International Journal of Finance and Investment
ISSN: 3078-6002 (Print), ISSN: 3078-2295 (Online) | Vol. 2, No. 1, 2025

Forecasting Financial Fortunes: Unveiling the Secrets of Stock Prediction Models

Yuanchun (Justin) Chang

Department of Applied Mathematics. School of National Chengchi University, Taipei,116011, Taiwan
changjustin.yc@gmail.com

Abstract: Stocks are the most representative area of modern financial markets, serving as a cornerstone of investment strategies for individuals and institutions. Over the years, analysts have devoted substantial effort to developing methodologies for predicting stock price movements more accurately. For seasoned and ordinary investors, buying and selling stocks represents a significant avenue for generating additional income. However, the inherent volatility of stock markets, driven by many economic, political, and psychological factors, poses a persistent challenge to achieving precise predictions. Mathematical models play a crucial role in forecasting stock price trends in this context. These models range from traditional statistical techniques like linear regression and Autoregressive Integrated Moving Average (ARIMA) to advanced machine learning frameworks such as Support Vector Machines (SVM) and Long Short-Term Memory (LSTM) networks. Each approach offers distinct advantages yet also exhibits limitations when applied to different market conditions, particularly in the face of sudden, unpredictable events. This paper examines the performance of various predictive models, focusing on their strengths and weaknesses in capturing trends across diverse stock types. This research aims to illuminate the potential and limitations of predictive modeling in stable and volatile market environments by analyzing its outcomes. It highlights the importance of aligning forecasting techniques with specific market dynamics and stock characteristics, offering valuable insights for analysts seeking to refine their strategies in the ever-evolving landscape of financial markets.

Keywords: Stock Market Prediction, Mathematical Models, Machine Learning, Time Series Analysis.

1. Introduction

This essay explores the role of various mathematical models in predicting stock prices, ranging from traditional linear regression to more complex techniques such as ARIMA, SVM, and LSTM networks. The model discussed in this essay is suitable for analyzing data from different time periods. It also examines whether the accuracy of predictions improves with more data, or if results are inaccurate due to outdated data not being relevant to the present.

The article also focuses on comparing stocks that are more stable and unstable than the market average to understand if they can keep up with today's turbulent financial markets.

While these models help us analyze past data to understand potential future outcomes, it's important to note that in volatile financial markets, stocks are influenced by various factors. Therefore, mathematical models cannot guarantee accurate predictions of stock performance in future trends. This will be a major challenge for mathematical models going forward.

2. Organization of the Text

2.1. Model Introduction, Process and Outcome

2.1.1. Linear Regression: Using the 2 years historical data

Linear regression is a statistical method to model the relationship between a dependent variable and one or more independent variables. It assumes a linear relationship between the input (independent variables) and the output (dependent variable). In stock market forecasting, this means using historical financial data to predict future stock prices. The model is mathematically represented by the equation

$$Y=\beta 0+\beta 1X1+\beta 2X2+\ldots+\beta n\text{-}1Xn\text{-}1+\beta nXn$$

Y: the future stock price

X1, X2, …. Xn: explanatory variables such as past stock prices, trading volumes, or macroeconomic variables

β0, β1, …βn: the coefficients that measure the impact of each independent variable on the stock price

How Linear Algebra works

The first step in applying linear regression to stock price prediction involves collecting and preparing data. This process involves gathering daily data for two years, including the open price, highest price, lowest price, close price, adjusted price, and volume.

After the data is prepared, it is divided into training and testing sets. The training set estimates the coefficients of the Linear Regression model by applying the least squares method, which aims to minimize the sum of squared residuals. This method effectively finds the 'line of best fit' that predicts the dependent variable with the slightest error.

Outcome and observation

(1) AAPL: According to historical stock data, the red line predicted by linear regression shows a stable upward trend, indicating a linear relationship with time. This method reduces inventory changes to an average rate of change, calculated to minimize forecast error from past data. However, this approach fails to capture significant biases, such as a company's announcement in May 2024 that significantly helped the stock price rise, highlighting its limitations in adapting to sudden market changes.

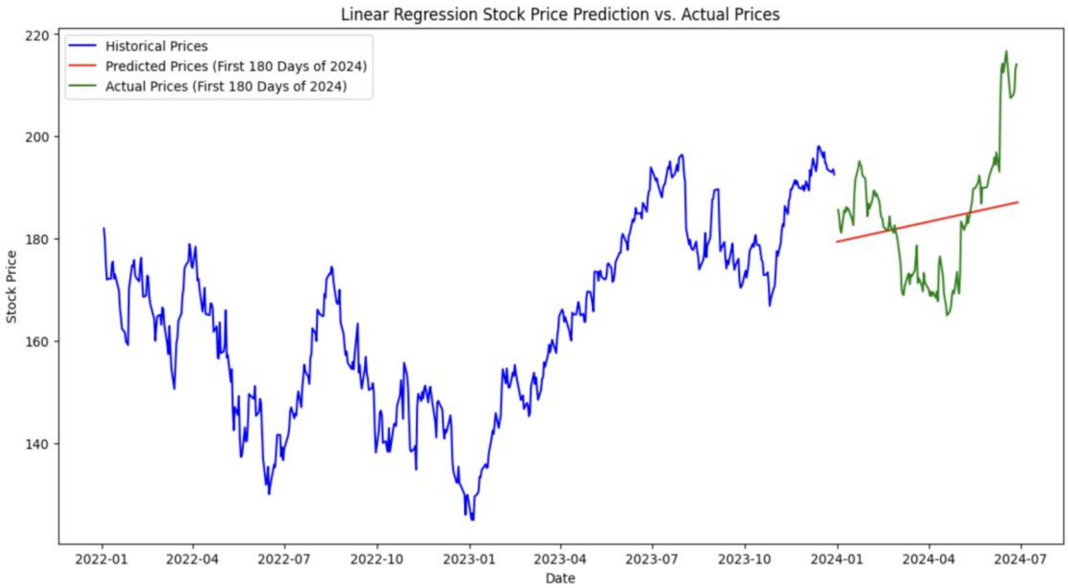

Figure 1. Linear regression of AAPL

(2) TSLA: The red line in the chart reflects the initial decline in Tesla stock prices in early 2024, reflecting a trend consistent with early data points on actual prices. The red line follows a slight, steady decline in the stock price, showing that the model's predictions were relatively accurate during this period. This indicates that for certain stages, especially those in which price changes are mild and gradual, the model's predictions agree with actual results, demonstrating its effectiveness in situations of lower volatility and fewer sudden changes in the market.

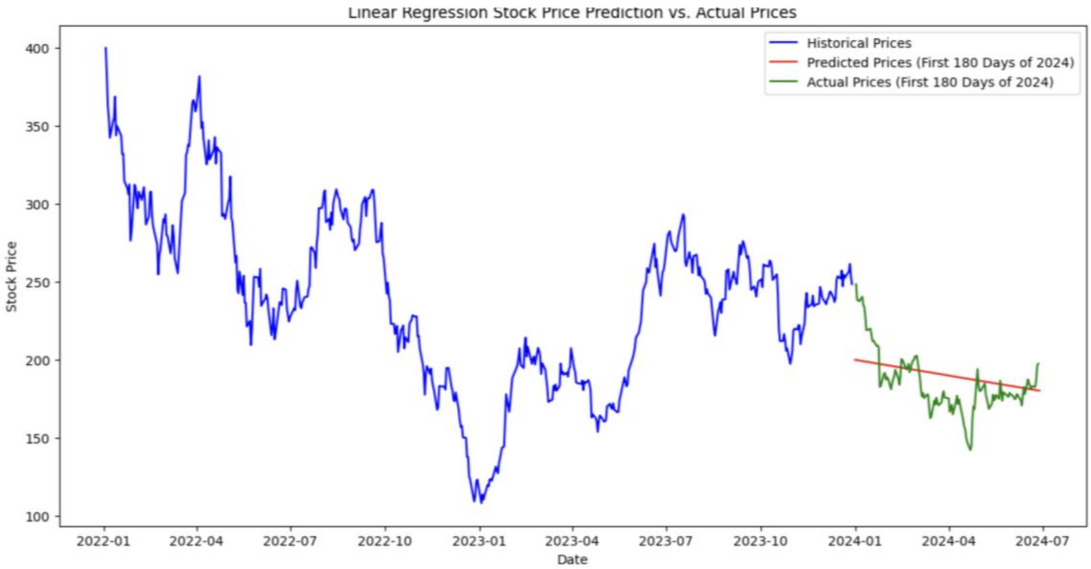

Figure 2. Linear regression of TSLA

(3) JNJ: While successfully predicting a downward trend in JNJ stock in early 2024, the red line on the chart doesn't fully capture the extent of the stock's actual decline. The exact price is significantly lower compared to the red line, showing that market conditions are more bearish for JNJ than the linear trend would predict. This discrepancy highlights the model's limitations in adjusting for larger-than-expected market moves, suggesting that linear models need to account for potentially negative news or market sentiment that is more influential in larger declines.

