Evaluating LSTM Neural Networks with Moving Average Features for Short-Term Stock Price Forecasting: A Comparative Study of the Technology and Pharmaceutical Sectors

Manav Patel

Abstract

This study evaluates whether historical daily trading data combined with simple moving average (SMA) features can improve the accuracy of Long Short-Term Memory (LSTM) neural networks in forecasting short-term stock prices, and whether forecasting accuracy differs systematically between the technology and pharmaceutical sectors. Daily open, high, low, close, and volume data were retrieved through the yfinance Python library for six publicly traded firms: three technology companies (Apple, Alphabet, and NVIDIA) and three pharmaceutical companies (AbbVie, Eli Lilly, and Pfizer). Simple moving averages over 10-, 20-, and 50-day windows were computed from closing prices and supplied to the models as additional inputs. LSTM architectures were implemented in both TensorFlow and PyTorch, trained on chronologically ordered data using an 80/10/10 train-validation-test split, and evaluated with mean squared error, root mean squared error, and mean absolute error. Ten-day out-of-sample forecasts were generated for closing price, daily high, daily low, and trading volume beginning July 21, 2025, and were compared against realized market values. Across the evaluation window, the pharmaceutical sector produced a lower mean absolute percentage error (3.06 percent) than the technology sector (7.27 percent), with AbbVie the most accurately forecast security (2.01 percent) and NVIDIA the least (12.33 percent). Directional accuracy remained near or below chance for every security, indicating that low percentage error reflects proximity to a slow-moving price level rather than a reliable ability to anticipate day-to-day movement. The results indicate that moving averages help LSTM models track the general level and direction of price series in stable, low-volatility conditions, but that models trained exclusively on historical price information cannot account for the exogenous news, regulatory, and innovation shocks that drive a substantial share of price variance, particularly in the technology sector.

Keywords: Long Short-Term Memory, LSTM, Stock Price Prediction, Moving Averages, Deep Learning, Time Series Forecasting, Technology Sector, Pharmaceutical Sector, Financial Machine Learning, Sector Comparison

1. Introduction

The stock market allows individuals and institutions to buy and sell ownership shares of publicly traded companies. Shareholders may benefit from capital appreciation as a company's value increases and, in some cases, from dividend distributions. Stock prices are determined by supply and demand dynamics and are influenced by company-specific fundamentals, broader economic conditions, and market sentiment (Akana, Drayton, and Lee). In 2023, global stock market capitalization surpassed one hundred trillion

dollars, underscoring the economic scale of the question of whether price movements can be anticipated at all.
Daily stock price data typically includes the opening price, the highest and lowest prices reached during the session, the closing price, and trading volume. These variables provide a structured summary of intraday market behavior and form the basis of most quantitative financial analysis. Historical price and volume series are widely used to identify trends, volatility, and momentum, and historical data may additionally be affected by corporate actions such as dividend distributions and stock splits (Beers).
One widely used technique for trend analysis is the moving average. Simple moving averages (SMAs) compute the mean price over a fixed window, whereas exponential moving averages (EMAs) place greater weight on recent observations (Monfared). Moving averages smooth short-term fluctuations and highlight the underlying direction of a series, and are therefore frequently used as technical indicators in financial modeling. A rising moving average is often interpreted as evidence of an upward trend, while a falling moving average is read as the opposite.
Although historical price patterns can reveal trends, stock markets are influenced by numerous exogenous factors, including macroeconomic conditions, geopolitical events, natural disasters, inflation, and firm-specific news, all of which limit predictability (Beers). Price forecasting models must therefore be evaluated cautiously and with appropriate performance metrics rather than presented as reliable trading instruments.
The technology and pharmaceutical sectors offer a useful contrast for this question. In the technology sector, firms such as Apple (AAPL), Alphabet (GOOGL), and NVIDIA (NVDA) are driven by innovation, consumer demand, and technological breakthroughs; their equities are often volatile, with rapid growth potential but also sharp declines caused by competitive pressure (Velasquez). The pharmaceutical sector, represented here by AbbVie (ABBV), Eli Lilly (LLY), and Pfizer (PFE), is oriented around drug development, regulatory approvals, and patents. Pharmaceutical equities tend to be more stable and to offer dividends, but are sensitive to clinical trial outcomes, decisions by the Food and Drug Administration, and patent expirations (Speights). During the COVID-19 pandemic, for example, pharmaceutical equities rose on vaccine development while technology equities benefited from the shift to remote work (Akana, Drayton, and Lee). Comparing these two sectors therefore reveals whether a single modeling approach behaves consistently across different market regimes.
This research addresses a practical need as well as a methodological one. Retail investors account for roughly one quarter of United States trading volume and increasingly transact through mobile applications (Mikulic), yet surveys indicate that only about 43 percent of Americans are stock-market literate (Akana, Drayton, and Lee). Understanding both the capability and the limits of machine learning models applied to public price data is therefore directly relevant to financial literacy.
This study accordingly investigates the following research question: can historical trading data and moving averages improve LSTM model accuracy in predicting future stock prices of technology and pharmaceutical companies? The objective is not to claim precise price prediction, but to evaluate whether adding moving average features improves short-term forecasting accuracy relative to models trained on raw price and volume data alone, and to determine whether that improvement is sector-dependent.

2. Background and Related Work

Long Short-Term Memory networks are a class of recurrent neural network designed to model sequential data by maintaining an internal memory state (Sayah). This architecture allows LSTMs to capture temporal dependencies across multiple time steps while mitigating the vanishing gradient problem that affects standard recurrent neural networks (Monfared). Because financial time series are ordered, autocorrelated, and non-stationary, LSTMs are frequently selected for price forecasting in preference to feedforward architectures, which treat each observation independently and therefore discard temporal structure.

Prior applied work has established the general workflow adopted here. Monfared demonstrates an LSTM applied to Alphabet closing prices using scaled historical inputs and rolling sequence windows, and Sayah presents a comparable pipeline combining exploratory market analysis with LSTM-based prediction. Both studies establish that LSTMs can reproduce the shape of a historical price series, but neither isolates the marginal contribution of technical indicators such as moving averages, and neither compares performance across sectors with structurally different volatility profiles. The present study addresses those two gaps.

The literature also establishes clear boundaries on what such models can achieve. Beers documents the extent to which price changes are driven by news events that are, by construction, absent from historical price data, and Alberto et al. note that scaling and preprocessing decisions materially affect neural network convergence when input features differ by orders of magnitude, as price and volume do. These constraints inform both the methodology and the interpretation of results presented below.

3. Methodology

The approach adopted in this study involves retrieving historical stock data, conducting exploratory data analysis, computing movement metrics, and applying LSTM models to generate forecasts. Each stage is described below.

3.1 Data Collection

Historical daily stock data were retrieved using the yfinance Python library for the period January 1, 2010, through July 20, 2025. Where a company's listing history began later than 2010, data were retrieved from the listing date onward; AbbVie, for example, contributes data from 2013. The dataset includes the open, high, low, and close prices together with trading volume for each session. Dividend and stock split information were retrieved but were not used as model inputs.

3.2 Exploratory Data Analysis

Before model training, the dataset was examined through statistical summaries, correlation heatmaps, pair plots, violin and box plots, line plots of daily, weekly, and monthly attributes, candlestick charts, yearly mean bars, histograms, SMA overlays, and risk-return scatter plots. This analysis was used to identify relationships among stock variables and to confirm the strong collinearity among open, high, low, and close prices before those variables were passed to the models.

3.3 Feature Engineering

Simple moving averages with window sizes of 10, 20, and 50 trading days were computed from closing prices. These windows were selected to capture short-, medium-, and longer-term trend components respectively. The features were included to reduce high-frequency noise and to direct the model toward momentum and direction rather than isolated daily fluctuations.

3.4 Data Preprocessing and Scaling

All continuous input features were normalized using Min-Max scaling to improve numerical stability during training, since neural networks are sensitive to differences in feature magnitude; without normalization, features such as trading volume can dominate the loss function and impair convergence (Alberto, Einhorn, Fisch, Le, and Sautter). Scaling parameters were computed exclusively on the training set and then applied to the validation and test sets in order to prevent data leakage (Sayah). Because the scaling parameters are derived from historical training data, values falling outside the training range may be clipped or extrapolated, which represents a known limitation of this preprocessing approach.

3.5 Train-Validation-Test Split

The dataset was split chronologically to preserve temporal order: the first 80 percent of observations formed the training set, the following 10 percent the validation set, and the final 10 percent the test set. This temporal split ensures that each model is evaluated on data occurring after its training window, which more accurately reflects real-world forecasting conditions than a randomized split.

3.6 Sequence Construction

Input sequences consisted of rolling windows of 10 consecutive trading days, with the model predicting the closing price on the eleventh day. This window size was chosen to balance short-term temporal context against computational efficiency.

3.7 Model Architecture

LSTM models were implemented in both TensorFlow and PyTorch so that results would not depend on a single framework implementation. The TensorFlow model consisted of two stacked LSTM layers with 64 and 32 hidden units respectively, followed by two fully connected layers with 25 and 1 units. It received open, high, low, close, and volume as inputs. The PyTorch model consisted of a single LSTM layer with 64 hidden units, a dropout layer to reduce overfitting, and a linear output layer; it received the same five price and volume variables together with the 10-day and 20-day SMA features, either individually or in combination, allowing the marginal contribution of the moving average features to be assessed. Data for the PyTorch models were scaled to the interval between zero and two, and mini-batches of size 32 were supplied through PyTorch DataLoaders.

