FORECASTING CBOE OPENING VALUES: AN EFFECTIVE ARIMA APPROACH WITH TRANSFORMATION AND MODEL DIAGNOSTICS

Porawat Visutsak[1,*], Ratchanon Wongpanti[2], and Ponrudee Netisopakul[2]

Received: May 23, 2024; Revised: July 08, 2024; Accepted: August 13, 2024

Abstract

This study investigates the use of an ARIMA model, coupled with Monte Carlo simulation, to forecast the opening value of a Volatility Index (VIX) time series. The data obtained from the Chicago Board Options Exchange (CBOE) for the years 1992–2019 have been transformed into stationary data using a detrend method and first-order difference. The Augmented Dickey-Fuller (ADF) test is used to ensure the data are adequately transformed. The autocorrelation function (ACF) and partial ACF (PACF) are then used to identify series with serial correlation and determine whether an autoregressive (AR) model is appropriate. Significant moving average (MA) lags are also determined for model identification. The Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC) are used to find the best-fit ARIMA model. Among the evaluated ARIMA models, which ranged from a random walk ARIMA(0, 1, 0) to ARIMA(2, 1, 2), the ARIMA(2, 1, 2) model was found to be the most optimal, exhibiting the lowest AIC and BIC values. This model was then used to forecast the opening value for the year 2014, using 2013 data as the real data. The generated ARIMA(2, 1, 2) model demonstrates reasonable alignment with the actual 2014 data, suggesting its potential for forecasting in this context. Additionally, Monte Carlo simulations are employed to assess the model's robustness by generating a range of potential outcomes, providing a more comprehensive understanding of the uncertainty associated with the forecasts.

Keywords: **ARIMA forecasting; Model selection; Monte Carlo simulation; Stationarity transformation; Time series analysis**

---

[1] *Department of Computer and Information Science, Faculty of Applied Science, KMUTNB, Bangkok, Thailand. E-mail: porawatv@kmutnb.ac.th[*]*
[2] *Faculty of Information Technology, KMITL, Bangkok, Thailand. E-mail: 65076055@kmitl.ac.th; ponrudee@it.kmitl.ac.th*
[*] *Corresponding author*

DOI: https://doi.org/10.55766/sujst-2024-04-e04932

Suranaree J. Sci. Technol. 31(4):030218(1-11)

Introduction

Forecasting financial time series, particularly stock market indices, is a complex yet crucial endeavor with implications for investors, financial institutions, and policymakers. The inherent volatility and non-stationary nature of these time series pose significant challenges for accurate predictions. Traditional forecasting models often struggle to capture the intricate patterns and dependencies within financial data. The efficacy of the Autoregressive Integrated Moving Average (ARIMA) model, a widely used statistical method in time series forecasting, is investigated in this study for predicting the opening value of a specific stock market index.

This study focuses on the Chicago Board Options Exchange (CBOE) data from 1992 to 2019. The data are transformed into a stationary series using detrending and differencing techniques. The Augmented Dickey-Fuller (ADF) test is employed to ensure the adequacy of data transformation. Subsequently, autocorrelations and partial autocorrelations are identified to assess the appropriateness of an autoregressive (AR) model, and the presence of significant moving average (MA) lags. The Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) guide the selection of the best-fit ARIMA model, which is then used to forecast the opening value for the year 2014, using 2013 data as the benchmark.

Stock Price Forecasting: A Review of Models and Methodologies

The ability to forecast stock prices accurately is a coveted goal in finance, with implications for investors, analysts, and policymakers alike. Numerous quantitative models have been proposed and tested for this purpose, each exhibiting their individual strengths and weaknesses. The literature review focuses on two prominent approaches: ARIMA models and Artificial Neural Networks (ANNs).

1. ARIMA Models: A Statistical Approach

ARIMA models, a class of statistical models widely used for time series forecasting, have found extensive application in stock price prediction. Their appeal lies in their ability to capture linear relationships and temporal dependencies within time series data. Several studies have explored the effectiveness of ARIMA models in predicting stock prices across different markets. Adebiyi *et al.* (2014) investigated the comparative performance of ARIMA and ANN models using New York Stock Exchange data, ultimately highlighting the superiority of ANN in capturing complex patterns and non-linear relationships. In their study, they utilized published stock data from the New York Stock Exchange and, through empirical analysis, unequivocally demonstrated that the ANN model surpassed the ARIMA model in predicting stock prices. Mohankumari *et al.* (2019) studied the application of ARIMA models for predicting daily stock trends, contributing to a deeper understanding of market behaviour. This research focuses on the application of this statistical method within the stock market context, aiming to assess its predictive capabilities using data from the stock market to analyze trends and develop a predictive model. Khanderwal and Mohanty (2021) conducted an in-depth analysis of ARIMA models for stock price prediction, demonstrating their potential for short-term forecasting while acknowledging limitations related to linearity and stationarity assumptions. The research methodology involved building an ARIMA model using historical stock price data and evaluating its predictive performance. The study’s findings revealed that the ARIMA model demonstrated promising results in forecasting stock prices, especially in short-term investments. Mashadihasanli (2022) focused on the Istanbul Stock Exchange, employing ARIMA models to forecast future price movements with varying degrees of accuracy. This study employed historical stock market data from Istanbul to train and test the ARIMA model’s predictive capabilities. By analyzing past price patterns and trends, the model aimed to forecast future price movements in the Istanbul market. Alotaibi (2022) further explored the application of ARIMA models in stock market trend prediction, emphasizing their value as a tool for informed investment decisions. The study highlighted the ARIMA model as a statistical tool used to analyze and forecast time series data, particularly in the context of financial markets. The research investigated the effectiveness of the ARIMA model in capturing and predicting stock price fluctuations.

