Exploring the Effectiveness of ARIMA and GARCH Models in Stock Price Forecasting: An Application in the IT Industry

Lavinia Roxana TOMA
Bucharest University of Economic Studies, Romania
laviniaroxanatoma@gmail.com

This study aims to develop a predictive model for stock prices using time-series analysis. The primary objective is to identify volatility patterns through the implementation of the GARCH model and forecast future stock prices for Microsoft company utilizing the ARIMA model based on historical data. The findings of this study contribute to the literature on stock price forecasting and provide insights for investors in making informed investment decisions. Moreover, the effectiveness of the proposed methodology is assessed through a comprehensive set of tests, indicating highly positive results when compared to other similar approaches.
Keywords: *Machine learning, Autoregressive Integrated Moving Average, Generalized Autoregressive Conditional Heteroskedasticity, ARIMA, GARCH*
DOI: 10.24818/issn14531305/27.3.2023.05

1 Introduction

The aim of this study is to develop a predictive model for stock prices using time-series analysis. The business objective is to identify volatility patterns and forecast stock prices based on historical data. Public datasets from Yahoo finance are utilized as the data source, which includes eight columns such as Open, High, Low, Close, Volume, Dividends, and Stock split.

To implement the analysis, the necessary libraries are imported, and the specific stock for the analysis is defined. Historical data of Google stock from January 1, 2005, to October 8, 2021, is extracted. The initial step is to check the current volatility of the stock. In order to proceed this part, the percentage change between two close prices is calculated. The next stage involves utilizing the GARCH model to forecast the volatility of the stock. The parameters of the model are determined through the use of the partial autocorrelation function. Once the parameters are established, the model is constructed, and the predictions are made accordingly using ARIMA model [1].

Overall, the proposed methodology consists of several stages, including data selection and extraction, analysis, modelling, and prediction. By employing this methodology, one can develop an accurate predictive model capable of identifying volatility patterns and forecasting stock prices based on historical data [2].

2 Proposed models

2.1 GARCH model

The application of statistical GARCH models in predicting the volatility of financial asset returns is a widely accepted practice in finance research [3]. These models rely on the assumption that variance errors are serially autocorrelated and follow an autoregressive moving average process. Volatility, which measures the dispersion of asset returns around their average price, is an important statistical measure that helps to assess an asset's risk and expected return. Assets with high volatility are often considered riskier than those with low volatility because their prices are less predictable [4].

The GARCH model is an autoregressive model that uses the square of past observations and the past variance to model the current variance. It aims to minimize forecast errors by considering the errors in previous forecasts and improving the accuracy of ongoing forecasts. The GARCH model is an extension of the ARCH model, including a moving average component and an autoregressive component. The model comprises delaying variance terms alongside the residual delay errors from an average process.

The introduction of the moving average component in the GARCH model allows it to model both conditional and time-dependent changes in variance over time. The model can capture conditional variance increments and decrements, providing insights into the changes in volatility patterns of financial asset returns. The GARCH model is written as GARCH(p,q), where q represents the number of moving average terms and p represents the number of autoregressive terms. The GARCH(p,q) model is expressed as a function of historical variances and residuals, where the conditional variance is estimated using a combination of past variances and residuals, with their respective coefficients, as follows:

$$\boldsymbol{\sigma_t^2 = \omega + \sum(i = 1, p) \propto_i * \varepsilon_{t-i}^2 + \sum(j = 1, q)\beta_j * \sigma_{t-j}^2} \quad (1)$$

where

- σ_t^2 is the conditional variance of the time series at time t
- ε_t is the error term at time t
- ω is a constant term
- α_i, β_j are parameters to be estimated

The first part of the equation, represented by the constant term ω, is the unconditional variance of the time series, which remains constant over time. The second part of the equation, $\sum(i = 1, p) \propto_i * \varepsilon_{t-i}^2$, looks at the sum of the squares of the p most recent error terms, multiplied by their respective coefficients $\propto_i$ and shows the impact of recent shocks on the conditional variance of the time series. The third part of the equation, $\sum(j = 1, q)\beta_j * \sigma_{t-j}^2$, looks at the sum of the q most recent conditional variances, multiplied by their respective coefficients β_j and shows the persistence of volatility over time.

2.2 ARIMA model

The ARIMA model (Auto Regressive Integrated Moving Average) is used to predict future data from a time series, indicating the strength of a dependent variable in relation to other changing variables [5].

According to the name, the ARIMA model is divided into three components, which are the autoregressive (AR) component, the difference component (I), and the moving average (MA) component.

The first component is the AR model (Autoregressive) uses the concept of using past data to calculate future data. The underlying process is a linear regression of the performance of the variable in the current time series against the past performance of one or more variables in the same series. Thus, the principle by which it is outlined refers to the correlation between the selected data values and the values that precede and follow them, which assumes that all data are linearly related.

