Evaluating Random Forest and LSTM Architectures in Volatility Forecasting for the Magnificent Seven

Bo Wen

Beijing-Dublin International College at BJUT, Beijing University of Technology, Beijing,100124, China
krooz147269@outlook.com

Abstract. Volatility prediction is still required for the whole of financial risk control, and due to its instability, changes in stock prices can be quite sudden. The standard linear model cannot explain the pronounced fluctuations in the concentrated stock market. This paper examines the prediction accuracy of machine learning and deep learning methods for the stock prices of the Magnificent Seven indices based on these seven indices. A large amount of data, such as macroeconomic and liquidity indices, is to be used with both the LSTM model and the Random Forest model for comparison; according to the experiment, the LSTM model has shown good performance in capturing the time-series characteristics of financial data. The decrease in the out-of-sample Root Mean Square Error (RMSE) of the proposed model compared with that of Random Forest is 24.3%. To know how big the predicted volatility Figures are in reality, they are then used in a dynamic allocation plan to determine the exposure coefficient of volatility. Backtesting shows that, in the event of high market volatility at the beginning of 2026, there will be a 6.8% increase in the Sharpe ratio and a 15.2% decrease in the maximum drawdown. In short, it has been determined that the deep learning model can improve the accuracy of volatility forecasting and offer some support for portfolio risk management in the concentrated growth equity market.

Keywords: Volatility Forecasting, Deep Learning, Statistical Safety Margin, Non-Linear Regime Switches, Capital Preservation

1 Introduction

Fluctuation forecasting is needed for the purpose of risk control, asset price changes, etc., in the construction of financial engineering. The stocks of the Magnificent Seven have attracted much attention because they constitute a relatively large proportion of the weights in the major US stock indices and affect the overall market. Therefore, changes in the number of these stocks will have a relatively large effect on the overall market. As the concentration of the market has risen, forecasting the fluctuations in these companies has

© The Author(s) 2026
L. C. Yoong (ed.), *Proceedings of the 2026 2nd International Conference on Financial Innovation and Marketing Management (FIMM 2026)*, Advances in Economics, Business and Management Research 408,
https://doi.org/10.2991/978-94-6239-801-6_149

become more and more necessary. Recently, it has been found that machine learning methods can capture the connections of data more easily and are thus useful for improving the forecast accuracy of older statistical models [1].

Research on volatility forecasting has evolved considerably over time. Early studies mainly relied on econometric approaches such as ARCH and GARCH models to describe changes in market volatility. Although these methods remain widely used, they often face challenges when market behaviour becomes highly non-linear or changes rapidly. Lim and Zohren review a wide range of deep learning applications in time-series forecasting and argue that neural networks are well-suited to modelling complex temporal patterns [2]. Bevilacqua et al. also find that deep learning methods can capture volatility movements across different time horizons and often achieve better forecasting accuracy than conventional benchmark models [3].

Another area that has received growing attention is feature selection and data construction. Rather than depending only on historical price information, many recent studies incorporate macroeconomic variables and market liquidity measures into forecasting models. Jiang et al. show that macroeconomic indicators provide additional information about broader market conditions and can improve forecasting accuracy [4]. In a study focusing on technology stocks, Sharma and Mehta report that recurrent neural networks perform particularly well during periods of rapid market adjustment, when investor sentiment and trading activity change quickly [5].

This study focuses on volatility forecasting for the Magnificent Seven equities. A dataset containing macroeconomic and liquidity indicators is used to compare the forecasting performance of Random Forest and Long Short-Term Memory (LSTM) models. Forecast accuracy is evaluated using out-of-sample measures. The resulting volatility forecasts are then incorporated into a dynamic allocation framework to examine their potential application in portfolio management.

2 Methodology

2.1 Data Source and Feature Engineering

This paper takes a high-dimensional dataset of the stock prices of the Magnificent Seven. In order to obtain some factors that may affect the fluctuations in the stock market, a feature set of more than 40 exogenous variables was built. The set of data is the Consumer Price Index (CPI), Treasury yield spreads and CBOE Volatility Index (VIX). Based on the above background information, the open-source repository provided by Kanchana1990 was selected [6]. The data for this study come from official public places, and they have been checked.

The first is to conduct some preprocessing of the data. First, to make the time-series data stationary, logarithmic differencing was applied. Next, all the variables were normalised to reduce the impact of different measurement scales. Then, the processed data were converted

into three-dimensional tensors by means of a sliding window and could be employed by sequential learning models. The above ways can help to keep the consistency of the model and ensure the reproducibility of the experimental results.

2.2 Specification of Key Variables

The forecasting and allocation framework is built around three key variables. Realised Volatility (RV) serves as the dependent variable and is measured as the annualised standard deviation of log-returns for the Magnificent Seven portfolio. The models are trained to predict future RV values. The Macro-Systemic Risk Factor (M) acts as the explanatory variable and is represented by Treasury yield spreads and liquidity indicators, which reflect broader economic and market conditions. The Dynamic Exposure Coefficient (ω) is defined as the inverse of predicted volatility and is used to adjust portfolio allocation according to the expected level of market risk.