2. Artificial Neural Networks: A Machine Learning Approach

ANNs, inspired by the biological neural networks in the human brain, offer a powerful framework for modeling complex non-linear relationships. ANNs exceptional capacity to discern intricate, non-linear relationships within extensive datasets, coupled with their inherent ability to adapt to shifting market dynamics, have solidified their position as the preferred methodology for stock price forecasting. Chandar, (2021) introduced nine new hybrid models for intraday stock price forecasting, combining three ANNs (BPNN, RBFNN, and TDNN) with nature-inspired optimization algorithms (GA, PSO, ABC). These models utilized technical indicators from historical data as input and were evaluated on four NSE datasets (National Stock Exchange of India) using root mean square error (RMSE), Hit Rate, Error Rate, and prediction accuracy metrics. According to the results, PSO-BPNN achieved the highest accuracy, while other models demonstrated varying degrees of success in predicting intraday stock prices. Chhajer *et al*. (2022) explored the application of machine learning techniques, specifically ANNs, Support Vector Machines (SVMs), and Long Short-Term Memory (LSTM) networks, for stock market prediction. The stock market's volatile and dynamic nature makes accurate prediction a challenge, but machine learning offers potential solutions due to its ability to learn from data and adapt to changing patterns. The paper discussed the strengths and weaknesses of each technique, highlighting their effectiveness in recognizing trends and patterns within stock market data. It also presented case studies and experiments conducted using these models on various stock markets, demonstrating their potential for accurate forecasting. While acknowledging the challenges and limitations, the paper emphasized the growing importance of machine learning in the financial sector and its potential to revolutionize stock market prediction and analysis. More complicated works on ANNs and Genetic Algorithms (GAs) for stock forecasting were also reviewed. Sharma *et al*. (2022) proposed a hybrid model combining ANNs and GAs for improved stock market forecasting. The aim of the model was to overcome the limitations of individual techniques, such as local minima issues in ANNs and difficulty in finding suitable fitness functions in GAs. By using GAs to optimize the weights of ANNs, the hybrid model sought to enhance prediction accuracy and minimize errors. The model was tested on DOW30 and NASDAQ100 indices, including data from the volatile COVID-19 period. The results demonstrated that the hybrid GANN model outperformed the traditional Backpropagation ANN (BPANN) in both short-term and long-term forecasting accuracy. The researchers concluded that the integration of GA with ANN effectively captured the fluctuating behavior of stock data, offering a promising approach for stock market prediction. Senthil *et al*. (2023) compared ARIMA and RNN-LSTM models in the context of the Indian stock market, highlighting the strengths and weaknesses of each approach. A case study approach was utilized to compare the accuracy and effectiveness of each model. The researchers explored the strengths and weaknesses of both ARIMA and RNN-LSTM in forecasting stock prices and market movements. By analyzing the results, it can be concluded that the ARIMA model outperformed the RNN-LSTM model in forecasting stock market trends in the short run. Specifically, the ARIMA (4, 1, 5) model was found to be more efficient in predicting the direction of change in the stock market indices based on the root mean square error (RMSE), mean absolute deviation (MAD), and mean absolute percent error (MAPE) metrics. These findings suggest that the ARIMA model is better suited for short-term forecasting in this specific context.

Challenges and Considerations

While both ARIMA and ANN models have shown promise in stock price forecasting, their application is not without challenges. For instance, ARIMA models require careful consideration of stationarity assumptions and may struggle to capture non-linear relationships. ANN models, while adept at handling non-linearity, can be computationally intensive and prone to overfitting. Furthermore, both approaches rely heavily on the quality and availability of data, with issues like missing data, outliers, and non-stationarity potentially impacting their performance.

The Way Forward

Despite these challenges, the ongoing research in this field continues to refine and enhance both ARIMA and ANN models. For example, hybrid models that combine the strengths of both approaches have been proposed, and advancements in machine learning techniques are constantly pushing the boundaries of ANN capabilities. Additionally, research on model diagnostics, such as that undertaken by Roy and Guria (2004), Liu and Zhang (2018), Hickey *et al*. (2019), Amin *et al*. (2022), and Lin and Liu (2022), has been crucial for ensuring the reliability and validity of model predictions. The work by Roy and Guria (2004) focused on the challenges posed by autocorrelation in regression analysis and discussed methods to identify and address these issues. Liu and Zhang (2018) introduced a novel surrogate approach for residuals and diagnostics in ordinal regression models. Hickey *et al*. (2019) provided a step-by-step guide on how to conduct regression diagnostics, and Amin *et al*. (2022) explored methods to assess the model's assumptions and identify potential issues like outliers or influential observations. They discussed diagnostic measures such as residuals, leverage values, and Cook's distance, explaining how these can be used to evaluate the model's fit and identify problematic data points. Lin and Liu (2022) introduced a novel approach, combining the strengths of various existing methods, offering a more comprehensive and flexible tool for assessing the adequacy of discrete regression models.

