alphanumeric
https://alphanumericjournal.com/

Volume 14
Issue 01
Year 2026

Research Article

Global technology company stock price forecasting using long short-term memory (LSTM) architecture under macro financial variables

Aynur İncekırık [1]

[1] Asst. Prof., Department of Econometrics, Faculty of Economics and Administrative Sciences, Manisa Celal Bayar University, Manisa, Türkiye

Abstract

This study aims to analyze the price dynamics of technology stocks such as Apple (USA), Samsung Electronics (South Korea), Xiaomi (China), Sony (Japan), LG Electronics (South Korea), and Nokia (Finland) using deep learning models. The analysis also includes exchange rates such as the US Dollar Index (DXY) along with the Chinese Yuan (USD/CNY), Japanese Yen (USD/JPY), and South Korean Won (USD/KRW) against the US Dollar. The study uses daily opening and closing prices for the period 09.06.2018–09.02.2026; relationships between variables were examined using Pearson correlation analysis. Stock prices, dollar index, and exchange rate data were obtained from the "Yahoo Finance" website. The results show a strong positive correlation between Apple and Sony, Apple and Samsung, and Samsung and Sony. In contrast, Xiaomi has a moderate positive correlation with Apple and Samsung, while the relationship between Apple and LG and Apple and Nokia remains weak. These findings reveal that sector-based common factors are influential in pricing. In the forecasting process conducted with three different layered LSTM models, the data was divided into training and test sets while maintaining chronological integrity; the models were evaluated using RMSE, MAE, and MAPE metrics. The results show that the performance of LSTM models varies by variable; while the single-layer LSTM architecture produced lower errors in most stock and currency series, two- or three-layer structures proved superior in some series. Overall, the study demonstrates that deep learning approaches are effective in modeling the price behavior of global technology companies and offer a strategic decision support tool in sustainable finance.

Keywords Macro Financial Variables, Deep Learning, Long Short-Term Memory (LSTM), Global Technology Company, Stock Price Prediction, Machine Learning, Time Series Analysis

Jel Codes G17, C45, E44

2026.14.01.MIS.03

Correspondence

A. İncekırık
aynur.incekirik@cbu.edu.tr

Timeline

Submitted	Mar 02, 2026
Revision Requested	Jun 03, 2026
Last Revision Received	Jun 03, 2026
Accepted	Jun 15, 2026
Published	Jun 30, 2026

Citation

İncekırık, A. (2026). Global technology company stock price forecasting using long short-term memory (LSTM) architecture under macro financial variables, *alphanumeric journal*, 14 (1), 41-64. https://doi.org/10.17093/alphanumeric.1901302

ORCID

A. İncekırık 0000-0002-5029-6036

Contents

1. Introduction

With the rapid development of the economy, stock markets or equity markets play an important role today. More and more people are choosing to invest in stocks as time goes by. Stock prices reflect each company's performance, and fluctuations in prices are of vital importance as they are closely related to investors' profits. Businesses and investors aiming to maximize profits want to predict future trends in stock prices; this allows investors to react in advance, keep pace with market opportunities, and achieve large financial gains with relatively low risk. Therefore, in the capital market, accurately predicting stock prices is important for investors' investment decisions and risk management. However, this is quite a difficult task because tons of transactions occur every second; these transactions are fraught with politics, unexpected events, the global economic situation, volatility, and uncertainty (Yang, 2022). Applying traditional forecasting methods to today's increasingly complex capital market is often difficult, making it challenging to capture the complex dynamics of stock prices. In particular, algorithms known as machine learning provide more accurate prediction results than were previously possible.

In recent years, rapid advances in machine learning and deep learning techniques have created a remarkable development in market forecasting by providing innovative methods for analyzing financial trends and market behavior. These techniques can examine voluminous historical data, identify trends, and use these predictions to forecast future stock values. These algorithms work by discovering patterns and features among countless data points through advanced computational processes. Predicting stock prices is an important area of research in finance, and selecting appropriate machine learning models is crucial for improving prediction accuracy. Researchers have explored the potential of machine learning algorithms to predict stock prices as their needs have grown (Aggarwal, 2019). Information about the dynamics and direction of the relationship between macroeconomic variables and stock prices is very important for policymakers. Accurate prediction models are critical in the financial sector for navigating the complex structure of market behavior and maximizing returns.

With technology constantly evolving, the importance of technology companies is increasing day by day. The current financial situation of technology companies is as important to investors as it is to managers. Investors need to benefit from analysis results showing the financial performance of the technology companies they are interested in to make the right decisions and profitable investments. Xiaomi, a global technology company, is a well-known Chinese technology company focused on electronic products such as smartphones and showing high growth. However, it has shown unstable performance due to market competition and challenges. Headquartered in the United States, Apple is known as a multinational company in computer software, technological products, and mobile phone manufacturing. This company is a global technology giant and is best known for its iPhone and other Apple products. It also benefits from a strong brand, a stable customer base, and diversified revenue streams. Therefore, it offers solid growth rates. Samsung, headquartered in South Korea, is a multinational company in information and communication technologies and mobile phone manufacturing. This company operates in various sectors, including electronics and telecommunications. Samsung's diversified portfolio makes it a balanced choice for risk-averse investors (Na, 2025).

The primary objective of this study is to comparatively examine the performance of three different LSTM models in predicting stock prices of global technology companies using macro-financial variables and to reveal the extent to which these architectures can capture different time series dynamics. Within the scope of the study, the explanatory power of the LSTM model, one of the unidirectional recurrent neural network architectures, on the nonlinear structure of stock prices is empirically evaluated. Thus, it is analyzed which model can more effectively represent the complex and dynamic structure of the technology market, and it is aimed to contribute to the literature from a methodological perspective. The findings are expected to be instructive for both academic studies and investment decisions.

2. Literature Review

In recent years, stock price forecasting using deep learning techniques has been extensively researched and has shown a marked increase in the academic literature, attracting widespread interest from academics, investors, and government agencies. Modeling and forecasting financial time series is among the topics intensively researched in the machine and deep learning literature. Reddy et al. (2020) achieved some success in their work using traditional techniques for predicting the next lag of time series data, such as statistical algorithms like Exponential Smoothing and Autoregressive Integrated Moving Average (ARIMA). With the emergence of deep learning architectures and advanced computing processors, they analyzed the performance of these techniques in stock market forecasting. This study presents a performance comparison of Exponential Smoothing, ARIMA, Vanilla LSTM, and Stacked LSTM models. Empirical analysis reveals the superior performance of deep learning techniques with an RMSE score as low as 3.208 on daily closing price stock data over a ten-year period. In this study, LSTM was used to predict stock market activity. The findings show that advanced versions of LSTM provide more detailed results than standard algorithms. They have indicated that this paradigm is efficient for both private and institutional investors.

Kang (2022) the opening prices of three Asian countries' stock market indices (KOSPI, Hang Seng, and Nikkei) using US financial market indicators. Exchange rates against the dollar in each country were also included in the model. The architectures considered were linear regression, artificial neural network, recurrent neural network, long short-term memory, support vector machine, and random forest. For ANN and LR, Korea's KOSPI and Japan's Nikkei showed the least performance of 19.9% in both RMSE and MAE. Hang Seng's performance was 3.8%. When US financial indicators were included, the performance of all three countries improved. LR and ANN showed superior performance. In the KOSPI index, ANN performed better than linear regression only in RMSE, while in the Nikkei index, ANN performed better than LR in MAE. The ANN architecture showed the best performance in predicting stock market indices.

Xiao et al. (2022) used MAE, MSE, and RMSE performance indicators to analyze the performance of different stocks predicted by LSTM and ARIMA models in their study. In the analysis, they took 50 publicly traded company stocks from finance.yahoo.com. The dataset used in this study consists of the highest prices on trading days corresponding to the period between January 1, 2010, and December 31, 2018. For the LSTM model, data from January 1, 2010, to December 31, 2015, was selected as the training set, data from January 1, 2016, to December 31, 2017, as the validation set, and data from January 1, 2018, to December 31, 2018, as the test set. For the ARIMA model, data from January 1, 2016, to December 31, 2017, was selected as the training set, and data from January 1, 2018, to December 31, 2018, was selected as the test set. For both models, 60 days of data were used to predict the next day. In conclusion, both the ARIMA and LSTM models were found to be consistent with both predicted and actual values. LSTM was found to perform better than ARIMA in predicting stock prices.

In their work, Saracık & İncekırık (2023) applied LSTM and GRU methods to predict the stock prices of five companies in the XELKT index on the Istanbul Stock Exchange using the Google Colab software program. The data set used covers the period from January 2, 2013, to December 30, 2022. The most successful predictions were determined in analyses conducted over four different days. Since the model performance metrics obtained were below 1 for MSE and below 5% for MAPE, it was concluded that both techniques performed successfully. When these two techniques were compared, it was observed that the LSTM technique was slightly more successful than the GRU technique.