3.8 Training Procedure

Models were trained using the Adam optimizer and mean squared error loss. Training was conducted for between 50 and 100 epochs depending on convergence behavior, with training and validation losses recorded at regular intervals so that overfitting could be detected.

3.9 Evaluation Metrics

Model performance was evaluated using mean squared error (MSE), root mean squared error (RMSE), and mean absolute error (MAE), computed on the held-out test set and on the ten-day out-of-sample forecast window. Mean absolute percentage error (MAPE) is additionally reported so that securities trading at very different price levels, such as Eli Lilly near 760 dollars and Pfizer near 24 dollars, can be compared on a common scale. Directional accuracy, defined as the proportion of consecutive day pairs in which the predicted change and the realized change share the same sign, is reported as a complementary measure, since a model may achieve low absolute error while failing entirely to anticipate the direction of movement.

3.10 Forecast Generation

Two forecast horizons were produced. Short-horizon forecasts covered the ten trading days from July 21 to July 30, 2025, and included closing price, daily high, daily low, and trading volume; internal consistency constraints were applied so that the predicted low did not exceed the predicted open or close and the predicted high represented the session maximum. Long-horizon forecasts covering 100 days beginning July 21, 2025 were also generated, incorporating historical volatility and sector-specific event simulation such as drug approvals for pharmaceutical firms and product announcements for technology firms. The long-horizon forecasts are included solely for qualitative visualization and are not interpreted as reliable predictions, because error accumulation in recursive time-series forecasting makes long-term projections increasingly unstable.

4. Results

Results are reported in four parts: closing price forecasts over the ten-day evaluation window, quantitative error metrics for those forecasts, long-horizon closing price simulations, and forecasts of daily high, daily low, and trading volume.

4.1 Ten-Day Closing Price Forecasts

Table 1 reports the predicted and realized closing prices for all six securities across the ten-day evaluation window beginning July 21, 2025.

Table 1. Predicted and actual closing prices, July 21–30, 2025 (US dollars)

Date	AAPL Pred.	AAPL Act.	GOOGL Pred.	GOOGL Act.	NVDA Pred.	NVDA Act.	ABBV Pred.	ABBV Act.	LLY Pred.	LLY Act.	PFE Pred.	PFE Act.
7/21/2025	204.06	212.48	194.20	191.15	186.50	171.38	189.00	184.85	747.00	762.18	23.77	24.26
7/22/2025	202.11	214.40	189.80	192.11	188.00	167.03	189.00	187.11	755.00	776.44	23.35	25.14
7/23/2025	195.75	214.15	190.90	191.51	190.00	170.78	192.00	190.55	760.00	798.89	23.55	25.36
7/24/2025	207.48	213.76	186.10	193.20	197.00	173.74	190.50	190.83	769.00	805.43	23.70	25.35
7/25/2025	200.28	213.88	184.30	194.08	194.00	173.50	190.00	190.28	774.00	812.69	23.83	24.79
7/26/2025	196.94	214.05	186.00	193.42	190.00	176.75	190.80	188.52	778.00	808.11	24.02	24.31
7/27/2025	192.32	214.05	184.30	193.42	191.50	176.75	191.80	188.52	762.00	808.11	23.80	24.31
7/28/2025	200.57	214.05	185.80	196.43	197.00	176.75	195.50	188.52	757.00	808.11	23.72	24.31
7/29/2025	205.59	211.27	186.80	196.43	211.00	175.51	196.80	191.22	762.00	762.95	23.57	24.30
7/30/2025	202.30	209.05	185.40	197.44	211.50	179.27	201.00	189.31	752.00	760.08	23.92	23.81

Forecast accuracy varied substantially across securities. For Apple, the predicted closing price on the first day of the window was 204.06 dollars against a realized value of 212.48 dollars, an error of approximately eight dollars. AbbVie, by contrast, was forecast within two dollars of the realized closing price on several days, including July 23, 2025. Eli Lilly forecasts tracked the general movement of the price series but were consistently offset in level. These differences in error magnitude reflect differing sensitivity to market conditions and price volatility rather than differences in model specification, since an identical architecture and training procedure were applied to every security.

4.2 Quantitative Error Metrics

Table 2 summarizes forecast error for each security across the ten-day window. Mean absolute error and root mean squared error are expressed in dollars and are therefore not directly comparable across securities trading at different price levels; mean absolute percentage error and directional accuracy permit that comparison.

Table 2. Forecast error metrics by security and sector, July 21–30, 2025

Security	Sector	MAE ($)	RMSE ($)	MAPE (%)	Mean Bias ($)	Directional Accuracy (%)
AAPL	Technology	12.37	13.45	5.79	−12.37	22.2
GOOGL	Technology	7.17	8.06	3.68	−6.56	11.1
NVDA	Technology	21.50	22.58	12.33	+21.50	44.4
ABBV	Pharmaceutical	3.79	5.06	2.01	+3.67	33.3
LLY	Pharmaceutical	28.70	32.77	3.58	−28.70	55.6
PFE	Pharmaceutical	0.89	1.08	3.58	−0.87	22.2
Technology mean	—	13.68	14.70	7.27	—	25.9
Pharmaceutical mean	—	11.13	12.97	3.06	—	37.0

Three findings emerge from Table 2. First, the pharmaceutical sector was forecast more accurately than the technology sector on a percentage basis, with a mean absolute percentage error of 3.06 percent against 7.27 percent, a difference of more than a factor of two. Second, the range of accuracy within each sector was wide: AbbVie was forecast to within 2.01 percent while NVIDIA, in the same evaluation window and under the same model configuration, was forecast to within only 12.33 percent. Sector membership therefore constrains but does not determine forecast accuracy. Third, and most importantly for interpretation, directional accuracy was at or below the 50 percent chance level for five of the six securities, with only Eli Lilly exceeding it at 55.6 percent. A model can therefore achieve a low percentage error by remaining close to a slowly moving price level while providing essentially no information about the direction of the next day's movement. This distinction is not visible in error metrics alone and is central to the interpretation offered in Section 5.

The mean bias column indicates a further systematic pattern. Errors were not randomly distributed around zero; each security exhibited a consistent directional offset across the entire window. Apple, Alphabet, Eli Lilly, and Pfizer were systematically under-forecast, while NVIDIA and AbbVie were systematically over-forecast. NVIDIA displayed the largest bias in magnitude, with predictions exceeding realized values by 21.50 dollars on average, consistent with a model extrapolating an established upward trend into a period when the realized series did not sustain it.

4.3 Ten-Day Closing Price Forecast Visualizations

Figures 1 through 6 present the ten-day closing price forecasts for the six securities, ordered by sector.

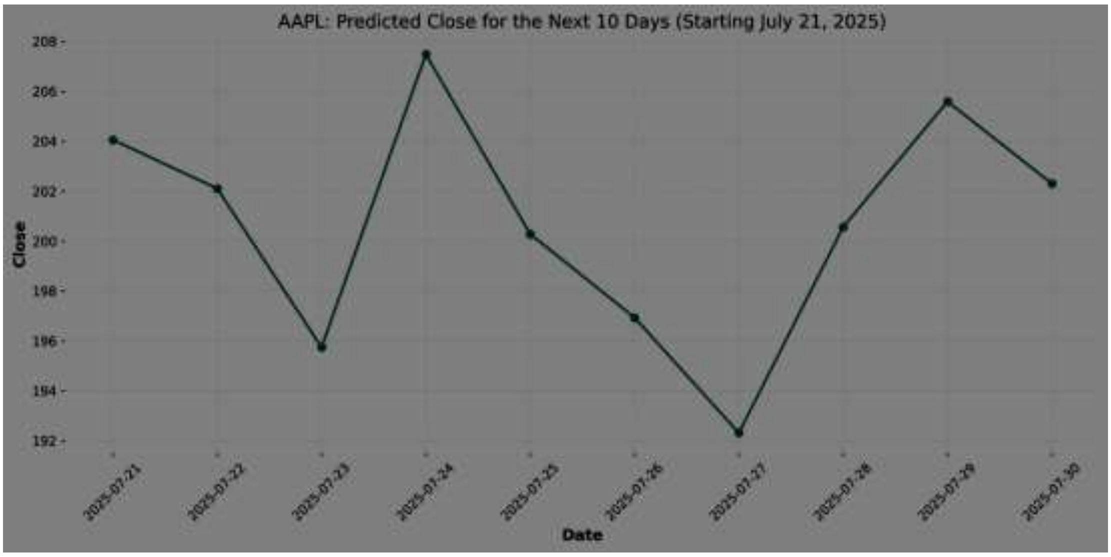

Figure 1. Apple (AAPL): predicted closing price, ten-day horizon beginning July 21, 2025.