In brief, this review highlights the significant contributions of both ARIMA and ANN models to the field of stock price forecasting. While each approach has its own merits and limitations, their continued development and refinement hold promise for further advancing one's ability to predict and understand the complex dynamics of financial markets. Building on this foundation, this research specifically investigates the efficacy of the ARIMA model for predicting the opening value of a specific stock market index. Given its established popularity and success in time series forecasting, the ARIMA model serves as a relevant benchmark for analysis in this study.

Materials and Methods

This study employs CBOE data from 1992 to 2019. This dataset provides a treasure trove of historical options data. It encompasses comprehensive information on option prices, including opening, high, low, and closing prices for various contracts. Additionally, it contains details on trading volume, the number of contracts traded each day, and open interest, reflecting the total number of outstanding contracts. The dataset also includes Greeks (delta, gamma, theta, vega, and rho), which are crucial measures of option price sensitivity to various factors. Furthermore, it is expected to feature implied volatility, representing the market's anticipation of future price fluctuations. The dataset comprises 7317 records of daily Volatility Index (VIX) data. Each record contains the following information for a specific date: 1) date: the date of the observation; 2) vix_open: the opening value of the VIX on that date; 3) vix_high: the highest value reached by the VIX during that day; 4) vix_low: the lowest value reached by the VIX during that day; and 5) vix_close: the closing value of the VIX on that date. MATLAB R2023b is implemented in this study to transform the VIX data from a CSV file into a timetable format in MATLAB. Firstly, the CSV file is imported into a table, and the date column is then converted into a datetime format for compatibility. Next, a timetable is created using the datetime column as the row index and the remaining columns (vix_open, vix_high, vix_low, and vix_close) as variables. Finally, the data are visualized by plotting the VIX closing values over time, with the dates as the x-axis. It is important to label the plot clearly with a title, y-axis label (VIX value), and x-axis label (Timeline) for better understanding. Figure 1 illustrates the proposed method. The details are further explained in the following section.

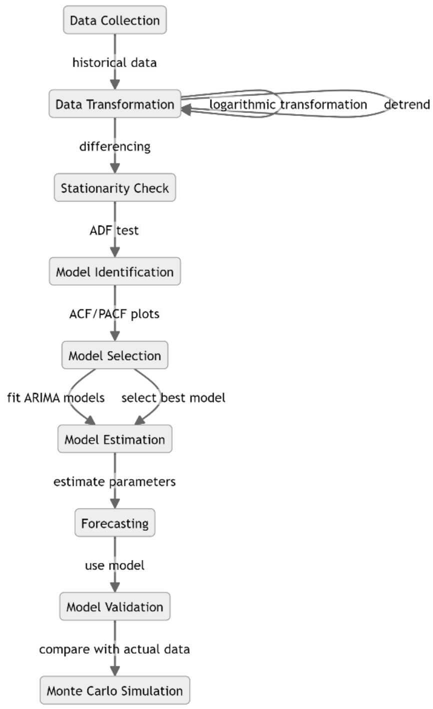

Figure 1. System block diagram

1. Data Collection: Gather historical CBOE (VIX) data from 1992 to 2019.

2. Data Transformation:
- Apply logarithmic transformation to stabilize variance.
- Detrend the data to remove linear trends.
- Apply first-order differencing to achieve stationarity.

3. Stationarity Check:
- Use the Augmented Dickey-Fuller (ADF) test to verify stationarity.

4. Model Identification:
- Analyze Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) plots to identify potential AR and MA terms.

5. Model Selection:
- Fit various ARIMA models (e.g., ARIMA(0, 1, 0) to ARIMA(2, 1, 2)).
- Use AIC and BIC to select the best-fitting model (ARIMA(2, 1, 2) in this case).

6. Model Estimation:
- Estimate the parameters of the selected ARIMA(2, 1, 2) model.

7. Forecasting:

- Use the estimated model to forecast VIX opening values for 2014.

8. Model Validation:

- Compare forecasted values with actual 2014 data to assess model performance.

9. Monte Carlo Simulation:

- Enhance analysis by simulating multiple VIX scenarios using the ARIMA(2, 1, 2) model.

In Figure 2, the graph illustrates the closing prices of the VIX over a period of 28 years, from 1992 to 2019. The x-axis represents the year, and the y-axis represents the VIX value. The line chart shows the fluctuations in the VIX closing prices over time, highlighting periods of high and low market volatility.

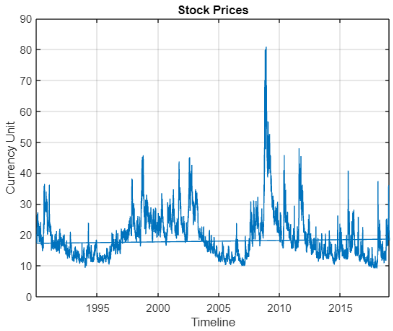

Figure 2. VIX closing prices from 1992 to 2019

In the realm of time series analysis, achieving stationarity is often a crucial step. Stationary time series exhibit consistent statistical properties like mean, variance, and autocorrelation across time, simplifying the modeling and forecasting process. A common approach for addressing non-stationarity in financial data, such as stock prices or volatility indices, involves a series of transformations.

First, a logarithmic transformation is often applied to compress the data range and stabilize variance, particularly when values tend to display exponential growth or high volatility. Subsequently, detrending is employed to remove any linear trend component, as trends can introduce non-stationarity by causing the mean to shift over time. Finally, differencing, a powerful technique for handling non-stationarity arising from unit roots, calculates changes between consecutive time points and can often transform non-stationary series into stationary ones.