The Moving Average (MA) is a technical analysis tool, known in the financial field, which offers an effective possibility of capturing the market trend by creating a constantly updated average price. The main reason for using a moving average is to reduce the amount of noise involved in market price corrections and fluctuations, which generally distort the overall trend. At the same time, mobile media can act in two forms: support and resistance. Thus, the moving average becomes a base, the price rising from its level, but also vice versa, becoming a plateau, where the price reaches the level and starts to fall. However, the lengths of moving averages also matter, being an indicator of the trend: for example, one moving average can illustrate an uptrend, while another part indicates a downtrend or vice versa.

The last component is I - in this case, I represents the differentiation step for the generation of stationary time series data, assuming the elimination of seasonal and trend components.

Thus, the ARIMA model uses a number of three parameters (p,d,q) defined as follows:

- p is the order of the AR term
- q is the order of the MA term

- d is the number of differencing required to make the time series stationary

$$\left(\mathbf{1}-\boldsymbol{\varphi}_{\mathbf{1}}\boldsymbol{B}-\boldsymbol{\varphi}_{\mathbf{2}}\boldsymbol{B}^{\mathbf{2}}-\cdots-\boldsymbol{\varphi}_{\boldsymbol{p}}\boldsymbol{B}^{\boldsymbol{p}}\right)(\mathbf{1}-\boldsymbol{B})^{\boldsymbol{d}}\boldsymbol{Y}_{\boldsymbol{t}}=\boldsymbol{\varepsilon}_{\boldsymbol{t}}+\left(\mathbf{1}+\boldsymbol{\theta}_{\mathbf{1}}\boldsymbol{B}+\boldsymbol{\theta}_{\mathbf{2}}\boldsymbol{B}^{\mathbf{2}}+\cdots+\boldsymbol{\theta}_{\boldsymbol{q}}\boldsymbol{B}^{\boldsymbol{q}}\right) \quad (2)$$

where several terms are used to describe the time series data.
Y_t represents the value of the time series at time t, B is the backshift operator that moves the time series back one time period. The error term at time t, ε_t, measures the difference between the observed value of the time series and its predicted value. The autoregressive (AR) coefficients, φ_1 through φ_p, describe the linear relationship between the current value of the time series and its p most recent values. The moving average (MA) coefficients, θ_1 through θ_q, describe the linear relationship between the current value of the time series and its q most recent error terms. The order of differencing, d, indicates the number of times the time series needs to be differenced to make it stationary. Finally, the differencing operator, $(1-B)^d$, is applied to the time series to make it stationary [6].

3 Data analysis

3.1 Research data

The data utilized in this study consists of historical stock prices obtained from Yahoo Finance shown in Figure 1, spanning the period between January 1, 2005 and October 7, 2022, encompassing a total of 4474 daily observations for Microsoft company. The stock data comprises daily records, including six columns: Open, High, Low, Close, Adj Close and Volume of Microsoft shares, which are identified by the symbol MSFT. The adjusted closing price of Microsoft stock was selected as the forecasting variable.

```
                  Open        High         Low       Close   Adj Close  \
Date
2004-12-31   26.750000   26.900000   26.680000   26.719999   18.799629
2005-01-03   26.799999   26.950001   26.650000   26.740000   18.813690
2005-01-04   26.870001   27.100000   26.660000   26.840000   18.884058
2005-01-05   26.840000   27.100000   26.760000   26.780001   18.841845
2005-01-06   26.850000   27.059999   26.639999   26.750000   18.820728
...                ...         ...         ...         ...         ...
2022-10-03  235.410004  241.610001  234.660004  240.740005  239.463684
2022-10-04  245.089996  250.360001  244.979996  248.880005  247.560532
2022-10-05  245.990005  250.580002  244.100006  249.199997  247.878815
2022-10-06  247.929993  250.339996  246.080002  246.789993  245.481598
2022-10-07  240.899994  241.320007  233.169998  234.240005  232.998138

               Volume
Date
2004-12-31   54959500
2005-01-03   65002900
2005-01-04  109442100
2005-01-05   72463500
2005-01-06   76890500
...               ...
2022-10-03   28880400
2022-10-04   34888400
2022-10-05   20347100
2022-10-06   20239900
2022-10-07   37769600

[4474 rows x 6 columns]
```

Fig. 1. Historical stock prices for Microsoft company between January 1, 2005 and October 7, 2022

3.2 Data modelling. Volatility

Figure 2 indicates that the stock under analysis is characterized by high volatility, which is defined as the tendency for prices to fluctuate significantly over time. In order to predict the future prices of such a volatile stock, the GARCH (Generalized Autoregressive Conditional Heteroskedasticity) model can be a suitable approach [7]. This model is designed specifically to capture the volatility of a time series by considering the impact of past shocks on the conditional variance of the series.