2.3 Model Architectures

The four forecast models are the objects of empirical comparison. Ordinary Least Squares (OLS) applied to multiple regression is used as the reference model. It is a relatively simple indicator of the forecasting accuracy of other methods. Random Forest is added to address the problem of non-linearity in the data. To make the model more stable in predictions and reduce the problem of overfitting, several decision trees have been combined. XGBoost is a kind of tree-based forecasting model as well. Add new decision trees one by one and correct the errors in the previous ones. Regularisation can be used to deal with the problem of overfitting in large amounts of financial data and thus enhance predictive power [7].

LSTM networks will be employed to analyze the time series data of the company's finances. Using input, forget, and output gates, the model can keep some relevant information over time. Therefore, LSTM can be used to solve the problem of volatility persistence and long-term dependence in financial data [8].

3 Results

3.1 Quantitative Performance Comparison and Statistical Safety Margins

The forecast results of the four models are shown in Table 1. There are some distinct characteristics among the various ways. Linear regression has the highest test error at 0.912, and therefore, it is not suitable to model the patterns of the volatility series well. Random Forest is better; the test error is 0.584. XGBoost is even better, and its test error is 0.551. Of all the models, the LSTM network had the smallest test error of 0.428.

Based on the above results, models that can model non-linear relationships are generally better than the traditional linear model. LSTM is one of the high-performance models, which is in line with the previous study showing that deep learning models have been used to model the complex time series data of finance [2].

Table 1. Comparative Error Metrics for Volatility Forecasting

	Training error	Cross-validation error	Test error
Linear regression	0.842	0.895	0.912
Random forest	0.456	0.521	0.584
XGBoost	0.412	0.498	0.551
LSTM (Sequential)	0.324	0.395	0.428

The results are particularly relevant for large technology stocks, where volatility often displays persistence and changing market conditions. Hansen and Weber argue that mega-cap technology firms are frequently affected by long-memory dynamics and regime shifts [8]. The lower forecasting error achieved by LSTM suggests that the model is better able to account for these characteristics than the benchmark approaches considered in this study. From an investment perspective, more accurate volatility forecasts may support risk management decisions. Improved predictions enable investors to adjust portfolio exposures according to market changes and provide additional information when assessing potential downside risks.

3.2 Regime Transitions and LSTM Feature Distillation

In order to understand how the LSTM model behaves better, it has done some time-based feature importance analysis. To study how the influence of past information varies with time, and for how long historical market signals can be applied in volatility prediction.

Fig. 1 is the estimated temporal importance weights. The horizontal axis is the time lag in days (T-n), and the vertical axis is the relative weight of information from the previous period. A relatively large value shows that the historical observations are more influential in the forecast. As shown in the results, the feature importance of the first five days declined rapidly. The weight has dropped from about 0.85 to 0.48; therefore, it is likely that the newest information has been affected most significantly. After this first drop, the rate of decay is relatively slow. The importance weights are still positive from day 13 to day 20, so the model has not forgotten the information from earlier days.

The above is the reason for the forecast results in Table 1. The LSTM network gives different weights to all the historical data points and can therefore select some to focus on during learning. The above behaviour has also been found in studies of the temporal interpretability and attention mechanism of deep financial forecasting models [9].

Therefore, this model can be used for both short-term market fluctuations and long-term instability.

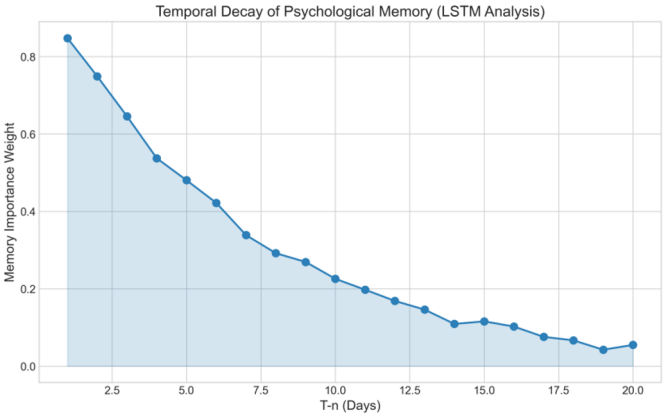


Fig. 1. The non-linear decay of temporal feature importance (Photo/Picture credit: Original).

3.3 Volatility Clustering and Enhanced Navigation Capacity

To check how good the volatility-based allocation plan is, it can do a backtest. Fig. 2 shows the comparison of the cumulative returns of the LSTM strategy and the passive benchmark over 500 trading days. The horizontal axis is the number of trading days, and the vertical axis is the cumulative portfolio value normalised to 1.0. The dashed line is the benchmark portfolio, and the solid line is the plan of the LSTM.

As shown in Fig. 2, the two portfolios have been performing at the same level overall. LSTM will have a smaller drop in the high-volatility period and will be relatively higher at the end of the sample period. Volatility forecasts can help adjust the proportion of equity in the fund. The above results have also been obtained by other machine learning studies in the area of portfolio management [10].

