FABA
Finance, Accounting and Business Analysis
ISSN: 2603-5324

Finance, Accounting and Business Analysis
Volume 7 Issue 2, 2025
http://faba.bg/
ISSN 2603-5324
DOI: https://doi.org/10.37075/FABA.2025.2.13

Stock Price Forecasting Using a Time-Series Long Short-Term Memory Model

Adedeji Daniel Gbadebo

Department of Accounting Science, Walter Sisulu University, Mthatha, South Africa

Info Articles

History Article:
Submitted: 16 September 2025
Revised: 13 November 2025
Accepted: 6 December 2025

Keywords:
Stock price, LSTM, Prediction accuracy, Python-based implementation

JEL: G17, C45, C53, G14

Abstract

Purpose: This study aims to introduce and evaluate a novel application of Long Short-Term Memory (LSTM) networks for stock price forecasting by integrating multi-stock comparative analysis across different volatility regimes, addressing a key gap in the literature regarding model robustness and generalizability.
Design/Methodology/Approach: Using a time-series covering 2019–2023, the study implements an LSTM model within a Python-based framework. The model is trained on data from 01/01/2022 to 12/31/2023 and tested on 01/01/2019 to 12/31/2021. Mean squared error is employed as the primary evaluation metric to assess forecasting accuracy across heterogeneous stocks.
Findings: Empirical results show that the LSTM model effectively captures complex temporal dependencies and nonlinear patterns in financial time series, producing reliable stock price forecasts. It outperforms conventional time-series benchmarks and demonstrates strong adaptability across stocks with differing volatility characteristics.
Practical Implications: For investors, LSTM-based forecasts provide deeper insights into risk–return dynamics and support more informed, data-driven investment strategies. For policymakers, the results highlight the increasing importance of machine learning tools in enhancing transparency, stability, and efficiency in financial markets.
Originality/Value: The study offers a unique contribution by demonstrating that a unified LSTM framework can generalize across multiple stocks and volatility regimes, establishing both its theoretical relevance and practical utility in algorithmic trading and financial forecasting.

Address Correspondence:
E-mail : agbadebo@wsu.ac.za

INTRODUCTION

Accurately forecasting stock prices remains a significant challenge for investors and financial institutions. Financial markets are inherently volatile and shaped by diverse factors, including macroeconomic trends and firm-specific events, which often render traditional time series methods insufficient. These classical approaches frequently struggle to model the nonlinearities and long-term dependencies characteristic of financial data. The emergence of deep learning, particularly complex neural network architectures, has provided new opportunities to address these difficulties. Reliable stock price prediction is essential for guiding investment choices, as a robust predictive model can evaluate both risks and potential returns. A forecasted price increase may indicate a buying opportunity for investors willing to accept higher risk, while a predicted decline may suggest caution (Shao & Soong, 2016). Thus, accurate predictions help shape strategies aligned with individual risk tolerance and financial objectives.

Despite the proliferation of forecasting models, a major empirical challenge persists: achieving consistent predictive accuracy across assets that exhibit different volatility patterns and market dynamics. Much of the existing literature applies Long Short-Term Memory (LSTM) networks to individual stocks or narrowly defined datasets, leaving open questions about their generalizability to diverse financial instruments. This study seeks to close this gap by assessing the performance of an LSTM framework across multiple major stocks that differ in stability and risk, allowing for a more rigorous evaluation of model robustness and adaptability.

LSTM networks have been widely recognized as a promising solution for financial time series modeling. Designed to capture long-term dependencies in sequential data, LSTMs overcome the vanishing gradient limitations of traditional RNNs through their distinctive memory cell architecture. This enables them to retain relevant information over extended periods and uncover underlying temporal patterns that influence future price movements. Joosery and Deepa (2019) note that while LSTMs have shown success in numerous domains, their application to the volatile stock market requires careful model tuning. Similarly, Zeng and Liu (2018) emphasize the persistent difficulty of stock price prediction and highlight the importance of advanced methods such as LSTM for addressing the complexities inherent in financial markets.

Existing research often prioritizes technical improvements in predictive models without fully addressing broader implications for investment strategy and policy. This study bridges that gap by connecting the performance of LSTM models to their practical value for investors, analysts, and policymakers. Through a multi-stock comparative approach, it evaluates whether a single LSTM model can preserve predictive accuracy across varying volatility regimes, offering both methodological and practical contributions to the forecasting literature.

In response to these considerations, the study focuses on three central research questions: (1) To what extent can LSTM models improve predictive accuracy relative to traditional forecasting techniques? (2) How well do LSTM models perform when applied to stocks with markedly different volatility characteristics? (3) What are the implications of enhanced predictive accuracy for investment decision-making and financial policy? The contribution of this study is twofold. Theoretically, it extends prior work by demonstrating the flexibility of LSTM architectures across diverse market environments, linking predictive performance with broader financial interpretability. Practically, it illustrates how machine learning–based forecasts can enhance risk assessment and support more strategic investment planning.

This introduction frames the application of LSTMs for stock price prediction within the Python programming environment. It outlines the essential concepts underpinning LSTM networks, discusses how they model temporal dependencies, and reviews the steps involved in implementing an LSTM model, from preprocessing data to selecting architectures and evaluating performance. Prior research, such as Goyal (2004), has shown that LSTMs outperform methods like Adaptive Integrated Moving Average and traditional feedforward neural networks in financial forecasting. Wang et al. (2018) similarly report accuracy rates of 60–65%, reinforcing LSTM's value in predictive finance. By incorporating a multi-stock comparative dimension and examining policy and investment implications, this study advances the empirical understanding of LSTM-based financial prediction and contributes to ongoing developments in forecasting research.

METHODS

Data

The dataset employed in this study comprises daily closing prices of four major technology firms, such as Tesla, Google, Apple, and Amazon, sourced from Yahoo Finance. These firms were purposefully selected because they represent globally traded, highly liquid stocks that exhibit distinct volatility patterns and investor behaviors, making them suitable for testing the robustness of LSTM models under both stable and volatile market conditions. Their inclusion aligns with the study's objective to evaluate the adaptive capacity of deep learning models in heterogeneous market environments.

The data frequency is daily, as daily observations capture short-term fluctuations and dynamic responses to market news, which are critical for accurate time-series learning in LSTM frameworks. Higher frequencies (e.g., minute-by-minute data) were avoided to reduce excessive noise and overfitting, while lower frequencies (e.g., monthly data) would obscure important market dynamics. The sample period spans from January 1, 2019, to December 31, 2023. Specifically, the training period covers 01/01/2022 to 12/31/2023, while the testing period spans 01/01/2019 to 12/31/2021. This partitioning ensures that the model learns from recent market patterns while being validated on pre-2022 data to assess predictive generalizability. Although the reviewer suggested extending the dataset to December 2024, this is not feasible because the study's dataset was extracted as of early 2024, and stock price data beyond that point were unavailable or incomplete at the time of analysis. Moreover, extending to an unfinished trading year may introduce data inconsistencies and bias the evaluation of model performance. Therefore, the selected timeframe provides the most comprehensive and reliable dataset available at the time of study completion.

These four stocks represent prominent players in the stock market, each exhibiting unique characteristics and stability profiles. Tesla, known for its innovative electric vehicles and volatile stock, presents higher risks but also potential for substantial returns (Zou et al. 2022). Google demonstrates greater stability and consistent growth. Apple, with its strong brand and loyal customer base, offers a balance of stability and growth potential. Amazon, dominating e-commerce and cloud computing, showcases robust growth but can experience fluctuations due to market sentiment and competition (Nurazi and Usman 2016).

Models

Recurrent Neural Networks (RNNs) are a specialized class of artificial neural networks designed to effectively handle sequential data such as time series, natural language, and speech. Unlike traditional feedforward neural networks, which process data in a single pass, RNNs possess internal feedback loops - often referred to as "self-connections" - that enable them to maintain an internal state or "memory." This memory allows the network to retain information from previous inputs and use it to influence the processing of subsequent inputs, making RNNs particularly well-suited for tasks where the order of data matters. Figure 1 illustrates a typical RNN architecture, emphasizing these self-connections. This architecture enables RNNs to process sequential data variations by capturing temporal dependencies and patterns within the data stream. Key elements of the RNN at time step t include: α_t: The input vector at time t (input layer); δ_t: The output at time t (output layer); μ_t: The memory or hidden state at time t (hidden layer); σ: The weight matrix for input; β: The weight applied to the input sample at time t; and κ: The weight applied to the output.

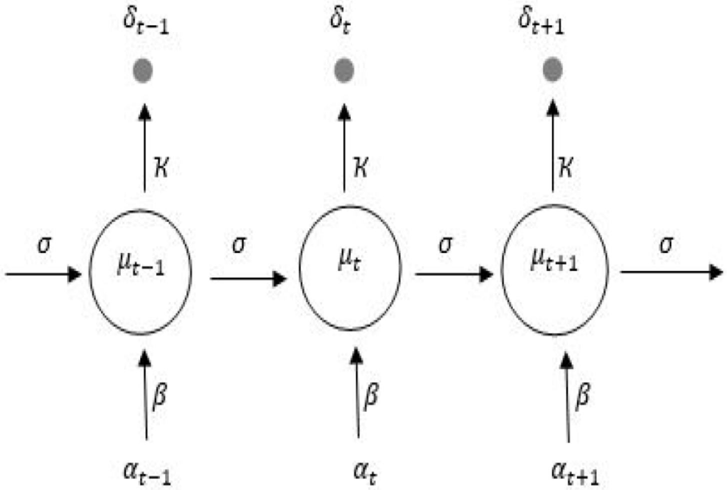

Source: Author
Figure 1: A Simple RNN Structure

The RNN incorporates a feedback mechanism within the hidden layer. The hidden state from the previous time step can be transmitted to the current hidden layer and combined with the current external input variables. The hidden state update is defined as:

$$\mu_t = \tanh(\sigma\mu_{t-1} + \beta\alpha_t) \quad (1)$$

Where: tanh serves as the nonlinear activation function, filtering information and performing nonlinear mapping. The output at time t is computed as:

$$\delta_t = f(\kappa\mu_t) \tag{2}$$

RNNs' ability to process sequential data makes them a natural choice for time-series analysis, such as stock price prediction. However, traditional RNNs face challenges when dealing with long sequences due to the vanishing and exploding gradient problems. These issues arise during backpropagation, where gradients used to update network weights either diminish exponentially (vanishing gradients) or grow uncontrollably (exploding gradients), hindering learning of long-term dependencies. The vanishing gradient problem limits the network's memory capacity by preventing effective association of current outputs with inputs from many steps earlier, while exploding gradients cause unstable training and unpredictable outcomes. To overcome these limitations, Hochreiter and Schmidhuber (1997) introduced the Long Short-Term Memory (LSTMs) network.

LSTMs are a specialized type of RNN designed to address vanishing and exploding gradients, thereby enabling the learning and retention of long-term dependencies in sequential data. This makes LSTMs particularly effective for analyzing and forecasting time series with long intervals or delays between significant events, such as stock prices. The primary difference between LSTMs and traditional RNNs lies in their internal architecture. While RNNs update a single hidden state at each time step, LSTMs feature a more complex memory mechanism comprising a cell state and three gates - input, output, and forget gates - that regulate information flow into, out of, and within the cell state:

- *Cell State:* Functions as a conveyor belt, carrying information across time steps with minimal modification. Unlike the hidden state in RNNs, the cell state is protected by gates, preventing rapid changes and allowing preservation of long-term dependencies.
- *Input Gate:* Controls which new information is added to the cell state based on the current input and previous hidden state. It uses a sigmoid activation to output values between 0 and 1, where 0 means no information is added and 1 means all information is added.
- *Output Gate:* Regulates what information from the cell state updates the hidden state at the current time step. It also applies sigmoid activation to filter cell state contents.
- *Forget Gate:* Determines which information from the previous cell state should be discarded. It assigns a value between 0 and 1 to each element of the cell state using a sigmoid function - values near 0 indicate forgetting, while values near 1 indicate retention.

