Predicting Stock Returns and Optimizing Portfolios: An Analysis of 15 Technology Companies Based on ARIMA, GARCH and Monte Carlo Simulation

Yifan Liu

International Economics and Trade, Southwestern University of Finance and Economics, Chengdu, 611139, China

42219166@smail.swufe.edu.cn

Abstract. Traditionally, portfolio management focused on using historical averages and qualitative assessments of market trends. However, with more data available and improvements in computational power, more complex methods have been developed. This study aims to analyze the stock prices of 15 technology companies spanning from 2016 to 2023, creating a training set to calculate returns. Utilizing ARIMA and GARCH models, this study predicted stock returns for the subsequent years, 2023-2024. These models are used to capture the trends and volatility in the stock prices, providing valuable insights into potential future performance. To further enhance the analysis, this study employed Monte Carlo simulation to evaluate various portfolio combinations. This approach enabled to assess the risk and return characteristics of different investment strategies, ultimately identifying the most representative strategy for the given dataset. The findings contribute to the field of financial forecasting and portfolio optimization, highlighting the potential of advanced statistical techniques in predicting stock returns and informing investment decisions.

Keywords: Returns Prediction, Portfolio Optimization, ARIMA and GARCH model, Monte Carlo simulation.

1 Introduction

Portfolio design, a cornerstone in the intricate landscape of financial markets, aims to maximize returns while mitigating risk. Traditionally, portfolio management relied heavily on historical averages and qualitative assessments of market trends. However, with the increasing availability of data and advancements in computational capabilities, more complicated methods have emerged.

The field of portfolio management has a long history, with various theoretical frameworks guiding investment decisions. The Capital Asset Pricing Model (CAPM) and the Black-Litterman model, for example, have provided insights into asset pricing and the incorporation of investor views into market expectations [1, 2]. In recent years, quantitative finance has introduced a new pattern, where mathematical models

J. Lu (ed.), *Proceedings of the International Workshop on Navigating the Digital Business Frontier for Sustainable Financial Innovation (ICDEBA 2024)*, Advances in Economics, Business and Management Research 315,
https://doi.org/10.2991/978-94-6463-652-9_66

and statistical techniques play a key role in portfolio construction and risk management.

The Autoregressive Integrated Moving Average (ARIMA) model and its extensions have been extensively used in finance for forecasting stock prices and returns [3]. These models capture the temporal dynamics of financial data, allowing for more accurate predictions of future behavior. Similarly, the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model has become a staple in modeling the volatility of financial time series, a crucial aspect of risk management [4]. Monte Carlo simulation, another powerful tool in quantitative finance, enables the generation of numerous scenarios for portfolio performance based on assumed distributions of asset returns [5]. This technique allows for the optimization of portfolios under different risk-return objectives, providing investors with a comprehensive understanding of potential outcomes.

While these models and techniques have individually demonstrated their utility in portfolio management, there remains a notable gap in research regarding their integrated application. Specifically, the combination of time series analysis (ARIMA-GARCH) for return and volatility forecasting with Monte Carlo portfolio optimization has not been extensively explored. This gap limits the investors and portfolio managers to fully use the predictive power of these models in a unified framework. Most existing research in this field tends to focus on specific aspects of portfolio management, such as return forecasting or risk assessment, rather than providing a holistic approach that integrates both [6, 7]. This fragmentation can lead to suboptimal portfolio designs, as the interplay between return prediction and risk management is not fully accounted for.

To fill this research gap, this study aims to design portfolios with integrated time series analysis for return and volatility forecasting with Monte Carlo portfolio optimization. By doing so, this study seeks to provide a comprehensive framework that not only predicts future returns but also optimizes portfolios based on these predictions, taking into account the inherent volatility of financial assets. The proposed approach involves several key steps. First, this study forecasted stock returns and performed tests to find the best ARIMA model. Next, this study will employ the ARIMA-GARCH model to estimate the volatility of these returns, a crucial input for risk management. Finally, Monte Carlo simulation will be employed to optimize portfolios based on the forecasted returns and volatilities, exploring different risk-return trade-offs.