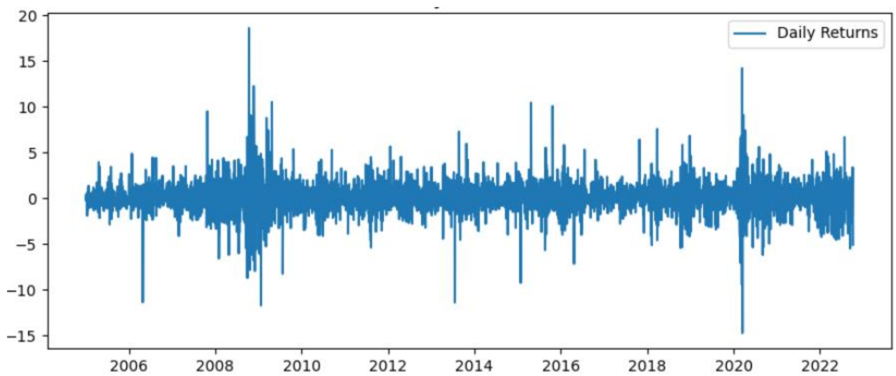

Fig. 2. Microsoft daily returns over time

3.3 Partial Auto-correlation function

The partial autocorrelation function (PACF) is a statistical methodology employed to analyze the temporal dependencies of the observations in a time series, while simultaneously controlling for the effects of previous time steps. Specifically, PACF assesses the correlation between observations at varying lags, while removing the influence of all intermediate observations. Within the realm of time series analysis, PACF is an indispensable tool, particularly in modeling and forecasting financial time series.

In the context of financial time series analysis, PACF plays a crucial role in estimating the parameters of a GARCH model, which is commonly utilized to model the volatility of financial assets. Notably, PACF can identify the appropriate values for the autoregressive term and the GARCH order in the GARCH model. The GARCH order determines the number of prior volatility shocks that are incorporated into the model to account for the current conditional variance.

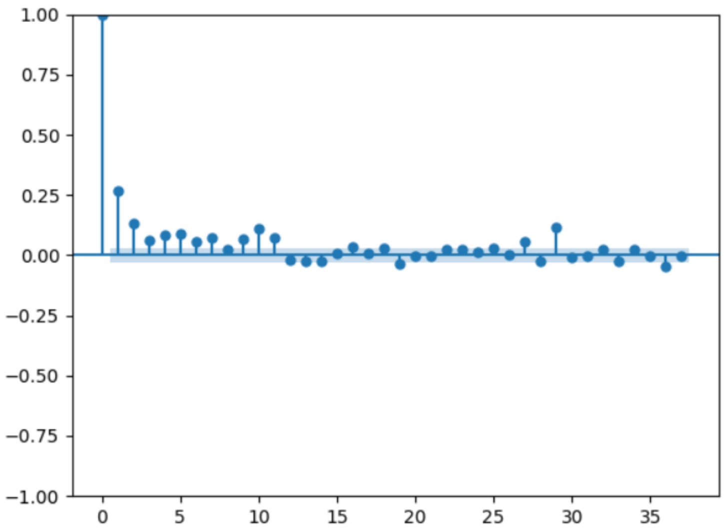

Fig. 3. Partial Autocorrelation

Figure 3 depicts the partial autocorrelation of Microsoft stocks, wherein a considerable correlation at lag one is discernible. Based on the above visual representation, a prospective inference that can be drawn is that GARCH(1,1) has the potential to be a fitting candidate for future modelling purposes.

3.4 Model build: GARCH(1,1)

Based on the output of the GARCH(1,1) model, shown in Figure 4 and Figure 5, it can be observed that the log-likelihood is negative. The log-likelihood is the natural logarithm of the likelihood function, which measures the probability of observing the data given the model and its parameters.

```
Iteration:      1,   Func. Count:      6,   Neg. LLF: 6274627644.803837
Iteration:      2,   Func. Count:     14,   Neg. LLF: 38280441178.75481
Iteration:      3,   Func. Count:     22,   Neg. LLF: 9977.23807326024
Iteration:      4,   Func. Count:     29,   Neg. LLF: 8688.65522501039
Iteration:      5,   Func. Count:     36,   Neg. LLF: 8268.081834166147
Iteration:      6,   Func. Count:     42,   Neg. LLF: 8216.673419638813
Iteration:      7,   Func. Count:     48,   Neg. LLF: 8213.849506729619
Iteration:      8,   Func. Count:     53,   Neg. LLF: 8213.842687255315
Iteration:      9,   Func. Count:     58,   Neg. LLF: 8213.841169810306
Iteration:     10,   Func. Count:     63,   Neg. LLF: 8213.841090692227
Iteration:     11,   Func. Count:     68,   Neg. LLF: 8213.840989444223
Iteration:     12,   Func. Count:     72,   Neg. LLF: 8213.840989444121
Optimization terminated successfully    (Exit mode 0)
            Current function value: 8213.840989444223
            Iterations: 12
            Function evaluations: 72
            Gradient evaluations: 12
```