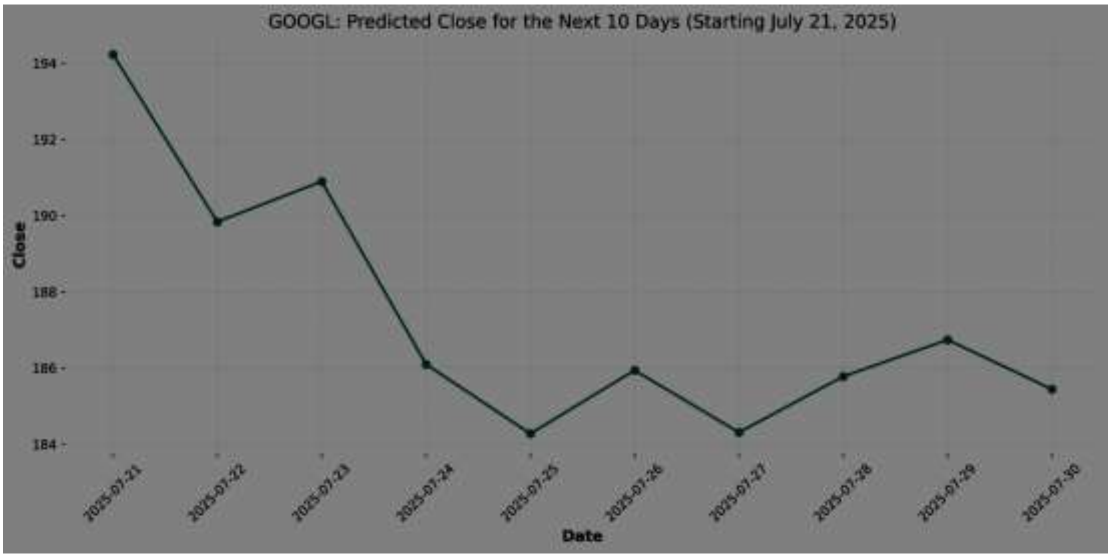

Figure 2. Alphabet (GOOGL): predicted closing price, ten-day horizon beginning July 21, 2025.

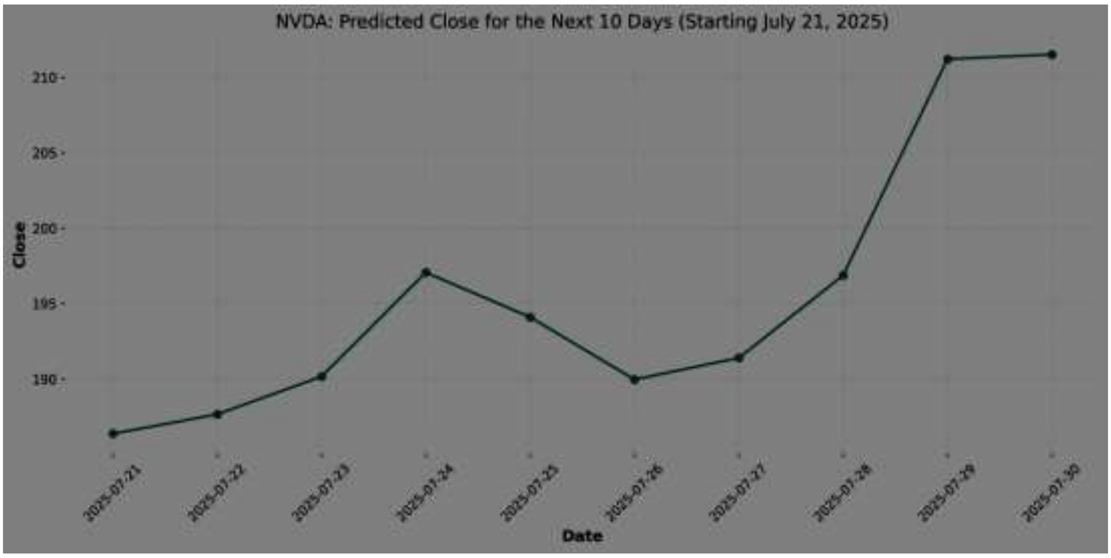

Figure 3. NVIDIA (NVDA): predicted closing price, ten-day horizon beginning July 21, 2025.

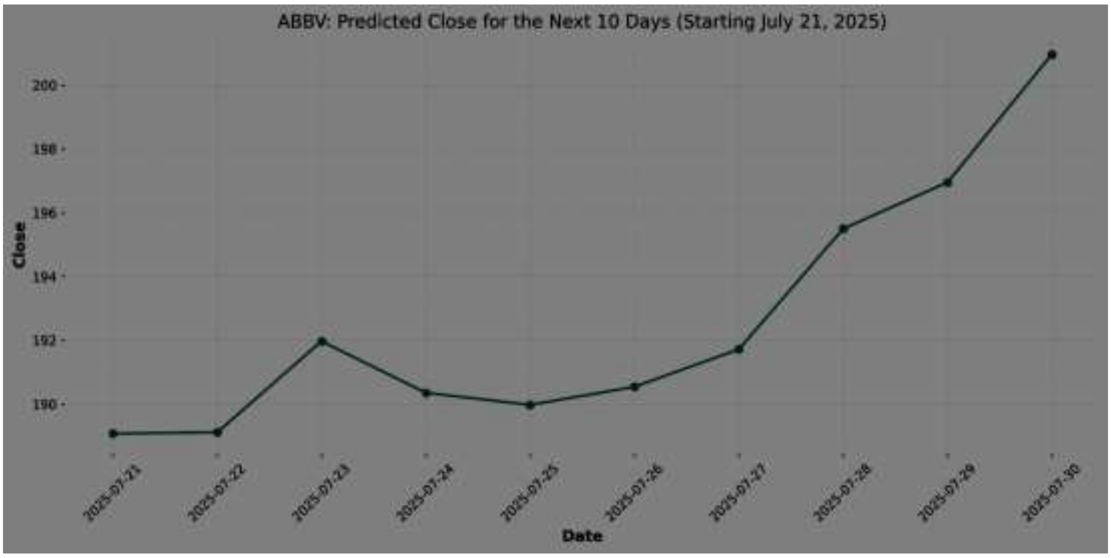

Figure 4. AbbVie (ABBV): predicted closing price, ten-day horizon beginning July 21, 2025.

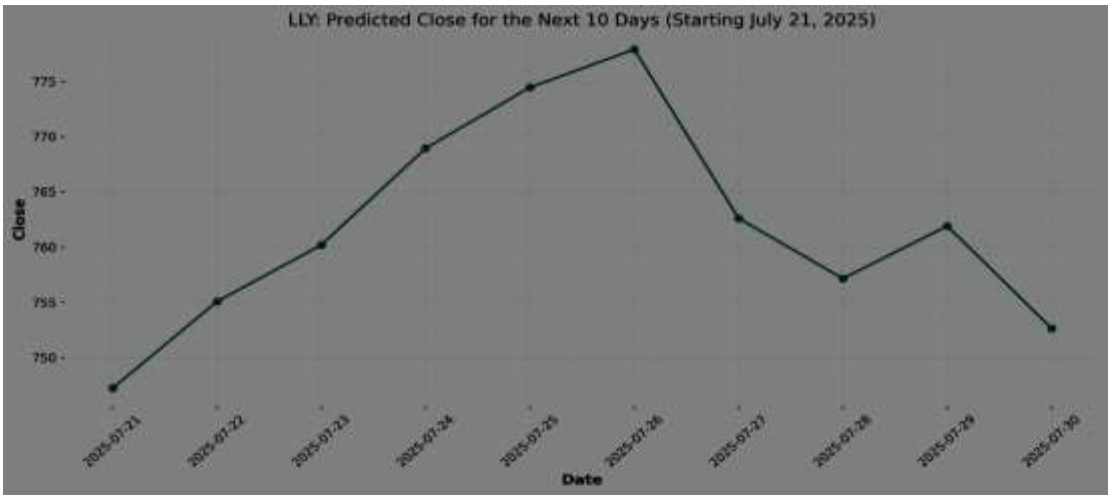

Figure 5. Eli Lilly (LLY): predicted closing price, ten-day horizon beginning July 21, 2025.

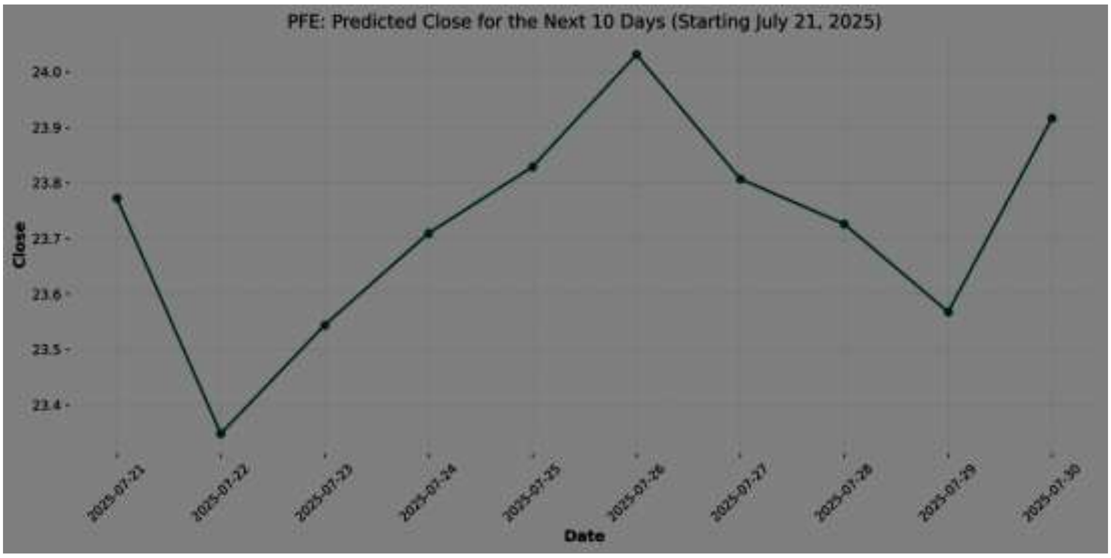

Figure 6. Pfizer (PFE): predicted closing price, ten-day horizon beginning July 21, 2025.

4.4 Long-Horizon Closing Price Simulations

Figures 7 through 12 present 100-day simulated closing price trajectories for the same six securities. These trajectories incorporate historical volatility and simulated sector-specific events and are presented for qualitative visualization only. Given that the ten-day forecasts already exhibit percentage errors of between 2 and 12 percent, and that recursive forecasting compounds error at each successive step, the 100-day trajectories should be read as illustrations of plausible volatility structure rather than as forecasts of price levels.