While this transformation pipeline is effective, it is crucial to always verify stationarity after transformation using tests like the ADF. Furthermore, the order of integration (the number of differencing operations required) should be determined as this influences the selection of appropriate time series models like ARIMA. Additional transformations like the Box-Cox (to stabilize variance for data not following a log-normal distribution) or seasonal decomposition (to address seasonal patterns) can also be considered when necessary.

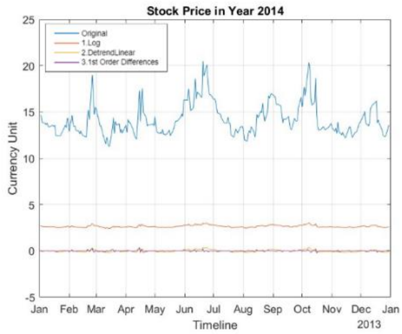

Figure 3. Results of data transformation

Figure 3 shows the effects of applying three transformations on a stock price time series:

1) Log Transformation (orange line): Taking the logarithm of the original stock prices (blue line) to compress the range of values. This is evidenced by the large fluctuations in the original data being reduced in the log-transformed data. The log transformation helps to stabilize variance and make the data more amenable to further analysis.

2) Detrending (yellow line): The detrended series removes any linear trend that may be present in the log-transformed data. In the graph, it appears that the slight upward trend in the log prices has been removed, resulting in a series oscillating more closely around zero. This is crucial because trends introduce non-stationarity, and removing them is a step toward making the time series stationary.

3) Differencing (purple line): The differenced series represents a change in detrended log prices from one time step to the next. As can be observed, the scale of the y-axis changes dramatically. The differenced series has a much smaller range than the previous transformations. Differencing is often effective in removing non-stationarity caused by unit roots (also known as random walks), where the current value depends on the previous value plus some random shock. This final transformation results in a series that appears to be more stationary, as it fluctuates around a constant mean without any apparent trends or persistent changes in variance.

To determine if time series data is stationary, hypothesis testing is employed. The null hypothesis posits that the data is non-stationary (has a unit root), while the alternative hypothesis suggests it is stationary (lacks a unit root). The ADF test is a common method for calculating a test statistic and a *p*-value. A small *p*-value (often below 0.05) implies stronger evidence against the null hypothesis, favoring stationarity.

While the ADF test is frequently used after transforming data (e.g., log transformation, detrending, and differencing), it can also be applied directly to the original data after differencing. This simpler approach is valid, especially if the additional transformations are deemed unnecessary.

It is important to note that the ADF test is not foolproof and depends on certain assumptions. A low *p*-value does not guarantee stationarity but offers more substantial evidence. Conversely, a high *p*-value does not confirm non-stationarity; it indicates the test is inconclusive. If in doubt, alternative tests should be considered like the KPSS. Generally, if direct differencing achieves stationarity and passes the ADF test, it is often preferred for its simplicity. However, there are scenarios where additional transformations may be required for specific modeling purposes.

Before performing ARIMA modeling, it is necessary to check for autocorrelation in the model residuals. In time series analysis, the autocorrelation function (ACF) and partial autocorrelation function (PACF) are essential tools. The ACF measures the correlation between a time series and its past values at different time lags. This helps identify serial correlation, assess the suitability of an AR model, and pinpoint MA terms for potential inclusion in the model. Significant spikes at specific lags in the ACF plot suggest potential MA terms.

The PACF, on the other hand, measures the correlation between a time series and its past values after accounting for the effects of intervening lags. In similarity to the ACF, it helps identify serial correlation, but also helps to assess the suitability of a MA model and determine potential autoregressive (AR) terms for inclusion. Significant spikes at specific lags in the PACF plot suggest potential AR terms.

Interpreting both the ACF and PACF plots involves looking for significant spikes outside the confidence interval, which indicates potential AR or MA terms to include in the model. A gradual decay in the ACF plot might suggest the need for further differencing to achieve stationarity. In contrast, a sharp cutoff after a few lags in the PACF plot suggests the suitability of an MA model.

Together, the ACF and PACF plots provide crucial information for selecting the appropriate order of an ARIMA model. The ARIMA model is defined by three parameters: *p* (number of AR terms), *d* (number of differencing operations), and *q* (number of MA terms). The significant lags in the PACF plot suggest the value for *p*, while the significant lags in the ACF plot guide the value for *q*. If the data has already been differenced, the value of d is usually 1.

For instance, if the ACF plot shows a significant spike at lag 1 and the PACF plot cuts off after lag 2, an ARIMA(2, 1, 1) model might be suitable. Figures 4 and 5 show the results of ACF and PACF in this study, respectively.

