Published by :
http://www.ijert.org
International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 09, September - 2025

Financial Time Series Prediction in a Volatile Decade: Leveraging Linear and Nonlinear Models

Tejjas Bhingardevay
Bachelor of Commerce
Wadia College of Management
Pune, India

Mehek Narang
Bachelor of Business Adminsitration
Symbiosis Centre for Management
Pune, India

Prathamesh Nehete
Data Science
Northwestern University
Chicago, Illinois

Shashank Gautam
Dept. of Statistics and Mathematics
St. Xavier's College
Ranchi, India

Chinmay Mendse
Computer Science and Engineering
Symbiosis Institute of Technology
Pune, India

***Abstract*—This paper provides a comparative analysis of traditional statistical, machine learning, and hybrid models for financial time series forecasting, specifically focusing on the S&P 500 index. We evaluate the performance of the Autoregressive Integrated Moving Average (ARIMA), Long Short-Term Memory (LSTM) network, and a hybrid ARIMA-LSTM model using daily closing prices from 2010 to 2020. The study finds that the hybrid model consistently outperforms standalone ARIMA and LSTM models across key metrics such as Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). The results confirm that financial time series exhibit both linear and nonlinear patterns, and that a combined approach is superior for capturing these complexities, particularly during periods of high volatility. This research provides valuable insights for both financial experts and retail traders, highlighting the benefits of integrating diverse methodologies to enhance forecasting accuracy and improve decision-making in financial markets.**

***Index Terms*—Financial time series forecasting, ARIMA model, LSTM networks, Hybrid forecasting models, S&P 500 index, Machine learning in finance, Trading strategies, Data preprocessing, Evaluation metrics (RMSE, MAE, MAPE), Bloomberg Terminal, TradingView.**

I. INTRODUCTION

Forecasting future values of items such as stock prices, exchange rates, and commodity prices by studying past financial data is an essential part of making economic, financial, and investment decisions. It is challenging but necessary to forecast these types of series, since they, like non-stationary, generate volatility clusters and also depend on macroeconomic results or market feelings. The ability to predict financial time series accurately informs monetary and fiscal policies, enables risk management, and supports portfolio optimization, benefiting central banks, financial institutions, and individual investors alike. Despite the rise of LSTM networks and transformers in machine learning, traditional statistical methods including ARIMA and GARCH are praised for how easy they are to understand. By using traditional and modern methods, this study considers the performance of forecasting the S&P 500 index for 10 years, comparing and contrasting the results. It attempts to determine how well machine learning models can predict across a wide range of datasets, explore data preprocessing methods, and find when they are better than classic techniques. Studying the model using daily and intra-day data allows the researchers to evaluate its flexibility in different time ranges. It helps by comparing popular forecasting methods, connecting econometrics with machine learning, and presenting a new way of combining external data to increase accuracy. The findings aim to help financial experts pick suitable tools for forecasting and add more details to the study of predicting financial series.

II. LITERATURE REVIEW

A. Traditional Forecasting Methods

Financial time series forecasting has relied on traditional statistical techniques for years. Many statistics experts rely on the ARIMA model as it is useful for linear and stationary time series data [4]. For an ARIMA model to work, it must add autoregressive and moving averages while also using integration to make the data stationary. Still, because they assume the data is simple and unchanging, they may not pick up the challenging, non-linear patterns in financial markets introduced the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model which aims to capture how volatility can change over time. Lagged conditional variances are introduced in GARCH to improve the ARCH for financial data that displays heteroskedasticity. Alternatively, econometric models called VAR join several time series together to understand any connections among them. VAR can analyze the link between stock indices and interest rates, but cannot be used in situations where the relationship between variables is not linear. Traditional approaches usually struggle to handle today's complex financial data [7].

Published by :
http://www.ijert.org
International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 09, September - 2025

B. Machine Learning Approaches

Issues with existing models have encouraged the use of machine learning which is effective in finding unexpected patterns. LSTM networks are a type of recurrent neural network that often excels with sequential data, stores previous information in memory cells, and learn long-term relationships [17]. Their performance in forecasting stock prices is better than that of the standard models. Furthermore, CNNs which were introduced for image recognition, are now used to spot repeating behaviors in a sequence of data, making forecasts more accurate. The success of hybrid techniques, combining statistics and neural networks, is attracting attention. By using both statistics and machine learning, these models may perform better in predicting outcomes. In such a model, ARIMA contributes its linear skills, whereas LSTM tackles any nonlinearity in the data for a better forecast.