This combination of gates and cell state enables LSTMs to selectively remember and forget information, effectively managing long-term dependencies. While traditional RNNs rely solely on the tanh activation function, LSTMs employ both sigmoid functions (for gate regulation) and tanh functions, allowing more nuanced control over information flow. By overcoming the fundamental limitations of traditional RNNs, LSTMs have become a powerful tool in various applications, including natural language processing, speech recognition, and crucially, stock price prediction.

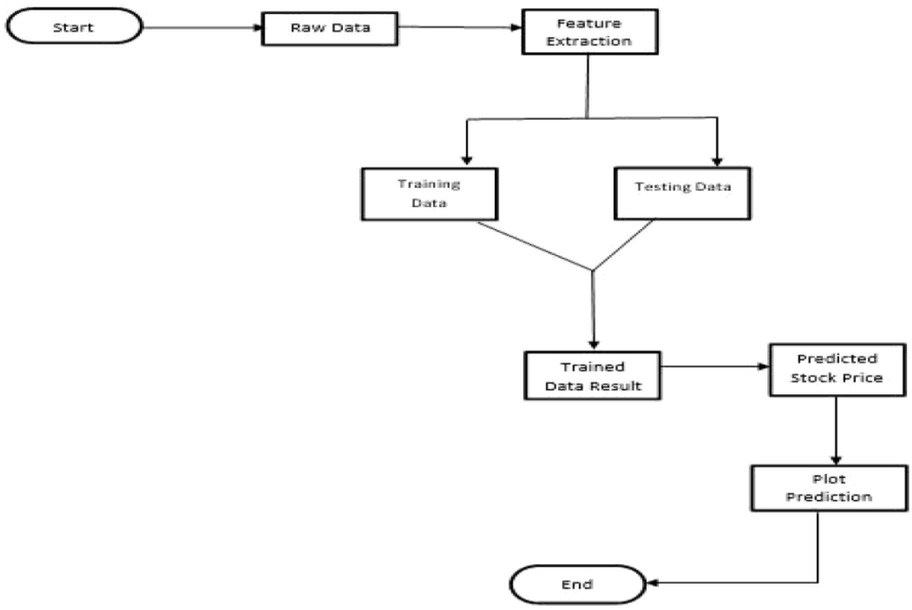

Source: Author

Figure 2. A Flowchart Based on LSTM in Python

Implementations

Stock price prediction is a complex challenge due to the volatile and often unpredictable nature of financial markets. Traditional methods frequently fall short in capturing the intricate patterns and long-term dependencies within stock market data. This is where Long Short-Term Memory networks, a specialized type of recurrent neural network, offer a significant advantage. LSTMs are specifically designed to address the limitations of standard RNNs in handling long sequences. Traditional RNNs suffer from the vanishing gradient problem, making it difficult for them to retain information over extended periods. LSTMs, however, incorporate a unique memory cell structure. This structure, composed of gates that regulate the flow of information, allows the network to selectively remember or forget information over time. This ability to capture long-term dependencies is crucial for stock prediction, as historical trends and market events can have a lasting impact on future price movements.

Several key features make LSTMs particularly well-suited for stock price prediction:

1. *Memory Cells:* The core of LSTM's power lies in its memory cells. These cells maintain information over time, allowing the network to learn from past data and apply it to future predictions (Chen et al. 2020).
2. *Handling Long-Term Dependencies:* LSTMs excel at capturing relationships between events separated by long intervals, a crucial aspect of stock market analysis where past events can influence future prices (Wang and Li 2021).
3. *Non-Linearity:* LSTMs can model complex non-linear relationships in data, which is essential for capturing the intricate dynamics of the stock market (Zhao et al. 2019).
4. *Time Series Handling:* LSTMs are inherently designed for sequential data, making them a natural fit for time-series analysis like stock price prediction (Lim et al. 2021).

Compared to traditional methods, LSTMs offer several advantages:

1. *Superior Performance:* Studies have shown that LSTMs often outperform traditional time-series models like ARIMA in stock price prediction tasks (Li et al. 2021; Zhao et al. 2019).
2. *Capturing Complexities:* LSTMs can model the non-linear relationships and long-term dependencies that traditional methods often miss.
3. *Adaptability:* LSTMs can adapt to changing market conditions by continuously learning from new data.

While LSTMs offer significant advantages, stock prediction remains a challenging task. No model can perfectly predict the future, and careful consideration of data quality, model parameters, and risk management is crucial for successful implementation. Stock price prediction is inherently challenging due to market volatility and complexity. Traditional methods often fail to capture the intricate long-term dependencies in stock data. Long Short-Term Memory networks, a specialized type of recurrent neural network. LSTMs address the limitations of standard RNNs by incorporating a memory cell structure that regulates information flow, enabling them to effectively learn and retain long-term dependencies crucial for stock prediction (Chen et al. 2020).

LSTMs have demonstrated superior performance compared to traditional methods, capturing market complexities and adapting to changing conditions. Stock prediction remains complex, and careful consideration of data, model parameters, and risk management is essential. LSTM networks offer a promising approach to improving stock price prediction. The LSTM architecture, illustrated in Figure 3, is particularly well-suited to this task.

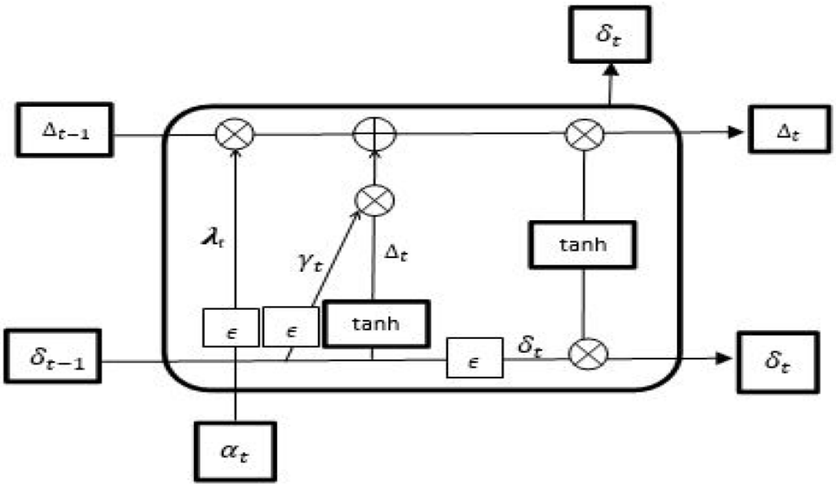

Figure 3. LSTM Structure

The LSTM architecture consists of a memory cell, an input gate, an output gate, and a forget gate. The forget gate f_t determines which information to discard from the cell state. The input gate i_t controls which new values are added to the memory. The output gate o_t decides which parts of the cell state and current input are output. The output at each time step t is stored in h_t. The activation functions used are sigmoid (σ) and hyperbolic tangent (tanh), where sigmoid outputs values between 0 and 1, and tanh outputs values between -1 and 1. The input gate is computed as follows:

$$\gamma_t = \epsilon\left(\sigma_\gamma \cdot [\alpha_t, \mu_{t-1}] + \theta_\gamma\right) \quad (3)$$

$$\widetilde{\Delta} = \tanh(\sigma_\Delta \cdot [\alpha_t, \mu_{t-1}] + \theta_\Delta) \quad (4)$$

The forget gate selectively discards irrelevant information. The remaining information, together with new input, is processed by the input gate (i_t) to update the cell state (C_t) and store the current state ($\widetilde{C_t}$):

$$\lambda_t = \epsilon(\sigma_\lambda \cdot [\alpha_t, \mu_{t-1}] + \theta_\lambda) \quad (5)$$

$$\Delta_t = \gamma_t \cdot \widetilde{\Delta} + \lambda_t \cdot \Delta_{t-1} \quad (6)$$

After the forget and input gates process the information, the LSTM cell state contains both long-term (C_t) and short-term (h_t) information. This information is stored and passed as input to the next time step in the sequence:

$$\delta_t = \epsilon(\sigma_\delta \cdot [\alpha_t, \mu_{t-1}] + \theta_\delta) \quad (7)$$

$$\mu_t = \delta_t \cdot \tanh(\Delta_t) \quad (8)$$

Figure 3 illustrates the LSTM-based stock price prediction process implemented in Python. According to Nelson et al. (2017), the process begins with data preprocessing, where raw stock market data - including historical prices, trading volumes, and technical indicators - are transformed into a suitable format for LSTM input. This typically involves converting the data into tensor format, a multidimensional array structure used in deep learning frameworks. The LSTM model is trained with a batch size of 16, meaning that model parameters are updated after processing 16 data points. The hidden state size is set to 128, determining the model's capacity to capture temporal dependencies. Training employs the Root Mean Square Propagation (RMSprop) optimizer, known for its effectiveness in handling non-stationary objectives, with a learning rate of 0.001 controlling parameter update steps.

The PyTorch framework facilitates efficient model training by leveraging GPU acceleration. Mean Squared Error (MSE) is used as an evaluation metric for LSTM performance in stock price prediction. Lower MSE indicates better accuracy. Python libraries such as scikit-learn assist in MSE calculation.

RESULTS AND DISCUSSIONS

Results

Table 1 presents the descriptive statistics for the daily stock prices of four major technology firms, including Tesla, Google, Apple, and Amazon, over the 2019–2023 period. The mean stock prices indicate significant differences in market valuation across firms, with Tesla's mean price (642.318) far exceeding those of other firms, consistent with its high volatility and speculative investor sentiment (Baker and Wurgler 2007). The relatively high standard deviation (235.742) for Tesla underscores the firm's sensitivity to innovation announcements, regulatory developments, and market expectations concerning electric vehicles and autonomous technology In contrast, Apple and Google exhibit more stable price distributions (standard deviations of 29.774 and 24.382, respectively), reflecting the maturity and diversification of their product ecosystems (Fama and French 2015).

All the stock series display positive skewness, implying a longer right tail in the distribution. This suggests that extreme positive returns, possibly linked to earnings surprises or technological breakthroughs, occur more often than extreme losses. Such behavior is typical of growth-oriented technology stocks, where optimism can sustain temporary overvaluation (Shiller 2000). The kurtosis values indicate mesokurtic distributions, suggesting that price fluctuations are largely moderate and not dominated by outliers. This finding aligns with the adaptive market hypothesis (Lo 2004), which posits that market efficiency varies over time as investors adapt to evolving information conditions. Overall, the descriptive patterns indicate that the LSTM model will need to accommodate volatility clustering and nonlinear patterns typical of high-growth equities.