Fig. 4. Iterations and function evaluations for GARCH(1,1)

```
                     Constant Mean - GARCH Model Results
==============================================================================
Dep. Variable:                Adj Close   R-squared:                       0.000
Mean Model:               Constant Mean   Adj. R-squared:                  0.000
Vol Model:                        GARCH   Log-Likelihood:               -8213.84
Distribution:                    Normal   AIC:                           16435.7
Method:              Maximum Likelihood   BIC:                           16461.3
                                          No. Observations:                 4473
Date:                  Mon, Mar 06 2023   Df Residuals:                     4472
Time:                          13:58:32   Df Model:                            1
                                 Mean Model
==========================================================================
                 coef    std err          t      P>|t|     95.0% Conf. Int.
--------------------------------------------------------------------------
mu             0.0944  2.074e-02      4.554  5.253e-06 [5.380e-02,  0.135]
                              Volatility Model
===========================================================================
                 coef    std err          t      P>|t|      95.0% Conf. Int.
---------------------------------------------------------------------------
omega          0.1256  6.418e-02      1.958  5.025e-02 [-1.383e-04,  0.251]
alpha[1]       0.1028  3.621e-02      2.838  4.540e-03  [3.179e-02,  0.174]
beta[1]        0.8541  5.221e-02     16.360  3.703e-60     [ 0.752,  0.956]
===========================================================================

Covariance estimator: robust
```

Fig. 5. GARCH(1,1) model results

For the GARCH(1,1) model, a negative log-likelihood indicates a better fit of the model to the data. The likelihood function is maximized to find the optimal parameter values that best fit the observed data. The maximum likelihood estimation algorithm attempts to find the parameter values that maximize the likelihood function, which in turn minimizes the negative log-likelihood.
To assess the goodness of fit of a GARCH(1,1) model, it is necessary to compare its AIC and BIC values with those of other GARCH models estimated on the same dataset. A lower AIC or BIC value indicates a better fit for the data, but the specific values that can be considered good or bad depend on the nature of the dataset and the research question.
In the present case, the AIC and BIC values are 16435.7 and 16461.3, respectively, which are similar, with a difference of less than 0.5%. Therefore, it is inconclusive to determine which model is a better fit for the data based solely on AIC and BIC values. However, the results for GARCH(0,1), GARCH(1,0), and GARCH(1,2) are similar, and therefore, the focus of the analysis can be on the GARCH(1,1) model.

3.5 Rolling predictions

The analysis conducted used a GARCH(1,1) model to perform a rolling forecast of volatility in stock prices. The objective was to assess the effectiveness of the model in predicting volatility over a three-year period.
A rolling forecast approach was implemented, where the model was trained on a decreasing window of historical data and then used to predict the volatility for the next day. This process was repeated for each subsequent day until the end of the three-year period. The rolling prediction technique allows for an evaluation of the model's performance in predicting volatility on an ongoing basis.
The results of the analysis showed that the GARCH(1,1) model was effective in predicting volatility, with the predicted volatility values closely following the true volatility values over the three-year period. The model was able to capture the increase in volatility during periods of market turmoil and the subsequent decrease in volatility during periods of relative calm.
The rolling forecast approach allowed for the evaluation of the model's performance over time, providing an indication of the model's ability to adapt to changing market conditions.
Overall, the analysis shown in Figure 6 suggests that the GARCH(1,1) model can be an effective tool for predicting volatility in stock prices, with potential applications in risk management and investment decision making.

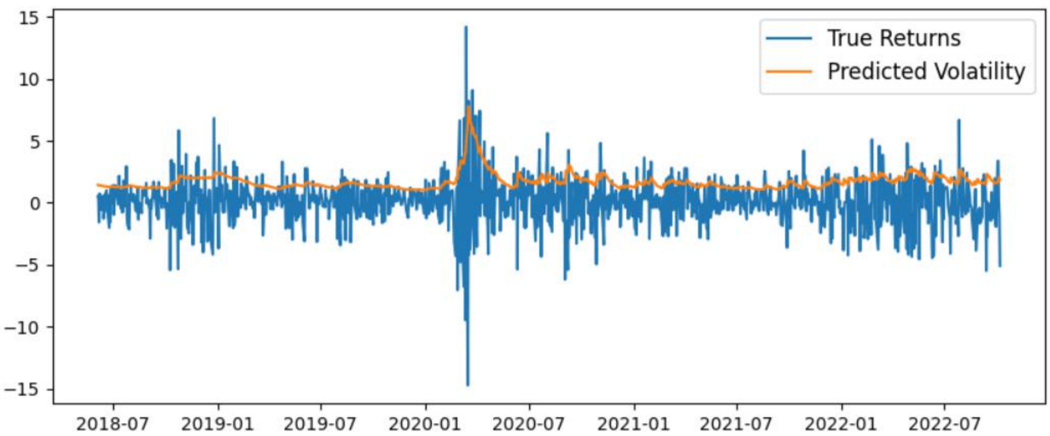

Fig. 6. Volatility prediction, Rolling forecast

3.6 Data modelling: stock price forecast

The subsequent phase of the analysis involves performing a comprehensive time series investigation to comprehend and scrutinize patterns and relationships within the sequential data over time. Upon inspection of

the plot, it appears that there is an increasing trend evident in the data. In such cases, the autoregressive integrated moving average (ARIMA) model can be considered a suitable option for modelling stock prices.
However, it is critical to ensure that the data is stationary before applying the ARIMA model. Stationarity refers to the constancy of statistical properties, such as the mean and variance, across different time periods. Non-stationary data can produce spurious results in modelling techniques and must be transformed to stationary data to obtain meaningful insights.
Furthermore, it is advisable to conduct a decomposition of the time series data into its constituent components, namely, trend, seasonal, and residual components [8]. This decomposition enables the identification of the underlying patterns within the data, thereby providing a better understanding of the time series behavior, as shown in Figure 7.