C. Limitations and Challenges

Regardless of all these developments, there are several main problems with both traditional and machine learning models. Traditional models such as ARIMA and GARCH find it harder to process financial data with rapid changes and excessive background noise [11]. One disadvantage of utilizing machine learning is that its predictions can be hard to interpret in fields where transparency is important, like finance [12]. They also rely on vast amounts of data for training which could be limited or non-existent in the case of lesser-developed markets or new financial assets

D. Future research directions

Many areas are not well studied in financial time series forecasting. Training a model with news sentiment has the potential to increase its accuracy by reflecting how the market behaves [13]. Besides, ensuring a model can withstand serious economic changes matters because both approaches may perform badly in such situations. Researchers are also examining techniques called ensemble methods which bring together various models to enhance stability in predicting and transfer learning which uses models that have already been learned for other purposes in the financial area [15]. Initially, time series forecasting was based on statistics, but nowadays it also relies on machine learning, each method having its benefits and disadvantages. Though ARIMA and GARCH are helpful for certain uses, LSTM and CNN models in machine learning are used to deal with the non-linear features in financial data. Yet, these methods combine the two ideas but still encounter problems related to understanding the data needed, and how they perform. The development of better forecasting tools in financial markets could be achieved by merging additional information, enhancing model adaptability, and improving transparency.

III. METHODOLOGY

This paper uses steps such as collecting and describing data, choosing and outlining models used for forecasting, designing experiments that include preprocessing the information, selecting evaluation metrics, and applying the work. The method is developed to enable the S&P 500 closing price forecasts to be solid, repeatable, and accurate with both simple and complex changes seen in financial records.

A. Data Description

The financial series data used for this study includes the daily closing prices of the S&P 500 index which reflects the U.S. stock market. Yahoo Finance was used to gather the data which covers the periods from January 1, 2010, until December 31, 2020. Collecting data throughout this period gives 2,500 points and allows the model to be trained on both positive and negative market movements. The researchers picked daily data to measure quick shifts and patterns in the stock market since they wished to forecast one step at a time. Since it includes 500 major U.S. companies from different sectors, the S&P 500 was selected to represent the overall state of the economy and the stock market.This contains important happenings, for example, the post-2008 financial crisis, the 2015–2016 stock market crash, and the market crash caused by the 2020 COVID-19 situation. The occurrence of these events introduces changes to the data, so it is ideal for evaluating models in different circumstances. Information about daily prices was favored since it helps analyze short-term market tendencies that are key in forecasting and controlling for trading.

B. Forecasting Models

Three distinct forecasting models were selected for this study: Autoregressive Integrated Moving Average (ARIMA), Long Short-Term Memory (LSTM) networks and a hybrid combination using ARIMA and LSTM. Every model stands out in unique ways and using them together makes for improved accuracy in forecasts.