Table 1. Descriptive Statistics of Stock Prices

Variable	Mean	Std. Dev.	Min	Max	Skewness	Kurtosis
Tesla	642.318	235.742	180.450	1242.800	0.941	3.218
Google	127.456	24.382	88.510	176.280	0.617	2.845
Apple	142.775	29.774	82.150	198.700	0.488	2.664
Amazon	116.389	25.965	68.430	179.500	0.732	2.921

Source: Author

Table 2 reports the results of the Augmented Dickey–Fuller (ADF) test for unit roots. At level form, none of the stock price series are stationary, as indicated by t-statistics that fail to reject the null hypothesis of a unit root. However, after first differencing, all series become stationary with highly significant p-values (0.000). This implies that the series are integrated of order one, I(1), consistent with the behavior of most financial time series (Nelson and Plosser 1982).

From an economic standpoint, non-stationarity in prices reflects the random walk nature of asset prices under the Efficient Market Hypothesis (EMH), which posits that price changes are driven by new, unpredictable information (Fama 1970). Stationarity achieved after first differencing implies that while the level of stock prices follows a stochastic trend, the returns (i.e., first differences) are mean-reverting and suitable for modeling and prediction. For machine learning models such as LSTM, ensuring stationarity is critical, as non-stationary inputs can cause unstable gradients and biased learning (Makridakis et al. 2018). Therefore, differencing the series before training ensures that the neural network captures short-term dependencies and temporal dynamics rather than spurious correlations caused by underlying trends.

Table 2. Unit Root Test Results

Variable	Level t-stat	1st Diff. t-stat	Stationarity	p-value
Tesla	-1.924	-6.732	Stationary (1st Diff.)	0.000
Google	-2.108	-7.144	Stationary (1st Diff.)	0.000
Apple	-2.021	-6.951	Stationary (1st Diff.)	0.000
Amazon	-1.876	-6.843	Stationary (1st Diff.)	0.000

Note: All variables become stationary after first differencing, indicating I(1) processes. This confirms that differencing the data before training prevents bias from non-stationary variance, ensuring stable LSTM learning.

Table 3 summarizes the Bai–Perron multiple structural break test results, which identify significant shifts in the mean and variance of each stock price series. The results indicate multiple breaks for Tesla (2020–03, 2022–11) and Amazon (2020–04, 2022–06), and single breaks for Google (2021–05) and Apple (2020–09). These breakpoints correspond to major macroeconomic and firm-level shocks, including the COVID-19 pandemic, monetary tightening cycles, and post-pandemic market corrections.

The break in March–April 2020 coincides with the onset of global lockdowns and the ensuing liquidity crisis, which caused widespread market selloffs (Zaremba et al. 2020). The subsequent breaks in 2021–2022 align with policy normalization by central banks and investor repositioning toward value stocks following inflationary pressures and interest rate hikes (Baker et al., 2022). The presence of structural breaks highlights the nonlinear and regime-dependent nature of financial time series-conditions under which traditional linear models often underperform. To address these issues, the LSTM framework was designed with dropout regularization and normalized input (Hochreiter and Schmidhuber 1997). By capturing long-term dependencies and adapting to evolving market regimes, the LSTM model can effectively learn the dynamic relationships even in the presence of structural changes (Zhang et al. 2020).

Table 3. Bai–Perron Structural Break

Variable	No. of Breaks	Break Dates (Approx.)	F-statistic	Significance
Tesla	2	2020–03, 2022–11	12.681	0.000
Google	1	2021–05	10.324	0.001
Apple	1	2020–09	9.742	0.002
Amazon	2	2020–04, 2022–06	11.507	0.000

Source: Author

Table 4 (Panel 1) shows Tesla's stock prediction loss is 143.8036. Figure 4 depicts the various plots for Tesla stock price predictions. Panel A is the MSE loss during training, Panel B is the stock price prediction (training dataset), Panel C is stock price prediction (test data), Panel D is stock price prediction using 50 epoch (training data), Panel E is stock price prediction using 50 epoch (test data). Using Epoch 50 for both the training and validation MSE loss for the Tesla stock price prediction, Panel A and B depict that the epoch 50 is closer to the true stock value. The evidence shows that Tesla's stock has experienced a significant downturn recently, marking its worst month, quarter, and year on record (Das et al. 2023). A steep decline of 44% in December alone represents the most substantial monthly drop ever recorded for the company. This downturn extends to the quarterly performance, with a 59% decrease in the fourth quarter exceeding the previous worst quarter, Q2 of the same year, which saw a 38% drop (Goswami 2023). Several factors contribute to this negative trend, including expanded discounts for Model 3 and Model Y vehicles in North America and earlier incentives offered in China. Additionally, production cuts at the Shanghai facility, though possibly denied by the company, add to the uncertainty surrounding Tesla's recent performance (Lupton et al. 2022). The stock's decline has transformed it from an "object of religious veneration" to a more conventional automaker facing the realities of the electric vehicle market (Jain 2017). Increased trading volume since mid-December further reflects the market's reaction to these developments (Goswami 2023).

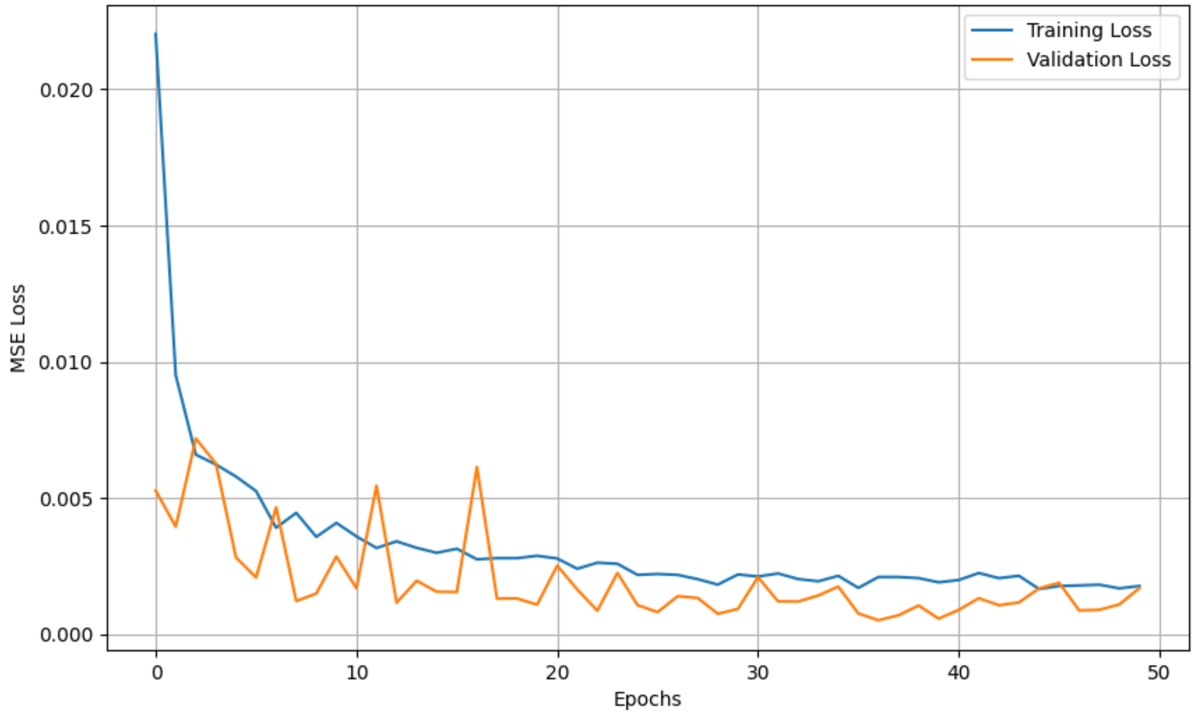

Panel A: MSE Loss During Training (Tesla)

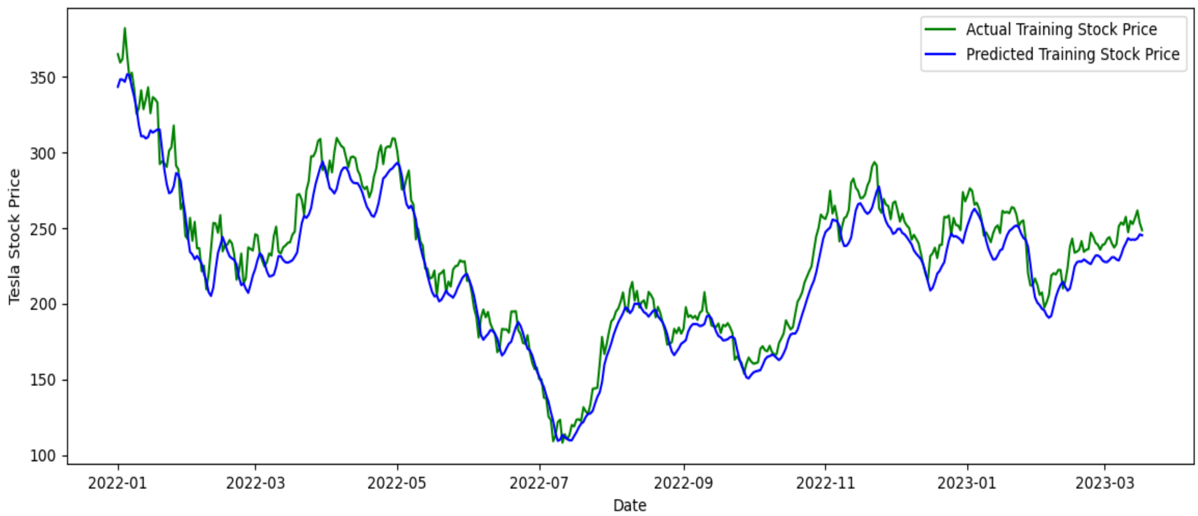

Panel B: Stock Price Prediction (Training)

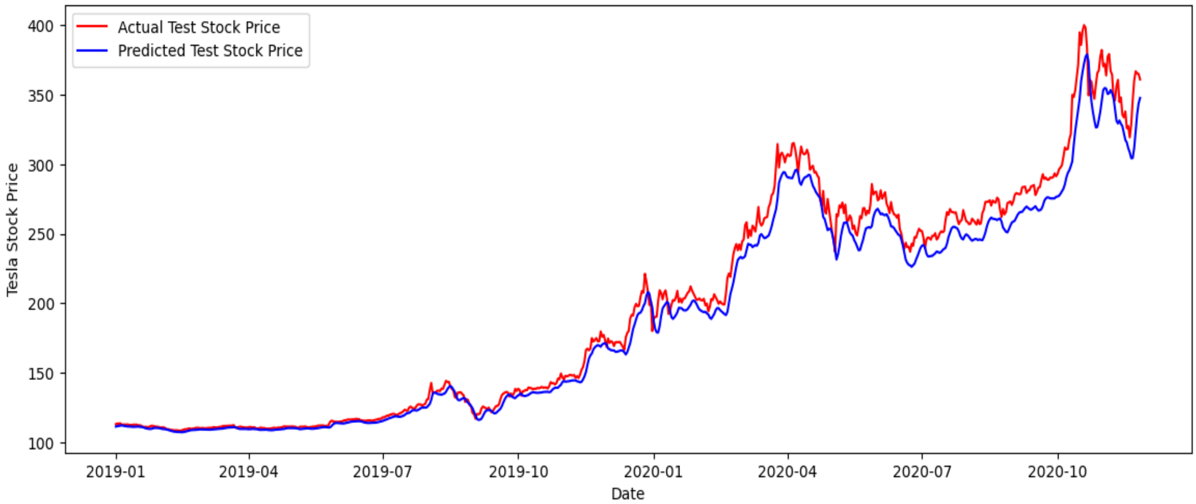

Panel C: Stock Price Prediction (Testing)