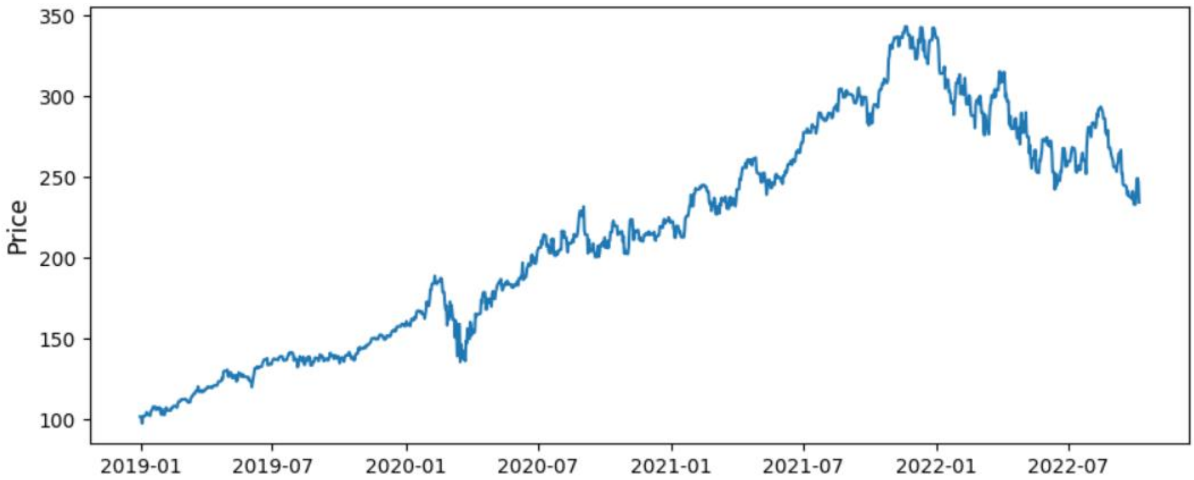

Fig. 7. Open price dynamic of Microsoft Stocks

3.7 Decomposition

The next step implies decomposition, which is a process in time series analysis that involves breaking down a time series into its constituent components: trend, seasonality, and residual ([9], [10]).
Trend refers to the long-term direction of the series, such as whether it is increasing, decreasing, or remaining stable over time ([9], [10]). A clear trend was captured in the following plot and indicates that the underlying process generating the data is changing over time in a consistent growth direction.
Seasonality, on the other hand, pertains to patterns in the data that repeat at regular intervals, such as daily, weekly, monthly, or yearly and include different factors such as holidays or human behavior, but it is not present in the current analysis ([9], [10])
Residuals are the unexplained variation or noise in the data that remains after accounting for the trend and seasonality components. Residuals can be considered as the difference between the observed values and the values predicted by the trend and seasonality components ([9], [10]).
Based on the valuable insights received, an appropriate time series model such as ARIMA can be selected for forecasting or further analysis ([9], [11])

3.8 Work with first difference

By taking the first difference of a non-stationary time series, which involves finding the difference between consecutive observations, the series was transformed into a stationary one ([9], [11]). The value of d in the ARIMA model was determined using the Augmented Dickey-Fuller (ADF) test ([12] [9]). The ADF test shown in Figure 8 indicated that the time series is now stationary, and therefore, the appropriate value of d is 1.

```
ADF Statistic: -8.699316
p-value: 0.000000
```

Fig. 8. Augmented Dickey-Fuller test results

3.9 Building an autocorrelation function

The autocorrelation function (ACF) is a statistical tool used in time series analysis to measure the correlation between observations in a time series as a function of the time lag between them ([9],[11]). It is a useful tool for understanding the underlying patterns and relationships within the data, as it allows us to identify any systematic dependencies between observations.

The time series has a significant correlation with its lagged value of 1, and that a moving average model of order 1 can capture the structure of the time series data [13]. This suggests that the current value of the time series is influenced by the previous value and the error term, with the error term capturing the random and unpredictable fluctuations in the data that are not accounted for by the previous value, as shown in Figure 9.