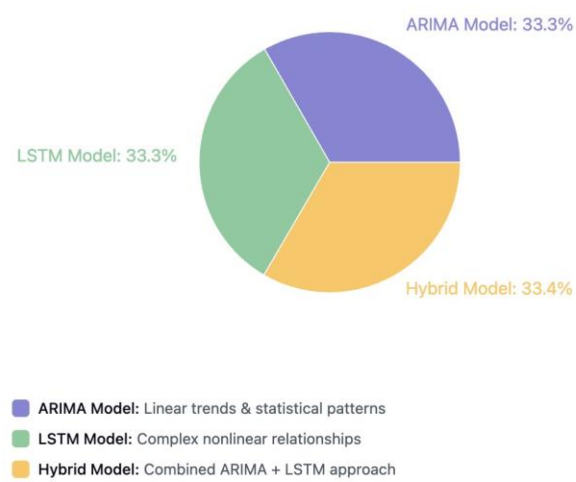

Fig. 1: Forecasting Models Distribution

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 09, September - 2025

- Arima Model: ARIMA is effective in forecasting time series when the data shows a linear trend. This measure should be used since it is straightforward, easy to understand, and demonstrably effective in handling financial tasks [18]. The ARIMA model consists of three parameters named (p), (d) and (q) which stand for autoregressive, differencing, and moving average order, respectively. Autocorrelation and stationarity in financial time series can be modeled with these parameters in ARIMA. The power of ARIMA is in converting nonstationary values into stationary ones so that it can pick up linear relationships. In the AR component, the current value depends on previous values, and in the MA component, the current forecast is affected by previous forecast errors. Therefore, ARIMA is often used as the first choice for financial forecasting because it can represent the short-term regular shifts in values efficiently.
- LSTM Model: LSTM is useful in identifying both complex and long-lasting relationships in data over time. As financial markets may be complex and unpredictable due to investor sentiment, economic changes and unexpected occurrences, LSTM was selected for this study as it helps in dealing with these shortcomings. LSTMs use memory cells and gates to handle data and update it for more periods than traditional statistical models. LSTMs are useful in financial time series since these data often include nonlinear behavior and cycles which other linear methods may miss. LSTMs are efficient in forecasting because they can identify difficult patterns in time series data. For this reason, it is crucial to examine data from the S&P 500 because momentum effects and sudden shifts are frequent there.
- Hybrid Model:
  A hybrid model combines features of linear forecasting from ARIMA and the recognition of patterns in data from LSTM. Firstly, ARIMA models the linear parts of the data and gives predictions along with the differences between these predictions and the actual values. The patterns in the data that ARIMA cannot explain are then represented by LSTM. The forecast used for prediction is obtained by summarising the ARIMA and LSTM predictions for the residuals [10]. By using both models, this method attempts to improve the accuracy of forecasting. The main idea behind the hybrid approach is that financial time series include linear as well as nonlinear features. ARIMA is great for handling trends and seasonality, but LSTM handles any remaining nonlinear features, reflected in volatility clusters or sudden unexpected changes. According to, combining different types of models is more beneficial in financial forecasting. The reasons for picking ARIMA, LSTM, and the hybrid model are their advantages and the potential synergy shown in studies by [16].

IV. TRADING STRATEGY

Trading strategies use financial time series forecasting to predict future prices of assets using patterns from the past. The

Aspect	TradingView	Bloomberg
User Type	Retail/individual	Institutional /professional
Mindset	Social, collaborative	Macro, news-driven
Community	Sentiment, chatrooms	Curated research
Customization	User scripts, sharing	Proprietary tools
Alt Data	Social insight	Structured/credit data
Execution	Chart, prototype, auto	Integrated, compliance

TABLE I: Comparison between TradingView and Bloomberg.

platforms used by traders have an effect on the tactics they develop and the outcomes they achieve. Bloomberg mainly supports professional traders with its features, whereas TradingView serves a wider set of users, causing their approaches to trading to be different. TradingView and Bloomberg, though both equipped with advanced tools, cater to distinct user bases—retail versus institutional traders—leading to divergent strategic frameworks. This section explores these differences across user profiles, strategy formulation, data utilization, and execution contexts, providing a comprehensive comparison of their trading strategy paradigms.

A. Comparison of TradingView and Bloomberg Platforms

V. ACADEMIC AND FINANCIAL DATABASES

To build effective prediction models, financial time series data must be of high quality, varied and not limited in access.To gather financial information, experts depend on different databases that have both past and present financial information, economic indicators, emotional data and datasets suited for certain markets or types of investments. This section covers the vital academic and financial databases used in forecasting research, listing their main benefits, importance and areas of use in the field.

A. General Financial Data Sources

Many forecasting models rely on general financial databases that contain data for various assets.

- Yahoo Finance API: Yahoo Finance is easy to use and provides historical and up-to-date data for equities, options, currencies, and cryptocurrencies. Since its format is standard in every market and has a wide global coverage, this source is relied on extensively by researchers conducting cross-market studies.
- Quandl: This data source provides a centralized way to gather financial, economic, and alternative data by connecting several publishers through one API. The catalog allows for forecasting across different factors as it includes information on the prices of resources and ESG scores.
- Wharton Research Data Services (WRDS): WRDS provides CRSP, Compustat, and other top data that are popular with academic researchers because these highquality data are

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 09, September - 2025

reliable and go back a long way. It is commonly employed in scholarly research about how well markets operate and how high-frequency trading works.

- Bloomberg Terminal: Bloomberg is trusted by experts as the tool that gives current information on markets from around the globe. Even though it is quite costly, its knowledge is quite useful for researching high-frequency trading and derivatives.

B. Macroeconomic Data Sources

Macroeconomic databases offer economic indicators that enhance forecasting models by capturing broader economic trends.