2 Methods

2.1 Dataset Preparation

This study analyzes 15 technology stocks as mentioned before, including companies such as Apple (AAPL), Microsoft (MSFT), and Google (GOOGLE). The data comes from Yahoo Finance and covers the period from January 2016 to July 2023. The study chooses historical data to predict the upcoming year. By using several years of historical data, these models can be trained effectively, and the risk of overfitting can be reduced. Finally, historical data includes different phases of the economic and business cycle, such as growth and recession, which can explain how assets will perform under different scenarios. This knowledge is critical to accurately predicting what might happen next in the market.

The daily logarithmic return was calculated for each stock in the dataset. Logarithmic returns are preferred in modeling because they simplify the mathematics involved in continuous compounding calculations and provide a more accurate picture of investment performance over time. By visualizing the daily logarithmic yield of each stock, this study gains insight into the volatility and overall performance of the stocks.

2.2 Stock Return Prediction

The ARIMA Model. This study first plots the ACF and PACF graphs of the fifteen stocks then defines the ADF function and performs an ADF test on the logarithmic return of each stock [8, 9]. Through the results of both, this study chooses to use the ARIMA model.

Then this study seeks the best orders of the ARIMA model based on the AIC principle. Akaike Information Criterion (AIC) is a statistical criterion for model selection that provides an assessment of the relative goodness of fit of different models. AIC principle is the basic principle of seeking a balance between minimizing information loss and model complexity [10]. According to it, given a set of candidate models, the lower the AIC value, the better the relative goodness of fit of the model. Therefore, this study defines a 'find_best_arima_model' function by programming that traverses sixteen combinations (for p,q in the range of 0-4) to find the one with the smallest AIC value.

The ARIMA-GARCH Model. After fitting the ARIMA model, there's one fact that the ARIMA model can only focus on the mean part of the modeling time series but cannot consider the volatility of the series. The GARCH model captures the volatility of time series, compensating for the shortcomings of the ARIMA model. Therefore, this study chooses to combine the ARIMA model with the GARCH model.

The Prediction and the Magnitude of Error Testing. The prediction is made by fitting the ARIMA model and the GARCH model together. The residual error of the ARIMA model is passed into the GARCH model as a parameter for fitting. Finally, RMSE and MAPE values were calculated to evaluate the predicted results.

The Root Mean Square Error (RMSE) involves calculating the square root of the average squared differences between the actual and predicted values [11]. However, because the stock returns are very small, the RMSE value may not be entirely convincing. The Mean Absolute Percentage Error (MAPE) assesses the accuracy of a model's predictions by calculating the average absolute percentage difference [12]. It helps to understand the error concerning the actual values, making it more suitable when the importance of errors depends on the size of the actual values.

According to the results, most stocks have a high MAPE value, suggesting that the predicted results may not be accurate. Therefore, this study intends to assess the error stability using the baseline method and further validate the prediction's reliability.

Forecast Stability Test. Firstly, use the actual and predicted daily return rates of each stock from 2022 to 2023, and that of data from 2023 to 2024, as the training and test sets separately. Second, calculate residuals for both the training set and test set.

With the "Normality of residuals" assumptions, this study assumes residuals of the training set are normally distributed [13]. By calculating the mean and standard deviation of the training set using the relevant formula, this study tries to figure out how many data points of the test set fall within the mean±3 standard deviations of the training set. To achieve this, "temp" is used as a variable to represent the quantity. According to the Empirical rule, if 99.7% of data points fall within mean±3 standard deviations, then it is normally distributed, which refers to stable forecast error[14]. From the results, all test sets' residual data points fall less than 99.7% of the range, which implies unstable errors.