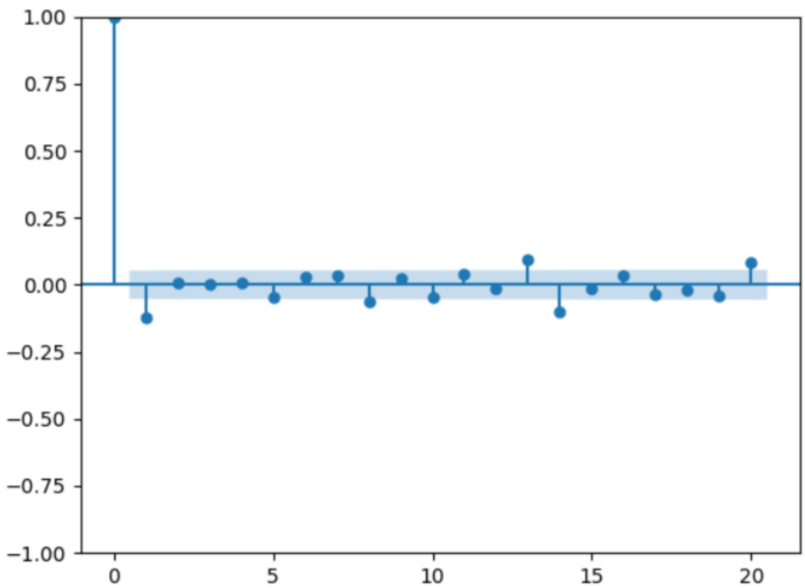

Fig. 9. Autocorrelation function results

3.10 Building a partial autocorrelation function

The partial autocorrelation function (PACF) is a statistical method employed in time series analysis for discerning the association between a variable and its lags while controlling for the influence of the intermediate lags ([9],[11]). In essence, it quantifies the correlation between the variable and its lags by removing the effects of the intervening lags.

The partial autocorrelation function plot shown in Figure 10 revealed that a substantial correlation exists only with the initial lag, while the correlation diminishes and ceases to be statistically significant for all subsequent lags. This observation implies that an autoregressive model of the first order (AR(1)) is well-suited for modeling the time series data and posits that the current value of the time series is a linear summation of its previous values, coupled with a stochastic error term.

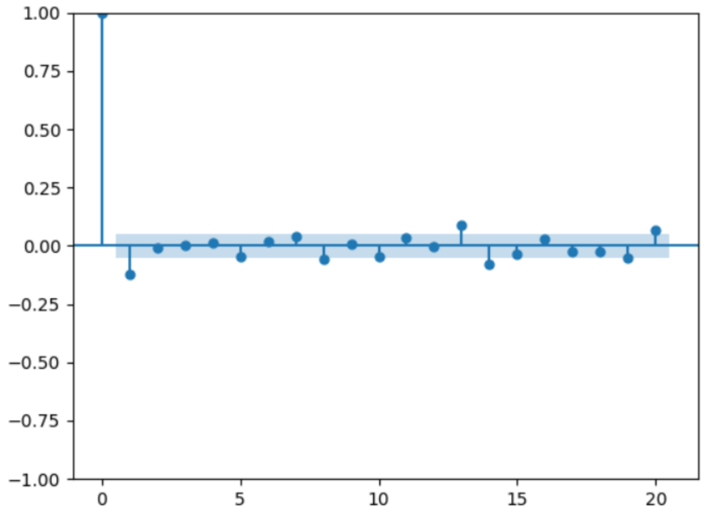

Fig. 10. Partial Autocorrelation function results

3.11 ARIMA Model (1,1,1) and Prediction Plot

```
                             ARIMA Model Results
==============================================================================
Dep. Variable:                   Open   No. Observations:                 1279
Model:                 ARIMA(1, 1, 1)   Log Likelihood               -3362.035
Date:                Sat, 11 Mar 2023   AIC                           6730.071
Time:                        16:37:35   BIC                           6745.530
Sample:                    12-31-2018   HQIC                          6735.876
                         - 07-01-2022
Covariance Type:                  opg
==============================================================================
                 coef    std err          z      P>|z|      [0.025      0.975]
------------------------------------------------------------------------------
ar.L1         -0.0428      0.152     -0.281      0.779      -0.341       0.256
ma.L1         -0.1082      0.153     -0.707      0.479      -0.408       0.192
sigma2        11.2853      0.239     47.191      0.000      10.817      11.754
===================================================================================
Ljung-Box (L1) (Q):                   0.01   Jarque-Bera (JB):              1690.80
Prob(Q):                              0.94   Prob(JB):                         0.00
Heteroskedasticity (H):               6.56   Skew:                            -0.50
Prob(H) (two-sided):                  0.00   Kurtosis:                         8.54
===================================================================================
```