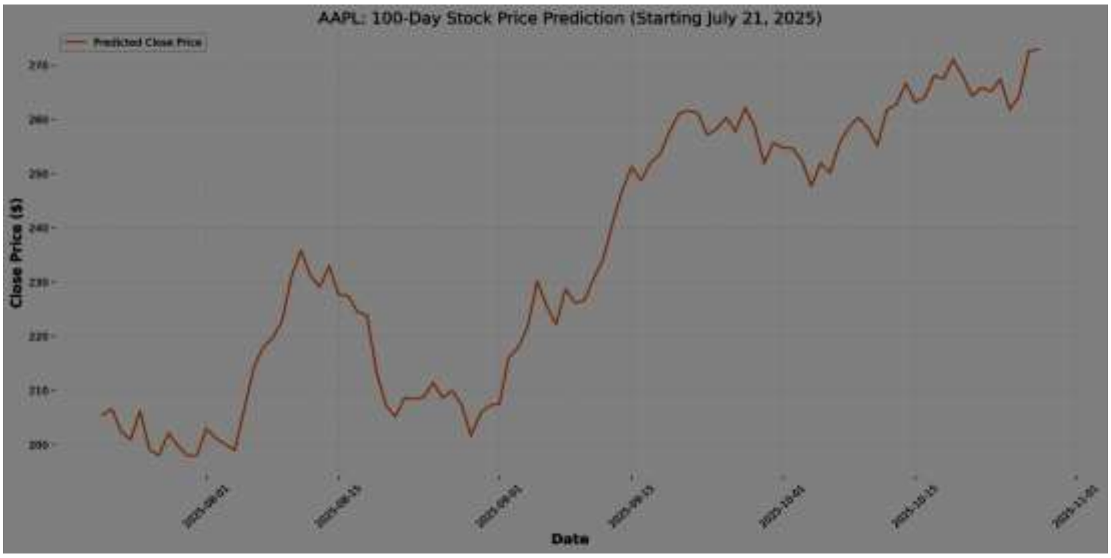

Figure 7. Apple (AAPL): 100-day simulated closing price trajectory.

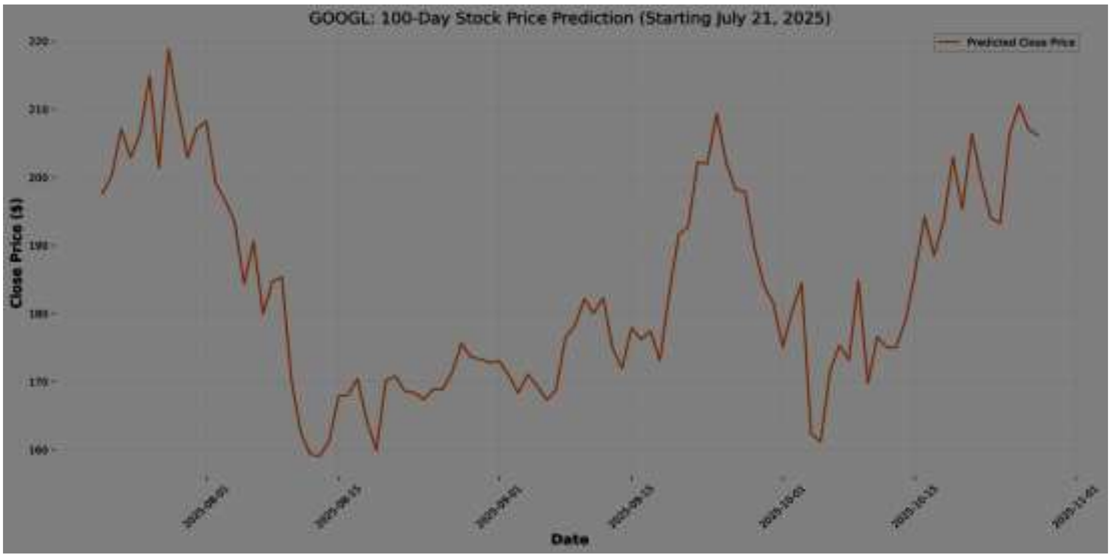

Figure 8. Alphabet (GOOGL): 100-day simulated closing price trajectory.

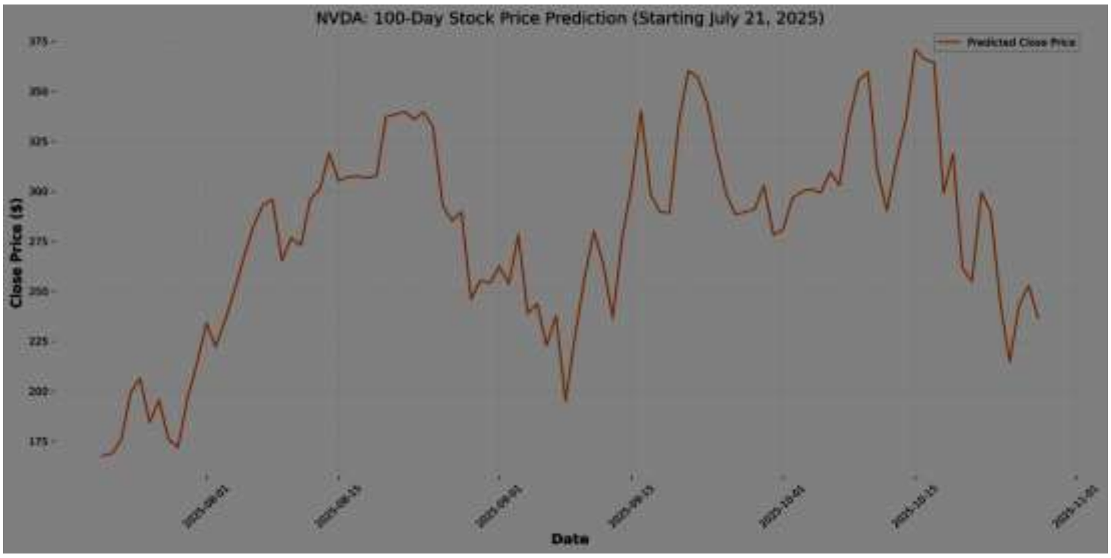

Figure 9. NVIDIA (NVDA): 100-day simulated closing price trajectory.

Figure 10. AbbVie (ABBV): 100-day simulated closing price trajectory.

Figure 11. Eli Lilly (LLY): 100-day simulated closing price trajectory.

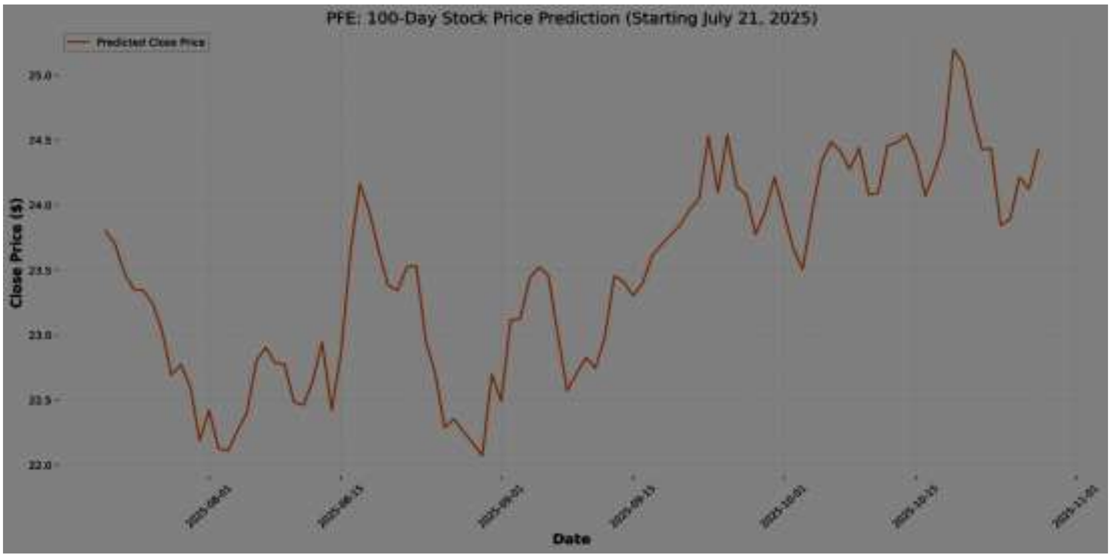

Figure 12. Pfizer (PFE): 100-day simulated closing price trajectory.

4.5 Daily High and Low Forecasts

Forecasts were additionally generated for the daily high and daily low of each session. Figures 13 through 24 present these results. Both sectors show wide predicted ranges between high and low, and in several cases those ranges are not realistic relative to the observed trading range. For Apple on July 21, 2025, the realized high and low were 215.33 and 211.19 dollars respectively, while the model predicted approximately 210 and 199 dollars, an implied intraday range roughly two and a half times the realized range. Numerous single-day spikes appear in these series. These discrepancies are consistent with a model that reproduces historical volatility patterns without access to the real-time information that determines how much of that volatility is actually realized on a given day.

