Volatility Forecasting of Stock Index via GARCH-Family and LSTM Models

Zihao Li
School of Statistics and Mathematics, Central University of Finance and Economics, Beijing, China

Abstract: In financial markets, reliable forecasting of stock index volatility constitutes a fundamental component of risk management and strategic investment decisions. Traditional economic models struggle to capture complex financial market relationships. For example, Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models fail to fully account for nonlinear dependencies and long-term memory. As a result, accurate forecasting remains a challenging task. In response to the challenges, the LSTM–GEM hybrid framework is proposed in this study, integrating GARCH-family models with Long Short-Term Memory (LSTM) networks. Conditional volatility predictions from GARCH-family models serve as input features to the LSTM network. This allows the hybrid model to model the linear and nonlinear patterns underlying financial time series. To assess the contribution of high-frequency information, we compare model performance when using only low-frequency inputs versus combining both data types. The experimental findings indicate that the LSTM-GEM model consistently achieves superior performance compared to both standalone GARCH-family and LSTM models. Furthermore, incorporating high-frequency data improves forecasting accuracy. The findings demonstrate that the LSTM-GEM model attains lower prediction errors, with the inclusion of high-frequency data further improving its accuracy.

Keywords: LSTM; GARCH; Volatility Prediction

1. Introduction

Financial risk management is crucial for ensuring market stability and promoting sustainable development, with volatility forecasting serving as a core tool for quantifying uncertainty and managing risks[1]. Volatility reflects the degree of fluctuation in asset prices and is especially important for asset pricing, derivative valuation, portfolio risk evaluation, and hedging activities. Accurate volatility forecasts are also essential for risk measurement frameworks, such as value-at-risk (VaR), which rely on reliable estimates to evaluate potential losses under normal market conditions. Consequently, volatility prediction has become a central focus of financial research and practice.

Despite extensive research, forecasting financial market volatility remains a challenging task. Traditional econometric models provide statistical rigor and interpretability but often rely on restrictive linear assumptions that fail that inadequately reflect nonlinear patterns and long-term dependencies present in the financial data. Deep learning approaches, including recurrent neural networks (RNNs) and LSTM models, offer strong capabilities in modeling nonlinear and temporal dynamics but tend to overlook established financial theories and may suffer from overfitting when applied independently. Moreover, the increasing availability of high-frequency data offers valuable intraday information; however effectively integrating such data into volatility forecasting frameworks remains a complex challenge.

Considering the drawbacks of conventional approaches, a new hybrid framework named LSTM–GEM is established by coupling the strengths of GARCH-family models with the temporal learning capability of LSTM networks. Specifically, the LSTM network takes as input low-frequency financial data, realized volatility derived from high-frequency data, and parameter estimates generated by GARCH-family models. By combining these inputs, the model captures both brief intraday variations and extended temporal dependencies in volatility, effectively integrating statistical modeling with data-driven deep learning.

This study makes several key contributions. From a theoretical perspective, the proposed hybrid model integrates statistical volatility modeling with deep learning, providing a framework capable of capturing complex market dynamics and multi-scale volatility patterns. From a practical perspective, empirical analyses indicate that the framework surpasses both standalone conventional econometric models and deep learning approaches in volatility prediction, enhancing forecast reliability and robustness, and serving as a dependable instrument for financial risk management and investment decision-making.

The remainder of this paper is structured as follows. Section II provides a review of existing research on volatility forecasting, including both econometric and deep learning approaches. Section III outlines the research methodology, including the overall modeling framework and general procedures for data preparation and model construction. Section IV describes the experimental design, evaluation metrics, and comparative analysis of model performance. Finally, Section V presents the conclusions, summarizes key findings, highlights limitations, and suggests directions for future work.

2. Related Work

This section is structured into three parts: the first part provides a comprehensive review of the evolution and extensions of GARCH-family models for volatility forecasting; the second part examines LSTM networks and their capacity to capture complex nonlinear temporal dependencies; and the third part examines hybrid frameworks that combine statistical volatility modeling with deep learning techniques.