Cao (2024) used a deep learning-based formulation for stock price prediction in his study. By applying three popular and advanced models—CNN, RNN, and LSTM—he trained them on data from six companies obtained from Yahoo Finance to predict their future prices. He applied the Sliding Window model to make future predictions in time series. The results of the models were calculated using MSE, MAE, and MAPE. It was found that the network architectures could capture the hidden factor and make accurate predictions. Looking at the findings, the RNN model showed the best performance in identifying changes in trends, as it predicts future prices by learning from short-term historical data. The LSTM model had lower accuracy because it evaluates long-term and short-term historical data together, and the stock market is a highly volatile sector. The CNN model performed well in processing high-dimensional data and gradient descent iteration; this facilitated local convergence rather than global convergence of results. Therefore, despite its good performance in the image and video domain, its performance in stock price prediction was not as good as the other two models.

Yousufi & İncekırık (2024) studied the fundamental components of the digital economy and the prediction of Ethereum's price movements in the cryptocurrency market using deep learning techniques. In the study, RNN, LSTM, and GRU models were used for Ethereum price predictions, and the performance of these models on time series data was measured. The data underwent various preprocessing steps such as scaling, sequence generation, and model training, and performance evaluation was analyzed using metrics such as MSE, MAE, and MAPE. The analysis results revealed the potential and reliability of deep learning algorithms in financial market predictions.

Farhan et al. (2025) used a deep learning-based LSTM model to predict the stock prices of three major Korean companies (Samsung Electronics, Naver, and SK Hynix) in their study. They obtained data collected from the Yahoo Finance database from the KOSPI market over 14 years. Their aim was to create a new feature by averaging the opening and closing prices and to evaluate its impact on prediction accuracy. Optuna was used for hyperparameter optimization to improve model performance. The results showed a significant improvement in prediction performance. Looking at Samsung Electronics, the RMSE decreased from 0.954 to 0.007 and the MAPE decreased from 2.25 to 1.75, while the R2 value increased from 0.949 to 0.962. The same was true for Naver and SK Hynix. The findings suggest daily stock price prediction by averaging the opening and closing prices.

Qi & Hu (2025) explored the potential of artificial intelligence in financial forecasting by applying a Long Short-Term Memory (LSTM) neural network to the task of predicting Apple Inc. (AAPL) stock prices in their study. The dataset used for training and testing consisted of daily closing prices from August 2020 to August 2023. The performance metrics used were Mean Squared Error (MSE), Mean

Absolute Error (MAE), and Root Mean Squared Error (RMSE). The results showed that the LSTM model performed quite well in predicting stock prices.

Saberironaghi et al. (2025) present a comprehensive analysis of various machine learning and deep learning approaches used in stock market forecasting, focusing on their methodologies, evaluation metrics, and datasets. Popular models such as LSTM, CNN, and SVM are examined, highlighting their strengths and weaknesses in predicting stock prices, volatility, and trends. Furthermore, ongoing challenges, including data quality and model interpretability, are addressed, and emerging research directions to overcome these obstacles are explored. This study aims to provide insights into the effectiveness of prediction models.

Saputra (2025) aims to compare the performance of two Long Short-Term Memory (LSTM) neural network models to predict the closing price of Nokia Corporation (NOK) stock: a single-variable model using only past closing prices and a multi-variable model incorporating opening, high, low, close, and volume (OHLCV) data. Daily data from 2015 to 2025 were used. The metrics used to measure the performance of the models were: Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Mean Absolute Percentage Error (MAPE). The findings show that the univariate LSTM model performs better than the multivariate model.

Tu (2024) aims to evaluate and compare the effectiveness of Logistic Regression (LR), Support Vector Machine (SVM), and Excessive Gradient Boosting (XGBoost) models in stock price prediction, emphasizing the advantages and limitations of imbalanced datasets. He obtained stock data from Yahoo Finance and evaluated the effectiveness of the three models. As a result, the LR model exhibited the highest performance with a very high accuracy rate. In contrast, the SVM and XGBoost models had lower performance than the LR model. When the findings are evaluated, it is emphasized that the study is a guide in model selection in the finance sector.

Wang et al. (2025) conducted time series analysis using a publicly available dataset from Kaggle containing Apple stock prices from the last 11 years to capture long-term dependencies and introduced a long-short term memory (LSTM) model. The aim was to predict Apple Inc.'s stock price. As a result, the LSTM model was successfully trained after 50 iterations of training. The LSTM model has a four-layer architecture consisting of two LSTM layers and two fully connected layers. Metrics such as mean squared error (MSE) and root mean squared error (RMSE) were used to evaluate the model's performance. The model performed well in terms of Mean Squared Error and other indicators and was able to more accurately reflect Apple's stock price trend. The results and related metrics demonstrate the effectiveness of the trained LSTM model. This provides investors with a practical reference value for forecasting and offers a new modeling perspective and methodology for stock market time series forecasting research.

Zhang (2025) selected Xiaomi's stock data (2020-09-01-2024-09-01) as the dataset for his study; the first 80% of the dataset was classified as the training set, and the last 20% as the test set. The "closing price" was selected as the single feature input, and a two-layer long short-term memory network (LSTM), a two-layer unidirectional recurrent neural network (RNN), and a one-dimensional convolutional neural network (CNN) were used to simulate the prediction of future stock prices. Prediction accuracy was evaluated using the Mean Absolute Percentage Error (MAPE). In this study, the MAPE values of the three models were compared across different batch sizes. It was found that the MAPE value of the LSTM model increased by 0.35%, the MAPE value of the RNN model increased by 0.23%, and the MAPE value of the CNN model increased by 1.14% as the batch size increased. According to the results, RNN is more accurate in capturing hidden changes in stock data fluctuations and shows more accurate results in predictions.

This literature review provides a comprehensive conceptual foundation on financial markets, price prediction models, investor sentiment, and sustainability-focused market dynamics. It aims to contribute to the literature both methodologically and practically with its extensive dataset covering the holistic modeling of multivariate long-term datasets using modern deep learning techniques, its multidimensional variable structure, and its prediction models.

3. Deep Learning Techniques

Deep learning techniques have recently gained popularity for time series analysis and stock prediction due to their ability to predict future states when given large historical data sets. In particular, the ability to model long-term dependencies in time series data makes these techniques critical for financial price prediction. Time series analysis is divided into two types: Univariate analysis is the analysis of a time series with a single input parameter at equal time intervals. Multivariate analysis is the analysis of a time series with multiple input variables that change over time. Stock market analysis can be considered multivariate because it involves multiple parameters that change over time. These parameters include time, opening price, closing price, and others, but are not limited to these. For an investor/trader, these parameters influence decision-making processes (Ojo et al., 2019).

In this study, the LSTM architecture, which belongs to the Recurrent Neural Networks (RNN) family, was used to predict the stock prices of global technology companies. The aim is to provide forecasts for financial analysts and investors by evaluating the performance of this model. The first stage of the research explains the LSTM model and its basic principles. Subsequently, the model is applied to daily stock price data, and its performance is compared using metrics such as MAE, RMSE, and MAPE. The results reveal the model's strengths and weaknesses, demonstrating its success in predicting stock prices. Different metrics can be used to evaluate the performance of time series prediction models; the most common ones are:

- Mean Absolute Error (MAE): Represents the average of the absolute values between the predicted and actual values.
- Root Mean Square Error (RMSE): The square root of the root mean square error better represents the scale of the error amount.
- Mean Absolute Percentage Error (MAPE): Represents the average absolute percentage between the predicted and actual values (Yousufi & İncekırık, 2024).

3.1. Long Short-Term Memory (LSTM)

Recurrent Neural Networks (RNNs) are a class of artificial neural networks designed to process sequential or time series data by capturing patterns over time. They maintain a hidden state that serves as an internal memory, allowing the current output to depend not only on the current input but also on previous inputs. Despite their advantages, RNNs face several limitations. When propagating errors backward over long sequences, they are sensitive to the vanishing or exploding gradient problem. This issue complicates training because vanishing gradients cause the model to focus primarily on short-term dependencies, making it difficult to learn patterns over longer time intervals (Martínez-Barbero et al., 2024).

LSTM is a special type of recurrent neural network (RNN). This architecture emerged in 1997 as a solution or alternative method to solve the problems of traditional recurrent neural networks (RNN). LSTM networks are faster and can solve complex problems that previous recurrent neural networks could not solve (Hochreiter & Schmidhuber, 1997). LSTM networks have the ability to learn long-term dependencies and are more effective at handling complex problems that standard RNN software struggles with. Unlike RNNs, LSTM networks do not suffer from the vanishing gradient problem, which is an important factor in stock price prediction; because the previous information processed in the neural network will affect future/subsequent information (Xie et al., 2018).