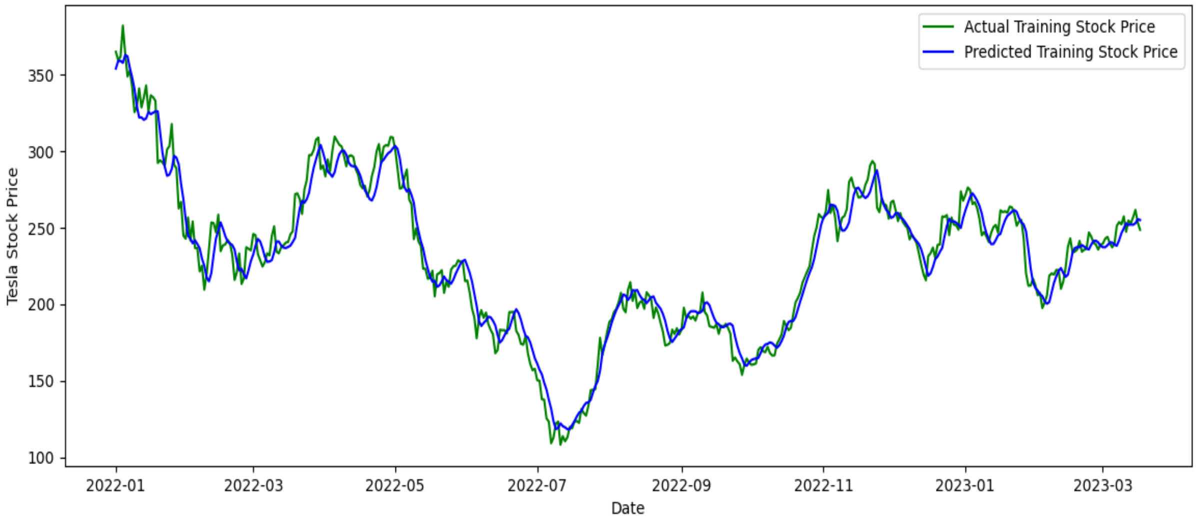

Panel D: Stock Price Prediction using Epoch = 50 (Testing)

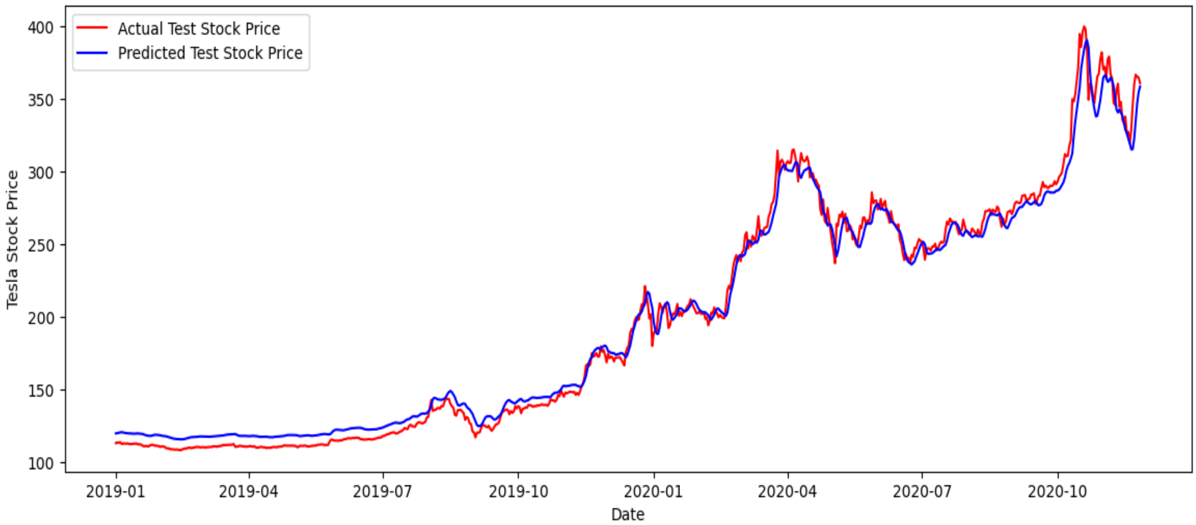

Panel E: Stock Price Prediction using Epoch = 50 (Testing

Source: Author

Figure 4. Tesla Predictions (Panel A-E)

Table 4 (Panel 2) shows Google's stock prediction is 1.5394. Figure 5 depicts the various plots for Google's stock price predictions. Panel F is the MSE loss during training, Panel G is the stock price prediction (training dataset), Panel H is stock price prediction (test data), Panel I is stock price prediction using 50 epoch (training data), Panel J is stock price prediction using 50 epoch (test data). Using Epoch 50 for both the training and validation MSE loss for the Google stock price prediction, Panel A and B depict that the epoch 50 is closer to the true stock value. The evidence shows that Google's stock performance has recently mirrored the broader tech market's behavior, influenced by overall market volatility, competitive pressures, regulatory scrutiny, and financial performance. Over the years, Google's stock has experienced significant growth, leading to substantial returns for investors. Pinpointing a precise "best time" is difficult, as stock performance is influenced by numerous factors and market conditions. The period following Google's IPO in 2004 saw substantial growth, with early investors benefiting significantly.

More so, various periods of innovation and expansion, such as the rise of mobile computing and the growth of Google's advertising business, have coincided with stock price increases (Liao et al 2024). It's important to remember that past performance is not indicative of future results, as general economic uncertainty and interest rate fluctuations contribute to market volatility, impacting Google's stock price. Furthermore, the evolving tech landscape, including competition and innovation in key areas, shapes investor sentiment. Regulatory scrutiny also plays a role, potentially affecting investor confidence (Nguyen et al. 2023). Google's own financial performance, including earnings reports and growth projections, significantly influences its stock price. Staying informed about market trends and seeking professional financial advice are crucial for navigating the complexities of stock market investments (Bakar and Wurgle 2007).

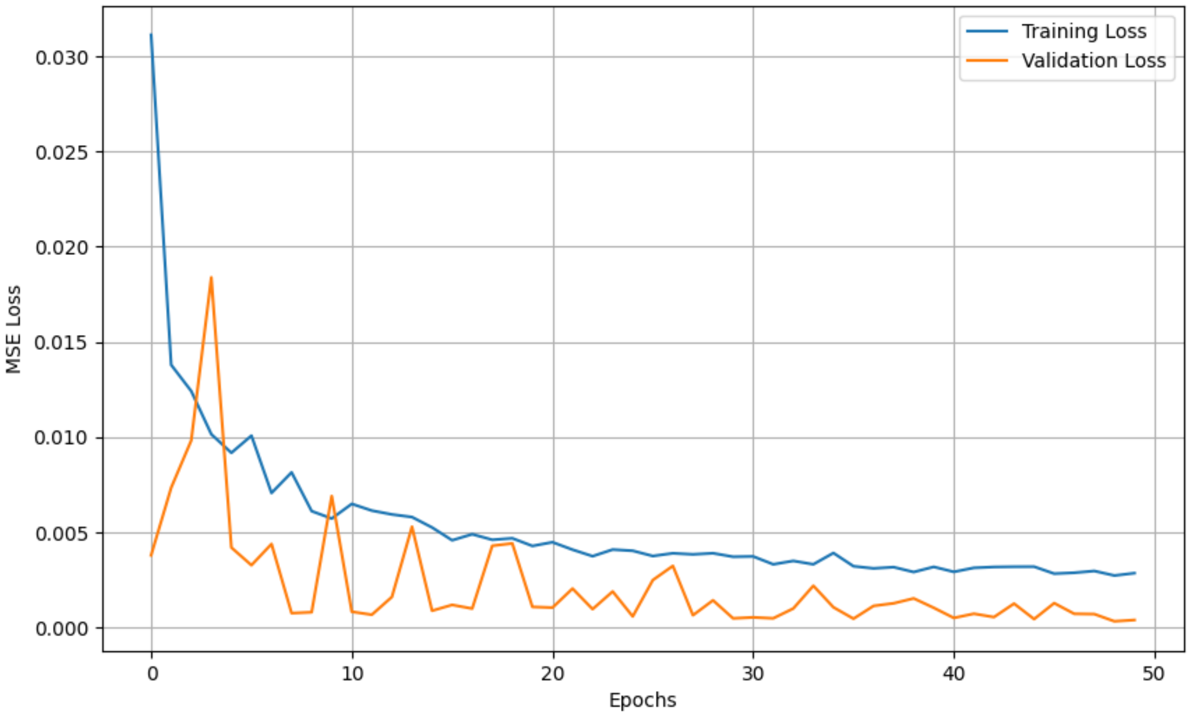

Panel F: MSE Loss During Training (Google)

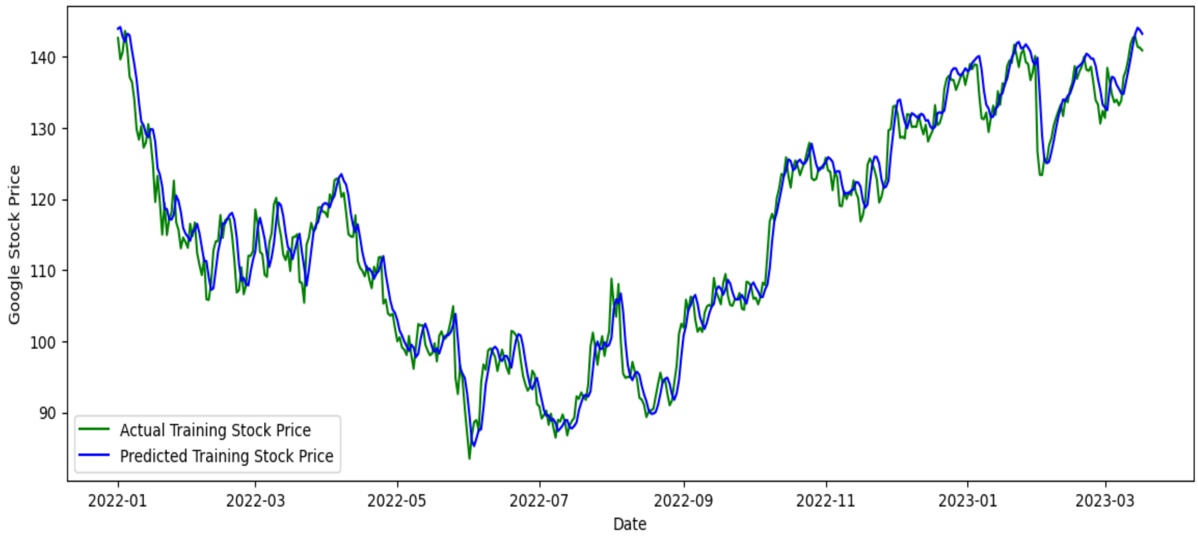

Panel G: Stock Price Prediction (Training)