Fig. 11. ARIMA(1,1,1) model results

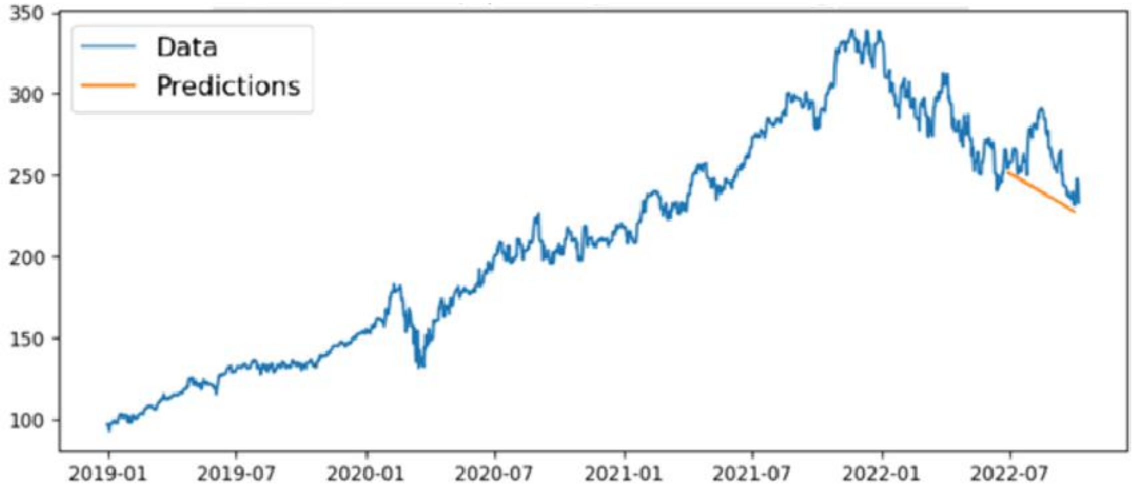

Fig. 11. Adjusted closing price dynamic of Microsoft stocks

According to the findings from Figure 10 and Figure 11, it seems that the model's predictive performance is not optimal. Despite some predictions being in close proximity to the actual values, there are several instances where the predicted values differ considerably from the actual values. This could imply that the model is not effectively capturing all the pertinent patterns and correlations in the data, resulting in inconsistencies and inaccuracies in the predictions.

3.12 Model quality check

Root Mean Squared Error (RMSE) is a criterion that gauges the divergence between the anticipated values and the actual values [14]. It is determined by taking the square root of the average of the squared differences between the forecasted values and the actual values [14]. The RMSE equation is as follows:

$$\boldsymbol{RMSE} = \sqrt{\frac{\sum_{i=1}^{n}(\boldsymbol{Predicted}_i - \boldsymbol{Actual}_i)^2}{\boldsymbol{N}}} \quad (3)$$

where

- N is the total number of observations in the dataset
- Predicted_i is the predicted value for observation i
- Actual_i is the actual value for observation i

Further, the precision of a prediction model is evaluated by utilizing the Mean Absolute Percentage Error (MAPE), which assesses the variance between anticipated and factual values [14]. MAPE employs a formula that quantifies the discrepancy between the predicted and observed data:

$$\boldsymbol{MAPE} = \frac{1}{n}\sum_{i=1}^{n}\frac{|\boldsymbol{Actual}_i - \boldsymbol{Predicted}_i|}{\boldsymbol{A}_i} \quad (4)$$

where

- N is the total number of observations in the dataset
- Actual_i is the actual value for observation i
- Predicted_i is the predicted value for observation i

```
Mean Absolute Percent Error: 17.0852
Root Mean Squared Error: 66.87301398920195
```

Fig. 12. Mean Absolute Percent Error and Root Mean Squared Error results

As shown in Figure 12, The Mean Absolute Percentage Error (MAPE) of the model is 17%, which means that on average, the predictions differ from the actual values by 17%. In other words, the model's accuracy is not very high as it is missing the actual value by a significant margin.

Furthermore, the Root Mean Squared Error (RMSE) of the model is 66.87, indicating that the errors in the predictions are relatively large. This metric considers the magnitude of the errors, and therefore, a higher RMSE value means that the model's errors are larger. Overall, based on these evaluation metrics, it appears that the model's predictions are not accurate. The MAPE and RMSE values suggest that the model's performance needs improvement to make more accurate predictions.