The LSTM architecture consists of three main gates: the input gate, the forget gate, and the output gate. These gates regulate the flow of information entering and exiting the cell state. Specifically, the input gate controls which new information is stored, the forget gate determines which information is discarded, and the output gate decides which part of the cell state will be used in the hidden state output (Gers & Schraudolph, 2002; Jansen, 2026). This enables the selective preservation of past information and the modeling of long-term dependencies. The numbered equations below illustrate the calculations related to each gate, cell state, and hidden state mentioned:

$$i_t = \sigma(W_i x_t + U_i h_{t-1} + V_i c_{t-1}) \tag{1}$$

$$f_t = \sigma\left(W_f x_t + U_f h_{t-1} + V_f c_{t-1}\right) \tag{2}$$

$$o_t = \sigma(W_o x_t + U_o h_{t-1} + V_o c_{t-1}) \tag{3}$$

$$\tilde{c}_t = \tan h(W_c x_t + U_c h_{t-1}) \tag{4}$$

$$c_t = f_t^i \odot c_{t-1} + i_t \odot \tilde{c}_t \tag{5}$$

$$h_t = o_t \odot \tan h(c_t) \tag{6}$$

Here, x_t represents the input vector, h_t represents the hidden state, and c_t represents the cell state. i_t, f_t, and o_t denote the input, forget, and output gates, respectively; σ represents the logistic sigmoid function, $\tan h$ represents the hyperbolic tangent function, and $\odot$ represents the element-wise multiplication. The forget gate controls how much of the previous cell information is added. The output gate regulates how the cell state is reflected in the hidden state and, consequently, in the model's output (Cao, 2024). Figure 1 shows how information flow is controlled through gate mechanisms in LSTM architectures. In an LSTM network, there are separate input, forget, and output gates, as well as a cell state.

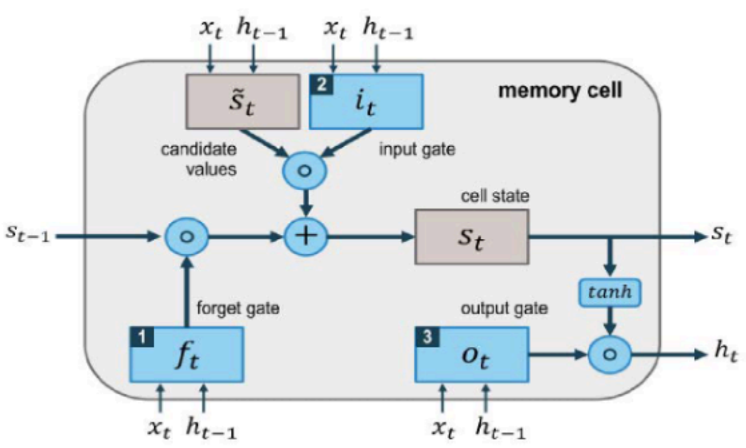

Figure 1. LSTM Network, Source: (Ojo et al., 2019)

3.2. Model Architectures

The study examined single-layer, two-layer, and three-layer LSTM model architectures. A single-layer LSTM model consists of a single LSTM layer, where the output of each layer serves as the input for the next layer. The two-layer LSTM model consists of two consecutive LSTM layers, with the output of each layer serving as the input for the next layer in the process. The three-layer LSTM model is constructed using three consecutive LSTM layers. This architecture aims to enable the model to learn deeper features.

4. Data Set and Application

In this study, deep learning-based time series models were used to predict the daily stock prices of six selected companies operating in the global technology sector. The analysis considered the daily opening and closing prices of the technology stocks of Apple (USA), Samsung Electronics (South Korea), Xiaomi (China), Sony (Japan), LG Electronics (South Korea), and Nokia (Finland). Thus, both companies that are leaders on a global scale and companies with different competitive and restructuring dynamics were evaluated together to model heterogeneous price behaviors. In order to control for macro-financial effects, the US Dollar Index (DXY) and the exchange rates of the Chinese Yuan (USD/CNY), Japanese Yen (USD/JPY), and South Korean Won (USD/KRW) against the US Dollar were included in the analysis. These variables were selected considering the economies of the countries in which the relevant companies operate and their sensitivity to global capital flows. This made it possible to examine the simultaneous effects of firm-specific dynamics and macroeconomic conditions.

The data set used in the study consists of daily frequency price series covering the period from June 9, 2018, to February 9, 2026. Stock prices, dollar index, and exchange rate data were obtained from the Yahoo Finance (2025) database. The preference for daily data allows for the modeling of short-term price movements and high-frequency market fluctuations. Both univariate and multivariate model structures were established in the analysis, and the contribution of macro-financial variables to the prediction performance was evaluated comparatively.

4.1. Descriptive Statistical Findings

Before proceeding to the modeling process, descriptive statistics were calculated to examine the basic distribution characteristics of the data set. This analysis reveals the central tendency and dispersion levels of the variables, enabling an assessment of the scale differences of the series, the volatility structure, and the price ranges observed throughout the sample period. Table 1 presents the mean, standard deviation, minimum and maximum values, median, and variance values for each variable. Reporting both opening and closing prices here allows for a comparative assessment of intraday price continuity and distribution characteristics.

Table 1. Descriptive Statistics

Variable	Mean	Std Deviation	Min	Median	Maximum	Variance
Apple Open	143.0674	64.9889	34.3238	146.2675	285.9325	4223.5635
Apple Close	143.1905	65.0568	35.1312	146.514	285.9225	4232.3868
Samsung Open	60,240.964	18,661.702	31,278.039	58,855.295	168,600	348,259,130.4
Samsung Close	60,210.4399	18,703.8203	31,278.0391	58,782.5059	169,100	349,832,895.3
Xiaomi Open	20.5419	12.8227	8.4	15.45	61.45	164.4222
Xiaomi Close	20.5263	12.8131	8.39	15.49	60.15	164.176
Sony Open	2254.1838	870.6761	870.2903	2,260.057	4650	758,076.81
Sony Close	2254.5422	869.8694	882.9143	2255.0254	4700	756,672.735
LG Open	92,738.9919	26576.4595	40,331.3902	89,997.8272	175,973.4444	706308200
LG Close	92610.7513	26,527.1142	39,855.2227	89,893.5117	177,896.625	703,687,788
Nokia Open	3.8613	0.6928	1.9684	3.9582	6.1657	0.4799
Nokia Close	3.8637	0.6947	1.9573	3.952	6.5575	0.4826
DXY Open	99.4032	5.0878	89.32	98.505	114.19	25.8858
DXY Close	99.4105	5.0888	89.44	98.515	114.11	25.8956
USD/KRW Open	1259.2739	111.2105	1083	1235.76	1486.13	12367.7839
USD/KRW Close	1259.2897	111.1619	1083.6801	1235.575	1486.13	12356.9591
USD/CNY Open	6.916	0.2866	6.3115	6.9721	7.35	0.0821
USD/CNY Close	6.9164	0.2866	6.3115	6.9721	7.35	0.0821
USD/JPY Open	127.6369	19.0497	103.028	125.444	161.607	362.8908
USD/PY Close	127.6377	19.0512	103.024	125.444	161.607	362,949

Table 1 shows that technology stocks exhibit significant scale differences in terms of price levels. For Apple, the average opening price is 143.0674 and the closing price is 143.1905, while for Samsung, these values are 60240.964 and 60210.4399, respectively. LG's average price is high at 92738.9919 (open) and 92610.7513 (close), while Nokia's average price remains quite low at 3.8613 and 3.8637. This situation reveals that there is serious scale heterogeneity between the variables and that normalization is necessary during the modeling phase. When standard deviations are examined as a volatility indicator, it is seen that the absolute volatility levels of Samsung (18661.702 open; 18703.8203 close) and LG (26576.4595 open; 26527.1142 close) shares are high. While the standard deviation for Apple is around 65, this value is around 870 for Sony. In contrast, Nokia's standard deviation remaining in the range of 0.6928–0.6947 indicates that absolute fluctuation is limited, parallel to the low price level. Xiaomi's standard deviation of around 12.82 points to a medium-scale volatility structure.

On the macro side, the DXY's average value is 99.4032 (open) and 99.4105 (close), with a standard deviation of approximately 5.08. The USD/KRW exchange rate averages 1259.27 with a standard deviation of around 111, while the USD/CNY standard deviation remains quite low at 0.2866. The USD/JPY averages 127.63 with a standard deviation of approximately 19 units, indicating significant exchange rate fluctuations throughout the period. When examining the minimum and maximum values, it is seen that Samsung traded in a wide range from 31,278 to 169,100, LG moved in a band from 40,331 to 177,896, and Apple rose from 34.32 to 285.93. These wide price ranges indicate that the analysis period included both upward and corrective phases.

The fact that the opening and closing statistics are quite close to each other reveals strong intraday price continuity in terms of both average and spread measures. This finding shows that the price series exhibit a consistent dynamic structure; however, it also implies that using the opening and closing variables together in the modeling phase may lead to information redundancy. Overall, descriptive statistics reveal that the dataset has a heterogeneous scale structure, varying volatility levels, and wide price ranges; this necessitates the application of appropriate data transformation and scaling techniques in the deep learning model.

4.2. Pearson Correlation Matrix

Before proceeding to the modeling stage, Pearson correlation analysis was applied to determine the linear relationships, co-movement dynamics, and possible multiple connection structures between stock prices and macro-financial variables. Within the scope of the analysis, both the opening and closing prices for each asset were evaluated as separate variables, thereby aiming to examine the intraday price continuity and the co-movement structure between variables together. The resulting correlation matrix is presented in Figure 2.