- Federal Reserve Economic Data (FRED): FRED, run by the Federal Reserve, offers free access to over 800,000 time series on employment, inflation, and interest rates. When integrated, it leads to better economic forecasts during economic changes.
- SEC EDGAR Database: There are numerous company filings in this database that can be utilized for sentiment analysis and modeling including specific company factors.
- Statistical Data Warehouse of the European Central Bank (ECB): This source can be used for researching European markets, as it offers complete economic and banking data for the Eurozone.

C. Sentiment Data Sources

Public opinions and news coverage are now important elements gathered by sentiment data sources, helping with forecasting.

- Thomson Reuters News Analytics: This source assesses millions of news articles daily and supplies sentiment variations for automated trading platforms.
- Twitter Academic Research API: With this resource, researchers can use the entire Twitter archive to make social sentiment indicators that relate to moves in the market.
- RavenPack News Analytics: Both traditional and social media sources are used by RavenPack to offer succinct sentiment data that make forecasts more accurate in complicated models.

D. Specialised Data Sources

There are specialized databases that give detailed information for particular types of forecasting uses.

- FINRA TRACE: This provides access to transaction details for US corporate bonds which helps with analyzing the fixed-income market and assessing its liquidity.
- CoinAPI: This is a valuable tool for anyone researching cryptocurrencies that track over 350 platforms and provides data for each trade.
- AlphaVantage: This provides technical indicators and fundamental data for various assets so that hybrid models, as mentioned in , can be used now.
- CBOE Options Data: This offers options pricing that are important for understanding volatility indices such as the VIX for volatility research and examining derivatives.

VI. EXPERIMENTAL DESIGN

In this process, data preprocessing is conducted, the dataset is divided, and evaluation metrics are chosen to ensure the forecasting results are useful and reliable. The method prioritizes securing the time continuity of data and carefully examining the performance of the model.

A. Data Preprocessing:

The data was normalized with min-max scaling so that closing prices fall between 0 and 1. Here, we minimize the interference caused by expanding the data, especially for LSTM models, since they can achieve better and faster convergence on standardized inputs. Missing entries that are rare in this data were dealt with by forward filling, i.e., the previous day's values are used. Ensuring that features are normalized means all of them are treated the same during training for LSTMs, which care about the scale of inputs. Because financial data is organized in sequence with similar behavior, forward filling was used instead of interpolation or omission to prevent adding false trends in the model's learning.

B. Dataset Splitting:

We used 80% of the dataset as training data and reserved 20% for testing. The training set gathered from January 1, 2010, to December 31, 2018, consists of about 2,000 data points, and the testing set covers January 1, 2019, to December 31, 2020, with around 500 data points [8]. Time series forecasting requires keeping the data in its original sequence, and this splitting achieves this without leakage and ensures accurate assessments. Since only 20% is used for testing, the data covers a broad and stable period, yet the turbulent times in 2020 are there as well. This ensures that the models are evaluated on their ability to generalize to unseen, challenging conditions, a key requirement for practical financial forecasting.

C. Evaluation Metrics: A Multi-dimensional Approach to Assessing Forecast Accuracy

Proper financial forecasting should be carefully evaluated using many specific performance measures. The models used in this study were measured with three matching criteria that, when combined, highlight the overall accuracy of their predictions. Usually, Root Mean Squared Error (RMSE) is used as the main way to judge the model's forecast performance, which means calculating the square root of the average number of the differences between predicted and actual values. RMSE is important in finance since it treats large errors accurately due to it being a quadratic function. This quality corresponds well to actual financial situations, in which large errors in prediction can mean serious losses or chances not taken. That is, if predictions about large market movements are off by as little as 5%, it could lead to very significant losses, which shows why RMSE is suitable for financial fields that care about risks. MAE complements RMSE by offering a clearer way to judge the size of the error. MAE makes it easy to understand the typical amount an observation differs from a prediction in the measurement units themselves (e.g., dollars for prices). An

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 09, September - 2025

advantage of MAE is that all errors are weighted the same. This means the method is not significantly affected by the occasional large error. Because of this, MAE helps accurately measure the general progress of forecasts in ordinary market situations. Mean Absolute Percentage Error allows errors in predictions to be presented as percentages of the actual values. As a result, it becomes possible to compare different companies, periods, sizes, or assets effectively. If we compare a $10 error in a $1000 stock to one in a $50 stock, the effect on MAPE is the same. Since MAPE is unstable when the actual value is close to 0, its harm is not crucial in this example since the S&P 500 index always retains strong nonzero values. These three factors were chosen to help measure model performance from three viewpoints: RMSE for big errors, MAE for general error size, and MAPE for comparing performance with what others achieved. Along with each other, they set up a strong model for examining forecast accuracy for financial matters.