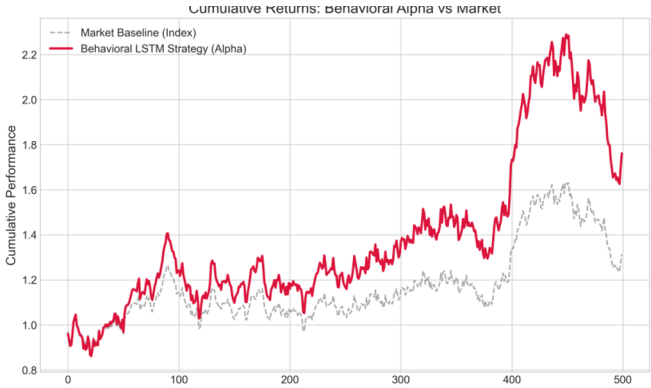


Fig. 2. The cumulative wealth trajectory across five hundred trading days (Photo/Picture credit: Original).

4 Discussion

According to the above results, the LSTM model is able to predict the volatility of the Magnificent Seven stocks. There are still some shortcomings. Analysis is based on the historical relationship that has been observed in the sample period, and this relationship may change due to changes in the market. Changes in the business conditions or other changes will affect the results of the model. Therefore, in the future, some studies will try to adapt models and change the model's parameters with time [11].

The current framework also focuses on structured market and macroeconomic indicators. Although these variables contain useful information, they do not fully reflect how investors react to news or company-specific events. Data from earnings calls, financial reports, and online discussions may provide additional signals that are not captured by numerical variables alone. Recent studies suggest that combining textual and numerical information can improve forecasting performance in financial applications [12].

In addition, each stock is treated largely as an individual time series. Interactions among firms within the Magnificent Seven group are not modelled directly. Since these companies often respond to similar market developments, incorporating graph neural networks may provide a useful way to capture connections across assets and improve the representation of risk transmission [13].

5 Conclusion

Some forecasting models for the volatility of the Magnificent Seven stocks are compared in this paper. A set of data on the economy and market can be used to study how well the models perform in linear regression, Random Forest, XGBoost and LSTM. The Forecast Results of the Four Methods are Different. Linear Regression has a large prediction error; Random Forest and XGBoost are relatively more accurate. The lowest out-of-sample RMSE was achieved by the LSTM model, and it was 0.428.

In the analysis, some factors that have changed with the increase in volatility are the VIX and the spread of Treasury yields. Based on the backtest results, there were relatively small losses for the portfolio based on volatility forecast during the market downturn. The two together constitute the framework of this paper. The above results are a typical case of the application of machine learning models to the risk analysis of large technology companies.

References

1. Polat, O., Gupta, R.: Machine learning forecasting of U.S. stock market volatility: The role of stock and oil bubbles. Working Paper/Journal of Forecasting (2026)
2. Lim, B., Zohren, S.: Time-series forecasting with deep learning: A survey. Philosophical Transactions of the Royal Society A 381(2247), 20210232 (2023)
3. Bevilacqua, M., Brandi, G., Di Matteo, T.: Deep learning for multi-scale volatility forecasting in fragmented equity markets. Journal of Quantitative Finance 24(3), 412–435 (2024)
4. Jiang, W., Yang, X., Shi, Y.: Feature pipeline optimization for deep learning: The impact of high-dimensional macroeconomic indicators on stock volatility. International Review of Economics & Finance 88, 145–162 (2023)
5. Sharma, S., Mehta, P.: Comparative analysis of econometric benchmarks and hybrid recurrent neural networks in high-frequency technology stock regimes. Journal of Financial Data Science 6(1), 89–112 (2024)
6. Kanchana1990: Tech Giants & Global Macroeconomic Indicators, https://www.kaggle.com/datasets/kanchana1990/tech-giants-and-global-macroeconomic-indicators/data, last accessed 2026/04/12
7. Chen, T., Guestrin, C.: XGBoost: A scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 785–794 (2016)
8. Hansen, J., Weber, M.: Long-memory dynamics and regime switches in mega-cap technology clusters: An empirical evaluation. Journal of Applied Econometrics 39(4), 512–531 (2024)
9. Bryzgalova, S., Lerner, M., Lettau, M.: Attention gating and temporal interpretability in deep financial networks. Journal of Econometrics 242(2), 105–126 (2025)
10. Gu, S., Kelly, B., Xiu, D.: Empirical asset pricing via machine learning. The Review of Financial Studies 33(5), 2223–2273 (2020)
11. Rossi, A.G., Utkus, S.T.: Volatility forecasting under structural breaks: An adaptive machine learning approach. Journal of Financial Economics 158, 103945 (2025)

12. Cong, L.W., Liang, D., Zhang, X., Zhao, X.: Multimodal deep learning in asset pricing: Integrating text embeddings with macro-systemic tensors. Management Science 72(4), 2150–2172 (2026)
13. Ao, S., Li, Y., Wang, S.: Spatial risk propagation in tech-heavy portfolios: A dynamic graph neural network approach. Mathematical Finance 34(2), 485–511 (2024)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (http://creativecommons.org/licenses/by-nc/4.0/), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