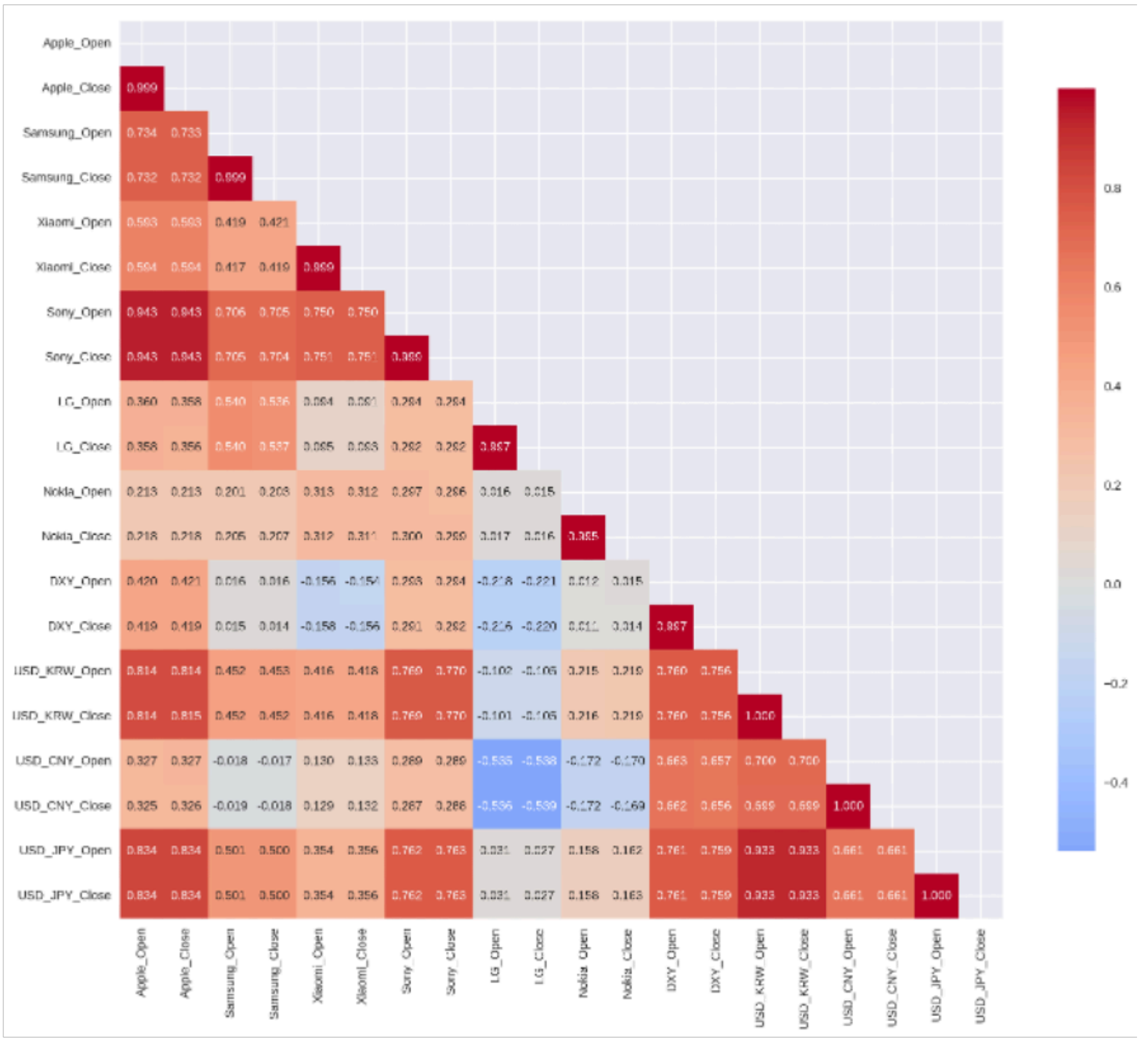

Figure 2. Pearson Correlation Matrix

Figure 2 shows that the Pearson correlation matrix reveals a clear tendency for global technology stocks to move in tandem. The high and positive correlation coefficients between Apple, Samsung, and Sony indicate that these companies respond similarly to global risk appetite, technology sector dynamics, and macroeconomic developments. The strong relationship between Apple and Sony, at 0.943, suggests that large-scale technology companies traded in developed markets move in sync with global capital flows. The correlation between Apple and Samsung in the range of 0.732–0.734 and the relationship between Samsung and Sony at approximately 0.705 indicate that sector-based common factors dominate pricing behavior. Xiaomi's correlation with Apple in the range of 0.593–0.594 and with Samsung in the range of 0.417–0.421 implies that local risks and regulatory factors specific to the Chinese market are reflected as an additional component to the global technology factor. LG and Nokia's lower correlation coefficients with other major companies suggest that company-specific dynamics may be more decisive for these companies. The fact that the relationship between Apple and LG remains at 0.358–0.360 and that between Apple and Nokia at 0.213–0.218 indicates that competitive strength and structural transformation processes may differentiate price behavior. In terms of macroeconomic variables, the positive relationship between DXY and Apple at 0.419–0.420 indicates that dollar appreciation can move in tandem with US-based stocks, while The negative relationship between DXY and Xiaomi in the range of (-0.154) to (-0.158) reveals that some Asian stocks may be negatively affected by dollar appreciation. High correlation coefficients between exchange rates show that global dollar dynamics create a common direction on regional currencies. In general, the matrix shows that common factors within the sector are strongly priced, while divergences emerge based on company characteristics.

The relationship between opening and closing prices is extremely high in the matrix. The Open–Close correlation is 0.999 for Apple, 0.999 for Samsung, 0.999 for Xiaomi, 0.999 for Sony, 0.997 for LG, and 0.995 for Nokia. These values indicate that intraday price formation shows a high degree of continuity and that the closing price is shaped as a natural extension of the opening price. Similarly, the opening-closing correlations of 0.997 for DXY, 1.000 for USD/KRW, 1.000 for USD/CNY, and 1.000 for USD/JPY reveal that daily price adjustments in the foreign exchange market occur with a high degree of parallelism. Such a strong correlation between opening and closing series indicates that daily price changes are shaped around the same information set and that short-term volatility largely follows the same trend. This finding reveals that the variables used in the analysis have high internal consistency and clearly reflects the intraday continuity of price dynamics.

These findings offer important methodological implications for deep learning models. First, the high correlation structure within the sector, particularly the strong co-movement observed between Apple, Samsung, and Sony, suggests that models can learn a common global technology factor; this points to the potential for cross-market information to improve prediction performance in multi-variable LSTM architectures. In contrast, stocks with lower correlations, such as LG and Nokia, impose the requirement for the model to learn firm-specific patterns and make the representational power of heterogeneous structures critical. The divergent relationships between exchange rates and the DXY on one hand and stocks on the other suggest that incorporating macro-financial variables into the model can add not only information richness but also regime sensitivity. On the other hand, the near-perfect correlation between opening and closing prices suggests that using these two variables in their raw form carries the risk of replication rather than information gain. Therefore, separating the information set in the model design using either price difference or return-based transformations will enable the LSTM to learn temporal dependencies more effectively. Overall, the correlation structure highlights the need for a multi-variable and carefully structured deep learning framework in terms of both variable selection and architectural design.

4.3. Data Preprocessing Transformations and Model Architecture

In the modeling process, stock price series for Apple, Samsung, Xiaomi, Sony, LG, and Nokia, as well as the DXY index and the USD/CNY, USD/JPY, and USD/KRW exchange rate series, were analyzed separately. Each time series was modeled independently, and a separate LSTM model was created for the respective series. Accordingly, the study aims to compare the forecasting performance of various LSTM architectures across different financial time series. Thus, both firm-specific price dynamics and the potential effects of exchange rate movements related to the global dollar index and Asian markets on stock prices were analyzed within the model architectures. All time series were synchronized on a date basis, considering only common trading days; observation gaps due to official holidays in different markets were excluded from the data set. No imputation method was applied to missing observations, preserving the natural volatility structure of the financial series. To prevent scale differences between variables from affecting the model learning process, all series were transformed to the [0,1] range using the Min-Max method. Normalization parameters were calculated only on the training data set, and the same coefficients were applied to the test data set. This approach increases the reliability of performance evaluation and prevents the model from indirectly accessing forward information. To transfer the time series structure to the model, the data set was converted into input-output pairs using the sliding window approach, and window lengths were set to 15, 30, 45, 60, 75, and 90 days. Each model input contains multivariate price information for past days throughout the relevant window, while the output represents the price value for the next trading day. Systematic testing of different window lengths aims to analyze the extent to which short-term price movements and longer-term trends can be learned by the model. The generated dataset was split into two subsets, with 80% used for training and 20% for testing, while maintaining the temporal order. The training data was used to learn the model parameters, while the test data was used to evaluate the out-of-sample prediction performance.