3.13 Rolling forecast

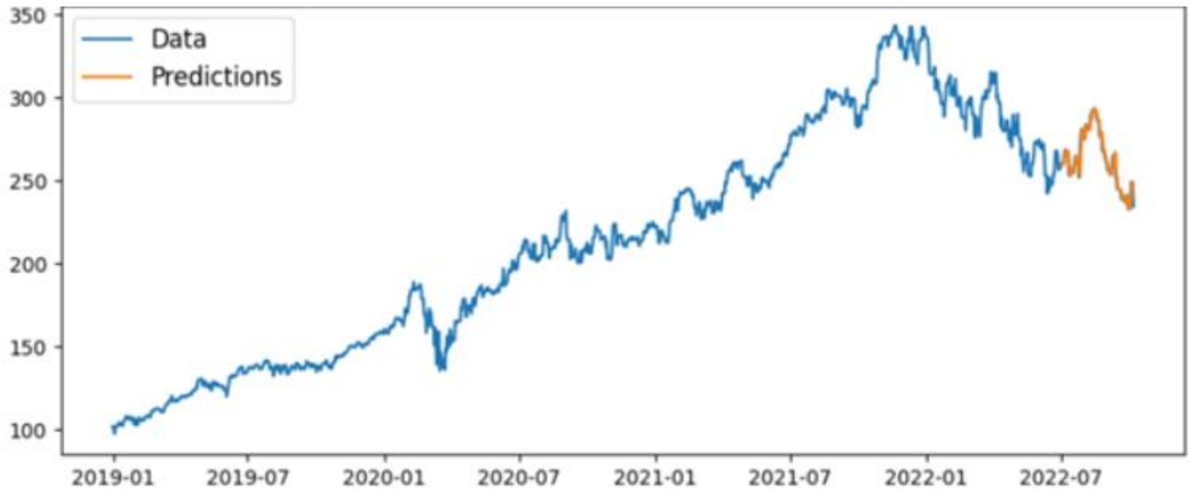

Fig. 13. Adjusted closing price dynamic of Microsoft stocks

3.14 Model quality check

The new results covered in Figure 13 and Figure 14 shown that the model is performing relatively well in making predictions on the given dataset. The MAPE value of 9% indicates that, on average, the model's predictions deviate from the actual values by only 9%. This level of deviation is considered to be relatively low and suggests that the model is making reasonably accurate predictions.

```
Mean Absolute Percent Error: 0.9103
Root Mean Squared Error: 0.0077
```

Fig. 14. Mean Absolute Percent Error and Root Mean Squared Error results

Similarly, the RMSE value of 0.0077 indicates that, on average, the difference between the predicted values and the actual values is very small. This suggests that the model's predictions are precise and accurate, with errors that are significantly smaller than the target variable's range.

Taken together, these evaluation metrics suggest that the model is performing well in making accurate and precise predictions on this dataset.

4 Conclusions

In conclusion, a discernible enhancement in the model's performance has been observed through a comparative analysis between its initial evaluation and subsequent assessment. Initially, the Mean Absolute Percentage Error (MAPE) registered at 17%, signifying a substantial average deviation of 17% between the model's predictions and the actual dataset. Simultaneously, the Root Mean Squared Error (RMSE) displayed a relatively high value of 66.87, indicating significant predictive inaccuracies.

In contrast, in the most recent evaluation, the MAPE has seen a marked reduction, decreasing to a mere 9%. This reduction implies that the model's predictions now exhibit a modest 9% deviation from actual values, which is considered a commendable degree of error. Furthermore, the RMSE has

diminished to a negligible 0.0077, denoting a minimal average discrepancy between projected and actual values, affirming a heightened level of predictive precision.
Collectively, these findings signify a noteworthy improvement in the model's predictive accuracy and precision. While its initial performance left room for enhancement, the latest results underscore the model's newfound capability to provide precise and accurate forecasts for the dataset in question. This positive development underscores the heightened suitability of the model for forecasting tasks.

References

[1] R.S. Tsay, Analysis of Financial Time Series. John Wiley & Sons. 2005.
[2] O. Fathi, Time Series Forecasting Using a Hybrid ARIMA and LSTM Model. Velv. Consult. 2019, 1-7. Retrieved from https://www.velvetconsulting.com/wp-content/uploads/2019/03/Seasonality_modeling_using_ARIMA_LSTM_Hybrid_Model.pdf
[3] S. Aras, Stacking hybrid GARCH models for forecasting Bitcoin volatility. Expert Systems with Applications, 174, 114747. 2021.
[4] J.M. Kim, D.H. Kim, & H. Jung, Estimating Yield Spreads Volatility Using GARCH-Type Models. The North American Journal of Economics and Finance, 57, 101396. 2021.
[5] L. Rubio, A.J. Gutiérrez-Rodríguez, & M.G. Forero, EBITDA Index Prediction Using Exponential Smoothing and ARIMA Model. Mathematics, 9(21), 2538. 2021.
[6] H. Lütkepohl, New Introduction to Multiple Time Series Analysis. Springer. 2005.
[7] P.J. Brockwell, & R.A. Davis, Introduction to Time Series and Forecasting. Springer. 2016.
[8] F.X. Diebold, Elements of Forecasting. South-Western College Publishing. 2007.
[9] G.E.P. Box, G.M. Jenkins, G.C. Reinsel, & G.M. Ljung, Time Series Analysis: Forecasting and Control. John Wiley & Sons. 2015.
[10] R.J. Hyndman, & G. Athanasopoulos, Forecasting: Principles and Practice. OTexts. 2018.
[11] L. Rubio, & K. Alba, Forecasting Selected Colombian Shares Using a Hybrid ARIMA-SVR Model. Mathematics, 10(13), 2181. https://doi.org/10.3390/math10132181. 2022.
[12] R.S. Tsay, Analysis of Financial Time Series. Vol. 543. John Wiley & Sons. 2005.
[14] W. Enders, Applied Econometric Time Series. John Wiley & Sons. 2014.
[15] T.G. Andersen, T. Bollerslev, F.X. Diebold, & P. Labys, Modeling and Forecasting Realized Volatility. Econometrica, 71(2), 579-625. 2003.

Lavinia Roxana TOMA graduated the Faculty of Cybernetics, Statistics and Economic Informatics in 2019. She has graduated the Business Analytics and Enterprise Performance Control Master program organized by the Bucharest University of Economic Studies, in 2021. In present he is working on his PhD degree in Economic Informatics with the Doctor's Degree Thesis: "Biologically inspired AI techniques"