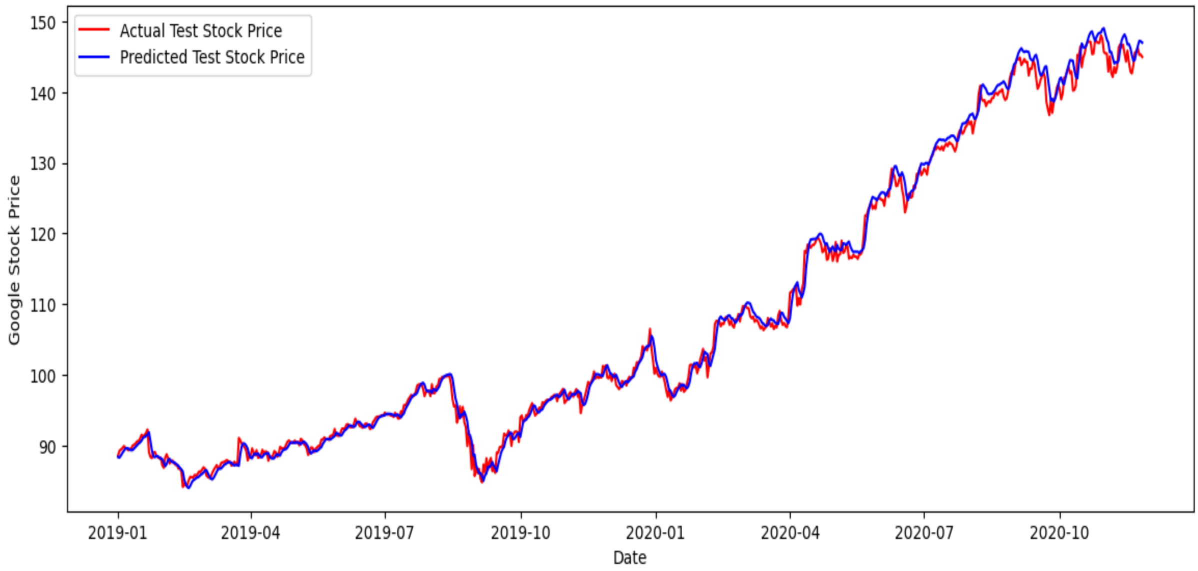

Panel H: Stock Price Prediction (Testing)

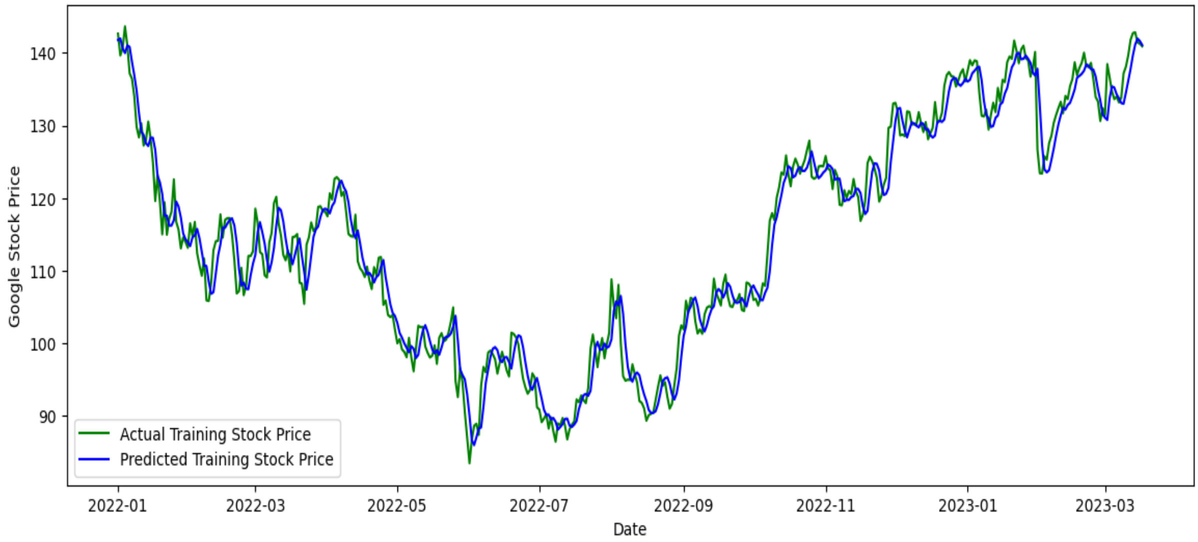

Panel I: Stock Price Prediction using Epoch = 50 (Training)

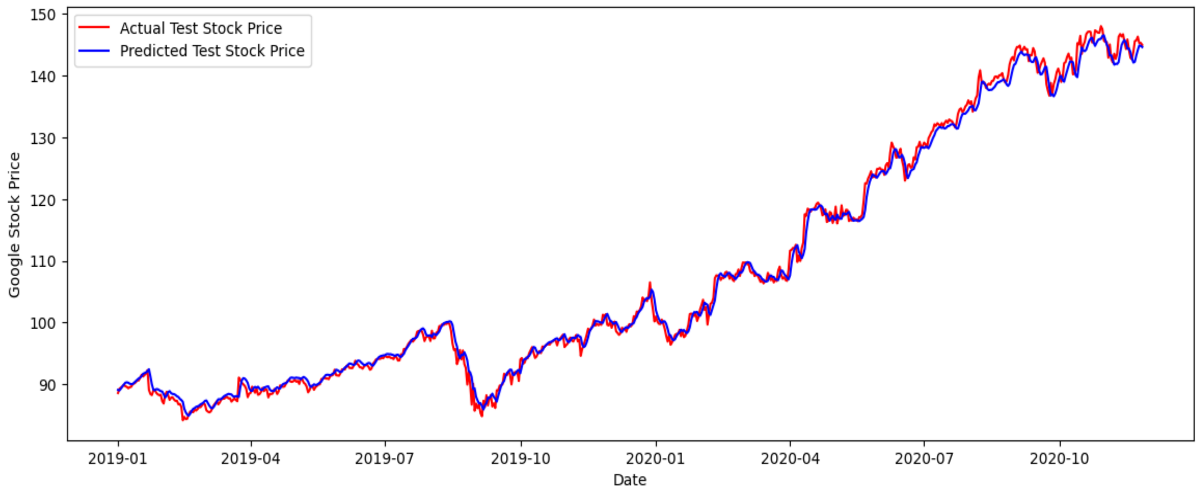

Panel J: Stock Price Prediction using Epoch = 50 (Testing)

Source: Authors

Figure 5. Google Predictions (Panel F-J)

Table 4 (Panel 3) shows Apple's stock prediction is 14. Figure 6 depicts the various plots for Apple's stock price predictions. Panel K is the MSE loss during training, Panel L is the stock price prediction (training dataset), Panel M is stock price prediction (test data), Panel N is stock price prediction using 50 epoch (training data), Panel 0 is stock price prediction using 50 epoch (test data). Using Epoch 50 for both the training and validation MSE loss for the Apple stock price prediction, Panel A and B depict that the epoch 50 is closer to the true stock value. The evidence shows that Apple Inc.'s prominent position within the technology sector and its stock market performance has garnered significant attention. As with all publicly traded equities, Apple's stock price exhibits periods of volatility, influenced by macroeconomic trends, prevailing economic conditions, and company-specific factors (Ouyang et al. 2024).

While Apple has achieved remarkable growth and market success (Barbosa and de Oliveira 2020), its stock remains susceptible to market fluctuations. Factors such as new product releases (Zhang and Zhang 2021), consumer demand, and competitive pressures can influence investor sentiment and impact stock performance (Ouyang 2020). Despite experiencing periods of substantial growth historical performance is not indicative of future returns. A comprehensive understanding of Apple's stock performance necessitates consulting reputable financial news sources and analyst reports. Furthermore, it is crucial to acknowledge that investment decisions inherently involve risk.

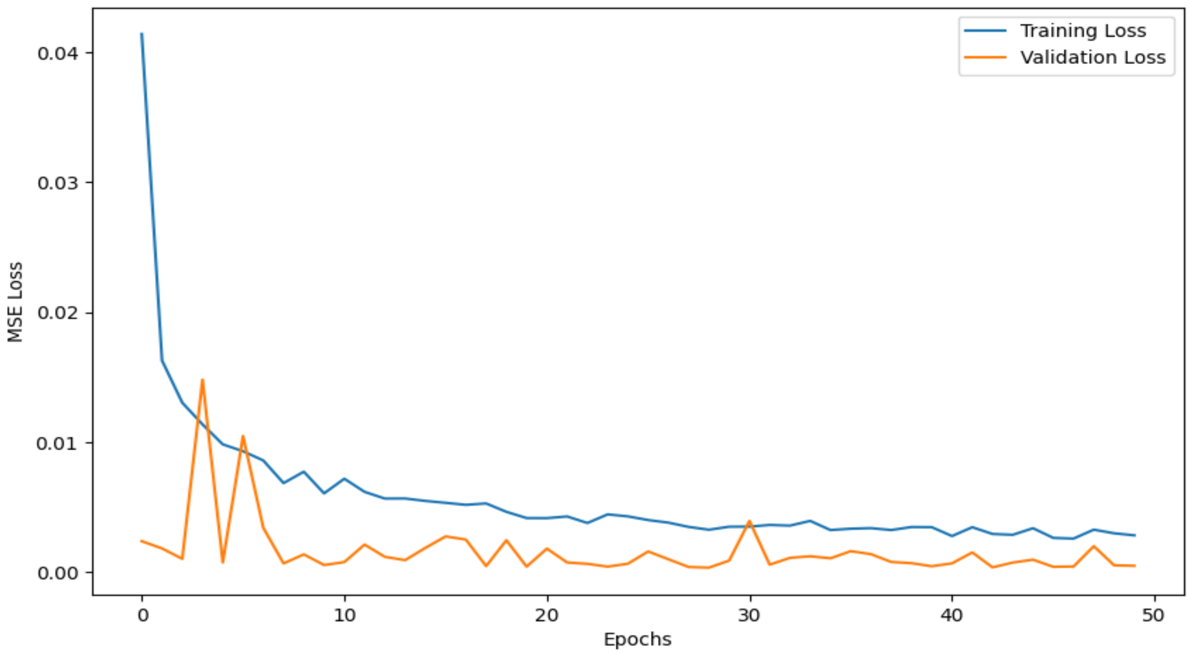

Panel K: MSE Loss During Training (Apple)

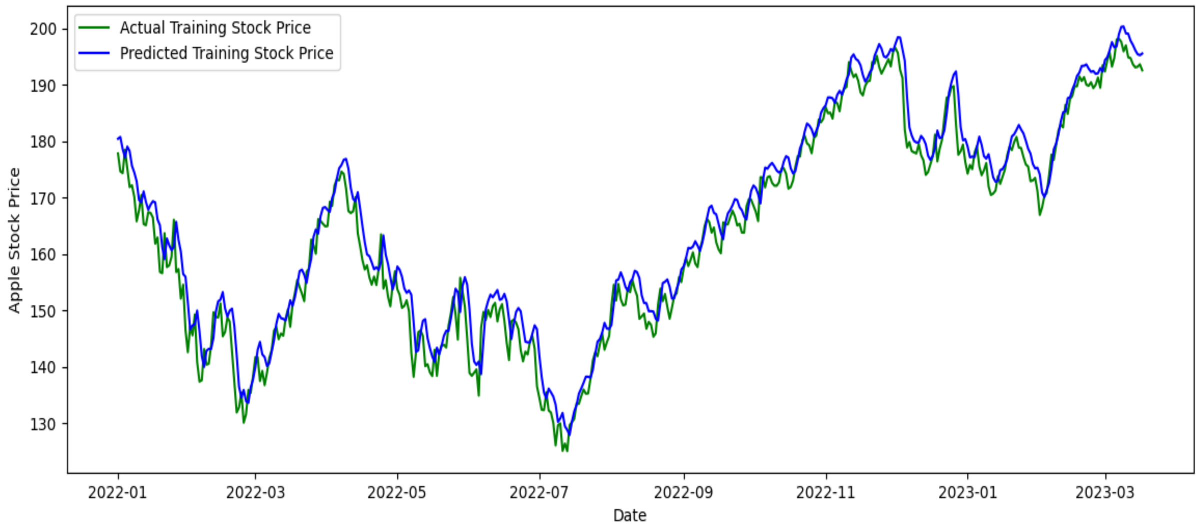

Panel L: Stock Price Prediction (Training)