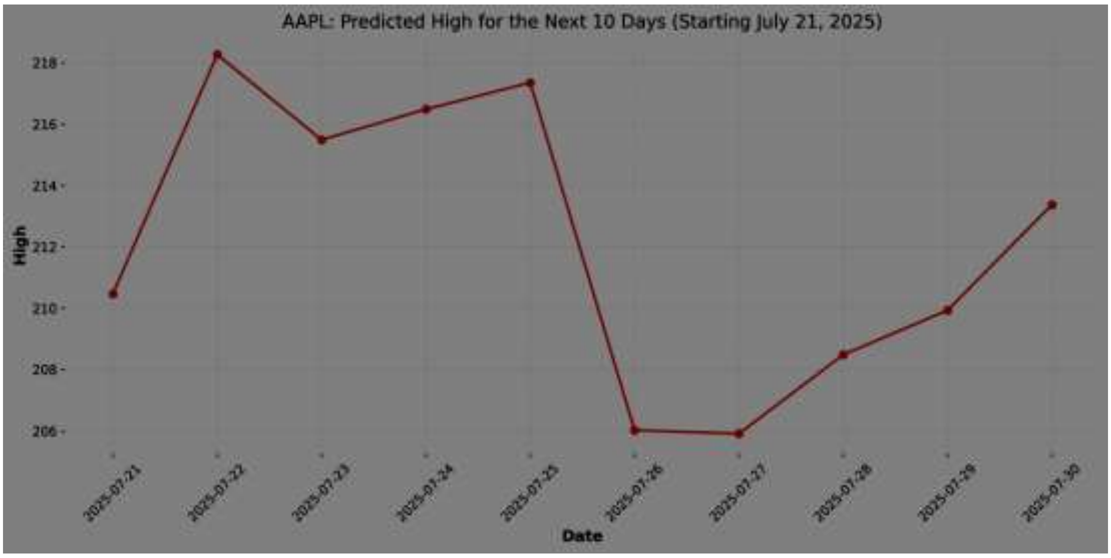

Figure 13. Apple (AAPL): predicted daily high, ten-day horizon.'

Figure 14. Apple (AAPL): predicted daily low, ten-day horizon.

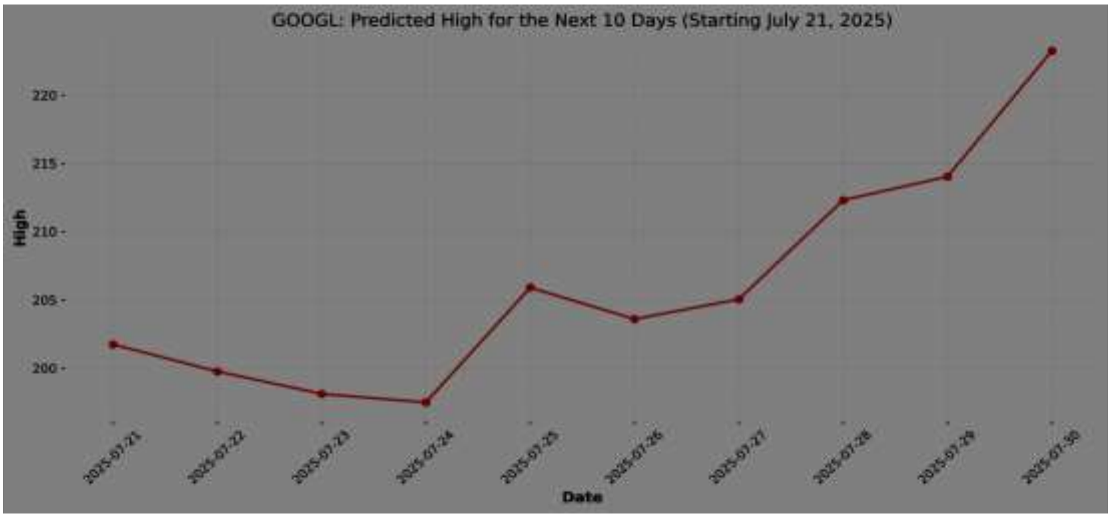

Figure 15. Alphabet (GOOGL): predicted daily high, ten-day horizon.

International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

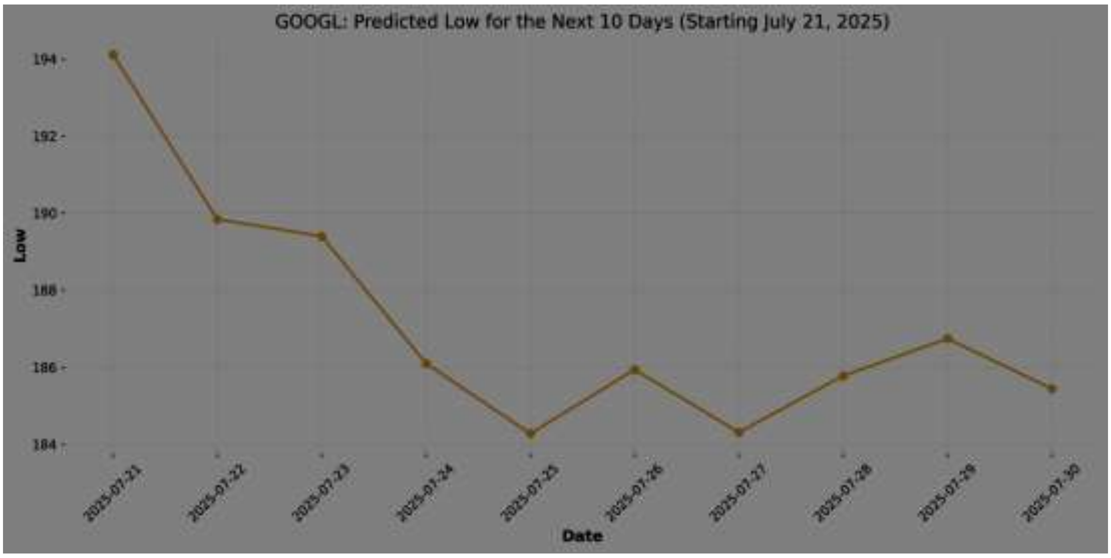

Figure 16. Alphabet (GOOGL): predicted daily low, ten-day horizon.

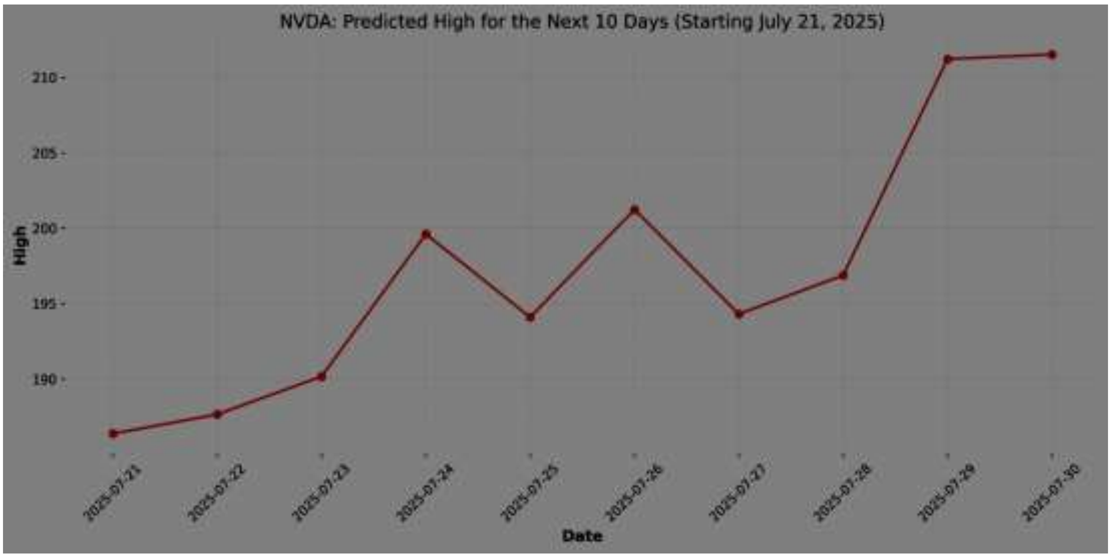

Figure 17. NVIDIA (NVDA): predicted daily high, ten-day horizon.

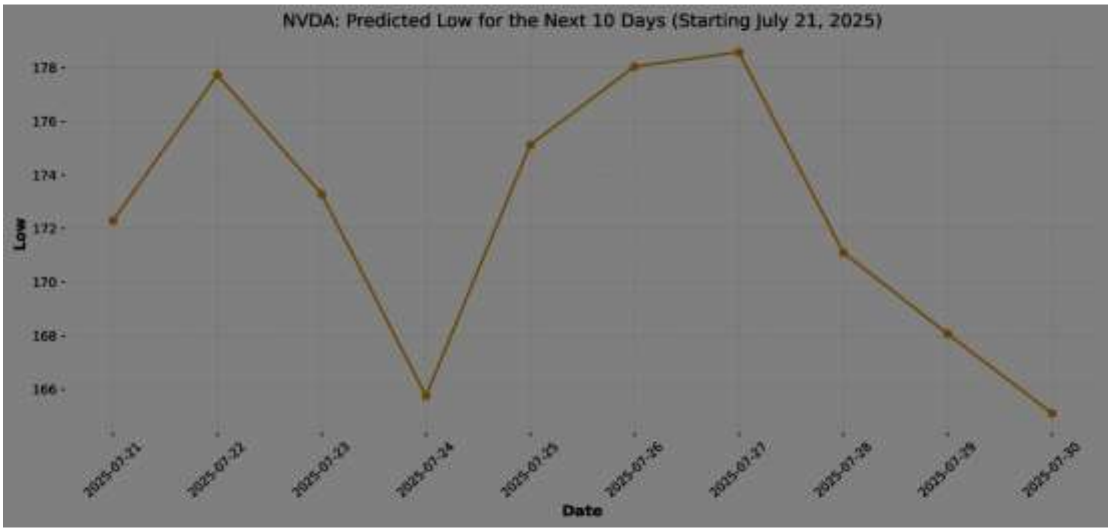

Figure 18. NVIDIA (NVDA): predicted daily low, ten-day horizon.