2.3 Portfolio Optimization

With portfolio optimization, this study will select the best combination of financial assets to maximize returns for a given level of risk or to minimize risk for a given level of expected returns. First, set the initial value and give each stock the same amount of investment and the same weight, whose weight is determined by the number of stocks. For example, if the total amount of investment held by the investor is 1, which means 100%, if the number of stocks is 5, each stock will get an initial weight of 20%, which is the total amount divided by the number of stocks.

Here, this study chose backtesting to evaluate the performance of the assets given the period from July 2023 to July 2024. The cumulative returns start at 1, indicating that the initial investment amount is normalized to 1. This makes it easier to track the percentage increase or decrease over time. Then this study used Monte Carlo to calculate the efficient frontier. The first step is to generate a large number of random portfolios. The number chosen is 10 million. For each portfolio, the weights are assigned randomly to the assets. And the weights must sum to 100%, meaning the entire portfolio is allocated among the selected assets. Second, calculate portfolio returns for each generated portfolio by taking the weighted sum of the expected returns of the individual assets based on the portfolio weights. Next, calculate the risk (variance or standard deviation) of each portfolio and plot each portfolio on a graph with risk (volatility) on the x-axis and return on the y-axis. Finally, this study backtested the portfolio with the highest Sharpe ratio and the portfolio with the lowest volatility. Through calculations, it leads to the conclusion of the optimal portfolios.

3 Results and Discussion

3.1 The Results of the Best Orders

From the results, in particular, for AMZN, the best order obtained is (0, 0, 0), which means that it is a simple average model itself, not including any autoregressive (AR) or moving average (MA) components. Such a model only predicts a constant value, resulting in a horizontal line without fluctuations.

So this study back to analyzing the PACF plot of AMZN shown in Fig. 1, it is clear that there is a significant partial autocorrelation at lag 1, but no significant partial autocorrelation at other lags, which indicates that using a lagged autoregressive term (AR (1)) may be a suitable choice. At the same time, analyzing the ACF plot of AMZN, it is also known that there is a significant autocorrelation at lag 1, but no significant autocorrelation at other lags, which indicates that using a lagged moving average term (MA (1)) may be more appropriate. Therefore, (1, 0, 1) was chosen as the best order of AMZN.

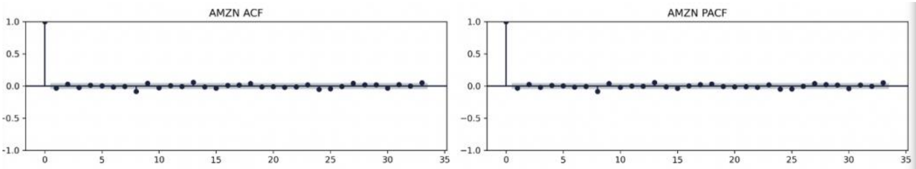

Fig. 1. The ACF and PACF plots of AMZN (Photo/Picture credit: Original).

3.2 The Prediction of the ARIMA-GARCH Model

Table 1. Mean Model for AAPL:(1,1)

	coef	std err	t
mu	1.0444e-03	1.7973-05	58.134

Table 1 is the Mean Model derived for APPL. This study selected the commonly used GARCH model with alpha and beta values of 1 [15]. From the results, the positive and highly significant constant mean (μ) indicates a small positive average return.

Table 2. Volatility Model for AAPL:(1,1)

	coef	std err	t
omega	7.0337e-06	7.571e-11	9.291e+04
alpha[1]	0.1000	2.710e-02	3.689
beta[1]	0.8800	2.350e-02	37.441

Regarding the Volatility Model shown in Table 2, the very small but highly significant ω suggests low baseline volatility. Overall, this output indicates that the GARCH (1,1) model fits the data well, with significant coefficients showing that past volatility and past shocks have a significant influence on current volatility. The high $\beta_1$value indicates high volatility persistence.