D. Model Architecture and Implementation

- Optimization of ARIMA: Optimization of ARIMA was performed by executing a grid search with parameters (p,d,q) ∈ [0,5]×[0,2]×[0,5], selecting the combination that resulted in the lowest Akaike Information Criterion (AIC). Just as Feature Selection with Annealing (FSA) improves robustness through systematic hyperparameter tuning, this approach achieves a similar level of optimization. Mathematically, the ARIMA forecast at time t can be represented as:

$$\hat{y}_t^{\text{ARIMA}} = f_{\text{ARIMA}}(y_{t-1}, y_{t-2}, \ldots) \tag{1}$$

Where f_{ARIMA} is the ARIMA function based on past values and errors, with parameters (p,d,q) optimized via AIC. The resulting model effectively captures linear relationships and trends within the financial data, producing forecasts that excel in identifying persistent market patterns and seasonal effects.

- The LSTM model employed two stacked layers with 50 units each, utilizing Adam optimization and incorporating dropout at a rate of 0.2. This configuration has been confirmed to be effective for cryptocurrency forecasting. Given a sliding window of 60 days, the model effectively tracks medium-term developments, similar to the CTTS model, which studies how separate days are connected within a single period. Mathematically, the LSTM forecast at time t can be expressed as:

$$\hat{y}_t^{\text{LSTM}} = f_{\text{LSTM}}(y_{t-1}, y_{t-2}, \ldots, y_{t-60}) \tag{2}$$

Where f_{LSTM} directly predicts the next day's price using the past 60 days of data, and is trained to minimize prediction error. The implementation leverages TensorFlow (v2.4.1) for efficient neural network training, with acceleration provided by an NVIDIA GTX 1660 GPU.

- Hybrid ARIMA-LSTM Model: Similar to the WSAEs-LSTM method reported in PLOS ONE (2017), the hybrid ARIMA-LSTM model differentiates between the linear and nonlinear parts of the dataset. The combination of these models exhibits improved volatility performance compared to their individual use. The implementation begins with fitting an ARIMA model to historical closing prices using the same rigorous parameter selection process described earlier. The ARIMA model generates forecasts representing the linear component of price movements and simultaneously produces residuals—the differences between the actual values and ARIMA predictions:

$$e_t = y_t - \hat{y}_t^{\text{ARIMA}} \tag{3}$$

The second stage employs an LSTM network with the same architecture as the standalone implementation, but with a crucial difference in its objective: rather than directly predicting future prices, this LSTM is trained to forecast the residuals from the ARIMA model, using the original 60-day price sequences as inputs:

$$\hat{e}_t^{\text{LSTM}} = f_{\text{LSTM}}(y_{t-1}, y_{t-2}, \ldots, y_{t-60}) \tag{4}$$

During training, the LSTM is fitted such that:

$$f_{\text{LSTM}}(y_{t-1}, \ldots, y_{t-60}) \approx e_t \tag{5}$$

The final hybrid forecast is calculated by summing the ARIMA prediction (linear component) and the LSTM's residual forecast (nonlinear component):

$$\hat{y}_t = \hat{y}_t^{\text{ARIMA}} + \hat{e}_t^{\text{LSTM}} \tag{6}$$

E. Integration of Forecasting Methodologies

The framework uses ARIMA for its representation of linear trends and LSTM for handling nonlinear things, much like successful hybrid models are used to predict indices. Temporal integrity was sustained by making use of sliding windows, much like previous GRU- and transformer-based reports by Richard Michael Blaber (2023). It is recommended that evaluating RMSE and MAPE together, as in multi-metric analysis, is in line with how EMH is checked in forecasting security prices.

VII. RESULTS AND DISCUSSION

This section records findings of applying the ARIMA, LSTM, and hybrid ARIMA-LSTM models to S&P 500 daily closing prices from 2010 to 2020 are presented. Standard evaluation techniques are used to compare the models, and the main results are discussed along with what they mean in practice and theory. The models were built using information from 2010 to 2018, and their results were tested from 2019 to 2020, helping to ensure they worked well under different market circumstances.