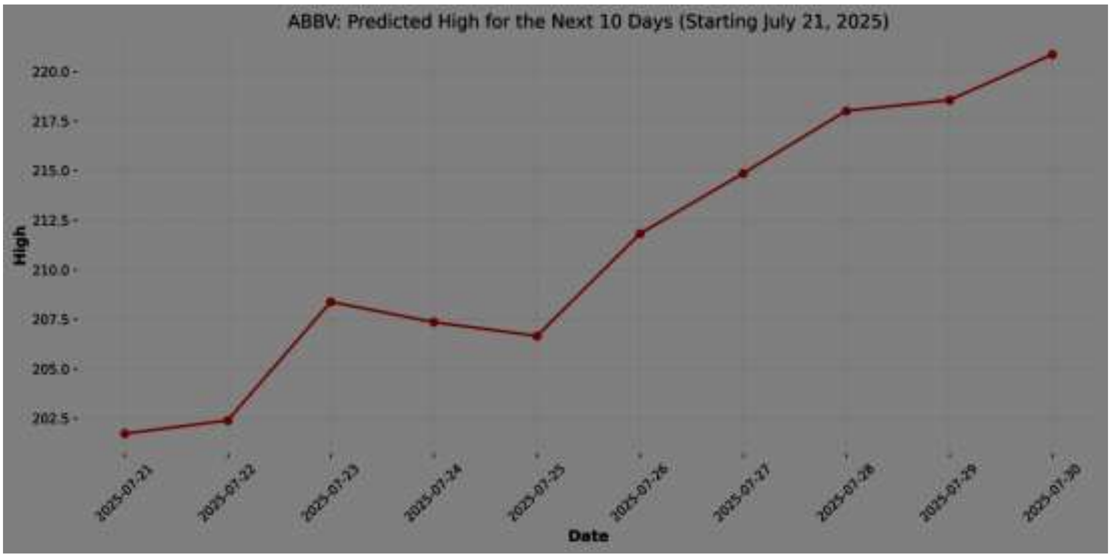

Figure 19. AbbVie (ABBV): predicted daily high, ten-day horizon.

Figure 20. AbbVie (ABBV): predicted daily low, ten-day horizon.

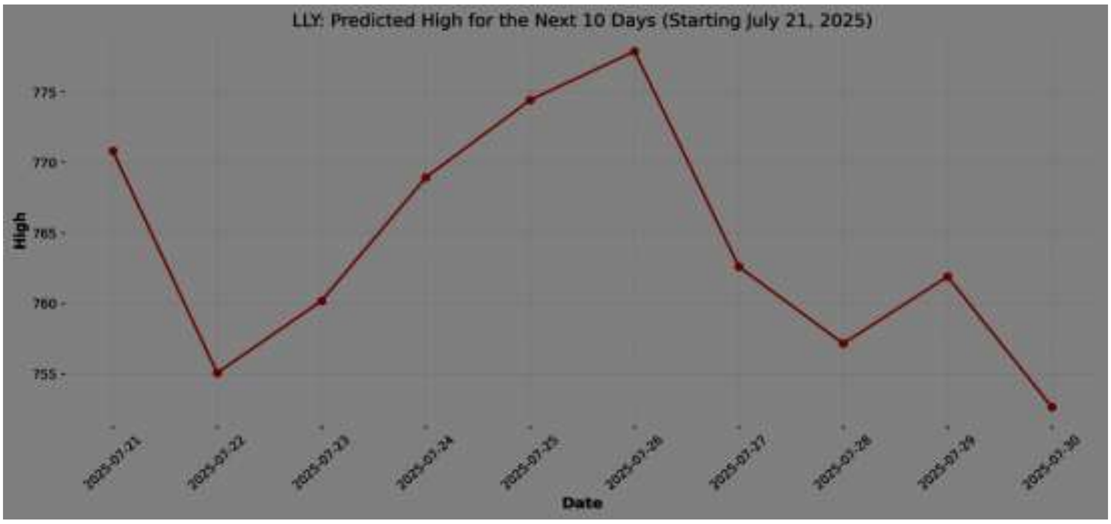

Figure 21. Eli Lilly (LLY): predicted daily high, ten-day horizon.

Figure 22. Eli Lilly (LLY): predicted daily low, ten-day horizon.

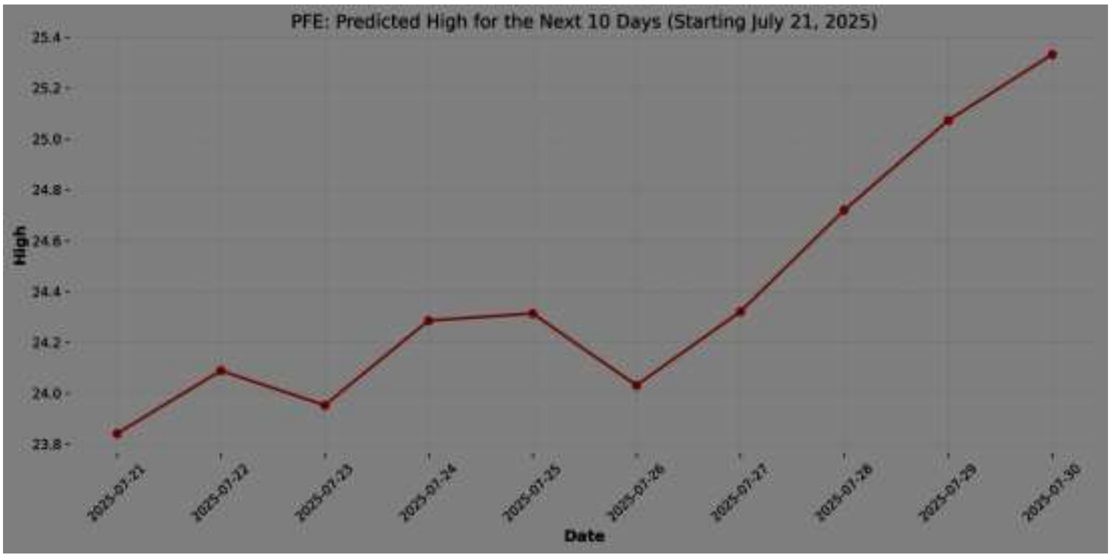

Figure 23. Pfizer (PFE): predicted daily high, ten-day horizon.

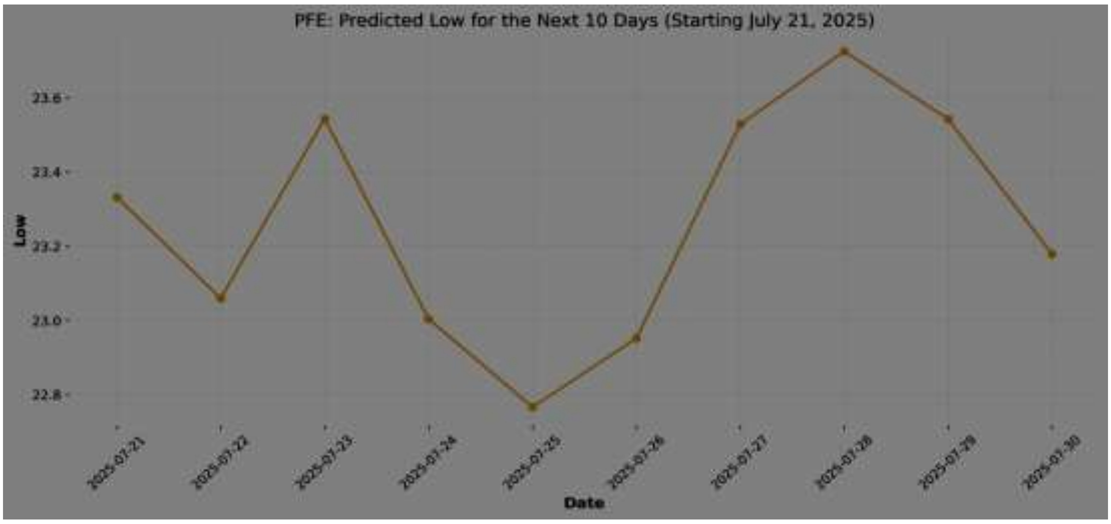

Figure 24. Pfizer (PFE): predicted daily low, ten-day horizon.

4.6 Trading Volume Forecasts

Figures 25 through 30 present ten-day trading volume forecasts. Volume forecasts for Eli Lilly, Apple, and NVIDIA contain pronounced single-day spikes, whereas those for Pfizer, AbbVie, and Alphabet follow smoother trajectories. Predicted technology sector volumes are substantially higher in absolute terms than pharmaceutical sector volumes, with NVIDIA forecast at a scale of approximately 1.4×10^8 shares against approximately 2×10^7 for AbbVie. This difference is consistent with greater investor interest and higher short-term trading activity in technology equities. Volume is driven predominantly by news flow and event timing, neither of which is represented in the input features, and the spikes visible in these series should therefore be treated as artifacts of the model's volatility structure rather than as anticipated trading events (Beers).