The study applied three different Long Short-Term Memory (LSTM) architectures. The LSTM Model Layer Structure and Parameter Distribution shown in Figure 3 consists of a single-layer architecture for the first model, a two-layer architecture for the second model, and a three-layer architecture for the third model. Fifty neurons were used in each LSTM layer. The effect of layer depth on the representation capacity of temporal patterns was examined comparatively, and dropout was applied at a rate of 25% in all architectures to reduce overfitting. Predictions were generated using a fully connected structure with a single neuron in the final layer. The models were trained using the MSE loss function and the Adam optimization algorithm; the number of epochs was set to 20, 50, 100, and 150, and the batch size values were set to 16, 32, 64, and 128.

LSTM 1 Katman

Model: "sequential"

Layer (type)	Output Shape	Param #
lstm (LSTM)	(None, 50)	10,400
dropout (Dropout)	(None, 50)	0
dense (Dense)	(None, 25)	1,275
dense_1 (Dense)	(None, 1)	26

Total params: 11,701 (45.71 KB)
Trainable params: 11,701 (45.71 KB)
Non-trainable params: 0 (0.00 B)

LSTM 2 Katman

Model: "sequential_1"

Layer (type)	Output Shape	Param #
lstm_1 (LSTM)	(None, 15, 50)	10,400
lstm_2 (LSTM)	(None, 50)	20,200
dropout_1 (Dropout)	(None, 50)	0
dense_2 (Dense)	(None, 25)	1,275
dense_3 (Dense)	(None, 1)	26

Total params: 31,901 (124.61 KB)
Trainable params: 31,901 (124.61 KB)
Non-trainable params: 0 (0.00 B)

LSTM 3 Katman

Model: "sequential_2"

Layer (type)	Output Shape	Param #
lstm_3 (LSTM)	(None, 15, 50)	10,400
lstm_4 (LSTM)	(None, 15, 50)	20,200
lstm_5 (LSTM)	(None, 50)	20,200
dropout_2 (Dropout)	(None, 50)	0
dense_4 (Dense)	(None, 25)	1,275
dense_5 (Dense)	(None, 1)	26

Total params: 52,101 (203.52 KB)
Trainable params: 52,101 (203.52 KB)
Non-trainable params: 0 (0.00 B)

Figure 3. Layer Structure and Parameter Distribution of LSTM Models

Model performances were evaluated using MAE, MAPE, and RMSE metrics, with lower error values interpreted as higher prediction accuracy. The results for closing prices are presented in Table 2, while the results for opening prices are presented in Table 3. Modeling opening and closing prices separately allows for a comparative analysis of the effect of intraday price formation and overnight information flow on prediction accuracy. This approach is based on the fact that, due to the variable-specific dynamic structures, different volatility levels, and non-linear relationships that financial time series may contain, it may not always be possible to achieve consistent model success across all series with a single fixed set of hyperparameters.

Table 2. Performance Comparison of Closing Price Models

Variable	Performance Criteria	Closing Performance Models					
		Day-Epoch-Batch		Day-Epoch-Batch		Day-Epoch-Batch	
Apple	MAE	45-100-64	3.730006	60-50-32	4.335903	30-150-128	4.635071
	MAPE		1.65445%		1.91338%		2.03928%
	RMSE		5.451564		6.200118		6.448527
DXY	MAE	45-100-64	0.426454	90-100-64	0.430593	60-150-128	0.42545
	MAPE		0.42066		0.42488		0.41841%
	RMSE		0.564038		0.569172		0.554295
LG	MAE	15-100-64	1697.12656	30-100-64	1688.903142	30-100-64	1650.311472
	MAPE		1.96831%		1.95926%		1.92795%
	RMSE		2401.04556		2403.778483		2322.133182
Nokia	MAE	30-50-32	0.073096	60-100-64	0.073132	30-100-64	0.07059
	MAPE		1.60920%		1.60826		1.55382%
	RMSE		0.126461		0.123181		0.121099
Samsung	MAE	30-50-32	1870.94275	90-100-64	2005.064774	60-100-64	2306.498936
	MAPE		2.42121%		2.49457%		3.00684%
	RMSE		2,818.29197		3,122.839183		3,303.18235
Sony	MAE	90-100-64	78.283846	30-50-32	77.423598	45-100-64	80.917716
	MAPE		2.19128%		2.17657%		2.33337%
	RMSE		107.81059		107.973422		110.210461
USD/CNY	MAE	15-100-64	0.013155	45-100-64	0.014436	75-150-128	0.01333
	MAPE		0.18321%		0.20154%		0.18555%
	RMSE		0.01998		0.020593		0.01897
USD/JPY	MAE	90-100-64	1.029614	60-100-64	1.095596	30-100-64	1.16739
	MAPE		0.68808%		0.73180%		0.77894%
	RMSE		1.316952		1.374043		1.455768
USD/KRW	MAE	60-100-64	7.551605	45-100-64	7.338781	60-100-64	7.148839
	MAPE		0.53588		0.52202%		0.50748%
	RMSE		10.122572		10.00346		9.620397
Xiaomi	MAE	45-50-32	1.006856	30-100-64	0.980954	15-100-64	1.035552
	MAPE		2.44139%		2.36963%		2.48628%
	RMSE		1.403755		1.413341		1.475805

The findings presented in Table 2 show that the performance of deep learning LSTM-based architectures varies significantly across variables in terms of closing prices in technology companies and exchange rate series. The superiority of single-layer, two-layer, and three-layer LSTM structures is not constant; there is no linear relationship between model depth and prediction accuracy. This clearly demonstrates that deeper architectures do not always produce lower errors in financial time series.

The lowest error values in the Apple series were obtained with a single-layer LSTM model using a 45-day window, 100 epochs, and a 64 batch configuration. This structure yielded MAE of 3.730006, RMSE of 5.451564, and MAPE of 1.65445%, indicating that error metrics increase as the number of layers increases. This result suggests that Apple closing prices have a relatively balanced and learnable temporal structure that does not require excessively deep structures. Similarly, in the Samsung series, the best performance was achieved by the single-layer LSTM model with a 30-day window, 50 epochs, and 32 batch configuration (MAE: 1870.94275, RMSE: 2818.29197, MAPE: 2.42121%). Error values are seen to increase significantly in two- and three-layer structures. This indicates that increasing model complexity may weaken generalization performance in Samsung closing prices.

In contrast, deeper structures provided an advantage in the LG and Nokia series. For LG, the lowest error was achieved with a three-layer LSTM model using a 30-day window, 100 epochs, and a 64 batch configuration: 1650.311472 MAE, 2322.133182 RMSE, 1.92795% MAPE. For Nokia, the three-layer model produced the lowest MAE (0.07059), RMSE (0.121099), and MAPE (1.55382%) values in the 30-100-64 combination. This finding shows that deeper architectures can represent temporal patterns more effectively in some series. For Sony and Xiaomi series, two-layer LSTM structures are relatively more successful. For Sony, a 30-day window, 50 epochs, and 32 batch configuration (MAE: 77.423598, MAPE: 2.17657%) produced the lowest error values. In Xiaomi, the 30-100-64 combination (MAE: 0.980954, MAPE: 2.36963%) provided the best performance. This indicates that medium model depth creates an optimal balance for some stocks.

It is noteworthy that the error values are relatively low and stable for exchange rates and DXY. The lowest RMSE value in DXY was obtained with the three-layer model with a 60-day window, 150 epochs, and 128 batch configuration (0.554295), but there were no significant differences between the models in terms of MAE and MAPE values. In the USD/CNY series, the three-layer model produced the lowest RMSE (0.01897) in the 75-150-128 configuration. In USD/KRW, the three-layer model achieved the lowest error values in the 60-100-64 combination (MAE: 7.148839, MAPE: 0.50748%). In contrast, in the USD/JPY series, the lowest error was achieved by the single-layer model in the 90-100-64 configuration (MAE: 1.029614, MAPE: 0.68808%).

Overall, it is clear that a single LSTM depth for closing prices is not superior across all series. While a single-layer structure produces lower errors in some series, two- or three-layer architectures yield more successful results in others. Furthermore, the optimal window length, epoch, and batch combinations also vary by variable. This finding clearly demonstrates that model success in financial time series depends on asset-specific dynamics and that universal success cannot be achieved with a fixed architecture or a fixed set of hyperparameters. These results show that increasing depth does not automatically mean increased performance in model selection; systematic model and hyperparameter optimization sensitive to the data structure is essential. The following Table 3 presents a performance comparison of opening price models.

The findings presented in Table 3 show that deep learning-based LSTM models generally performed consistently in predicting opening prices for technology companies and exchange rate series. When evaluated by variable, it is observed that single-layer LSTM produced lower MAE and MAPE values in some series, while two- and three-layer models showed relatively lower performance, especially in series with high volatility or large-scale prices. This indicates that deeper LSTM layers limit the ability to capture short-term fluctuations and that excessive parameter load increases the risk of overfitting in small datasets.