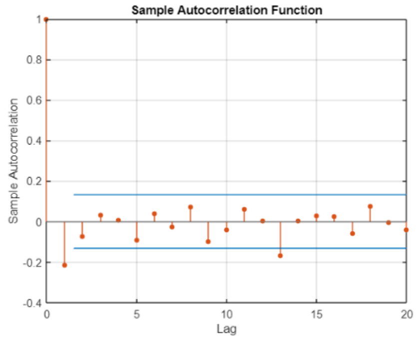

Figure 4. ACF plot

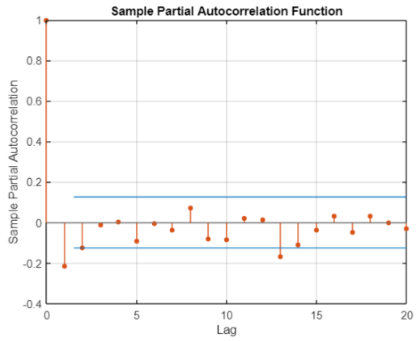

Figure 5. PACF plot

The ACF graph in Figure 4 reveals several key observations. As expected, the autocorrelation at lag 0 is perfect (ACF value of 1), indicating the time series is perfectly correlated with itself. However, for the remaining lags, the autocorrelation values mostly fall within the confidence intervals, suggesting no statistically significant correlation between the time series and its past values. This lack of significant autocorrelation points to a few possible interpretations. First, the time series is already stationary, meaning its statistical properties are consistent over time, which is ideal for many time series analysis techniques. Second, the time series seems to have limited short-term memory, meaning past values have minimal influence on future ones. Lastly, the observed pattern could be consistent with either a white noise process (completely random values) or a low-order MA model, potentially MA(1).

In Figure 5, the PACF graph reveals a very strong positive correlation between the differenced VIX close at a given time and its value at the immediate previous time, even after accounting for other lags. This is evidenced by the significant spike at lag 1, which extends beyond the significance bounds. However, there are no other significant spikes at higher lags, indicating no significant direct correlation between the differenced VIX close and its values at earlier lags once the lag 1 correlation is taken into account.

These observations strongly suggest that an autoregressive model of order 1 (AR(1)) is suitable for the data. This means the current value of the differenced VIX close is primarily determined by its value in the immediately preceding period, and this relationship is linear. Furthermore, the lack of significant spikes at higher lags indicates that there is no need to include additional autoregressive terms in the model.

The initial step of exploratory data analysis in building an ARIMA model suggested that the ARIMA(0, 1, 0) model could be a good fit for predicting the VIX. This type of model is defined by having no autoregressive terms, first-order differencing to achieve data stationarity, and no moving average terms. The ARIMA(0, 1, 0) model was then fitted to the actual stock market data (VIX closing prices for 2013). This process estimated the model's parameters, including the constant term (drift) and variance of the residuals (random fluctuations around the predicted values).

A summary of the model was then produced, revealing a constant term close to zero with a high *p*-value, suggesting that it was not statistically significant. The variance of the residuals was estimated to be 1.0606. The model's log-likelihood (-363.5336) was also calculated, providing a metric for comparing its fit to that of alternative models. Table 1 shows the parameters generated by ARIMA(0, 1, 0).

Table 1. Results of ARIMA(0, 1, 0)

	Constant	Variance
Value	-0.0044094	1.0606
Standard Error	0.067689	0.050316
T-Statistic	-0.065142	21.078
p-value	0.94806	1.263e-98

After fitting, the model was used to calculate residuals (differences between actual and predicted VIX values) and generate predictions. Figure 6 shows the actual data and prediction in 2013.

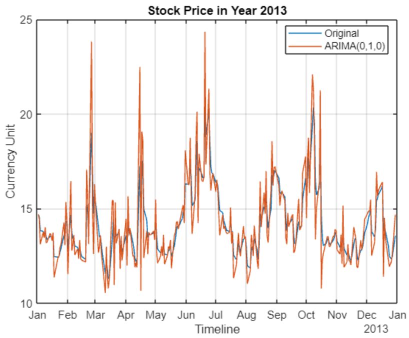

Figure 6. Prediction of VIX values in 2013

In essence, the ARIMA(0, 1, 0) model indicates that the best prediction for tomorrow's VIX value is today's value plus a small random change. The random changes (residuals) are assumed to follow a Gaussian (normal) distribution. This type of model is often referred to as a random walk with drift, implying that the VIX changes randomly from day to day, with a slight tendency to drift up or down over time (although the estimated drift, in this case, was very small).

In a later step, the balance between how well a model fits the data (goodness of fit) and its complexity is examined using the AIC and BIC. The AIC and BIC are calculated based on the model's log-likelihood (a measure of how well the model describes the observed data) and the number of parameters it includes. The AIC uses a simpler penalty term for model complexity compared to the BIC. Therefore, when dealing with larger datasets, the BIC tends to favour simpler models more strongly. Lower AIC or BIC values indicate a better trade-off between model fit and complexity.

In this study, the calculated AIC (731.0672) and BIC (738.1101) are relatively close, suggesting that the penalty for complexity is not a major deciding factor for this specific model. These values would be more meaningful when compared with the AIC and BIC of other models fitted to the same VIX data. The model with the lowest AIC or BIC would generally be preferred as it provides the best balance between fit and parsimony.

Beyond the AIC and BIC, examining the histogram of residuals is crucial for assessing model adequacy. Ideally, the residuals should be normally distributed (a bell-shaped curve), centered around zero, and show no clear patterns. A normal distribution suggests that the model's assumptions are reasonable, while a mean near zero indicates no systematic over- or under-prediction. The absence of patterns in the residuals further supports the model's ability to capture the underlying structure of the data. The generated residual histogram is shown in Figure 7.