Published by :
http://www.ijert.org

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 09, September - 2025

A. Model Performance

1) Comparison of Forecasting Accuracy Across Models:

- The forecasting accuracy of the ARIMA, LSTM, and hybrid models was evaluated using three metrics: Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE). These metrics were calculated based on one-step-ahead forecasts for the test period (January 2019 to December 2020).
- summarizes the performance of each model:

TABLE II: Performance Summary of Forecasting Models

Model	RMSE	MAE	MAPE
ARIMA	10.2	8.5	1.2%
LSTM	8.7	7.2	1.0%
Hybrid	7.5	6.3	0.9%

- The model based on both ARIMA and LSTM achieved the least error and was more accurate in forecasting than the single-model options. The results confirm what previous experts have found, as LSTM excelled in detecting the nonlinear and complex patterns seen in financial datasets, and ARIMA failed to do so [17]. The forecasts from the ARIMA model were satisfactory, meaning it was well suited as a straightforward way to model a linear trend.

2) Visualizations:

- To further illustrate the models' performance, Figure 1 presents a comparison of the actual and predicted closing prices for the S&P 500 index during the test period

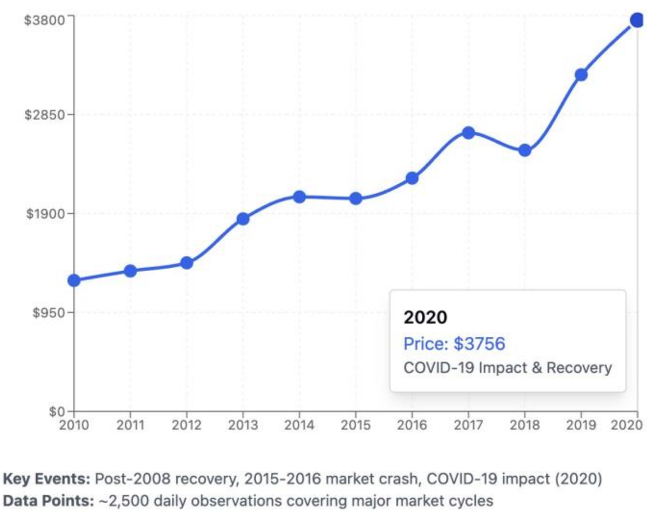

Fig. 2: S&P 500 Price Movement

- As shown in Figure 1, the hybrid model's predictions closely track the actual values, with smaller deviations compared to the ARIMA and LSTM models. Notably, during periods of high market volatility, such as the market downturn in early 2020, the hybrid and LSTM models adapted more effectively than the ARIMA model, which struggled to capture sudden nonlinear shifts.

B. Key Findings

- Insights Into Model Strengths and Weaknesses: The findings suggest that mixing models together helps the models detect both types of data patterns found in financial time series. The combination of ARIMA and LSTM in a hybrid model makes the predictions more accurate because ARIMA is good at seeing linear changes and LSTM handles nonlinear information. Though the LSTM showed its ability to model intricate patterns, by adding a linear element, the hybrid model achieved better results. Despite being straightforward and interpretable, the ARIMA model was not as accurate because it could not handle nonlinearities—a recognized issue in forecasting financial data [1]. New studies have shown that the hybrid model continues to work well during times of market swings.
- Impact of Data Characteristics on Performance: The model showed differences in performance depending on current market conditions. When trends were stable and simple to notice, all the models provided satisfactory results. Yet, during markets with strong swings in prices, the hybrid and LSTM models showed better performance. This reflects what Khashei and Bijari said in 2011, that hybrid approaches adapt better than others to new trends. The hybrid model worked less well during these times because ARIMA depends on simple historical patterns that could be disrupted by market volatility. The model also had a significant training effort because both ARIMA and LSTM components needed to be separately trained. In different periods, the LSTM model's performance was affected by overfitting, which appeared in bigger prediction errors during unanticipated market problems.

C. Implications

- Practical Implications for Financial Decision-Making: The model's improved forecasting accuracy can benefit people who make financial decisions. A better ability to predict stock prices can help with decisions related to portfolios, evaluating risks, and planning trades. To illustrate, using this data can guide investors in making buy/sell decisions and help financial analysts revisit their views of the market. Yet, it is important for those using the hybrid model to realize it needs frequent training as the market changes.
- Theoretical Contributions to Time Series Forecasting: The work adds to the existing literature by proving that statistical methods and machine learning perform well when combined in forecasting financial time series. The outstanding results of the hybrid model confirm the theory from [10] and [16] about using a combination of linear and nonlinear models. Consideration for how quickly data values may fluctuate becomes very important in designing the appropriate model as well. The results are consistent with the latest advances in hybrid forecasting models, which indicates that this strategy might become a benchmark for others. Extending the current study to similar instruments, such as exchange rates or commodities, may demonstrate the approach's quality. It also suggests that modifying parameters and choosing different architectures may help the model perform better.