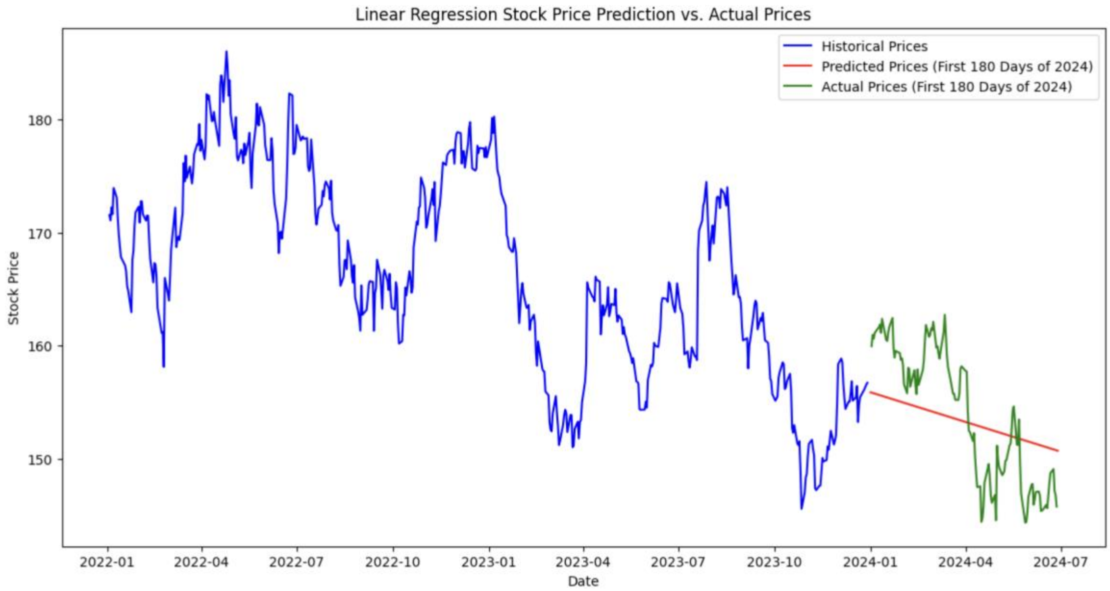

Figure 3. Linear regression of JNJ

(4) NVDA: The red line shows the modest upward trend in NVDA predicted based on past data, highlighting the typical expected performance based on a simple linear model. However, the actual stock price from early 2024 has shown a near-vertical surge, reflecting strong optimism about NVIDIA's prospects in the AI and GPU fields. This difference highlights the limitations of linear regression in predicting the intensity of stock price movements when the market reacts strongly to specific industry trends or technological advances. The model can predict the direction of a trend but not the steepness or rapid market dynamics driven by investor enthusiasm and specific industry developments.

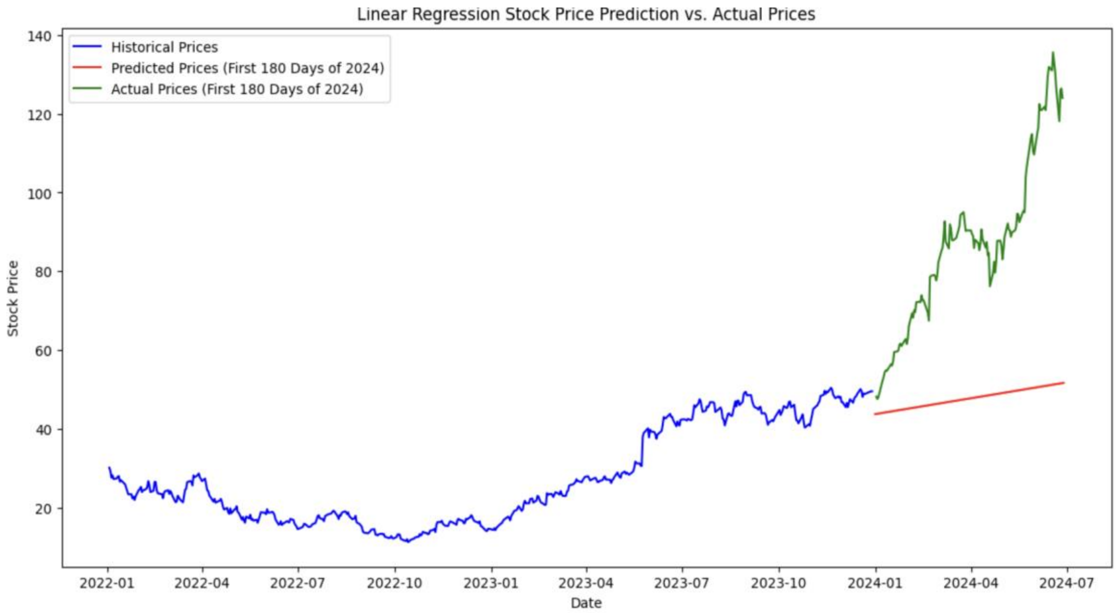

Figure 4. Linear regression of NVDA

(5) AVGO: As the chart shows, AVGO shares soared in early 2024, mainly due to the company's position in the communications chip market during the artificial intelligence boom. However, the red line is a linear regression forecast that reflects only a modest upward trend and begins to decline in early 2024, indicating its limitations in adapting to sudden market changes. This discrepancy highlights the model's inability to capture significant changes driven by industry-specific developments and rising market optimism, highlighting the challenges linear regression faces in volatile and fast-moving industries such as technology and artificial intelligence.

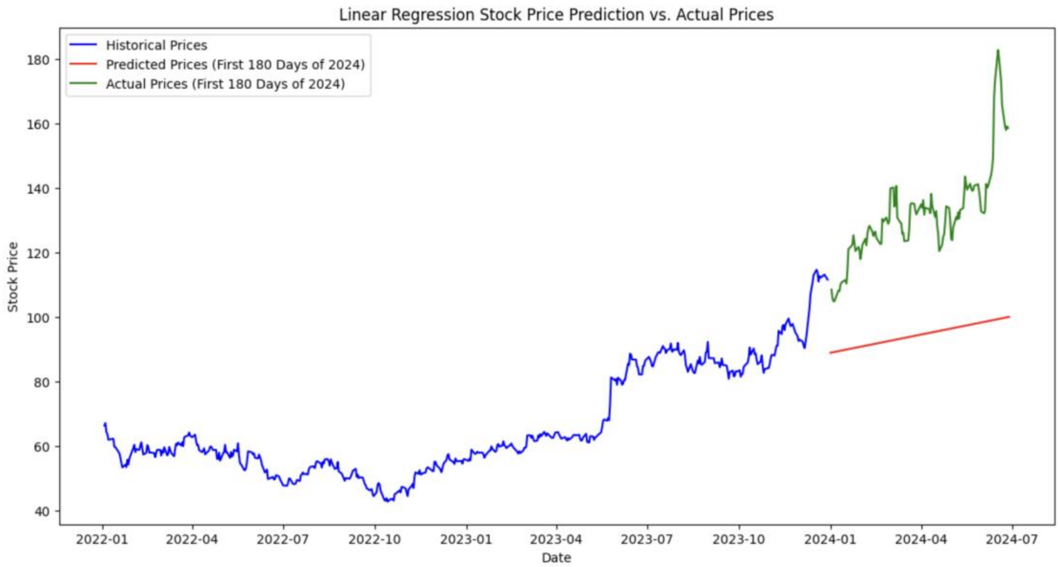

Figure 5. Linear regression of AVGO

(6) MSFT: MSFT's forecast red line shows a gradual upward trend. Still, it fails to capture the substantial share price gains in early 2024 driven by the company's deeper foray into artificial intelligence. The initial decline in forecasts for early 2024 reflects linear regression's limited ability to quickly adapt to the impact of sharp growth at the end of the year, highlighting its difficulty in handling sudden market swings related to technological advances and investor sentiment.