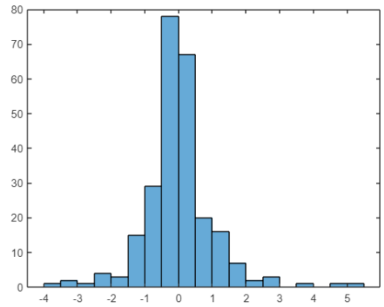

Figure 7. Residual histogram

In this step, the best-fit model is identified using the initial parameters generated by ARIMA (0, 1, 0). Figure 8 shows the pseudo-code for finding the best-fit ARIMA model.

```
// Initialization
model_name = "Arima(0,1,0)"
AIC = previous AIC value  // From your ARIMA(0,1,0) model
BIC = previous BIC value  // From your ARIMA(0,1,0) model
W = 2                // Counter for the models being evaluated

// Loop through different ARIMA model orders (p, d, q)
FOR i = 1 to 2 (p values)
   FOR j = 1 to 2 (d values)
      FOR k = 1 to 2 (q values)
         // Create and estimate a new ARIMA model
         ARIMA_model = create_arima_model(i, j, k)
         Estimate model parameters using StockData_TimeTable_2013.vix_close
         log_likelihood = get_log_likelihood(ARIMA_model)

         // Calculate AIC and BIC
         aic_value = calculate_aic(log_likelihood, 2, 250)
         bic_value = calculate_bic(log_likelihood, 2, 250)

         // Store model name and metrics
         model_name[W] = "Arima(i, j, k)"
         AIC[W] = aic_value
         BIC[W] = bic_value

         // Increment the model counter
         W = W + 1
      END FOR
   END FOR
END FOR

// Create a table for comparison
Create a table (TableComparison)
Fill the table with model names (Name), AIC values (AIC), and BIC values (BIC)
```

Figure 8. Pseudo-code of the proposed method

In this study, the pseudo-code begins by initializing the name and AIC/BIC values of the initial ARIMA(0, 1, 0) model, establishing a baseline for comparison. A counter, '*W*', is also initialized to keep track of the different models being evaluated. Next, the code iterates through various combinations of ARIMA orders (*p*, *d*, *q*), where '*p*' represents the autoregressive order, '*d*' is the differencing order, and '*q*' is the moving average order. These combinations are systematically explored, with the code considering models with up to two terms for each order (e.g., ARIMA(1, 1, 1), ARIMA(2, 1, 2)).

For each combination of (*p*, *d*, *q*), a new ARIMA model is created and fitted to the VIX data. The fitting process is performed quietly, suppressing any display output to improve efficiency. The AIC and BIC values are then calculated for the estimated model. The model name (e.g., 'ARIMA(1, 1, 2)'), along with its AIC and BIC values, are stored in respective arrays. The counter '*W*' is incremented to prepare for the evaluation of the next model in the sequence.

After all model combinations have been evaluated, a comparison table is created. This table summarizes the results, listing the model names alongside their corresponding AIC and BIC values. The ultimate goal of this code is to identify the best-fitting ARIMA model for the VIX data. By comparing the AIC and BIC values in the table, the model with the lowest value (for either AIC or BIC) is generally considered the most preferred, as it strikes the best balance between fit and complexity. Table 2 shows the results for the ARIMA models.

Table 2. Results for ARIMA based on the proposed method

Variable	*AIC*	*BIC*
ARIMA(0, 1, 0)	731.0672	738.1101
ARIMA(1, 1, 1)	715.5234	722.5663
ARIMA(1, 1, 2)	715.3272	722.3701
ARIMA(1, 2, 1)	723.3355	730.3784
ARIMA(1, 2, 2)	715.5734	722.6164
ARIMA(2, 1, 1)	715.3012	722.3441
ARIMA(2, 1, 2)	715.2912	722.3341
ARIMA(2, 2, 1)	720.7831	727.8260
ARIMA(2, 2, 2)	715.4597	722.5027

After comparing the various ARIMA models using AIC and BIC, the ARIMA(2, 1, 2) model was selected as the best fit for the VIX data due to its lower AIC and/or BIC values. This model is characterized by two autoregressive terms, first-order differencing, and two moving average terms. The chosen ARIMA(2, 1, 2) model is then fitted to the VIX data (closing prices for 2013), estimating its parameters while suppressing any display output for efficiency. A random number generator seed is set to ensure reproducibility of the results.

The model is then used to compute the residuals, namely the differences between the actual VIX values and the model's predictions. These residuals are then added back to the original data to generate the model's predicted values for the VIX. Finally, a plot is created to visually compare the actual VIX values and the model's predictions over time (Figure 9). The plot includes the original VIX data alongside the predictions from the ARIMA(2, 1, 2) model. It is properly labeled with a title, axis labels (dates on the x-axis and currency units on the y-axis), with a legend to distinguish the two lines. Gridlines are added for easier interpretation, and the plot is finalized.

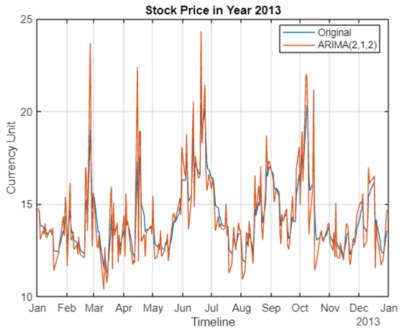

Figure 9. Prediction of VIX values according to the best-fit ARIMA model

This process aims to assess the performance of the best-fitting ARIMA(2, 1, 2) model by visualizing how well its predictions align with the actual VIX values over the course of 2013. The plot generated serves as a crucial tool for evaluating the model's accuracy and its ability to capture the dynamics of the VIX Index. The forecasting results are presented in the following section.