Table 3. Performance Comparison of Opening Price Models

Variable	Performance Criteria	Opening Performance Models					
		LSTM (Single-Layer)		LSTM (2-Layer)		LSTM (3-Layer)	
		Day-Epoch-Batch		Day-Epoch-Batch		Day-Epoch-Batch	
Apple	MAE	45-100-64	4.336819	60-100-64	4.900039	90-100-64	4.853536
	MAPE		1.91736%		2.15626%		2.14967%
	RMSE		6.121715		6.69884		6.822823
DXY	MAE	15-100-64	0.417574	30-100-64	0.433022	90-100-64	0.397987
	MAPE		0.41221%		0.42675%		0.39292%
	RMSE		0.539319		0.555231		0.520446
LG	MAE	45-100-64	1659.31302	15-100-64	1569.906781	15-100-64	1616.192522
	MAPE		1.91331%		1.80956%		1.86730%
	RMSE		2397.59599		2,298.485328		2365.03922
Nokia	MAE	30-100-64	0.070481	60-100-64	0.064656	75-150-128	0.06459
	MAPE		1.57794%		1.43808%		1.43763%
	RMSE		0.109545		0.10216		0.101842
Samsung	MAE	30-50-32	2073.04749	15-100-64	2091.86026	60-100-64	2344.311545
	MAPE		2.61824%		2.64218%		3.00643%
	RMSE		3,222.03606		3,196.44341		3,485.31007
Sony	MAE	30-100-64	74.975394	90-50-32	76.156588	60-100-64	72.193589
	MAPE		2.13035%		2.19613%		2.10497%
	RMSE		107.123884		108.740763		103.919882
USD/CNY	MAE	30-100-64	0.012907	60-100-64	0.013179	90-100-64	0.013926
	MAPE		0.17974%		0.18344%		0.19394%
	RMSE		0.019827		0.019194		0.019875
USD/JPY	MAE	45-100-64	1.054438	30-100-64	1.103041	30-100-64	1.199717
	MAPE		0.70467%		0.73802%		0.80188%
	RMSE		1.353616		1.400942		1.496817
USD/KRW	MAE	45-150-128	7.719622	75-150-128	7.80844	30-150-128	7.545757
	MAPE		0.54750%		0.55339		0.53605%
	RMSE		10.378602		10.287427		10.104031
Xiaomi	MAE	60-100-64	1.094266	45-100-64	1.08389	60-150-128	1.096861
	MAPE		2.67369%		2.64353%		2.65182%
	RMSE		1.496578		1.489512		1.528405

For Apple's opening price, a single-layer LSTM with a 45-day window, 100 epochs, and 64 batches configuration achieved the best performance with 4.3368 MAE and 1.917% MAPE; deeper layers showed a slight performance loss. In the DXY series, a three-layer LSTM with a 90-day window produced the lowest error with 0.3929% MAPE and 0.3979 MAE; this indicates that long-term dependencies are effective in currency indices and that deeper models may have an advantage in capturing this context. Error values for stocks with high price scales, such as LG and Samsung, are large compared to other series; however, among the model configurations, single and two-layer LSTMs offer relatively more balanced performance. For Nokia and Sony, series with medium volatility produced more consistent forecasts in short and medium-term windows; the three-layer LSTM provided only a small improvement in very long window and batch combinations. For exchange rate series such as USD/CNY, USD/JPY, and USD/KRW, single-layer LSTMs often provided the lowest MAPE values; this indicates that simple LSTM structures are more effective in foreign exchange markets dominated by short-term fluctuations. For Xiaomi and other technology stocks, model performance was closely related to window length, epoch, and batch combinations, highlighting the critical role of hyperparameter optimization.

The findings reveal that LSTM-based models are generally successful for opening prices, but the volatility and price scale structure specific to each series directly affects the choice of model depth, window length, epoch, and batch size. In this context, the variations in Table 2 demonstrate that a single LSTM architecture or a fixed set of hyperparameters cannot be optimal for all series in financial time series. Performance metrics related to closing prices are presented in detail in Table 2. Performance results for opening prices are presented in Table 3. In order to visually examine the model performance evaluated numerically in the relevant tables, comparisons of the opening and closing price predictions and actual values for the multi-layer LSTM model are presented in Figure 4 and Figure 5.

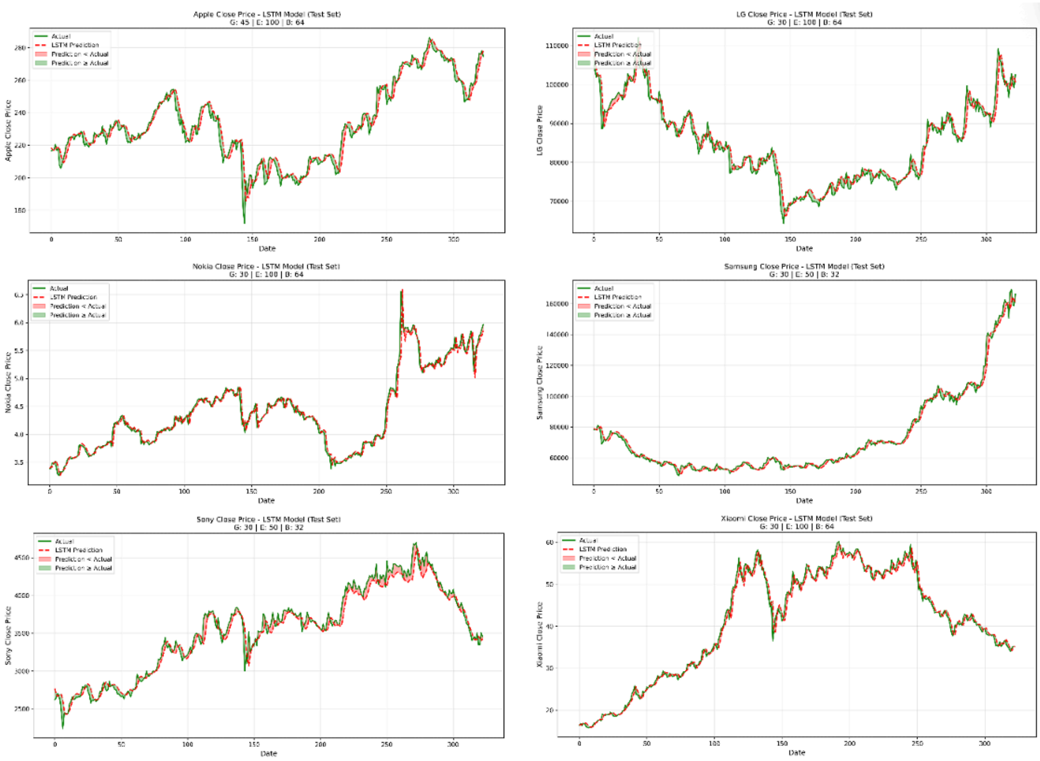

Figure 4. Comparison of Actual and Predicted Values for Closing Prices Using the Multi-Layer LSTM Model

When examining the graphs of closing prices provided in Figure 4, it is seen that the LSTM model successfully predicts the actual values to a large extent. The prediction curves are quite strong in capturing the general trend and accurately follow long-term upward and downward movements. The model's trend learning capacity is particularly high in the Apple, Samsung, and LG series. It is observed that the difference between actual and predicted values mostly remains low. However, minor deviations are seen in some stocks during sudden price breaks and sharp decline points. This can be explained by the fact that LSTM models may respond with a short delay during periods of high volatility. However, these deviations are not at a level that would impair overall performance. Overall, it can be said that the model demonstrates high accuracy and strong trend capture capability in predicting closing prices.

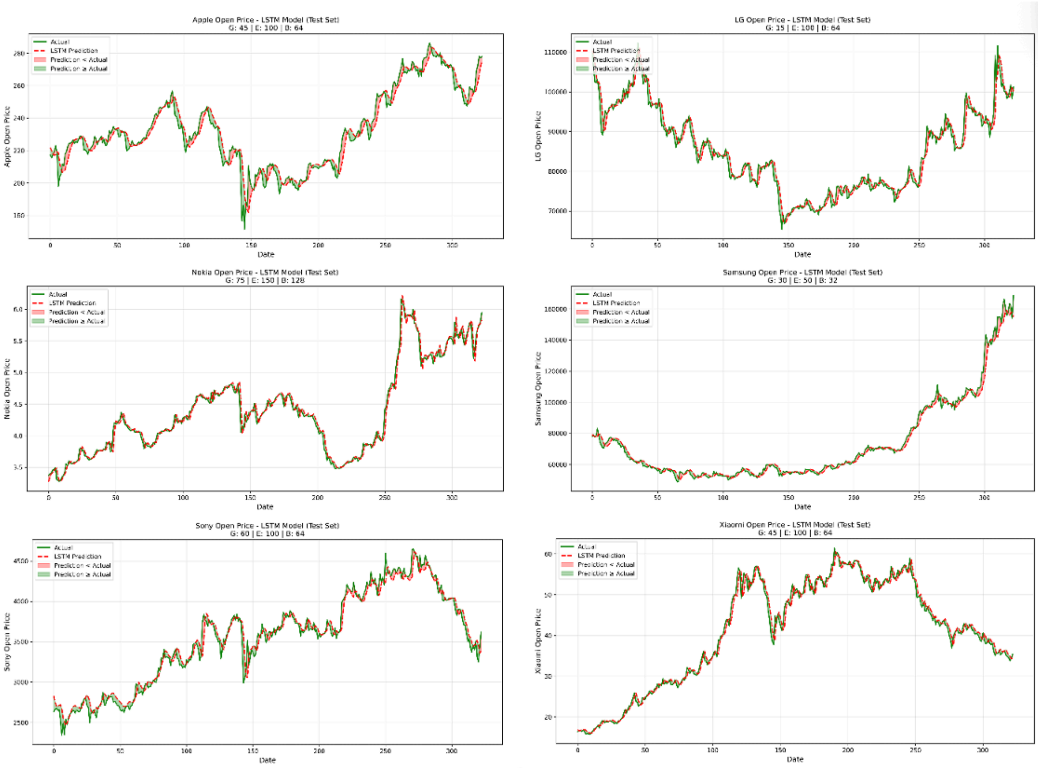

Figure 5. Comparison of Actual and Predicted Values for Opening Prices for the Multi-Layer LSTM Model