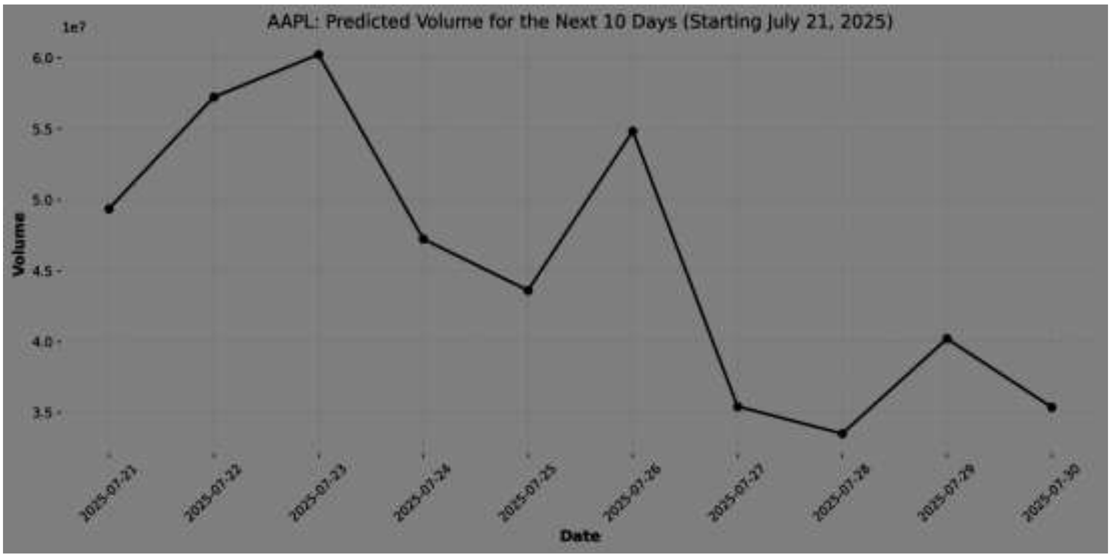

Figure 25. Apple (AAPL): predicted trading volume, ten-day horizon.

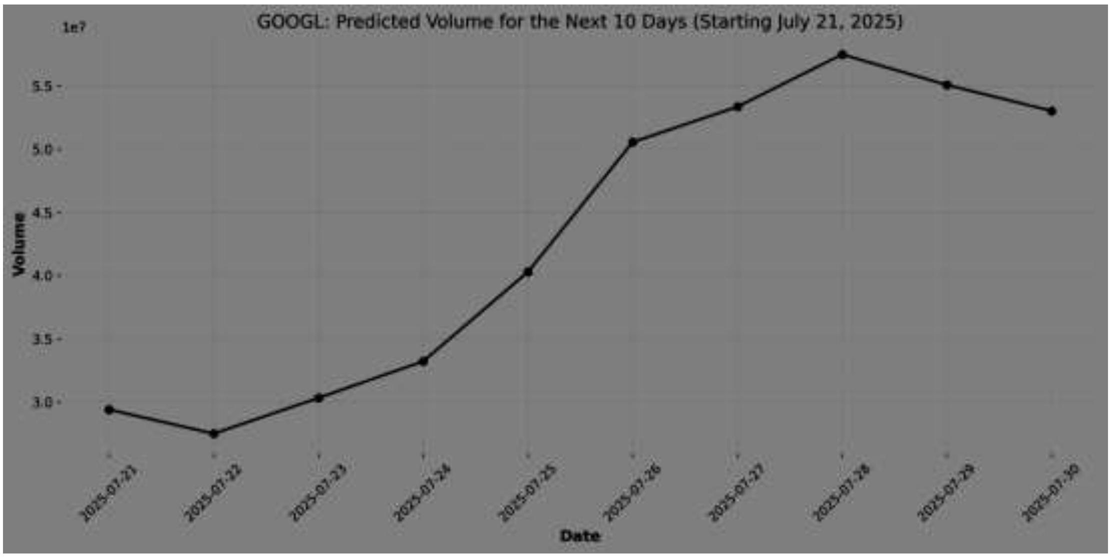

Figure 26. Alphabet (GOOGL): predicted trading volume, ten-day horizon.

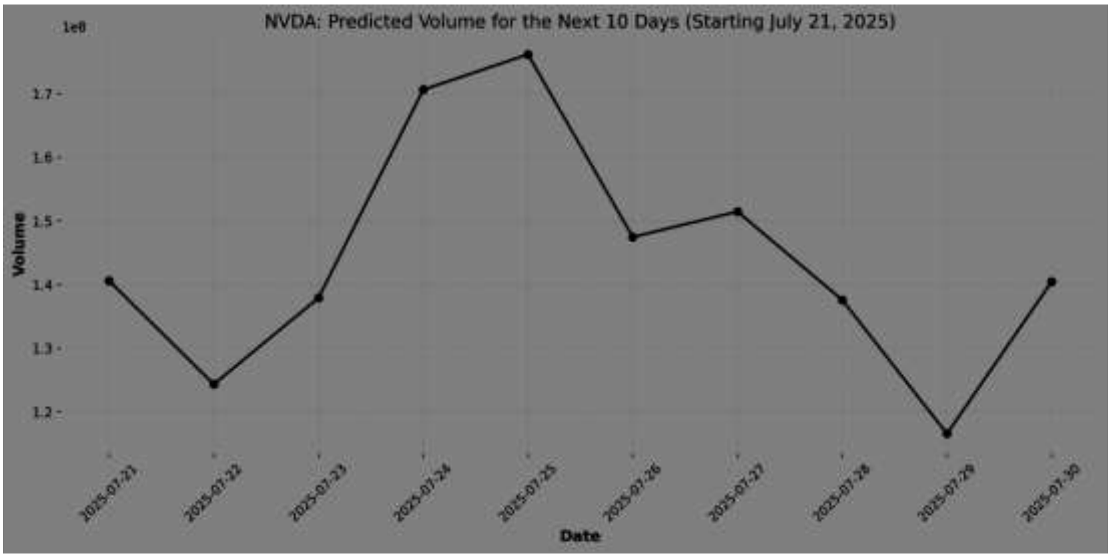

Figure 27. NVIDIA (NVDA): predicted trading volume, ten-day horizon.

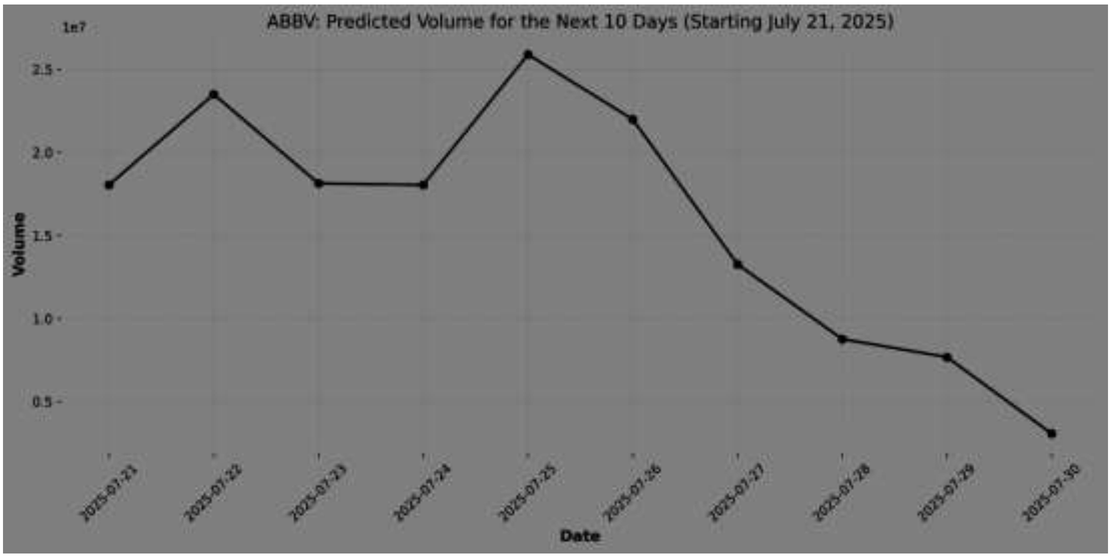

Figure 28. AbbVie (ABBV): predicted trading volume, ten-day horizon.

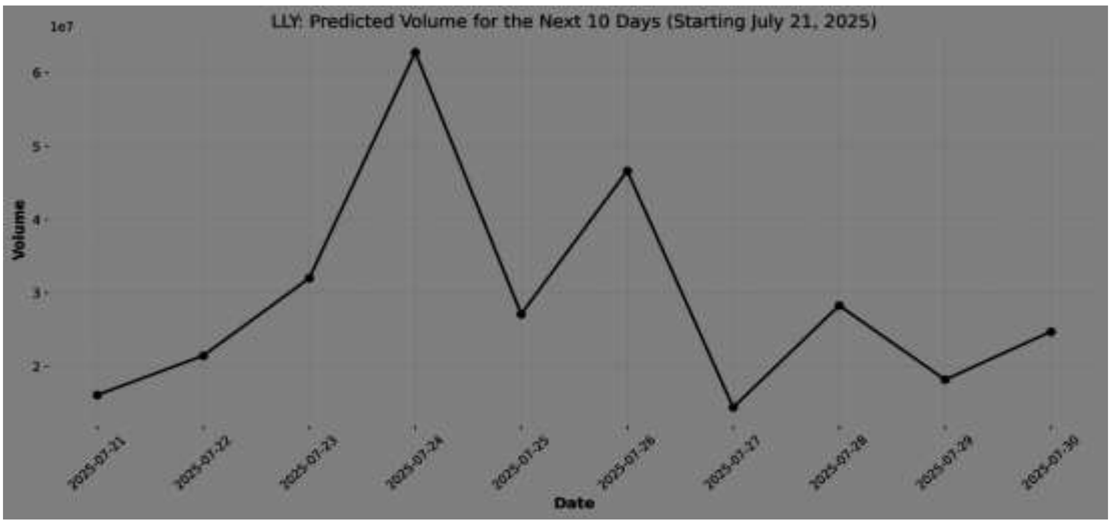

Figure 29. Eli Lilly (LLY): predicted trading volume, ten-day horizon.