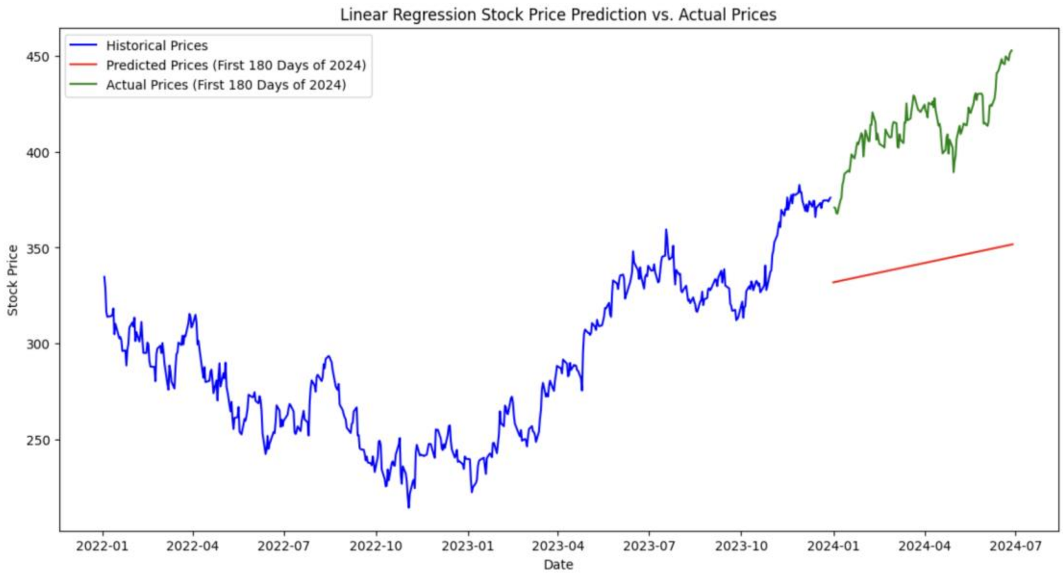

Figure 6. Linear regression of MSFT

Overall, linear regression tends to perform better when dealing with relatively stable stocks. However, it struggles when faced with sudden market crazes or excessive optimism from the public about a stock. Additionally, due to its algorithm, linear regression often fails to predict the stock price accurately on the first day of prediction. Furthermore, it cannot effectively predict sudden fluctuations and tends to oversimplify the stock's movement by using a simple straight line.

2.1.2. Support Vector Machine (SVM): Using 3 years data

The Support Vector Machine is a supervised machine learning algorithm used for both classification and regression tasks. In the context of stock price prediction, SVM functions as a regression model, learning to map historical stock data to corresponding stock prices.

How Support Vector Machine Works

Feature Space and Kernel Functions:

Support Vector Machines (SVM) work by converting the input data into a higher-dimensional feature space to identify a hyperplane that effectively separates the data points. When used for stock price prediction, input features usually consist of historical price data, trading volumes, and other pertinent financial metrics.

Regression with SVM:

In regression tasks such as stock price prediction, Support Vector Machine (SVM) aims to find a function that can predict continuous output values based on input features. This is accomplished by defining a loss function that penalizes deviations between predicted and actual prices.

Margin and Support Vectors:

SVM aims to maximize the margin between the hyperplane and the closest data points, which is important for defining the decision boundary. In stock prediction, the margin reflects the confidence level in the predicted prices compared to the historical data.

Kernel Trick for Non-linear Relationships:

SVM can utilize kernel functions such as the Radial Basis

Function to handle non-linear relationships between input features and output predictions, enabling it to capture complex patterns and relationships in stock data and enhance its predictive capability.

Prediction Process:

To predict future stock prices, a Support Vector Machine (SVM) uses the model parameters that have been trained using historical data. These parameters are then applied to new input data to make predictions. The predicted values reflect the anticipated future prices based on the patterns and relationships that have been learned.

Outcome and observation

(1) AAPL: The SVM model predicted the AAPL stock price in 2024 fairly accurately, reflecting its ability to forecast typical market changes. However, it failed to foresee a significant increase in share price following AAPL's announcement of new AI capabilities. This oversight highlighted the limitations of SVMs in adjusting to sudden market volatility events.

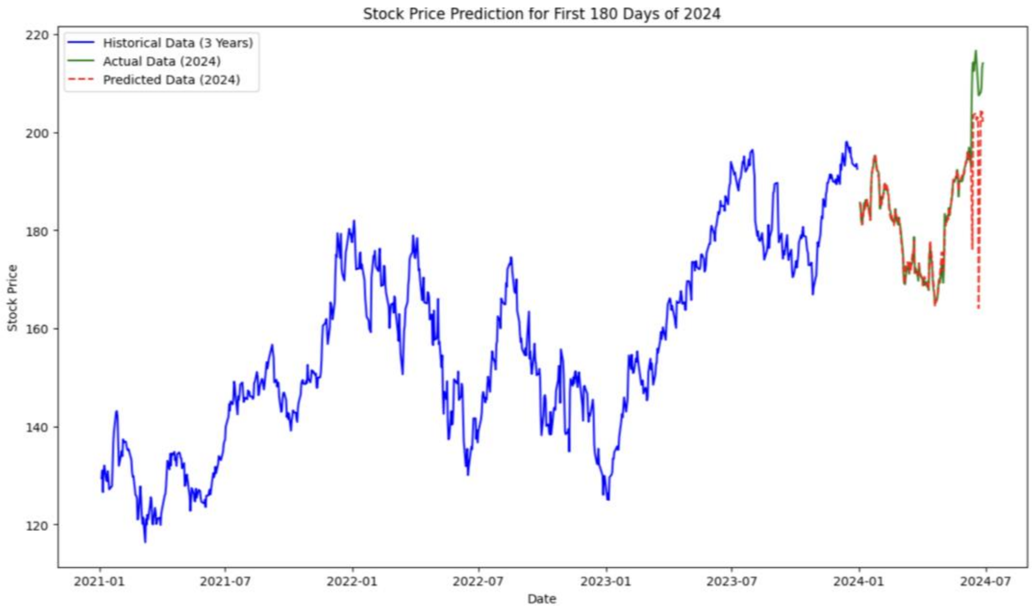

Figure 7. Support Vector Machine of AAPL

(2) TSLA: Despite TSLA facing negative news that can unpredictably affect its stock performance, the SVM model accurately predicted the stock's price movement during the first 180 days of 2024. This shows that the model can filter out noise and focus on essential patterns, even in the face of potential instability caused by external factors.

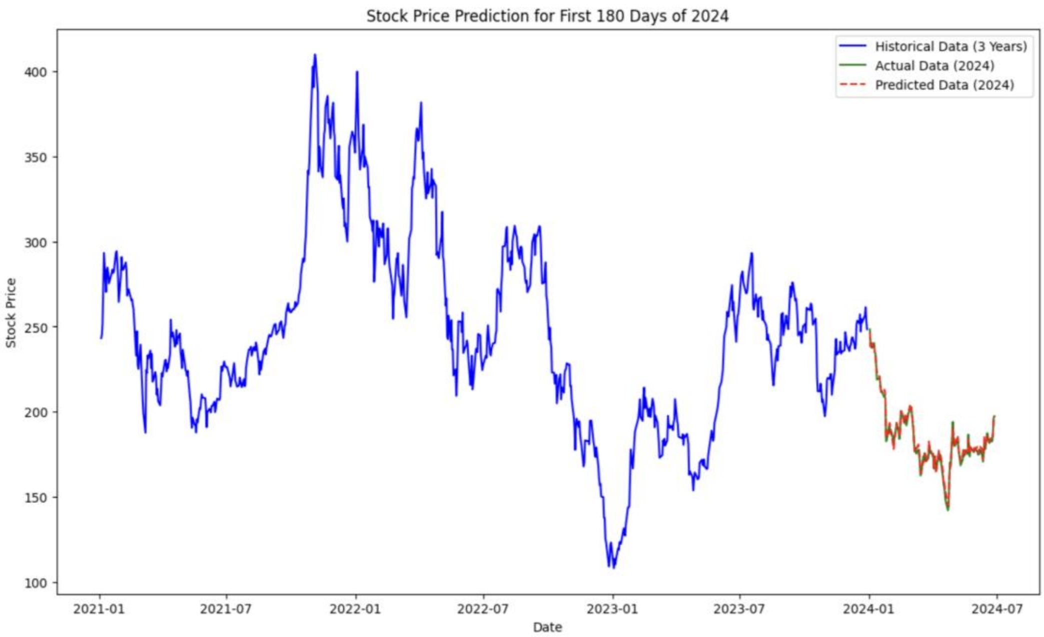

Figure 8. Support Vector Machine of TSLA

(3) JNJ: The SVM model accurately predicted stock price trends in the first half of 2024, closely matching actual data. It effectively captured stock volatility and trends, demonstrating its ability to simulate market complexity based on historical data. This accuracy shows that SVM is proficient at handling this stock's typical fluctuations and predicting its trajectory with impressive accuracy, even during usual market changes.