Results and Discussion

The proposed method in this study is designed to forecast the VIX for the first half of 2014 using historical VIX data from 2013. To achieve this, the data are prepared by extracting the relevant time period for analysis. Two different ARIMA models are then utilized, ARIMA(2, 1, 2) and ARIMA(0, 1, 0), to generate forecasts of the VIX's closing values for 2014. The results are then visualized by plotting the original VIX values from 2013, the forecasted values from both ARIMA models for 2014, and the actual VIX values for 2014. This visual comparison allows for an assessment of each model's performance in predicting the VIX.

The ARIMA(2, 1, 2) model suggests that the current VIX depends on its immediate past value, requires differencing twice to achieve stationarity, and incorporates previous forecasting errors. Conversely, the simpler ARIMA(0, 1, 0) model only requires differencing once to achieve stationarity and does not consider past values or forecasting errors.

Ultimately, the proposed method enables the user to visually evaluate the effectiveness of two different ARIMA models in forecasting the VIX and determine which model provides a better fit to the actual data for the first half of 2014. The forecasting results are shown in Table 3, and the plots in Figure 10.

Table 3. Forecasting results for 2014

Date	*vix_open*	*vix_high*	*vix_low*	*vix_close*
2014-01-02	14.3200	14.5900	14	14.2300
2014-01-03	14.0600	14.2200	13.5700	13.7600
2014-01-06	13.4100	14	13.2200	13.5500
2014-01-07	12.3800	13.2800	12.1600	12.9200
2014-01-08	13.0400	13.2400	12.8600	12.8700
2014-01-09	12.8300	13.2600	12.8300	12.8900
2014-01-10	12.6000	12.9000	12.1400	12.1400
2014-01-13	12.1800	13.6500	11.8200	13.2800
2014-01-14	12.8900	12.9000	11.9600	12.2800
2014-01-15	12.1500	12.4000	11.8100	12.2800
2014-01-16	12.3200	12.6600	12.2800	12.5300
2014-01-17	12.3400	12.9300	12.0400	12.4400
2014-01-21	12.6300	13.4200	12.6100	12.8700
2014-01-22	12.5700	13.1200	12.5500	12.8400
2014-01-02	14.3200	14.5900	14	14.2300

Figure 10. VIX forecasting values for 2014

The ARIMA(2, 1, 2) model is a more complex representation of the VIX, incorporating three key components. The autoregressive term indicates that today's VIX value is influenced by yesterday's value, capturing a degree of persistence in volatility. The two differencing terms are applied to transform the non-stationary VIX data into a stationary series, ensuring the model's assumptions are met. Lastly, the moving average term factors in past forecasting errors, allowing the model to adapt to changes in the underlying volatility dynamics.

In contrast, the ARIMA(0, 1, 0) model, also known as a random walk model, offers a simpler perspective. It suggests that the VIX changes randomly over time, and the best prediction for tomorrow's VIX is simply today's value. This model only applies one differencing term to achieve stationarity, making it less sensitive to short-term fluctuations than the ARIMA(2, 1, 2) model.

To enhance the robustness of the study analysis and account for the inherent uncertainty in the VIX, a measure of market volatility, additional tests were conducted using Monte Carlo simulation to generate a range of possible future VIX scenarios. While ARIMA models offer valuable insights through point forecasts, they cannot fully capture the range of potential outcomes. Monte Carlo simulation addresses this limitation by generating multiple possible future paths for the VIX, offering a more comprehensive understanding of risk and the potential range of future volatility.

In the financial world, risk management is paramount. Monte Carlo simulation allows the probabilities associated with different VIX levels to be assessed, providing valuable insights into the likelihood of extreme volatility events. This, in turn, enables more informed and strategic investment decisions.

Even the most sophisticated ARIMA model represents an approximation of reality. Monte Carlo simulation acknowledges this by accounting for model uncertainty. By simulating the fitted ARIMA model with random variations, a distribution of forecasts can be obtained that reflects the model's potential errors.

Monte Carlo simulation generates multiple simulated paths for the VIX, with each path representing a potential future scenario. These paths are visually represented in a plot, illustrating the range of possible outcomes for the VIX in the forecast period.

The value of Monte Carlo simulation lies in its ability to go beyond single-point forecasts, providing a richer picture of potential future volatility. By quantifying the probabilities of different volatility scenarios, risk can be managed effectively. Moreover, it enhances the robustness of forecasting by incorporating uncertainty and potential model errors. The evaluation results based on Monte Carlo simulation are presented in Figure 11.

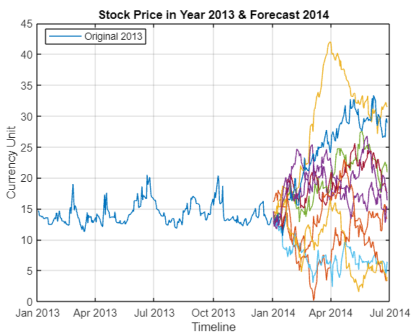

Figure 11. VIX forecasting values for 2014 based on Monte Carlo simulation

The graph in Figure 11 contains two key elements: 1) Original 2013 VIX Data: This line plot represents the actual historical VIX values for 2013. It serves as the basis for fitting the ARIMA model and provides a reference point for evaluating the forecast's accuracy. 2) Monte Carlo Simulations (2014 Forecasts): The multiple lines represent 10 different simulated paths for the VIX in 2014. Each path represents a potential scenario for the potential evolution of the VIX based on the fitted ARIMA(2, 1, 2) model and random variations. The spread of these lines indicates the range of uncertainty in the forecast.