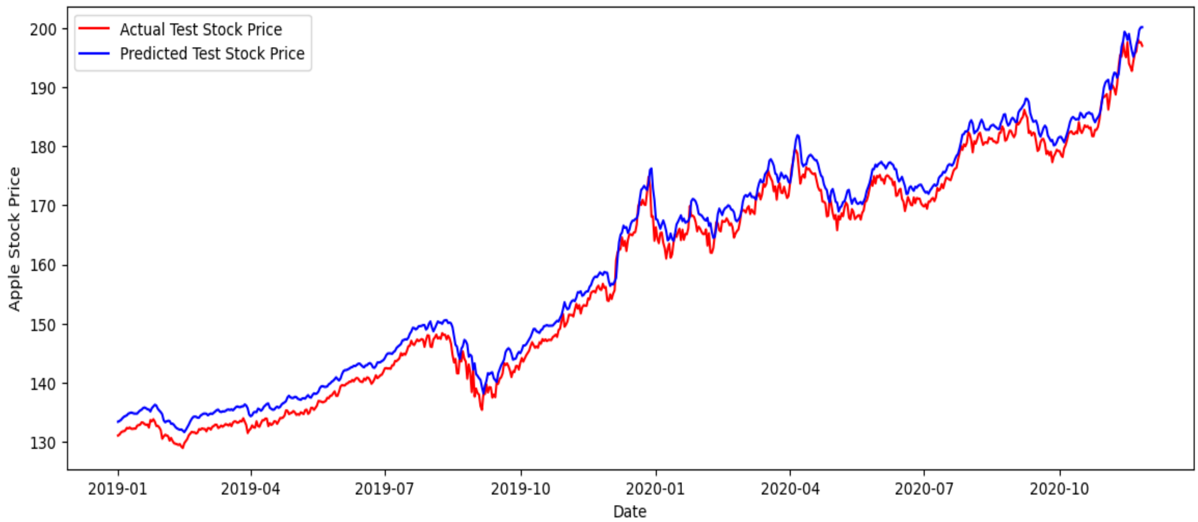

Panel M: Stock Price Prediction (Testing)

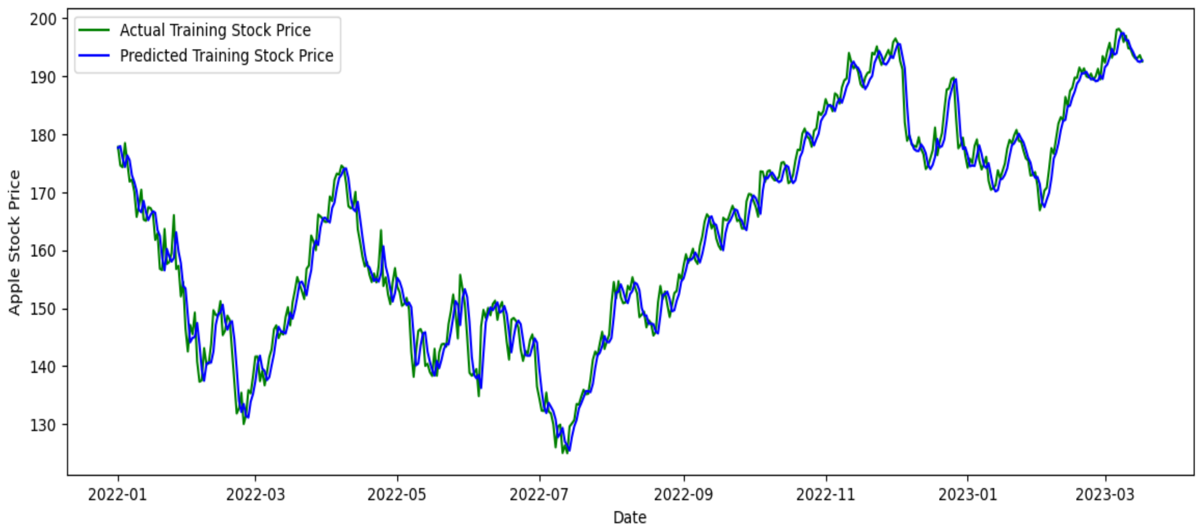

Panel N: Stock Price Prediction using Epoch = 50 (Training)

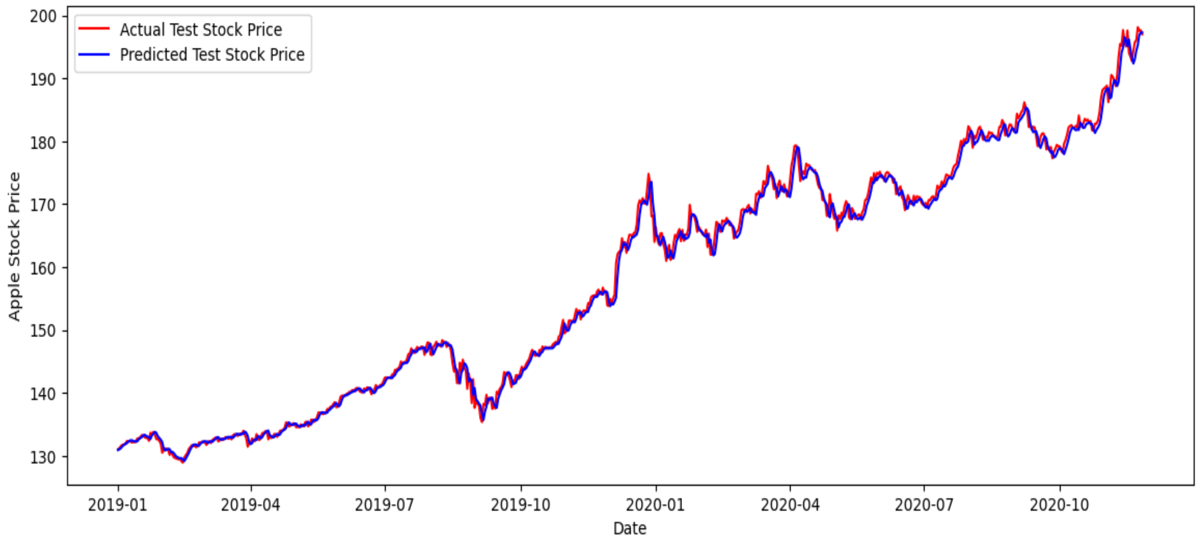

Panel O: Stock Price Prediction using Epoch = 50 (Testing)

Source: Author

Figure 6. Apple Predictions (Panel K-O)

Table 4 (Panel 4) shows Amazon's stock prediction loss is 2.4875. Figure 7 depicts the various plots for Amazon's stock price predictions. Panel P is the MSE loss during training, Panel Q is the stock price prediction (training dataset), Panel R is stock price prediction (test data), Panel S is stock price prediction using 50 epoch (training data), Panel T is stock price prediction using 50 epoch (test data). Using Epoch 50 for both the training and validation MSE loss for the Amazon stock price prediction, the Figures shows the epoch 50 is closer to the true stock value. Since its 1997 IPO, Amazon has evolved from an online bookstore to a dominant force in e-commerce, cloud computing, and digital streaming (Kumar and Ahuja 2020). Its customer-centric approach, innovative business model, and aggressive expansion fueled early growth, propelling its stock price upward (Das et al. 2023). While the dot-com bubble posed a challenge, Amazon's strong fundamentals enabled its recovery and continued market leadership. Despite recent market volatility, its strong performance and dominant position suggest a positive outlook, though challenges like competition and regulation remain (Wang et al 2022).

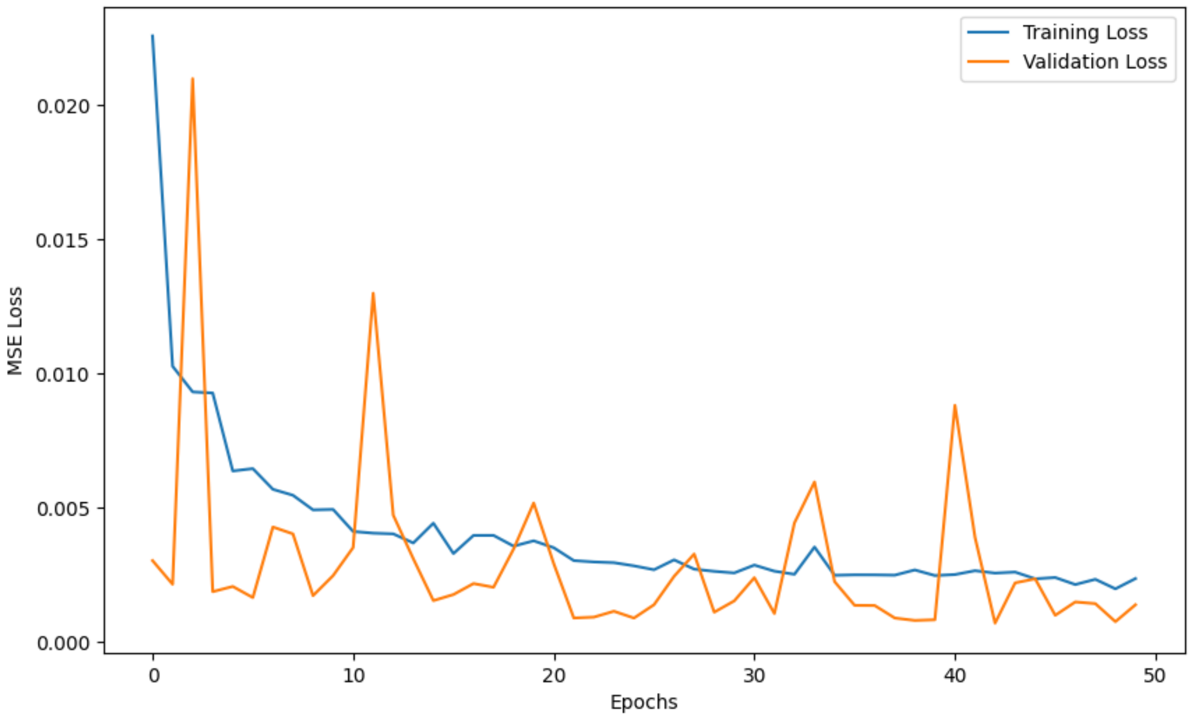

Panel P: MSE Loss During Training (Amazon)

