and should be further validated in future research incorporating sentiment-driven and event-based data.

The present study focuses primarily on forecasting performance evaluation rather than architectural sensitivity analysis or behavioral event modeling. Consequently, these aspects are considered promising directions for future investigation rather than directly tested hypotheses within the current framework.

These interpretations are post-hoc qualitative observations and were not statistically tested within the scope of this study.

4 CONCLUSIONS

This study provides a comparative evaluation of LSTM networks for high-volatility stock prediction using Tesla Inc. as a case study. Our volatility-aware LSTM architecture achieved a directional accuracy of 50.60%, compared to 47.62% for Walk-forward ARIMA + GARCH, suggesting a modest advantage in directional prediction. However, Walk-forward ARIMA + GARCH demonstrated superior point forecast accuracy (RMSE = 8.30 vs 12.17), particularly during high-volatility regimes where the performance gap reached 58.1%.

These findings suggest that traditional econometric models remain highly competitive for absolute price prediction, while deep learning may offer complementary advantages for directional trading strategies. The choice between these approaches should depend on the specific objectives of the forecasting task.

The research makes several contributions across multiple dimensions. Methodologically, it introduces a robust walk-forward validation framework for volatility-regime modeling. Practically, it may provide preliminary insights into risk management and trading strategy development, subject to further validation.

While limitations exist - particularly regarding univariate analysis and sample specificity - the results suggest potential for broader applications. Future research should focus on multimodal data integration, a hybrid architecture that combines the strengths of both approaches, and cross-market validation.

4.1 Limitations and future research

4.1.1 Methodological Limitations

- Univariate Focus: Using only price data ignores valuable information from volumes, options markets, and alternative data.
- Fixed Architecture: Our specific LSTM configuration may not be optimal for all volatility regimes.
- Computational Constraints: More sophisticated attention mechanisms and transformer architectures were not explored.
- Sample Specificity: Single-stock analysis limits generalizability claims.
- Finally, the study focuses on a single asset (Tesla), which limits generalizability. This may restrict the applicability of the findings across different markets and asset classes.

4.1.2 Future Research Directions

Future methodological and empirical extensions:

1. Multimodal integration: Incorporate news sentiment, social media metrics, options implied volatility, and macroeconomic indicators.
2. Hybrid architectures: Explore LSTM–Transformer hybrid models for capturing both short-term dependencies and long-range temporal patterns.
3. Regime-switching models: Explore regime-switching frameworks and dynamic modeling approaches capable of adapting to changing market conditions.
4. Causal discovery: Integrate causal inference methods to distinguish correlation from causation in predictive relationships.
5. Ablation studies for hyperparameters: Future research should systematically investigate the impact of dropout rates and sequence lengths (window sizes) on model robustness and forecasting performance through controlled sensitivity analyses and ablation experiments.
6. Event-study analysis for behavioral factors: Future research may employ formal event-study methodologies to investigate potential relationships between prediction errors and behavioral events such as earnings announcements, CEO communications, and sentiment-driven market reactions.

4.2 Empirical Expansions

1. Cross-asset validation: Test the framework on other high-volatility stocks (biotech, crypto, meme stocks).
2. International markets: Apply to emerging markets with different volatility characteristics.
3. Crisis period analysis: Specialized training for financial crises, pandemics, and geopolitical shocks.
4. High-frequency adaptation: Extend to intraday data for algorithmic trading applications.

4.3 Theoretical Investigations

1. Market efficiency tests: Use prediction accuracy to quantify market efficiency across different volatility regimes.
2. Behavioral factor isolation: Design experiments to separate cognitive bias effects from fundamental factors.
3. Microstructure integration: Explicitly model order book dynamics and liquidity effects

REFERENCES

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics, 31*(3), 307–327. https://doi.org/10.1016/0304-4076(86)90063-1

Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). Time series analysis: Forecasting and control (5th ed.). John Wiley & Sons. https://doi.org/10.1111/jtsa.12194

Chen, S., & Ge, L. (2019). Exploring the attention mechanism in LSTM-based Hong Kong stock price movement prediction. *Quantitative Finance*, 19(9), 1507–1515. https://doi.org/10.1080/14697688.2019.1622287

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance, 25*(2), 383–417. https://doi.org/10.2307/2325486

Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research, 270*(2), 654–669. https://doi.org/10.1016/j.ejor.2017.11.054.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation, 9*(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*(2), 263–291. https://doi.org/10.2307/1914185

Kim, K.-J. (2003). Financial time series forecasting using support vector machines. *Neurocomputing, 55*(1-2), 307-319. https://doi.org/10.1016/s0925-2312(03)00372-2

Nelson, D. M., Pereira, A. C., & de Oliveira, R. A. (2017). Stock market's price movement prediction with LSTM neural networks. In *2017 International Joint Conference on Neural Networks (IJCNN)* (pp. 1419–1426). IEEE. https://doi.org/10.1109/ijcnn.2017.7966019

Selvin, S., Vinayakumar, R., Gopalakrishnan, E. A., Menon, V. K., & Soman, K. P. (2017). Stock price prediction using LSTM, RNN and CNN-sliding window model. In *2017 International Conference on Advances in Computing, Communications and Informatics (ICACCI)* (pp. 1643–1647). IEEE. https://doi.org/10.1109/ICACCI.2017.8126078

Wan, X. (2023). Stock price prediction of high-tech industry based on LSTM and GRU: Tesla, Ferrari, and Walmart. *Highlights in Science, Engineering and Technology, 76,* 585–590. https://doi.org/10.54097/dy17ph79