Table 3. Forecasting returns for AAPL

time	prediction
2023-07-03	0.005245

2023-07-04	-0.000622
• • •	• • •
2024-07-09	0.031646
2024-07-10	-0.025499

Table 3 shows the forecasting returns for the full year from July 2023 to July 2024, which is the projected results.

3.3 The Results of the Monte Carlo Simulation

Through Monte Carlo simulation, 10 million portfolio weights were generated. The leftmost boundary of the shape if the efficient frontier, representing optimal portfolios when considering about sharp ratio and volatility (see Fig. 2).

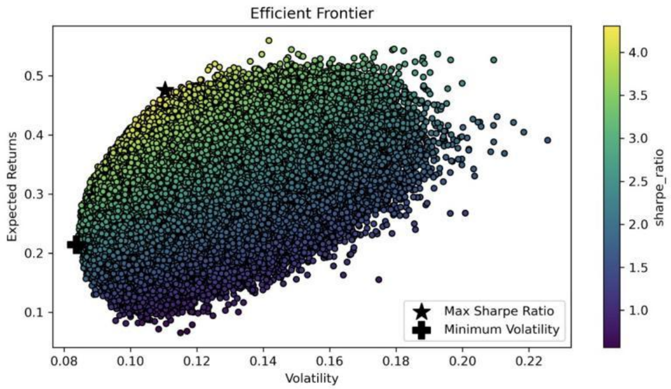

Fig. 2. The Monte Carlo simulation (Photo/Picture credit: Original).

3.4 The Backtesting and Results

Largest Sharp Ratio Outcome. Backtesting the portfolio with the highest Sharpe ratio showed a cumulative return of 59.83%, a Sharpe ratio of 4.24, and an annual volatility of 10.87%. Remarkably, the Sharpe ratio is decreasing constantly shown in Fig. 3, which may imply that the portfolio is becoming less effective.

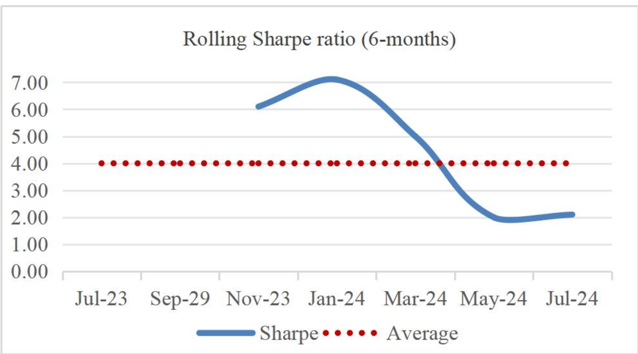

Fig. 3. Rolling sharpe ratio (6-months)-1 (Photo/Picture credit: Original).

Lowest Volatility. Through backtesting the portfolio with the lowest volatility, it can be found that the annual fluctuations slightly dropped to 8.239%, decreasing by approximately 2.5%, but the cumulative return and Sharpe ratio both decreased by half. It can be inferred that there is a marginal effect: Sacrificing returns does not achieve the expected low volatility. Also, the Sharpe ratio is continuously decreasing shown in Fig. 4, supporting the conclusion before.

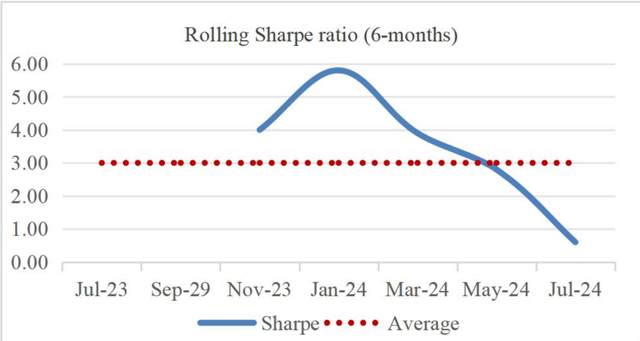

Fig. 4. Rolling sharpe ratio (6-months)-2 (Photo/Picture credit: Original).