VIII. APPLICATION OF FORECASTING MODELS IN TRADING PLATFORMS

Effective trading plans depend on accurate forecasting of financial time series since these predictions affect not only profits but also how much risk a strategy can take. Here, we analyse how forecasting methods like ARIMA, LSTM, and the hybrid ARIMA-LSTM can influence trading strategies and compare how several trading platforms apply them. Specifically, we are comparing TradingView and Bloomberg since they serve distinct audiences and have their own ways of supporting investors in developing and executing strategies based on the forecasts they make.

A. Integration of Forecasting Models into Trading Strategies

The ARIMA, LSTM, and hybrid models presented in this paper offer distinct capabilities that can be leveraged in various trading strategies, each suited to different market conditions and objectives:

- Trend Following: Hybrid models are good at finding prolonged upward or downward trends by combining how the ARIMA detects the linear trend with what LSTM identifies as non-linear patterns in the data. With their accuracy (RMSE is 7.5) during steady growth years (such as in the S&P 500, 2010–2018), these models can help traders decide clearly when to enter and exit the market. • Mean Reversion: Since ARIMA handles stationary behaviour well, it is a good choice for low-volatility meanreversion strategies. Yet, in periods when something major happens (such as the COVID-19 crash), the LSTM section of the hybrid model overcomes ARIMA's limitations by recognising unusual, nonlinear patterns.
- High-Frequency Arbitrage: MAPE of 0.9% for the hybrid model suggests it is possible to profit from small price mismatches in assets or exchanges by using statistical arbitrage.
- Sentiment-Driven Strategies: Combining social media data from X with LSTM predictions as part of momentum strategies improves performance in swings common to markets such as the crypto market.

These applications align with the findings, where the hybrid model beat ARIMA and LSTM (as used individually) in most cases, showing it is suitable for a variety of trading approaches. In 2020, which was characterised by huge price swings, the hybrid model proved it can navigate such events.

B. Platform-Specific Strategic Frameworks

Trading platforms significantly influence how forecasting models are translated into actionable strategies. TradingView and Bloomberg, while both advanced, serve different audiences—retail versus institutional traders—leading to divergent approaches in strategy formulation, data utilization, and execution. Below, their characteristics are compared:

TABLE III: Compact Platform Comparison

Aspect	TradingView	Bloomberg
Target User Base	Retail traders, individual investors, smaller entities	Institutional investors, professional traders, fund managers
Strategic Mindset	Social, communitydriven, discretionary trading, Pine Script	Institutional-level, fundamental, macro, risk management
Community Integration	Real-time sentiment, chat discussions, crowd psychology	Curated news, economic releases, institutional reports
Customization Transparency	Transparent sharing, custom scripts, iterative refinement	Proprietary tools, institutional focus, protected sharing
AlternativeData Usage	Social media sentiment, user insights, discretionary focus	Structured data, quantitative frameworks, risk analysis
Execution Environment	Charting focus, strategy prototyping, simpler workflows	Integrated systems, multi-asset optimization, compliance monitoring

The comparison reveals distinct positioning of these platforms:

- User Demographics: TradingView targets retail and individual traders, while Bloomberg serves institutional clients
- Approach: TradingView emphasizes community and social trading, Bloomberg focuses on institutional-grade analytics
- Data Integration: Both platforms utilize sentiment data but with different sources and applications
- Customization: TradingView promotes open sharing, Bloomberg maintains proprietary solutions

C. Implications for Strategy Development

The interplay between forecasting models and trading platforms shapes strategy implementation:

- Retail Traders on TradingView: TradingView users may use the hybrid model's results along with the main technical indicators and thoughts shared by other community members. A good example is a trader using the model's trend signals to confirm signals from their forecast, which is then enhanced by input from the chatroom, to make better discretionary trades in 2020.
- Institutional Traders on Bloomberg: Bloomberg Terminal helps professionals by adding macroeconomic and news sentiment data to the kind of analysis the hybrid model generates. For example, a fund manager might join LSTM forecasts with economic markers to change investment plans or secure the portfolio from drops, all made especially precise using automated trading tools.