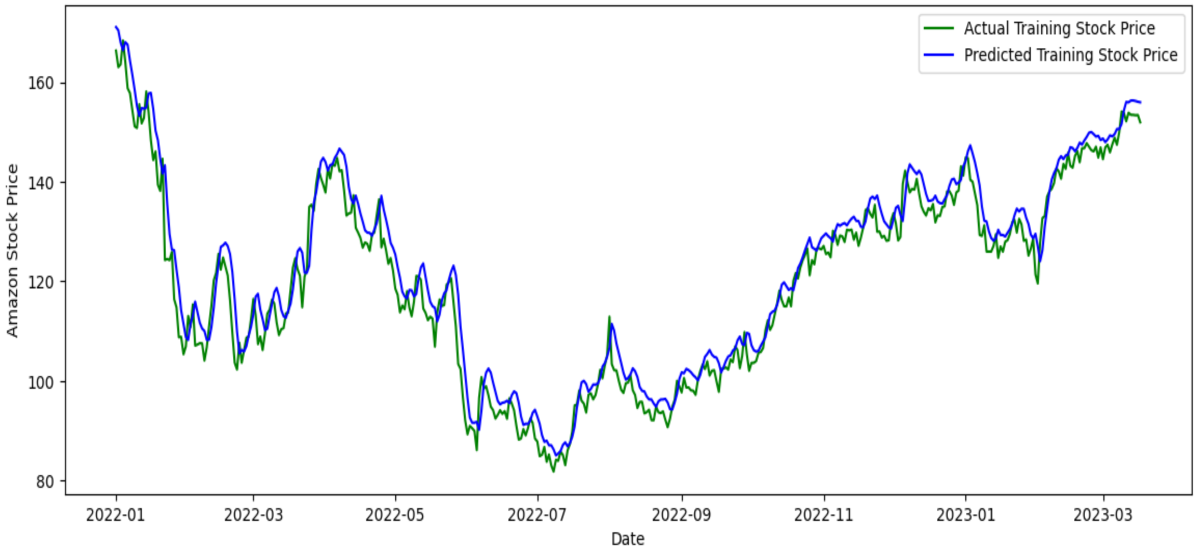

Panel Q: Stock Price Prediction (Training)

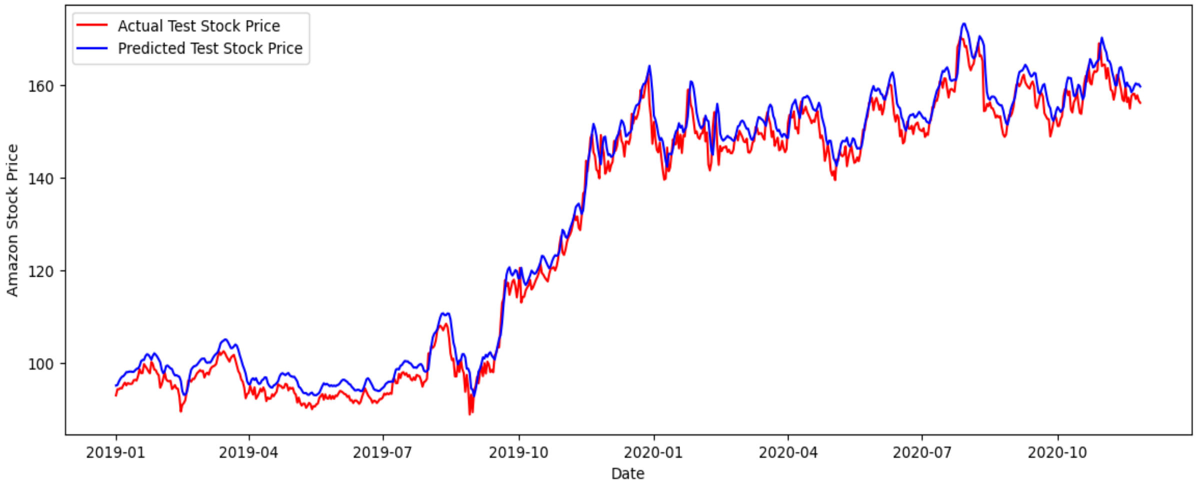

Panel R: Stock Price Prediction (Testing)

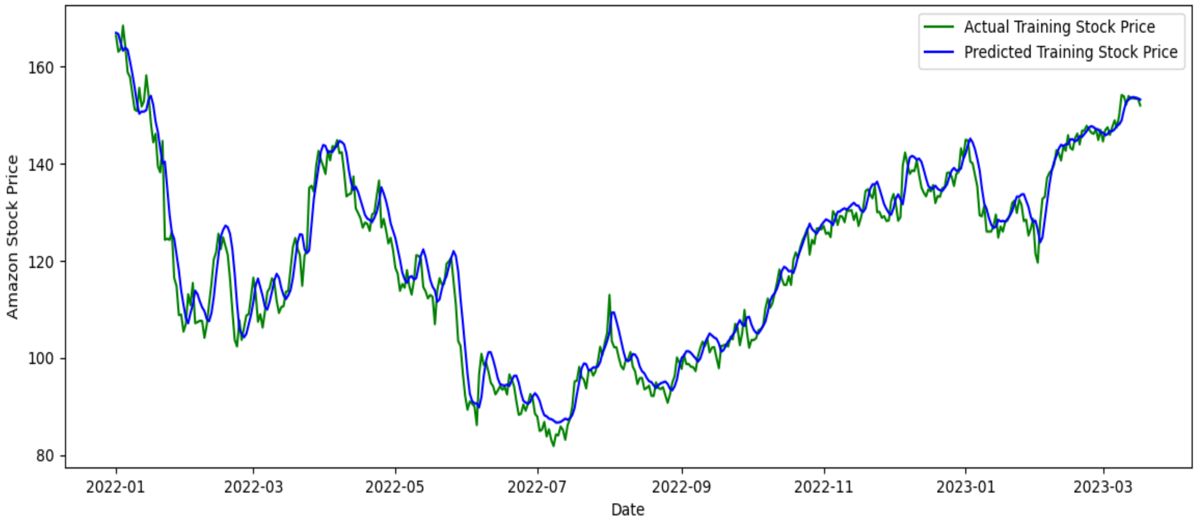

Panel S: Stock Price Prediction using Epoch = 50 (Training)

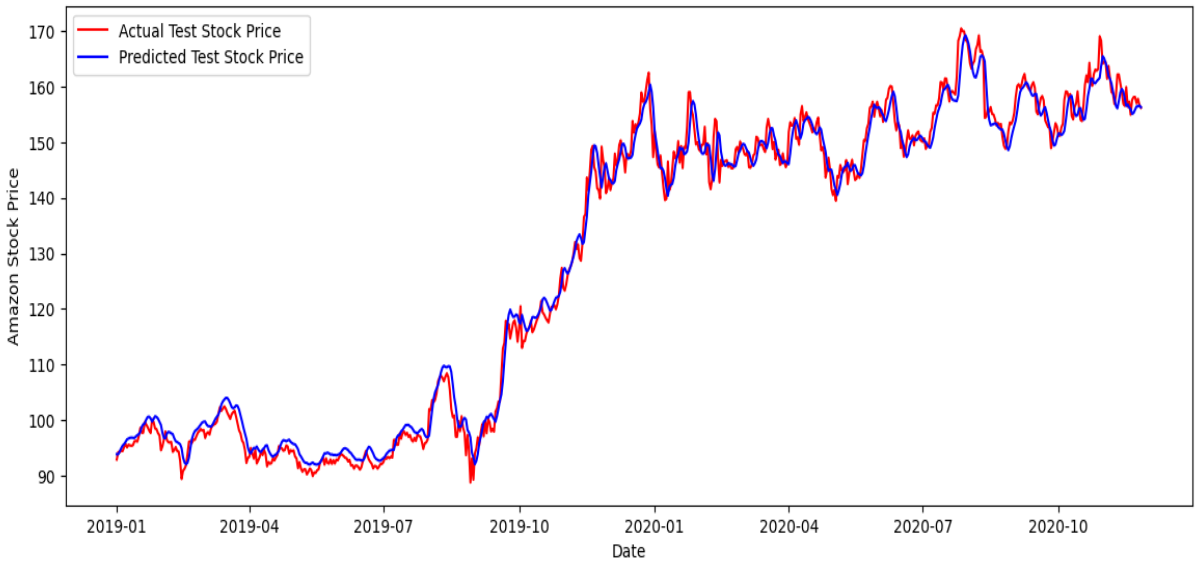

Panel T: Stock Price Prediction using Epoch = 50 (Testing)

Source: Authors

Figure 7. Apple Predictions (Panel P-T)

Table 4. Company Stock (Epoch and MSE)

	MSE	
Epoch	**Train**	**Validate**
Panel 1: Tesla		
1	0.0449	0.0053
50	0.0019	0.0017
Panel 2: Google		
1	0.0653	0.0038
50	0.0023	0.0004
Panel 3: Apple		
1	0.0908	0.0024
50	0.0026	0.0004
Panel 4: Amazon		
1	0.0383	0.0030
50	0.0022	0.0014

Source: Authors (2025)

Policy and Managerial Implications

The findings from the LSTM predictive framework reveal critical insights for policymakers, financial regulators, and firm managers operating in technology-driven capital markets. From a policy perspective, the high volatility and structural breaks detected, particularly during 2020–2022, underscore the sensitivity of equity markets to macroeconomic shocks such as pandemic disruptions and monetary tightening. Regulators must therefore strengthen macroprudential frameworks to mitigate systemic risks originating from speculative behavior in high-growth sectors (Borio 2014). By integrating real-time financial analytics powered by deep learning, central banks can enhance their surveillance systems to detect abrupt regime shifts or “flash crashes,” thus improving crisis response mechanisms (Adrian and Liang 2018).

Secondly, the presence of multiple structural breaks emphasizes the role of monetary policy in shaping investor expectations. The breakpoints around 2020–2022 correspond to interest rate adjustments and inflationary pressures, consistent with findings that liquidity tightening can induce valuation corrections in growth-oriented firms (Baker et al. 2022). Policymakers should therefore recognize that abrupt policy normalization can amplify volatility in technology equities that are priced heavily on future earnings potential. This calls for gradual policy signaling and enhanced forward guidance to allow market participants to recalibrate expectations smoothly, reducing the probability of herding and abrupt capital flight (Blinder et al. 2008).

From a managerial standpoint, the predictive evidence offers actionable guidance for strategic decision-making and risk management. The LSTM model’s superior predictive accuracy indicates that firm managers can leverage such models to anticipate short-term price movements, enabling better timing of share repurchases, capital issuance, or hedging activities. For example, the elevated MSE loss for Tesla signals the heightened uncertainty faced by firms engaged in emerging technologies with rapidly evolving market sentiments. Managers in such contexts should adopt flexible financing and inventory policies to accommodate abrupt demand or valuation shifts, in line with dynamic capability theory (Teece et al. 1997). On a broader economic level, the results carry implications for financial stability and innovation policy. The high volatility and episodic breaks in Tesla and Amazon prices mirror speculative tendencies in markets with incomplete information and innovation-driven narratives (Shiller 2000). Policymakers should thus design innovation-supportive yet stability-oriented interventions, including transparent disclosure requirements for emerging technologies and clearer regulatory oversight of AI-driven trading systems. Ensuring that innovation incentives do not foster unsustainable asset bubbles remains a central challenge for regulators in post-pandemic financial governance (Lo 2004).

Furthermore, the documented regime shifts reinforce the necessity of adaptive market regulations. As financial markets evolve through non-linear dynamics, static regulatory frameworks may become obsolete. Integrating AI-driven monitoring tools within stock exchanges and central depositories can enable the early detection of market anomalies, such as algorithmic mispricing or coordinated speculative trading (Zaremba et al. 2020). Such initiatives align with the adaptive market hypothesis, which emphasizes the need for continuous learning and adjustment in market institutions (Lo 2004). By adopting such adaptive mechanisms, regulators can promote more resilient and transparent markets.