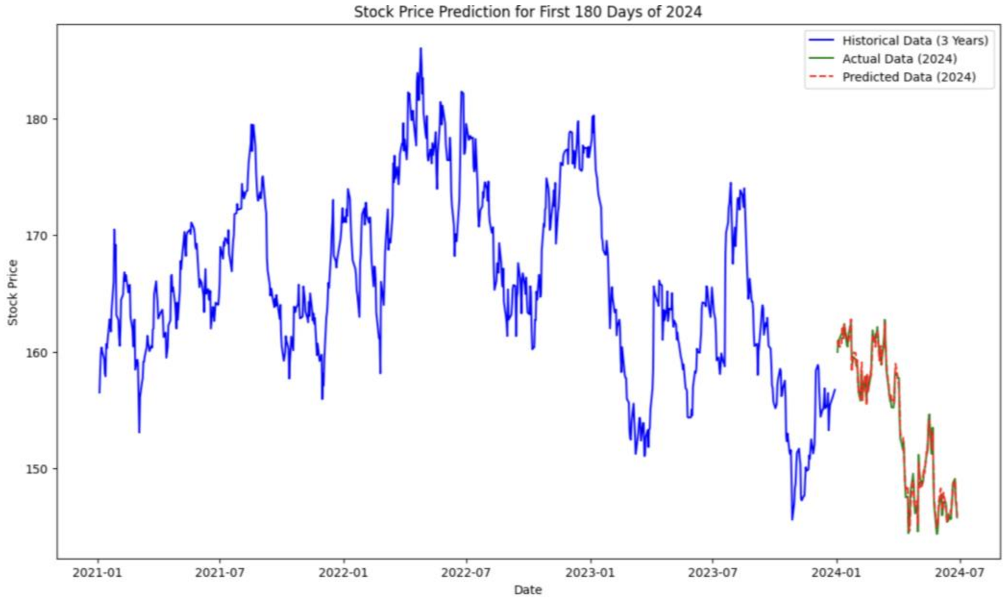

Figure 9. Support Vector Machine of JNJ

(4) NVDA: The SVM model underestimated the sharp rise in NVDA observed in early 2024. This was primarily due to SVM's reliance on historical data and its mathematical nature, which focuses on pattern continuity. Nvidia, a leader in GPU technology critical to AI, has seen its stock soar on rising demand and optimistic future forecasts for AI advancements. However, SVM's difficulty in adapting to rapid changes in market dynamics reflects its limitations in adapting to sudden investor enthusiasm that differs significantly from past trends. This highlights the challenges SVM faces in capturing changes in speculative markets driven by technological innovation.

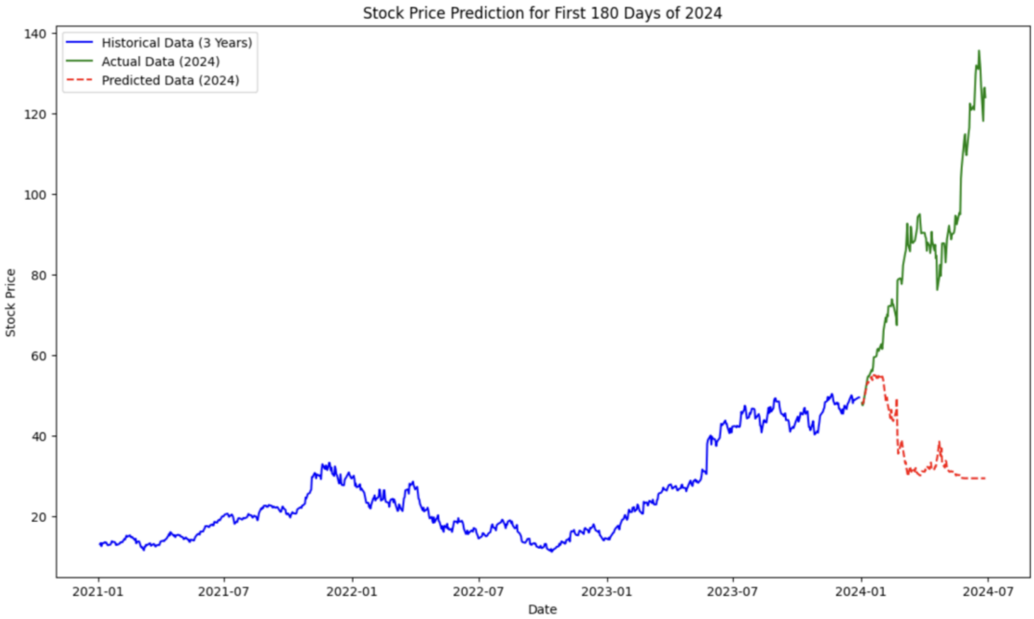

Figure 10. Support Vector Machine of NVDA

(5) AVGO: The SVM model incorrectly predicted a decline in AVGO's stock price despite the sharp rise in early 2024. This happened amidst rapid industry development and investor optimism driven by factors like artificial intelligence, fundamentally changing market conditions. The model's inability to predict AVGO's rise due to the artificial intelligence craze emphasizes the challenges traditional forecasting models face in fast-moving markets, where new technologies can significantly alter investment patterns and stock valuations.

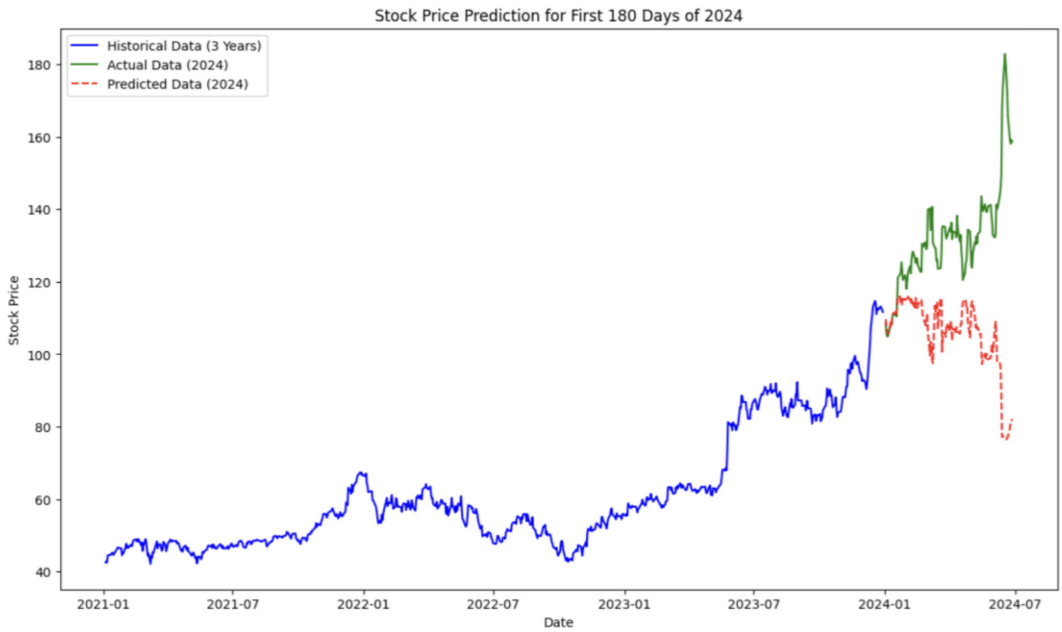

Figure 11. Support Vector Machine of AVGO

(6) MSFT: The chart shows that SVM predicted an increase in MSFT stock in 2023 and ongoing growth into 2024. However, it also anticipated a slight decrease in stock price early in the year. This forecast may be due to the model's tendency to expect a correction after significant gains, interpreting the rapid rise in 2023 as potentially unsustainable. This highlights the limitations of SVM in accurately predicting scenarios were continuous technological advancements and strong investor interest drive stock prices upward, going against traditional predictive models.