Through calculations, the S&P 500's index has increased by 30%. It is 8% higher than the returns of the lowest volatility portfolios, while nearly just half of that of the largest Sharpe ratio. This is due to the marginal effect mentioned beforehand. This leads to the conclusion that the optimal portfolios have relatively considerable returns.

4 Conclusion

In conclusion, this study successfully fulfilled its objective of generating a comprehensive training set through the meticulous analysis of stock prices for 15 technology companies spanning the years 2016 to 2023, coupled with the calculation of their respective returns. Utilizing advanced statistical models, the ARIMA and GARCH model, it was able to predict stock returns for the subsequent years, 2023-2024, with a focus on enhancing forecast accuracy. Furthermore, by implementing the Monte Carlo simulation, this study evaluated a multitude of potential portfolio combinations, ultimately identifying the most representative investment strategy.

While the findings contribute to financial forecasting and portfolio optimization, several limitations must be recognized. Firstly, the study assumes stationarity and normally distributed returns in the models, which may not always align with real-world market conditions. Secondly, the Monte Carlo simulation's accuracy depends on underlying assumptions and may not fully capture market complexity. To address these limitations, Future research can build on this foundation to deepen the understanding of financial markets and enhance investment decision-making processes.

References

1. Fama, E. F., & French, K. R.: The capital asset pricing model: Theory and evidence. Journal of Economic Perspectives 18(3), 25-46 (2004).
2. Idzorek, T.: A step-by-step guide to the Black-Litterman model: Incorporating user-specified confidence levels. In: Forecasting Expected Returns in the Financial Markets, pp. 17-38. Academic Press (2007).
3. Ling, S., & Li, W. K.: On fractionally integrated autoregressive moving-average time series models with conditional heteroscedasticity. Journal of the American Statistical Association 92(439), 1184-1194 (1997).
4. Francq, C., & Zakoian, J. M.: GARCH Models: Structure, Statistical Inference and Financial Applications. John Wiley & Sons (2019).
5. Harrison, R. L.: Introduction to Monte Carlo simulation. In: AIP Conference Proceedings 1204, p. 17. NIH Public Access (2010).
6. Yu, J. R., Chiou, W. J. P., Lee, W. Y., & Lin, S. J.: Portfolio models with return forecasting and transaction costs. International Review of Economics & Finance 66, 118-130 (2020).
7. Sunchalin, A. M., Kochkarov, R. A., Levchenko, K. G., Kochkarov, A. A., & Ivanyuk, V. A.: Methods of risk management in portfolio theory. Revista Espacios 40(16) (2019).
8. Mondal, P., Shit, L., & Goswami, S.: Study of effectiveness of time series modeling (ARIMA) in forecasting stock prices. International Journal of Computer Science, Engineering and Applications 4(2), 13 (2014).
9. Mushtaq, R.: Augmented Dickey-Fuller test (2011).
10. Akaike, H.: Factor analysis and AIC. Psychometrika 52, 317-332 (1987).
11. Hodson, T. O.: Root mean square error (RMSE) or mean absolute error (MAE): When to use them or not. Geoscientific Model Development Discussions (2022).
12. Goodwin, P., & Lawton, R.: On the asymmetry of the symmetric MAPE. International Journal of Forecasting 15(4), 405-408 (1999).
13. Schmidt, A. F., & Finan, C.: Linear regression and the normality assumption. Journal of Clinical Epidemiology 98, 146-151 (2018).
14. Wooditch, A., et al.: The normal distribution and single-sample significance tests. In: A Beginner's Guide to Statistics for Criminology and Criminal Justice Using R, pp. 155-168 (2021).
15. Jafari, G. R., Bahraminasab, A., & Norouzzadeh, P.: Why does the Standard GARCH (1, 1) model work well? International Journal of Modern Physics C 18(07), 1223-1230 (2007).