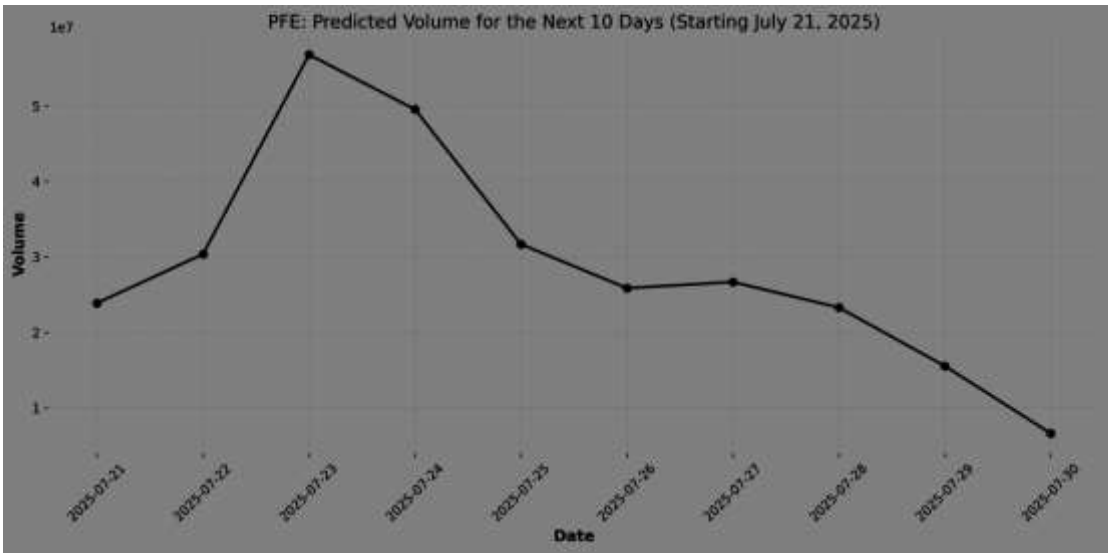

Figure 30. Pfizer (PFE): predicted trading volume, ten-day horizon.

5. Discussion

The results indicate that historical trading data and moving averages, used as the sole inputs to an LSTM model, do not yield reliably accurate short-term price forecasts, although they do allow the model to remain close to the prevailing price level in stable market conditions.

The sector comparison produced the clearest systematic result. Higher volatility in technology equities corresponded to higher forecast error, with a mean absolute percentage error of 7.27 percent against 3.06 percent for pharmaceutical equities. This pattern is consistent with the structural characteristics of the two sectors described in Section 1. Technology valuations respond rapidly to product launches, competitive developments, and shifts in expectations about innovation, none of which are represented in a feature set restricted to historical price and volume. Pharmaceutical equities such as AbbVie and Pfizer traded within narrower ranges during the evaluation window, and the smoother price paths that result are precisely the conditions under which moving average features are informative. The models therefore captured steady pharmaceutical patterns more easily than the more rapidly changing technology series.

Within-sector variation was nonetheless substantial and warns against generalizing from sector membership alone. NVIDIA was the least accurately forecast security in the study, at 12.33 percent, which is consistent with its rapid growth during the period and with the model extrapolating an established upward trend into a window in which that trend did not continue. Apple and Alphabet followed broader market trends and were forecast more closely, though still with errors of 5.79 and 3.68 percent respectively. Within the pharmaceutical sector, AbbVie tracked realized price movement most closely of any security in the study, Eli Lilly followed the general shape of the realized series while remaining offset in level, and Pfizer forecasts were mixed. Firms in the same sector can therefore exhibit materially different behavior depending on firm-specific events.

The directional accuracy results qualify these conclusions considerably and represent the study's most important finding. Five of the six securities achieved directional accuracy at or below chance, and the two securities with the lowest percentage error, AbbVie at 2.01 percent and Alphabet at 3.68 percent, achieved directional accuracy of only 33.3 and 11.1 percent respectively. This combination indicates that the models are effectively producing a smoothed continuation of the recent price level rather than anticipating movement. Low absolute error in this setting reflects the persistence of price levels over short horizons rather than genuine predictive skill. Any evaluation of financial forecasting models that reports only magnitude-based error metrics risks substantially overstating model capability, and the reporting of directional accuracy alongside MAE, RMSE, and MAPE should be regarded as necessary rather than supplementary.

Taken together, these findings suggest that moving averages are useful for identifying trends but should not serve as the only input to a forecasting model. Historical data and moving averages combined with LSTM architectures provide a workable means of describing the general pattern of a price series while remaining unsuitable as a basis for trading decisions.

6. Limitations

The principal limitation of this research is that the models rely exclusively on historical price data and moving averages (Monfared). Prices are influenced by news, social media activity, political developments, and innovation announcements, factors that are especially consequential in the pharmaceutical and technology sectors (Beers). Because these factors were not represented in the feature set, the resulting forecasts are not accurate enough to be relied upon, and even if such variables were incorporated, market randomness would place a ceiling on achievable accuracy.

Three further limitations apply. First, the evaluation window comprises ten trading days for six securities, which is a small sample; the reported error metrics describe model behavior during one specific period and may not generalize to other market regimes. Second, the Min-Max scaling parameters are derived from the training set, so values falling outside the historical training range may be clipped or extrapolated, a constraint that becomes more binding for securities in sustained trends such as NVIDIA. Third, the 100-day forecasts are recursive, and error accumulation renders long-horizon projections increasingly unstable; they are presented as visualizations only.

It is also worth noting that quantitative practitioners forecast markets successfully using real-time inputs and substantially more sophisticated models, but do so with considerable risk arising from market randomness. The gap between the results reported here and professional practice lies primarily in data access and feature breadth rather than in model architecture.

7. Conclusion and Future Work

This study evaluated whether historical trading data and moving averages improve the accuracy of LSTM models in forecasting short-term stock prices, and whether accuracy differs between the technology and pharmaceutical sectors. Across a ten-day evaluation window covering six securities, the pharmaceutical sector was forecast more accurately than the technology sector, at a mean absolute percentage error of 3.06 percent against 7.27 percent, and AbbVie was the most accurately forecast security at 2.01 percent while NVIDIA was the least accurate at 12.33 percent. Directional accuracy remained at or below chance for five of six securities.

It can therefore be concluded that historical trading data and moving averages help identify general price trends but cannot consistently predict exact price movements. Because both sectors depend heavily on innovation and news, and because those factors were not represented in the models, prediction accuracy remained insufficient for reliable decision-making. LSTM models trained on historical price data and moving averages capture short-term trends under certain market conditions, but their accuracy is limited by the absence of exogenous variables and by the inherently stochastic nature of financial markets (Beers).

Future work should extend this research in three directions. The feature set should be expanded to incorporate exogenous inputs, particularly news sentiment and macroeconomic indicators, since the sector difference observed here points directly to the absence of such information as the binding constraint. The evaluation design should be extended to multiple non-overlapping windows across different market regimes so that the sector comparison can be tested for stability rather than observed in a single period. Finally, a formal ablation study isolating the marginal contribution of each SMA window, and comparing LSTM performance against simpler baselines such as a naive random walk, would establish whether the added architectural complexity is justified by improved accuracy. Investigating these predictive models further connects directly to quantitative finance, where accurate forecasting carries substantial economic value.

Acknowledgements
The author thanks Vijay Krishnan for mentorship and guidance throughout this research.

References

1. Akana, Tom, et al. While About Half of Americans Report Owning Stocks Either Personally or Jointly with a Household. Federal Reserve Bank of Philadelphia, www.philadelphiafed.org/-/media/FRBP/Assets/Consumer-Finance/Briefs/Why-Some-Americans-Dont-Invest-in-the-Stock-Market.pdf.
2. Alberto, Sergio, et al. “The Retail Investor Report.” UMKC School of Law Institutional Repository, 2023, irlaw.umkc.edu/faculty_works/928/.
3. Beers, Brian. “Why Do Stock Prices Change Based on News Reports?” Investopedia, 2019, www.investopedia.com/ask/answers/155.asp.
4. Mikulic, Matej. “Global Pharmaceutical Industry.” Statista, 10 Jan. 2024, www.statista.com/topics/1764/global-pharmaceutical-industry/#topicOverview.
5. Monfared, Melissa. “Google’s Stock Price Prediction — LSTM.” Kaggle, 21 Sept. 2024, www.kaggle.com/code/melissamonfared/google-s-stock-price-prediction-lstm.
6. Sayah, Fares. “Stock Market Analysis and Prediction Using LSTM.” Kaggle, www.kaggle.com/code/faressayah/stock-market-analysis-prediction-using-lstm.
7. Speights, Keith. “Biotech vs. Pharma: What’s the Difference?” The Motley Fool, 6 Feb. 2026, www.fool.com/investing/stock-market/market-sectors/healthcare/biotech-vs-pharma/.
8. Velasquez, Francisco. “Tech Volatility Creates the Perfect Environment to ‘Nibble’ on Fintech, Market Veteran Says.” Yahoo Finance, 4 Feb. 2026, finance.yahoo.com/news/tech-volatility-creates-the-perfect-environment-to-nibble-on-fintech-market-veteran-says-172729443.html.