When examining the graphs of opening prices provided in Figure 5, it is seen that the model performs consistently, similar to closing prices. The predicted values closely follow the actual values and successfully reflect general price movements. The model produced quite strong results during periods when the trend direction was clear. The alignment between the prediction curve and the actual data is particularly striking in long-term upward trends. However, there are minor deviations in some periods with intense short-term fluctuations. Nevertheless, these deviations do not undermine the model's overall prediction success.

In conclusion, it can be stated that the model shows consistent and balanced performance in both opening and closing prices, shows no signs of overfitting, and offers a generalizable structure. The single-layer LSTM model prediction result for Apple's opening and closing prices is shown in Figure 6.

Figure 6 shows that the predictions generated by the single-layer LSTM model closely follow the actual Apple closing prices after the test start point. It is observed that the predictions generally align with the overall trend, especially during periods of accelerated upward momentum, while limited deviations occur during short-term fluctuations.

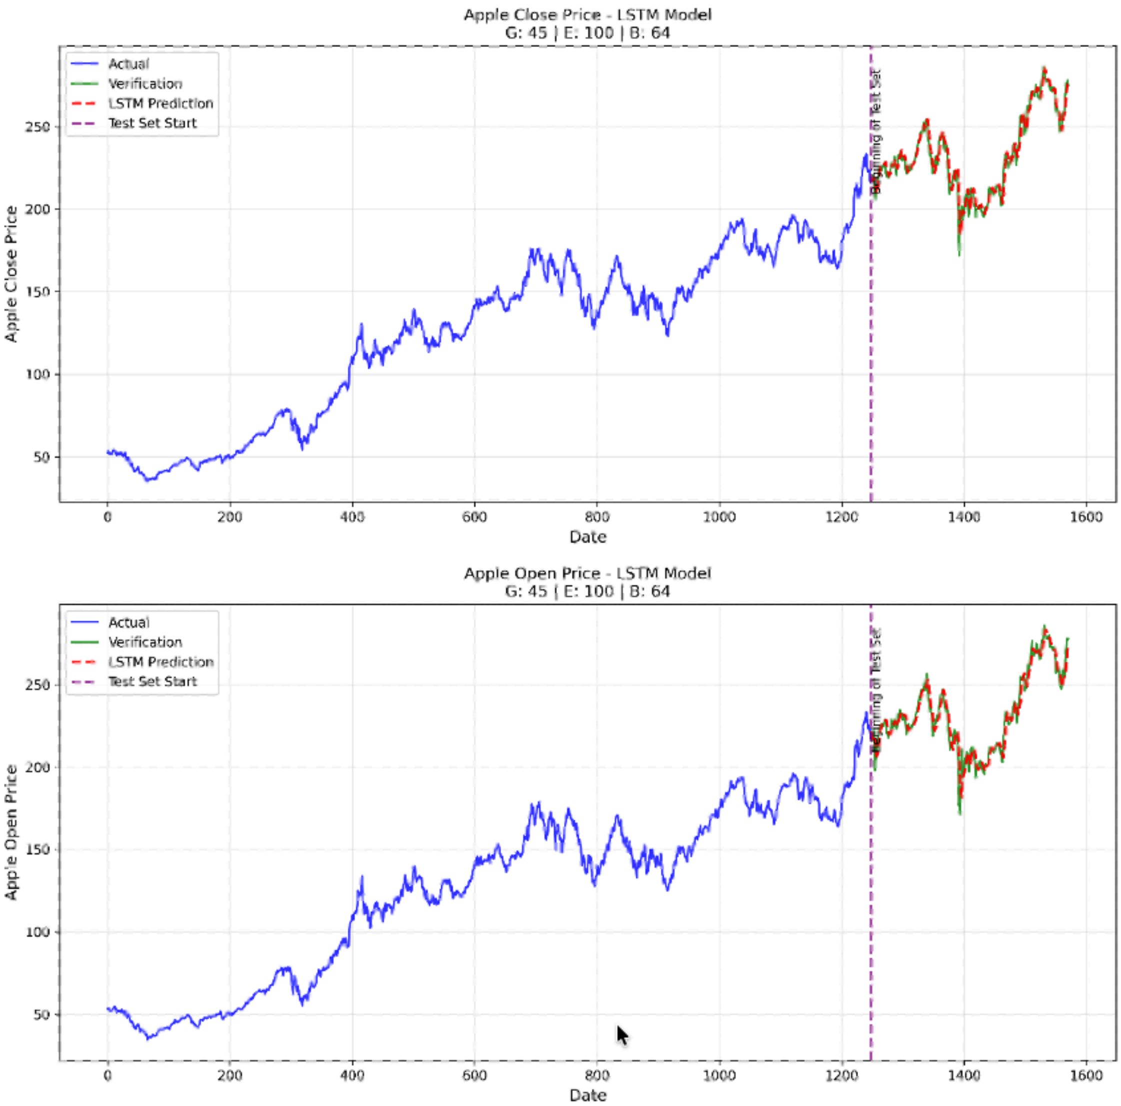

Figure 6. Single-Layer LSTM Model Prediction Result for Apple (Opening and Closing) Price Series

The findings of this study reveal that deep learning-based time series models have a high potential for accurately predicting opening and closing price movements in technology companies and foreign exchange markets. When visual and numerical results are evaluated together, it is seen that the models successfully track general trends, but prediction deviations relatively increase during periods of heightened volatility. Performance differences between model architectures are closely related to the temporal dependency structures and volatility characteristics of price series, indicating that a single model approach does not produce optimal results for all market conditions. The findings obtained in this context emphasize that model selection and hyperparameter tuning must be approached in a manner sensitive to the data structure.

5. Discussion and Conclusion

This study comprehensively examined the effectiveness of deep learning-based LSTM models in predicting daily opening and closing prices for technology companies and exchange rates. The analysis process began with descriptive statistics and correlation analyses performed prior to the modeling stage; then, single, two, and three-layer LSTM structures were trained under different window lengths, epoch numbers, batch sizes, and variable combinations, and their prediction performances were evaluated comparatively. This two-stage approach allowed for both a structural understanding of the relationships between variables and an empirical test of how these relationships translate into model prediction accuracy.

The analysis results revealed that descriptive statistics showed significant heterogeneity in technology stocks in terms of scale and volatility. Series such as Apple and Sony exhibited medium-scale volatility, while Samsung and LG had high absolute volatility; volatility was limited in small-scale stocks such as Nokia and Xiaomi. Among exchange rates, series such as DXY and USD/JPY exhibit medium- to long-term fluctuations, while USD/CNY remains quite stable. The high correlation between opening and closing prices (in the range of 0.995–0.999 for most series) indicates strong intraday price continuity and suggests that using these two variables together in the modeling stage may lead to information redundancy. Correlation analyses revealed strong positive relationships between Apple, Samsung, and Sony, indicating that global technology stocks respond simultaneously to similar market dynamics and macroeconomic factors. In contrast, LG and Nokia had lower correlations than other companies, suggesting that company-specific dynamics may be decisive. The divergent relationships between exchange rates and the DXY on one hand and stocks on the other suggest that incorporating macro variables into the model could enhance information richness and impart regime sensitivity.