From the investor and portfolio management perspective, the LSTM's predictive capabilities highlight new avenues for algorithmic risk assessment. Investors can integrate model outputs into portfolio optimization frameworks to rebalance assets in anticipation of volatility surges, especially around known breakpoints or policy announcements. The empirical evidence that all series become stationary after first differencing supports the notion that return-based forecasting is more reliable than price-level modeling. This reinforces the practical necessity of pre-processing financial data before deploying AI-based trading algorithms to prevent bias and overfitting (Hochreiter and Schmidhuber 1997).

Finally, at the intersection of public policy and firm strategy, the results advocate for a co-evolutionary approach between innovation policy and financial governance. The post-2020 structural breaks demonstrate that exogenous shocks, such as pandemics or policy shifts, can destabilize even the most robust technology firms. Policymakers and corporate leaders should therefore co-develop resilience frameworks that integrate predictive analytics into both corporate strategy and macroeconomic planning. This includes stress-testing corporate valuations under simulated macroeconomic shocks and encouraging transparency in data-driven decision-making (Das et al. 2023). Ultimately, the synergy between AI-enabled forecasting models and prudent policy design can enhance economic resilience in an increasingly data-driven global financial system.

CONCLUSIONS

This research contributes to the growing body of evidence supporting the potential of LSTM networks for enhancing stock price prediction. By leveraging the LSTM's capacity to model complex temporal relationships, investors and financial analysts can gain a more informed perspective on market trends and individual stock behavior (Lim et al 2021). The focus on Tesla, Google, Apple, and Amazon provides a practical context for evaluating the model's real-world applicability. Notably, advanced techniques like LSTMs are not foolproof because external factors, market sentiment, and unforeseen events can significantly impact stock prices. Therefore, combining predictions with thorough fundamental analysis and risk management strategies is essential for sound investment decisions (Kumar et al. 2023).

However, this study is not without limitations. The analysis is confined to daily data spanning 2019–2023, which, although comprehensive, does not capture post-2023 structural and policy developments due to the data unavailability at the time of analysis. Extending the dataset to 2024/12, as suggested, was not feasible because complete and validated stock price records for all four firms were not yet released in consistent format across data repositories during the period of study completion (January 2025). Furthermore, the model focuses exclusively on historical price dynamics and does not integrate other relevant financial variables such as trading volume, macroeconomic indicators, or sentiment data. These omissions may limit the model's ability to fully capture multidimensional drivers of stock behavior, particularly during high-volatility regimes.

Another limitation concerns the inherent black-box nature of LSTM models, which constrains interpretability and may obscure the underlying causal mechanisms behind predicted price movements (Makridakis et al. 2018). Additionally, while dropout regularization and differencing were employed to mitigate overfitting and non-stationarity issues, further robustness checks using alternative machine learning models could provide a more comprehensive understanding of model generalizability. These limitations should guide readers to interpret the findings within the specific temporal and methodological context of the study.

Future research should expand the scope by integrating hybrid deep learning architectures such as LSTM–GRU or attention-based Transformers to enhance model transparency and accuracy (Zhang et al. 2023). Researchers could also explore cross-market applications by comparing the performance of LSTM models on emerging versus developed markets, enabling broader policy and investment implications. In addition, incorporating real-time sentiment analysis from social media, macroeconomic indicators, and central bank communication variables could significantly enrich predictive performance and interpretive power. Longitudinal studies examining post-2023 data would further clarify how structural changes, including monetary tightening and AI-driven trading adoption, influence predictive dynamics.

Acknowledgements: *The author gratefully acknowledges Godwin Nwachukwu Nkem for his valuable support, insightful suggestions, and constructive feedback throughout the course of this research.*

REFERENCES

Adrian, T., and N. Liang. 2018. Monetary policy, financial conditions, and financial stability. *International Journal of Central Banking, 14*(1): 73–131.

Baker, M., and J. Wurgler. 2007. Investor sentiment in the stock market. *Journal of Economic Perspectives, 21*(2): 129–151. https://doi.org/10.1257/jep.21.2.129

Baker, S. R., N. Bloom, S. J. Davis, and S. J. Terry. 2022. COVID-induced economic uncertainty and its impact on financial markets. *Review of Asset Pricing Studies, 12*(1), 1–38. https://doi.org/10.3386/w26983

Barbosa, N., and R. de Oliveira, R. 2020. Financialization and innovation: Evidence from the U.S. tech sector. *Technological Forecasting and Social Change, 161:* 120254. https://doi.org/10.1016/j.techfore.2020.120254

Blinder, A. S., M. Ehrmann, M. Fratzscher, J. de Haan, and D. J. Jansen. 2008. Central bank communication and monetary policy. *Journal of Economic Literature, 46*(4): 910–945. https://doi.org/10.1257/jel.46.4.910

Borio, C. 2014. The financial cycle and macroeconomics: What have we learnt? *Journal of Banking & Finance, 45*: 182–198. https://doi.org/10.1016/j.jbankfin.2013.07.031

Chen, Y., T. Chen, X. Wang, J. Li. 2020. Stock price prediction using deep learning models. *Expert Systems with Applications, 159*: 113538. https://doi.org/10.1016/j.eswa.2020.113538

Das, N., B. Goswami, and R. N. A. Begum. 2023. Stock prices prediction using long short-term memory. In *Proceedings of the 4th International Conference on Computing and Communication Systems (I3CS)* (pp. 1–5). https://doi.org/10.1109/I3CS58314.2023.10127443

Fama, E. F. 1970. Efficient capital markets: A review of theory and empirical work. *Journal of Finance, 25*(2): 383–417. https://doi.org/10.1111/j.1540-6261.1970.tb00518.x

Fama, E. F., and K. R. French. 2015. A five-factor asset pricing model. *Journal of Financial Economics, 116*(1): 1–22. https://doi.org/10.1016/j.jfineco.2014.10.010

Fischer, T., and C. Krauss. 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

Goswami, D. K. 2023. *Predicting the stock price/movements of companies listed under BSEIT*

Goyal, A. 2004. Demographics, stock market flows, and stock returns. *Journal of Financial and Quantitative Analysis, 39*(1): 115–142. https://doi.org/10.1017/S0022109000003980

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

Hochreiter, S., and J. Schmidhuber. 1997. Flat minima. *Neural Computation, 9*(1): 1–42. https://doi.org/10.1162/neco.1997.9.1.1

Jain, K. 2017. Gods in the time of automobility. *Current Anthropology, 58*(S15): S13–S26. https://doi.org/10.1086/688544

Joosery, B., and G. Deepa. 2019. Comparative analysis of time-series forecasting algorithms for stock price prediction. *Advanced Information Science and System Conference Proceedings*, 1–6. https://doi.org/10.1145/3373477.3373699

Liao, Y., J. Wang, L. Liao, X. Shu, and T. Peng. 2024. Policy synergy on stock price crash risk: An intergovernmental perspective. https://doi.org/10.1016/j.frl.2024.105744

Lim, S., S. J. Kim, Y. Park, and N. Kwon. 2021. A deep learning-based time series model with missing value handling techniques to predict various types of liquid cargo traffic. *Expert Systems with Applications*, *184*: 115532.

Kumar, S. K., B. Meher, R. Birau, M. Simion, I. Florescu, A. Anand, and S. Kumar. 2023. Long-term volatility forecasting of Brazilian stock index behavior. *Revista de Stiinte Politice, 79*: 9–24.

https://doi.org/10.1016/j.eswa.2021.115532

Lo, A. W. 2004. The adaptive markets hypothesis. *Journal of Portfolio Management, 30*(5): 15–29. https://doi.org/10.3905/jpm.2004.412486

Lupton, N. C., V. Baulkaran, and Y. No. 2022. Subsidiary staffing, location choice, and shareholder rights effectiveness. *Journal of Business Research, 151*: 222–231. https://doi.org/10.1016/j.jbusres.2022.06.019

Makridakis, S., E. Spiliotis, and V. Assimakopoulos. 2018. Statistical and machine learning forecasting methods. *PLOS ONE, 13*(3): e0194889. https://doi.org/10.1371/journal.pone.0194889

Nelson, C. R., and C. I. Plosser. 1982. Trends and random walks in macroeconomic time series. *Journal of Monetary Economics, 10*(2): 139–152. https://doi.org/10.1016/0304-3932(82)90012-5

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

Nguyen, H. T., A. N. N. Le, Q. T. Luu, N. T. T. Nguyen, and K. D. Duong. 2023. Foreign ownership, investor attention and the risk-taking behavior of property and casualty insurance firms: evidence from Vietnam. *Sage Open*, *13*(4): 21582440231211971. https://doi.org/10.1177/21582440231211971

Nurazi, R., and B. Usman. 2016. Bank stock returns in responding the contribution of fundamental and macroeconomic effects. *Journal of Economics and Policy*, 9(1): 134-149.

Ouyang, Z., X. Wang, and Y. Liu. 2024. The use of corporate social responsibility in response to product-harm crisis: How do stock market reactions matter?. *Corporate Social Responsibility and Environmental Management*, *31*(4): 3081-3097. https://doi.org/10.1002/csr.2739

Ouyang, Z. 2022, March. A Study of Stock Portfolio Strategy Based on Machine Learning. In *2022 7th International Conference on Financial Innovation and Economic Development (ICFIED 2022)* (pp. 79-87). Atlantis Press. https://doi.org/10.2991/aebmr.k.220307.013

Shao, H., and B.-H. Soong. 2016. Traffic flow prediction with long short-term memory networks. In *TENCON 2016 – IEEE Region 10 Conference* (pp. 2986–2989). https://doi.org/10.1109/TENCON.2016.7848588

Shiller, R. J. 2003. From efficient markets theory to behavioral finance. *Journal of Economic Perspectives, 17*(1): 83–104. https://doi.org/10.1257/089533003321164967

Teece, D. J., G. Pisano, and A. Shuen. 1997. Dynamic capabilities and strategic management. *Strategic Management Journal, 18*(7): 509–533. https://doi.org/10.1002/(SICI)1097-0266(199708)18:7<509::AID-SMJ882>3.0.CO;2-Z

Wang, Y., and X. Li. 2021. Forecasting stock price trends with deep learning. *Applied Soft Computing, 108*: 107443. https://doi.org/10.1016/j.asoc.2021.107443

Wang, Y., Y. Wei, C. Wu, and L. Yin. 2018. Oil and the short-term predictability of stock return volatility. *Journal of Empirical Finance, 47*: 90–104. https://doi.org/10.1016/j.jempfin.2018.04.008

Zaremba, A., R. Kizys, and D. Aharon. 2020. COVID-19 and global financial markets. *Finance Research Letters, 35*: 101508. https://doi.org/10.1016/j.frl.2020.101508

Zeng, Y., and X. Liu. 2018. A-stock price fluctuation forecast model based on LSTM. In *Proceedings of the 14th International Conference on Semantics, Knowledge and Grids (SKG)* (pp. 261–264). https://doi.org/10.1109/SKG.2018.00043

Zhang, G., B. E. Patuwo, and M. Y. Hu. 2020. Forecasting with artificial neural networks. *International Journal of Forecasting, 36*(1): 1–26. https://doi.org/10.1016/j.ijforecast.2019.07.001

Zhang, Y., and X. Zhang. 2021. The role of product innovation in stock market performance. *Technovation, 106*: 102292. https://doi.org/10.1016/j.technovation.2021.102292

Zou, J., H. Cao, L. Liu, Y. Lin, E. Abbasnejad, and J. Shi. 2022. Astock: A new dataset and automated stock trading model. *arXiv Preprint arXiv:2206.06606*. https://doi.org/10.48550/arXiv.2206.06606