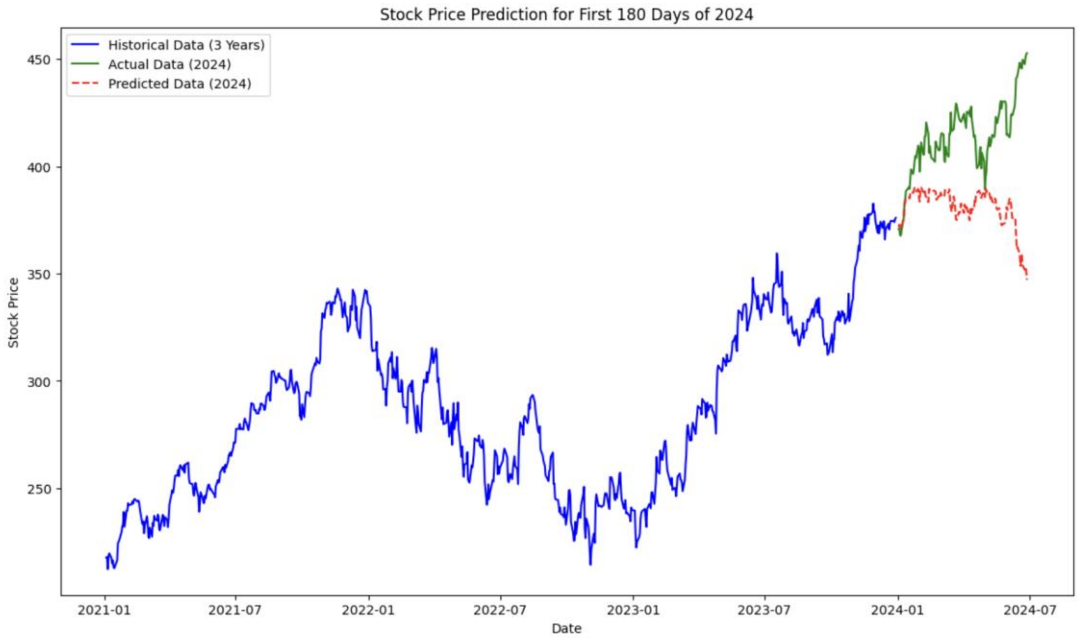

Figure 12. Support Vector Machine of MSFT

Overall, the SVM predictions for stocks such as TSLA, AAPL, JNJ, MSFT, NVDA, and AVGO in early 2024 demonstrate the model's strengths and weaknesses in financial forecasting. SVM performs well when market conditions are stable and historical data patterns are consistent, as observed with JNJ, where it accurately predicts price movements in line with past trends. However, the model encounters challenges adapting to rapid market transformations, mainly when influenced by technological advancements or significant external events. For example, in the tech sector with stocks like NVDA and AVGO, SVM failed to predict the sharp price increases driven by AI developments, highlighting its limitations in adjusting to new and fast-evolving market dynamics. This contrast emphasizes SVM's effectiveness in stable environments and its difficulties in capturing speculative momentum and investor reactions to new information in more volatile sectors.

2.1.3. Arima: Using 5 years of data

ARIMA stands for Autoregressive Integrated Moving Average. It is a statistical analysis model used to predict a series of future points based on past data. It provides the flexibility to model material where future points are influenced by past values. Its predictions are not the same as other models. It can only predict a confidence interval. The interval reflects the uncertainty in the predictions, acknowledging that the exact future values are unknown and could vary within a specified period.

How Arima works

Integration and Differencing:

ARIMA models address non-stationarity in the time series by differencing the data until it becomes stationary. This process involves subtracting previous observations from current observations, which helps stabilize the time series mean by eliminating changes related to trends or cycles.

Auto Regression and Moving Averages:

The autoregressive part of ARIMA predicts future values based on past values, establishing a direct relationship where past stock prices are used to forecast future prices. The moving average component, on the other hand, incorporates past forecast errors into the model to refine future forecasts. This dual approach allows ARIMA to adaptively improve its predictions based on errors observed in prior output.

Prediction Process

In the prediction phase, ARIMA uses the model parameters estimated from historical data to forecast future stock prices. These predictions are not just single-point estimates but are typically accompanied by confidence intervals, which provide a range of probable future values, reflecting the inherent uncertainties in predicting volatile data like stock prices.

Outcome and observation

(1) AAPL: The ARIMA model effectively forecasts AAPL stock prices by using past data trends and volatility, factoring in the autoregressive terms that account for previous price dependencies, integrated components for differencing to stabilize the mean, and moving average elements that smooth out noise and prediction errors. This combination allows ARIMA to adjust to sudden market changes like Apple's announcement, which positively impacts stock value, capturing changes that linear models often miss.

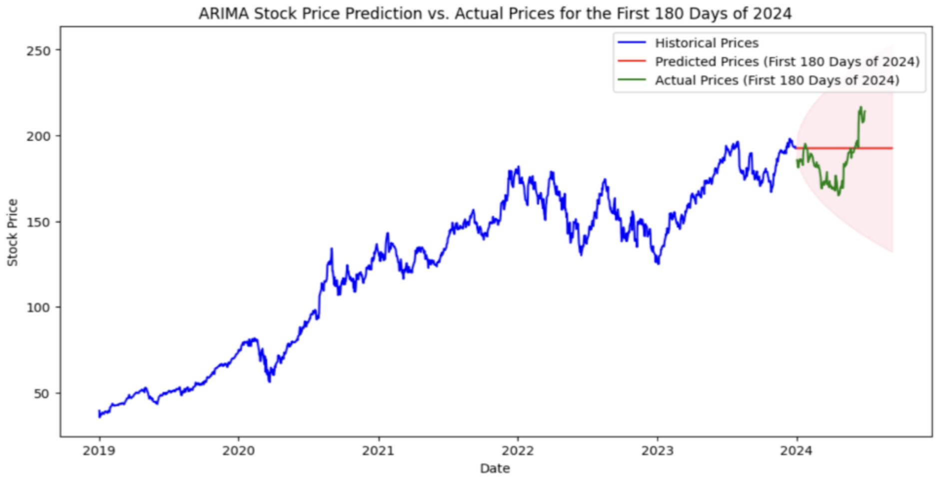

Figure 13. ARIMA of AAPL

(2) TSLA: The ARIMA model predicts TSLA's share price in 2024 in the chart. The wide confidence intervals indicate significant uncertainty about future price movements. Although I classify TSLA as relatively stable, ARIMA's broad forecast range reflects its response to the stock's historical volatility, which may include frequent price swings. This feature of ARIMA is very useful because it adjusts its forecasts based on past price action, predicting potential volatility even when a stock is experiencing stable periods. Therefore, the model suggests more uncertainty about TSLA's future prices, consistent with the inherent volatility observed in historical data.

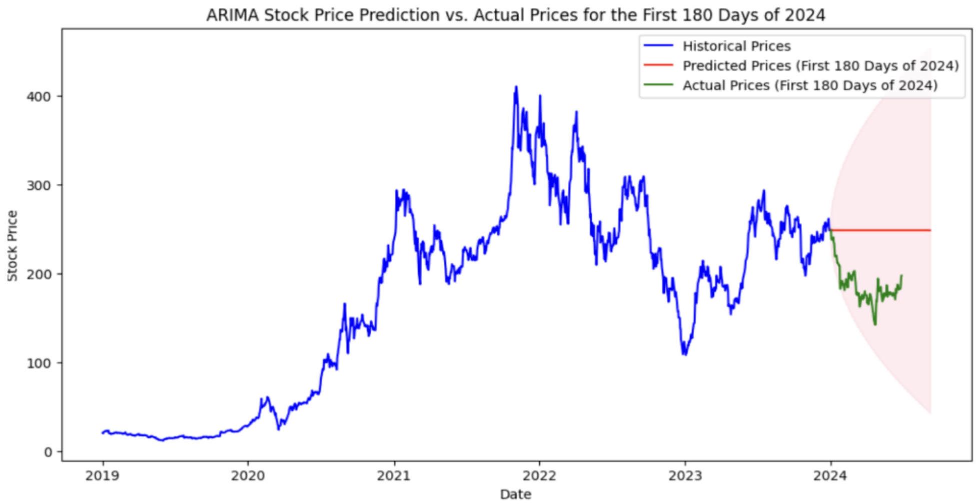

Figure 14. ARIMA of TSLA

(3) JNJ: The ARIMA model's forecast for JNJ is in the $80 range, while TSLA's forecast is in the $300 range. This shows that ARIMA closely tracks the actual price within the forecast confidence interval. It suggests that JNJ experiences mild price fluctuations, allowing the ARIMA model to maintain tighter confidence intervals. JNJ stock's consistent behavior highlights its stability and demonstrates ARIMA's effectiveness in capturing and predicting its price dynamics, even during market volatility.