Deep learning-based prediction results clearly demonstrate that model performance is closely related to time-dependent structures specific to the series and price volatility. For closing prices, single-layer LSTM models produced lower errors in most series, while two- or three-layer structures demonstrated superior performance in some series. For opening prices, single-layer models generally excelled at capturing short-term fluctuations, while deeper layers provided an advantage in series with high volatility or large scale. The optimal window length, epoch, and batch combinations varied by variable; for example, Apple and Samsung achieved the best performance with short windows and single-layer LSTM, while LG and Nokia reduced error metrics with long windows and deeper layers. The paired training method was applied for each LSTM layer; 480 training configurations were evaluated for single, two, and three-layer models, resulting in a total of 1,440 different models being examined. This comprehensive model scan made it possible to systematically test the performance of each series and hyperparameter combination and clearly demonstrated that achieving universal success with a fixed set of hyperparameters is not possible in financial time series. The performance of the models is highly sensitive to the data structure, volatility level, temporal dependence, and the nature of the relationships between variables; this situation reveals that the “one model-one setting” approach is no longer valid in deep learning applications.

When evaluating the findings, Table 2 and Table 3 show that the performance of LSTM-based deep learning models varies across variables for both closing and opening prices in technology company and exchange rate series. For closing prices, single-layer LSTM produced lower errors in most stock and currency series, while two- or three-layer structures proved superior in some series; for opening prices, single-layer models generally excelled in capturing short-term fluctuations, while deeper layers provided an advantage in series with high volatility or large scale. This reveals that there is no fixed linear relationship between model depth and prediction accuracy for both opening and closing prices, indicating that increasing depth does not automatically improve performance in financial time series. The optimal window length, epoch, and batch combinations vary depending on the series, playing a critical role in both datasets. For example, in Apple and Samsung series, short windows and single-layer LSTM produce the best performance for opening and closing prices, while in series such as LG and Nokia, deeper structures, longer windows, and higher epoch counts reduce error metrics. In exchange rates, the generally low error values and limited differences between models in DXY and USD/CNY indicate that short- and medium-term dependencies are dominant and that simple LSTM structures are sufficient for such series. Overall, both datasets clearly demonstrate that a single universal model ( ) cannot be valid for financial time series; model performance varies depending on the data structure, volatility, and temporal dependency characteristics.

In summary, the findings of this study indicate that deep learning-based time series models have high predictive potential in technology companies and foreign exchange markets; however, this potential can only be effectively utilized when a data-specific, model-specific, and hyperparameter-sensitive approach is adopted. Here, it is clearly demonstrated that there is no universal model architecture for financial time series and that model selection must be optimized according to the structural characteristics of each asset. In this respect, the study contributes methodologically to the academic literature and provides a guiding framework for practitioners in the design of deep learning-based forecasting systems. Future research could compare the LSTM model with other deep learning models to explore different model performances.

The findings of this study have yielded significant implications for both academic and financial sectors. From a researcher's perspective, the highly successful LSTM architecture in stock price prediction has increased the importance of deep learning models in financial analysis, leading to their widespread use. From an investor's perspective, the information obtained from these models enables them to make more accurate decisions, resulting in improvements in investment strategies and portfolio management. In conclusion, the study demonstrates the potential of three different layered LSTM architectures in stock market analysis and prediction, providing valuable insights for researchers and practitioners in the financial sector. As deep learning techniques evolve over time, it is expected that new avenues will open in financial analysis and in predicting and evaluating stock price movements.

Declarations

Conflict of interest The author(s) have no competing interests to declare that are relevant to the content of this article.

References

Aggarwal, D. (2019). Defining and measuring market sentiments: a review of the literature. *Qualitative Research in Financial Markets*, *14*(2), 270–288. https://doi.org/10.1108/qrfm-03-2018-0033

Cao, R. (2024). Stock Price Prediction Using Deep-Learning Models: CNN, RNN, and LSTM. *SHS Web of Conferences*, *196*, 2004. https://doi.org/10.1051/shsconf/202419602004

Farhan, M., Ahmed, A., Eesaar, H., Chong, K. T., & Tayara, H. (2025). A novel approach to stock price prediction: averaging open and close prices with LSTM. *Digital Finance*, *7*(3), 535–551. https://doi.org/10.1007/s42521-025-00151-6

Gers, F. A., & Schraudolph, & S. J., N. N. (2002). Learning precise timing with LSTM recurrent networks. *Journal of Machine Learning Research*, *3*(Aug), 115–143.

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

Jansen, S. (2026). *Machine Learning for Trading: Integrate GenAI, Causal Inference, and Reinforcement Learning into Real World Trading Systems* (3rd ed.). Packt Publishing.

Kang, W. I. (2022). *Predicting Asian Stock Market Index Using US Financial Market Indexes and Machine Learning Techniques* [Doctoral dissertation].

Martínez-Barbero, X., Cervelló-Royo, R., & Ribal, J. (2024). Portfolio Optimization with Prediction-Based Return Using Long Short-Term Memory Neural Networks: Testing on Upward and Downward European Markets. *Computational Economics*, *65*(3), 1479–1504. https://doi.org/10.1007/s10614-024-10604-6

Na, S. (2025). Portfolio Theory in Investment Decision Making: A Case Study of Xiaomi, Apple, and Samsung. In *Proceedings of the 2025 3rd International Academic Conference on Management Innovation and Economic Development (MIED 2025)* (pp. 482–487). Atlantis Press International BV. https://doi.org/10.2991/978-94-6463-835-6_51

Ojo, S. O., Owolawi, P. A., Mphahlele, M., & Adisa, J. A. (2019). Stock Market Behaviour Prediction using Stacked LSTM Networks. *2019 International Multidisciplinary Information Technology and Engineering Conference (IMITEC)*, 1–5. https://doi.org/10.1109/imitec45504.2019.9015840

Qi, R., & Hu, L. (2025). Stock Price Prediction of Apple Inc. Based on LSTM Model: An Application of Artificial Intelligence in Individual Stock Analysis. *European Journal of AI, Computing & Informatics*, *1*(3), 1–9. https://doi.org/10.71222/5hq2rh34

Reddy, S., Rao, S., & Sharma, D. (2020). Performance Analysis of Deep Learning and Statistical Models on Enhancing Stock Market Portfolio. *International Journal of Innovative Research in Engineering & Multidisciplinary Physical Sciences*, *8*(6), 20–29. https://doi.org/10.37082/ijirmps.2020.v08i06.003

Saberironaghi, M., Ren, J., & Saberironaghi, A. (2025). Stock Market Prediction Using Machine Learning and Deep Learning Techniques: A Review. *Appliedmath*, *5*(3), 76. https://doi.org/10.3390/appliedmath5030076

Saputra, R. (2025). A Comparative Analysis of Univariate and Multivariate LSTM Models for Nokia (NOK) Stock Price Prediction. *Jurnal Komtika (Komputasi Dan Informatika)*, *9*(2), 205–213. https://doi.org/10.31603/komtika.v9i2.15152

Saracık, Ö., & İncekırık, A. (2023). Stock Price Forecasting with Deep Learning Techniques. *Alphanumeric Journal*, *11*(2), 137–156. https://doi.org/10.17093/alphanumeric.1357466

Tu, W. (2024). Forecast and Analysis for Samsung Stock Price Based on Machine Learning. *Proceedings of the 1st International Conference on E-Commerce and Artificial Intelligence*, 234–238. https://doi.org/10.5220/0013213900004568

Wang, Y., Ge, Z., Jian, T., & Zhang, H. (2025). Predicting Apple's Stock Price with LSTM. In *Proceedings of the 2025 3rd International Academic Conference on Management Innovation and Economic Development (MIED 2025)* (pp. 954–960). Atlantis Press International BV. https://doi.org/10.2991/978-94-6463-835-6_102

Xiao, R., Feng, Y., Yan, L., & Ma, Y. (2022). *Predict stock prices with ARIMA and LSTM*. https://doi.org/10.48550/ARXIV.2209.02407

Xie, Q., Cheng, G., Xu, X., & Zhao, Z. (2018). Research Based on Stock Predicting Model of Neural Networks Ensemble Learning. *MATEC Web of Conferences*, *232*, 2029. https://doi.org/10.1051/matecconf/201823202029

Yahoo Finance. (2025). *Historical daily opening and closing prices for Apple, Samsung Electronics, Xiaomi, Sony, LG Electronics, Nokia, DXY, USD/CNY, USD/JPY, and USD/KRW*. https://finance.yahoo.com/

Yang, Y. (2022). Forecasting Apple Stock Closed Prices by LR and LSTM with Discrete Wavelet Transformation. In *Proceedings of the 2022 2nd International Conference on Economic Development and Business Culture (ICEDBC 2022)* (pp. 935–943). Atlantis Press International BV. https://doi.org/10.2991/978-94-6463-036-7_138

Yousufi, M. L., & İncekırık, A. (2024). Examining Ethereum price prediction in the digital economy using deep learning techniques. In *GENÇLERLE 360 10th International Student Congress / Recycling / Circular Economy in the Globalizing World* (pp. 58–66). Manisa Celal Bayar University. https://genclerle360.mcbu.edu.tr/wp-content/uploads/sites/11/2025/01/01.08.-Onuncu-bildiri-kitabi.pdf

Zhang, Z. (2025). Deep Learning in Stock Price Prediction. *2025 IEEE 3rd International Conference on Image Processing and Computer Applications (ICIPCA)*, 415–419. https://doi.org/10.1109/icipca65645.2025.11138890

This page is intentionally left blank