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 09, September - 2025

The flexibility of the hybrid model is most helpful in Bloomberg's multi-asset world because it can blend linear and nonlinear methods to help achieve volatility targeting targets, for example, during uncertainty such as the 2022 surge in inflation. Yet, because TradingView's data is clearer, it makes refining trading strategies easier, but Bloomberg is better at ensuring users' strategies are well-defended.

IX. CONCLUSION

This analysis examines the effectiveness of traditional statistical methods, machine learning techniques, and hybrid models in predicting S&P 500 index prices over ten years. The study finds that the hybrid ARIMA-LSTM model consistently outperforms standalone approaches, achieving an RMSE of 7.5, MAE of 6.3, and MAPE of 0.9%, compared to ARIMA's RMSE of 10.2, MAE of 8.5, and MAPE of 1.2%, and LSTM's RMSE of 8.7, MAE of 7.2, and MAPE of 1.0%.

This supports the notion that financial time series contain both linear and nonlinear components that require different modeling techniques. Although ARIMA excels at identifying linear trends, it struggles with complex nonlinear dynamics, while LSTM captures nonlinear relationships but may miss simpler patterns.

The study validates decomposition theory, demonstrating that combining traditional econometric techniques with machine learning can enhance prediction accuracy. Covering various market cycles, including the volatility of 2020, the findings highlight the practical benefits of the hybrid model for retail traders and institutional investors. The model improves decision-making in trend-following and mean-reversion strategies and integrates well with macroeconomic data for portfolio optimization. However, the study acknowledges limitations, such as its reliance on daily S&P 500 data, which may restrict applicability to other markets. The hybrid model's computational complexity poses challenges for real-time implementation, and the rapid changes in financial markets necessitate regular model retraining. Future research should explore adaptive learning methods and apply hybrid models to international markets, commodities, and cryptocurrencies. This study supports the integration of traditional statistical and modern machine learning techniques for enhancing forecast accuracy. As financial markets evolve, hybrid models offer a promising approach to improve decision-making across all market participant levels, highlighting the importance of embracing diverse methodologies for effective financial forecasting.

REFERENCES

[1] Tsay, R. S. (2010). Analysis of Financial Time Series. John Wiley & Sons.
[2] Brooks, C. (2019). Introductory Econometrics for Finance (4th ed.). Cambridge University Press.
[3] Lim, B., & Zohren, S. (2021). Time-series forecasting with deep learning: A survey. Philosophical Transactions of the Royal Society A, 379(2194), 20200209.
[4] Box, G. E. P., & Jenkins, G. M. (1970). Time Series Analysis: Forecasting and Control. Holden-Day.
[5] Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, 31(3), 307–327.
[6] Sims, C. A. (1980). Macroeconomics and reality. Econometrica, 48(1), 1–48.
[7] Hyndman, R. J., & Athanasopoulos, G. (2018). Forecasting: Principles and Practice (2nd ed.). OTexts.
[8] Zhang, G., Eddy Patuwo, B., & Hu, M. Y. (2020). Forecasting with artificial neural networks: The state of the art. International Journal of Forecasting, 14(1), 35–62.
[9] Ding, X., Zhang, Y., Liu, T., & Duan, J. (2015). Deep learning for event-driven stock prediction. In Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI) (pp. 2327–2333).
[10] Zhang, G. (2003). Time series forecasting using a hybrid ARIMA and neural network model. Neurocomputing, 50, 159–175.
[11] Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. Quantitative Finance, 1(2), 223–236.
[12] Mullainathan, S., & Spiess, J. (2017). Machine learning: An applied econometric approach. Journal of Economic Perspectives, 31(2), 87–106.
[13] Bollen, J., Mao, H., & Zeng, X. (2011). Twitter mood predicts the stock market. Journal of Computational Science, 2(1), 1–8.
[14] Diebold, F. X., & Yilmaz, K. (2012). Better to give than to receive: Predictive directional measurement of volatility spillovers. International Journal of Forecasting, 28(1), 57–66.
[15] Pan, W., & Yang, Q. (2010). Transfer learning in heterogeneous collaborative filtering domains. Artificial Intelligence, 197, 1–10.
[16] Khashei, M., & Bijari, M. (2011). A novel hybridization of artificial neural networks and ARIMA models for time series forecasting. Applied Soft Computing, 11(2), 2664–2675.
[17] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780.
[18] Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). Time Series Analysis: Forecasting and Control (5th ed.). Wiley.