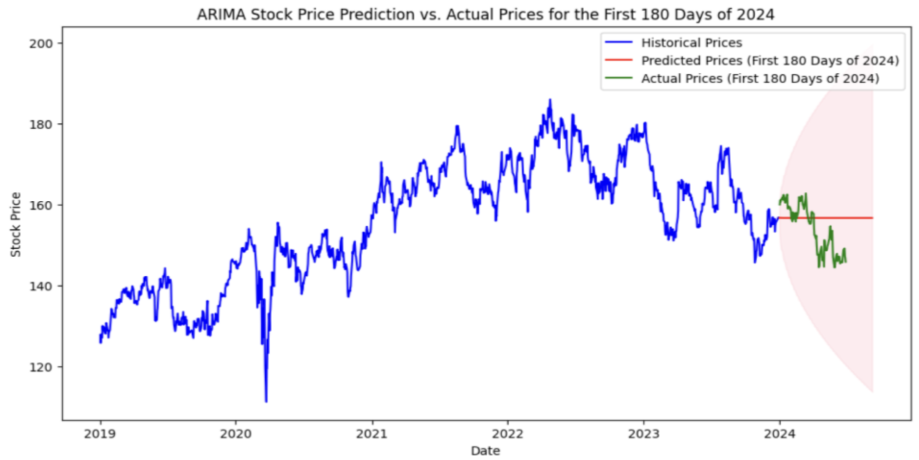

Figure 15. ARIMA of JNJ

(4) NVDA: The forecast plot generated by the Arima model for NVDA shows a significant difference between predicted and actual prices in early 2024. However, the recent surge in stock prices has far exceeded this forecast range due to the widespread belief that artificial intelligence represents a transformative future and NVDA's leadership in producing GPUs for artificial intelligence applications. This substantial increase illustrates the limitations of the ARIMA model: it struggles to predict extreme market moves driven by changes in investor sentiment and technological advances that sharply deviate from historical trends.

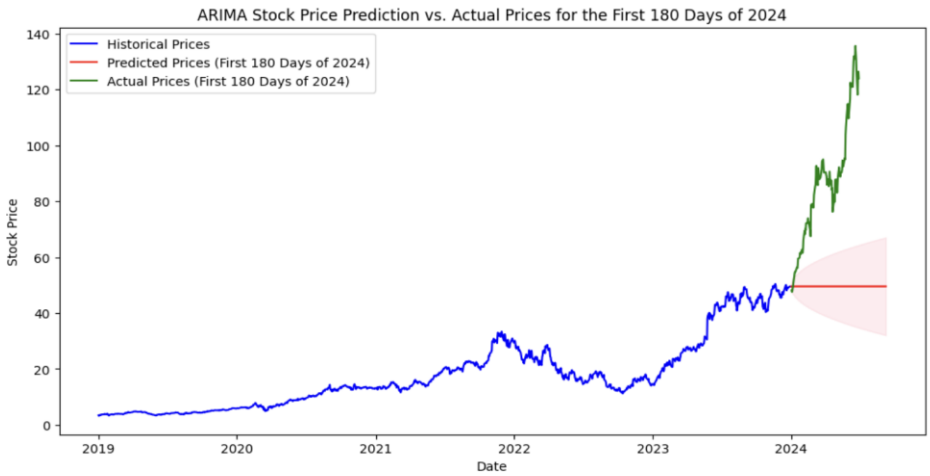

Figure 16. ARIMA of NVDA

(5) AVGO: The ARIMA model plot for AVGO stock also shows a significant difference between forecast and actual prices in early 2024, similar to what NVDA observed. Historically, AVGO has exhibited a stable and small upward trend, resulting in relatively narrow confidence intervals for ARIMA forecasts. However, due to the material impact of artificial intelligence on investor sentiment and expectations, AVGO's stock price experienced a significant increase that ARIMA did not foresee. This discrepancy highlights a standard limitation of the ARIMA model: its difficulty predicting sudden market changes influenced by emerging technologies and investor excitement that deviate significantly from established historical patterns.

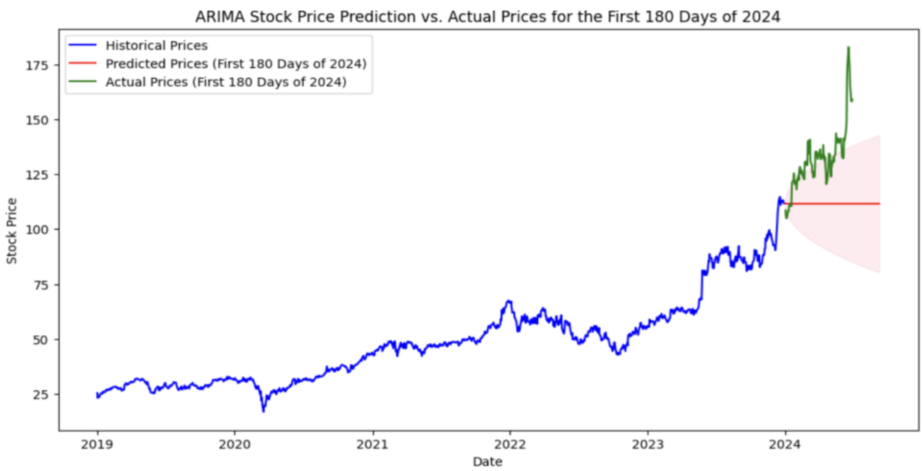

Figure 17. ARIMA of AVGO

(6) MSFT: The ARIMA model graph for MSFT stock effectively captures the stock trend, especially with the

substantial rise observed in 2023 that continued into 2024, though at a slightly moderated pace. This pattern likely contributed to the accuracy of ARIMA's predictions, which remained within the forecasted confidence interval. The model's ability to encompass such continuous growth within its predictions reflects the strength of ARIMA in tracking established trends over time. Despite historical fluctuations, the wide confidence interval accounts for MSFT's variability, ensuring the model is prepared for various potential outcomes. This shows how well ARIMA can adapt its forecasts to actual market behaviors, mainly when gradual rather than abrupt changes.

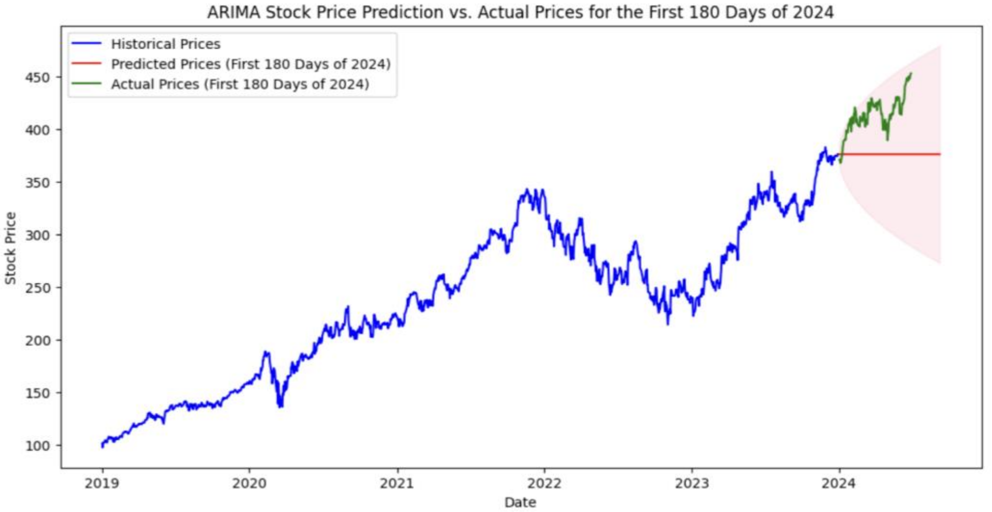

Figure 18. ARIMA of MSFT

Overall, The ARIMA model performs well in stable environments, accurately predicting trends for stocks with minimal fluctuations. It is effective for stocks unaffected by sudden market changes, providing precise forecasts within narrow confidence intervals. However, its effectiveness decreases with highly volatile stocks, leading to wider confidence ranges and reduced utility. Additionally, ARIMA struggles with stocks that are influenced by intense public sentiment shifts, failing to capture sudden spikes or drops. This highlights a limitation in ARIMA's ability to adapt to rapid market dynamics and external factors, suggesting the need to integrate more responsive or sentiment-aware analytical tools into forecasting models for stocks characterized by high volatility and emotional trading behaviors.