Conclusion

The challenge of accurately forecasting volatile and non-stationary financial time series is a persistent one. While traditional models often fall short, this investigation into the ARIMA model's applicability for predicting the opening value of the CBOE stock market index offers a potential solution. As revealed through the meticulous transformation of the data into a stationary series and the employment of rigorous model selection criteria, the ARIMA(2, 1, 2) model stands out as the most effective for predicting opening values in this specific context. According to the findings, among the evaluated ARIMA models ranging from a random walk (ARIMA(0, 1, 0)) to ARIMA(2, 1, 2), the ARIMA(2, 1, 2) model yields the most accurate forecasts, demonstrating the promising potential for predicting opening values in this context. This finding underscores the importance of tailoring statistical models to the unique characteristics of financial data and offers a promising avenue for further research and refinement in the field of financial forecasting.

In essence, the choice between these models depends on the underlying characteristics of the VIX data and the desired level of complexity. The ARIMA(2, 1, 2) model may be more suitable if the VIX exhibits persistence and responds to past forecasting errors, and ARIMA(0, 1, 0) for its simplicity and robustness to short-term noise if the VIX follows a random walk pattern.

Moreover, AIC and BIC offer a quantitative approach to model selection, guiding researchers toward models that strike a balance between accuracy and complexity. However, it is crucial to remember that these metrics are relative, not absolute, indicators of a model's quality. They are most useful when comparing multiple models to fit the same dataset. Additionally, graphical analysis of residuals complements these metrics, offering visual insights into a model's assumptions and overall performance.

In conclusion, while the ARIMA model identifies underlying patterns in VIX data, Monte Carlo simulation adds a crucial layer of realism by acknowledging the inherent uncertainty in financial markets. This combination equips decision-makers with a more informative and actionable foundation for navigating the complexities of volatility-related choices. Thus, Monte Carlo simulation can generate a range of possible VIX scenarios, making it valuable for risk assessment and decision-making in financial contexts.

References

Adebiyi, A.A., Adewumi, A.O., and Ayo, C.K. (2014). Comparison of ARIMA and Artificial Neural Networks Models for Stock Price Prediction. Journal of Applied Mathematics, 2014(1):614342. https://doi.org/10.1155/2014/614342

Alotaibi, R. (2022). ARIMA Model for Stock Market Prediction. In: Proceedings of the 2022 8th International Conference on Computer Technology Applications (ICCTA' 22). Association for Computing Machinery, New York, NY, USA, p. 1-4. https://doi.org/10.1145/3543712.3543723

Amin, M., Ullah, M.A., and Qasim, M. (2022). Diagnostic techniques for the inverse Gaussian regression model. Communications in Statistics - Theory and Methods, 51(8):2552-64. https://doi.org/10.1080/03610926.2020.1777308

Chandar, S.K. (2021). Hybrid models for intraday stock price forecasting based on artificial neural networks and metaheuristic algorithms. Pattern Recognition Letters, 147:124-133. https://doi.org/10.1016/j.patrec.2021.03.030

Chhajer, P., Shah, M., and Kshirsagar, A. (2022). The applications of artificial neural networks, support vector machines, and long-short term memory for stock market prediction. Decision Analytics Journal, 2:100015. https://doi.org/10.1016/j.dajour.2021.100015

Hickey, G.L., Kontopantelis, E., Takkenberg, J.J.M., and Beyersdorf, F. (2019). Statistical primer checking model assumptions with regression diagnostics. Interactive CardioVascular and Thoracic Surgery, 28(1):1-8. https://doi.org/10.1093/icvts/ivy207

Khanderwal, S. and Mohanty, D. (2021). Stock Price Prediction Using ARIMA Model. The International Journal of Human Resource Management, 2(2):98-107.

Lin, Z. and Liu, D. (2022). Model diagnostics of discrete data regression: a unifying framework using functional residuals. arXiv. 2022 Jul 4;2207.04299.

Liu, D. and Zhang, H. (2018). Residuals and Diagnostics for Ordinal Regression Models: A Surrogate Approach. Journal of the American Statistical Association, 113(522):845-54. https://doi.org/10.1080/01621459.2017.1292915

Mashadihasanli, T. (2022). Stock market price forecasting using the Arima model an application to Istanbul, Turkiye. İktisat Politikası Araştırmaları Dergisi, 9(2):439-454. https://doi.org/10.26650/JEPR1056771

Mohankumari, C., Vishukumar. M., Chillale, N.R. (2019). Analysis of Daily Stock Trend Prediction using Arima Model. International Journal of Mechanical Engineering and Technology, 10(1):1772-92.

Roy, S.S. and Guria, S. (2004). Regression Diagnostics in an Autocorrelated Model. Brazilian Journal of Probability and Statistics, 18(2):103-12.

Senthil Kumar, J.P., Sundar, R., and Ravi, A. (2023). Comparison of stock market prediction performance of ARIMA and RNN-LSTM model: A case study on Indian stock exchange. AIP Conf Proc., 2875(1):020010. https://doi.org/10.1063/5.0154124

Sharma, D.K., Hota, H.S., Brown, K., and Handa, R. (2022). Integration of genetic algorithm with artificial neural network for stock market forecasting. International Journal of System Assurance Engineering and Management, 13(2):828-841. https://doi.org/10.1007/s13198-021-01209-5