2.1.4. Long Short-Term Memory (LSTM): Using 10 years data

LSTM stands for Long Short-Term Memory. It is a type of recurrent neural network architecture designed to address the limitations of traditional RNNs in capturing long-term dependencies in sequential data. Unlike ARIMA, which relies on statistical modeling of time series data, LSTM networks are part of deep learning. They are specifically tailored for tasks involving sequential data, especially time series forecasting.

LSTM networks excel at learning and retaining patterns over long data sequences, making them well-suited for applications where understanding temporal dependencies is crucial.

How Long Short-Term Memory works

LSTM networks achieve their capability through a complex structure of memory cells and gates that regulate the flow of information. Key components include:

- Memory Cells: These maintain a memory state that can store information over time.
- Input, Forget, and Output Gates: These gates manage the flow of information into and out of the memory cell, regulating what information is stored, discarded, or outputted.
- Cell State: This represents the internal memory that runs along the entire chain of LSTM units, facilitating the learning of long-term dependencies.

Unlike ARIMA, which relies on historical statistical patterns and the adjustment of model parameters, LSTM networks autonomously learn and adapt to data patterns through their training process. This adaptability allows them to potentially capture and utilize intricate relationships in financial time series data, which may include non-linear dependencies and complex market trends.

Outcome and observation:

(1) AAPL: The LSTM model depicted in the graph predicted a slight decline in AAPL stock prices at the beginning of 2024. However, it didn't anticipate the subsequent increase after a noteworthy press conference. This highlights a standard limitation across various predictive models: the difficulty of integrating the impact of real-time, market-moving events into forecasts. Models like LSTM excel at analyzing historical data and identifying patterns over time. Still, they often struggle to predict the immediate effects of new information that hasn't been previously observed in the data. This example emphasizes the inherent challenge in forecasting stock prices influenced by unpredictable external events, which can result in significant deviations from predicted trends.

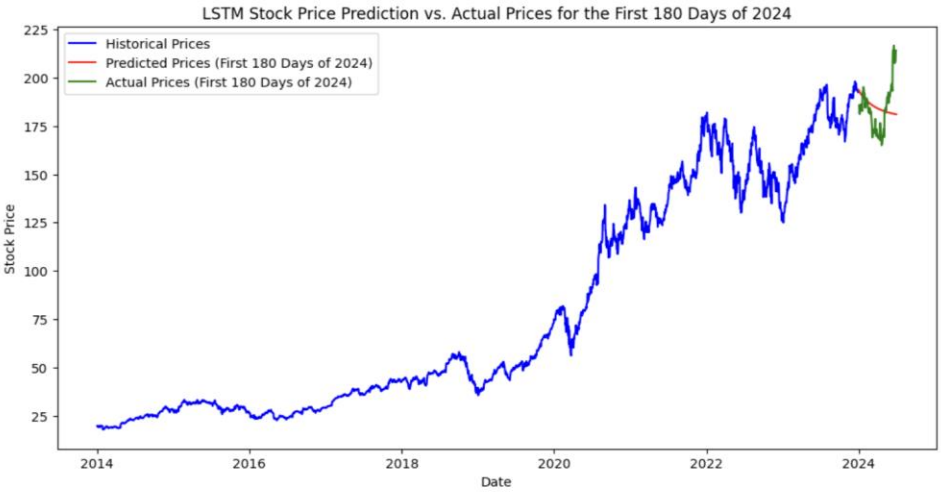

Figure 19. LSTM of AAPL

(2) TSLA: The LSTM model shown in the graph for TSLA stock indicates a clear difference between the model's predictions and the actual price movements during the first 180 days of 2024. The model expected TSLA to continue its upward trend from previous years, but instead, the stock price decreased. This difference likely stems from the model's inability to account for a sudden change in public sentiment or confidence, which can significantly impact stock prices. LSTM models are effective at identifying and forecasting patterns based on historical data. Still, they may struggle to adjust to rapid changes in market dynamics that are not reflected in the data they were trained on, such as sudden shifts in investor confidence or external economic factors.

Figure 20. LSTM of TSLA

(3) JNJ: The LSTM graph for JNJ stock shows that the model's predictions closely match market performance during the first 180 days of 2024. The LSTM model accurately forecasted both the continuation and the magnitude of the decline in JNJ's stock price. This success in prediction demonstrates LSTM's ability to effectively learn from and extrapolate historical trends, especially when the trends remain consistent over time. By analyzing patterns in past data, the LSTM model captured the ongoing downward trajectory of JNJ, showcasing its strength in situations where stock performance follows a clear and sustained trend. This accuracy is particularly valuable for predicting the performance of stocks like JNJ, where past behavior can indicate future performance.

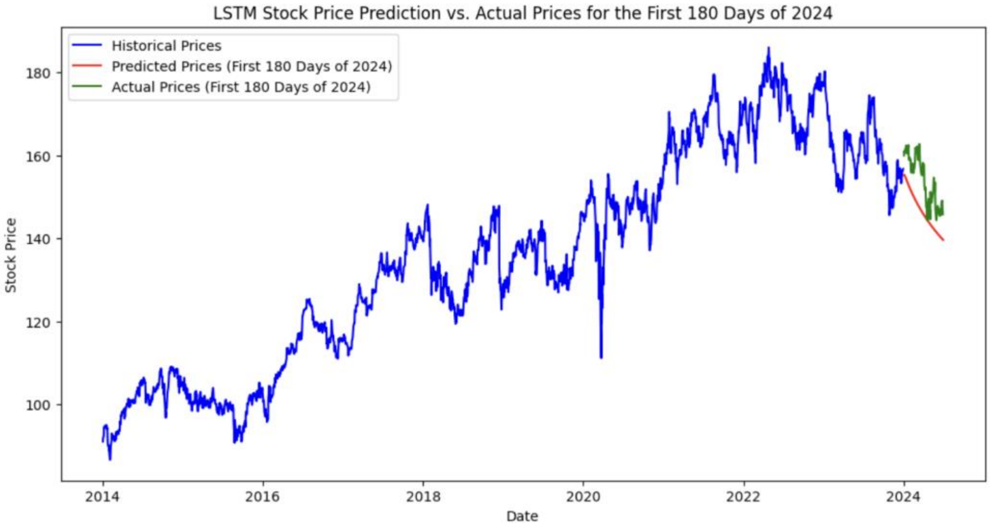

Figure 21. LSTM of JNJ

(4) NVDA: The LSTM graph for NVDA stock initially accurately predicted the stock's trajectory for the first 100 days of 2024, closely matching the actual prices during this period. However, after this initial phase, the model's forecast became more conservative, underestimating the continued upward trend that NVDA experienced. This difference highlights a standard limitation in LSTM models. While they are good at capturing data patterns over time, they can be cautious in their predictions when faced with sustained, aggressive growth trends. This conservative stance may be due to the model's reliance on historical data patterns, which, if not showing similar steep trends before, could lead the model to predict more moderate increases. These scenarios illustrate the challenges in machine learning-based predictions, especially in rapidly evolving sectors like technology, where investor sentiment and market dynamics can drive stock prices to unexpected heights.

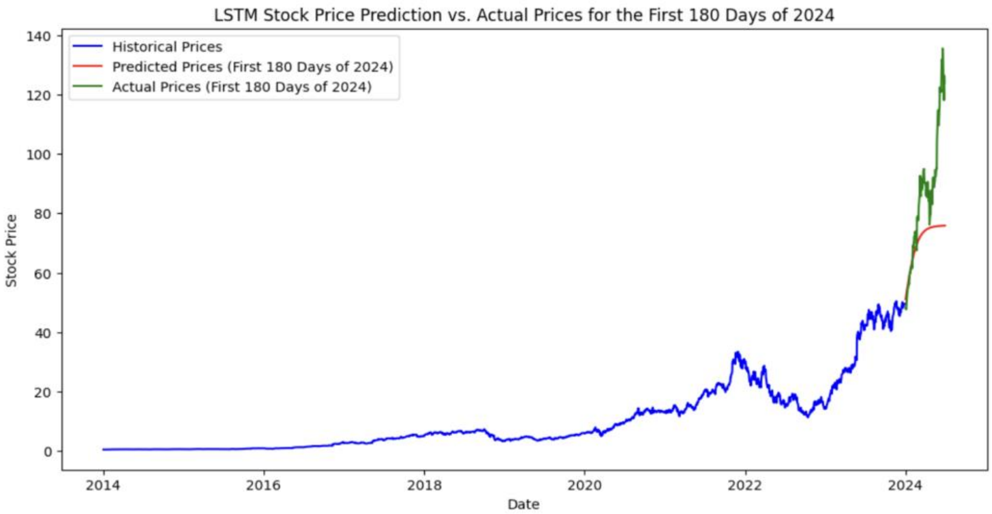

Figure 22. LSTM of NVDA

(5) AVGO: The LSTM model incorrectly predicted a sharp decline in AVGO's stock prices in early 2024, when there was a significant increase in reality. This mistake may have occurred because the model heavily relied on recent historical data, potentially skewed its predictive accuracy due to a downward trend at the end of 2023. LSTMs are sensitive to the most recent data trends, and this sensitivity can lead to projections that extend these patterns into the future. This case demonstrates the challenges LSTM models face in adjusting to sudden changes in market conditions that go against recent trends, revealing the inherent limitations of predictive analytics in volatile markets.

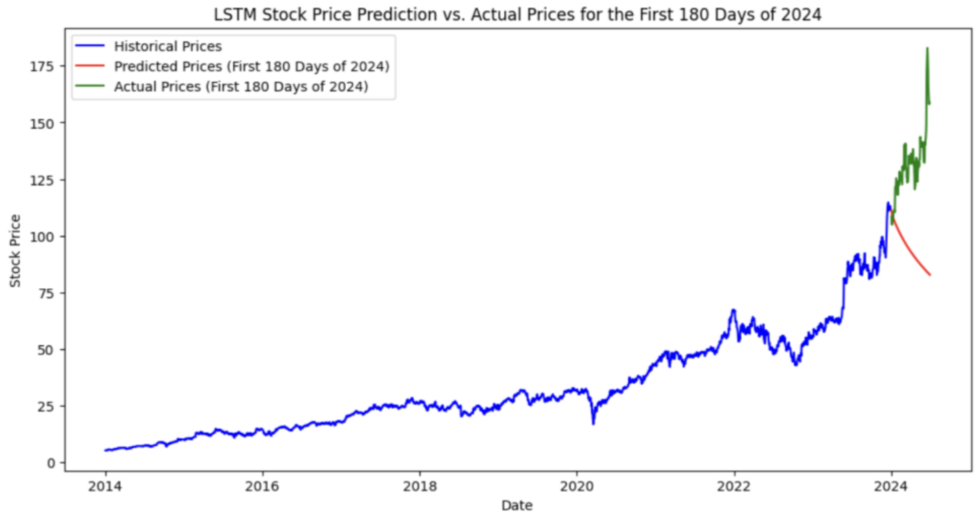

Figure 23. LSTM of AVGO

(6) MSFT: The LSTM model to predict MSFT stock showed a similar issue to the one observed with AVGO. The model predicted a significant decline for early 2024, but stock prices rose substantially. This prediction pattern suggests that the model might have been influenced by a minor decline in stock prices at the end of 2023, leading it to project this downward trend into the new year. LSTM models heavily weigh recent trends in the data they are trained on, which can result in forecast errors if these trends suddenly reverse. This scenario exemplifies the challenges in using LSTM for financial forecasting, especially when recent data might not fully capture future market dynamics or shifts in investor sentiment.

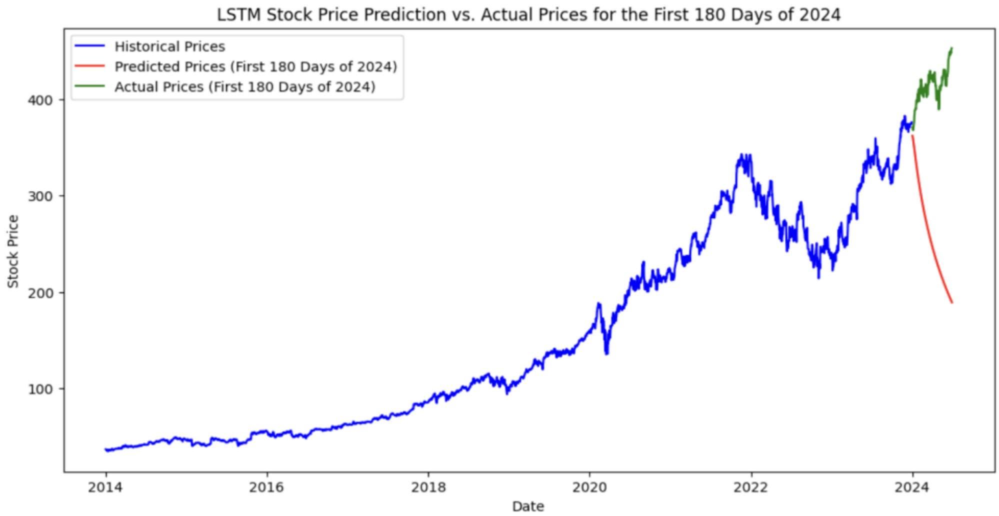

Figure 24. LSTM of MSFT

Overall, the prediction revealed that LSTM models perform well when market conditions align with historical trends, as demonstrated by JNJ. However, these models need help to make accurate predictions when faced with abrupt changes in trends, as seen with AVGO and MSFT. In these cases, the predicted declines did not match the sharp rises. This pattern highlights the limitations of LSTM models, which are proficient at tracking and extending established patterns but can be misled by sudden changes in short-term trends due to their reliance on recent historical data. These findings emphasize the need for models adapting to new market information and more complex investor behaviors. They also underscore the importance of further research and development in financial predictive analytics to address these challenges. This summary reflects on the broader challenges and necessary advancements in using machine learning techniques for financial forecasting in unpredictable market environments.

3. Conclusion

Predictive models play a crucial role in stock market forecasting, especially in handling the dynamics of stable and volatile stocks. For stable stocks like JNJ, models such as ARIMA and linear regression perform well by effectively extending historical trends into future forecasts. ARIMA excels at capturing time-based patterns and seasonal trends, making it practical for stocks with less volatile trends. On the other hand, linear regression works well in a consistent environment but struggles with sudden market movements and fails to predict abrupt changes in stock prices.

In contrast, volatile stocks experiencing sharp price swings due to technological innovation, market sentiment, or regulatory changes pose significant challenges. Models like LSTM and SVM are better equipped to handle these fluctuations because they can learn complex patterns and include longer-time dependencies. However, even these advanced models may struggle to predict extreme market behavior accurately. For instance, SVM may have difficulty predicting the rapid growth of technology stocks like NVDA and AVGO, while LSTM may miss sudden changes affected by external events or announcements.

The differences in model performance at different stock volatilities underscore the complexity of financial modeling and the importance of selecting appropriate models based on specific stock characteristics and market environments. Each model's underlying assumptions and mechanics determine its success rate, highlighting the need for a nuanced approach to stock market forecasting. Combining ARIMA, linear regression, LSTM, and SVM models can leverage their strengths to produce more accurate and comprehensive forecasts under various market conditions.

References

[1] H ochreiter S., & Schmidhuber J. Long short-term memory. Neural Computation, 1997, 9 (8): 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735.W.-K. Chen, Linear Networks and Systems (Book style). Belmont, CA: Wadsworth, 1993, pp. 123–135.

[2] Box G. E. P., Jenkins G. M., & Reinsel G. C. Time series analysis: forecasting and control. 4th ed. Wiley, 2008.

[3] Cristianini N., & Shawe-Taylor J. An introduction to support vector machines and other kernel-based learning methods. Cambridge University Press, 2000.

[4] Draper N. R., & Smith H. Applied regression analysis. 3rd ed. John Wiley & Sons, 1998.

[5] Stock J. H., & Watson M. W. Introduction to econometrics. 4th ed. Pearson, 2019.

[6] Goodfellow I., Bengio Y., & Courville A. Deep learning. MIT Press, 2016.

[7] Kim K. Financial time series forecasting using support vector machines. Neurocomputing, 2003, 55 (1–2): 307–319. https://doi.org/10.1016/S0925-2312(03)00372-2.

[8] Zhang G. P. Time series forecasting using a hybrid ARIMA and neural network model. Neurocomputing, 2003, 50:159–175. https://doi.org/10.1016/S0925-2312(01)00702-0.

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

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

[11] Yahoo Finance. Historical data for [specific stock, e.g., AAPL, TSLA, NVDA, etc.]. Retrieved from https://finance.yahoo.com.

[12] Bloomberg. NVIDIA stock rises on AI optimism. Retrieved from https://www.bloomberg.com.