2.1 Garch

Bollerslev[2] analyzed the Autoregressive Conditional Heteroskedasticity(ARCH) model proposed by Engle[3] and subsequently extended it by introducing the GARCH model, which addressed the overparameterization issue of the original ARCH model and has since been widely adopted. As research on volatility modeling progressed, scholars proposed various extensions to the standard GARCH framework to capture specific empirical characteristics of financial time series. These extensions retained the advantages of the original model while improving its ability to represent stylized facts such as periods of calm and turbulence, volatility clustering[4], and leptokurtic distributions. However, the standard GARCH model has limitations in representing asymmetric volatility responses to shocks, particularly the leverage effect, in which negative shocks influence volatility more strongly than positive shocks of equivalent magnitude. To overcome this limitation, Nelson[5] introduced the Exponential GARCH (EGARCH) model, effectively incorporating asymmetry. Similarly, Zakoian[6] systematically developed the Threshold GARCH (TGARCH) model, overcoming the symmetric effect assumption of standard GARCH models. Engle et al.[7] extended the ARCH framework to the ARCH-in-Mean (ARCH-M) model to account for situations where returns depend on both the mean and the conditional variance, while Engle and Kroner[8] proposed the multivariate GARCH model for jointly modeling multiple time series. Overall, these developments highlight the evolution of GARCH-family models in addressing the limitations of ARCH and capturing key features of financial time series for robust volatility modeling.

2.2 LSTM

Over the past few years, deep learning techniques have found growing application in financial time-series modeling, with LSTM networks showing strong capability in capturing nonlinear dependencies and long-term temporal patterns. Hochreiter and Schmidhuber[9] were the first to propose the LSTM model, introducing a gated recurrent architecture designed to model long-term dependencies in sequential data. SiamiNamin et al.[10] evaluated the predictive performance of ARIMA and LSTM models on economic and financial time series, with experimental results indicating that LSTM achieves superior predictive accuracy. Xiong et al.[11] employed an LSTM model incorporating domestic Google Trends data to forecast S&P 500 volatility, demonstrating performance superior to that of alternative benchmark models. Chen et al.[12] applied LSTM models to forecast Chinese stock market returns, confirming their superior performance over alternative methods. Fischer and Krauss[13] employed LSTM models for out-of-sample directional volatility prediction on S&P 500 constituents over the period 1992–2015, outperforming random forests, deep neural networks, and logistic regression models.

2.3 Hybrid Models

The effective integration of GARCH models with emerging methods has become a prominent research topic. Roh[14] introduced a hybrid framework that integrates financial time-series models (EWMA, GARCH, EGARCH) with a feed forward neural network using the KOSPI 200 index, showing that the combined approach outperforms individual models, with the EGARCH–network combination performing best. Kim and Won[15] proposed a hybrid framework integrating LSTM networks with multiple GARCH-family models for forecasting stock index volatility, demonstrating superior predictive performance. Hu et al.[16] developed an LSTM–ANN–GARCH hybrid model for copper price volatility prediction, which outperformed single models. Cao and Ren[17] investigated RMB exchange rate volatility prediction using a hybrid LSTM and GARCH-family model, demonstrating that the hybrid approach significantly outperforms individual models. Verma[18] developed a GARCH–LSTM hybrid model for crude oil futures volatility forecasting and demonstrated that the hybrid framework outperforms traditional GARCH-family models across multiple horizons. Pan et al.[19] constructed a MULTI-GARCH-LSTM hybrid model for gold futures prediction and showed it outperformed standalone GARCH-family models, LSTM, and other hybrids in predictive accuracy.

Existing volatility forecasting models often rely on traditional econometric methods and trading data, yet they underutilize high-frequency information and overlook the complementarity of statistical and deep learning approaches. This study tackles the aforementioned shortcomings by proposing a hybrid approach which merges GARCH-family models with LSTM networks, incorporating low-frequency indicators, realized volatility from high-frequency data, and GARCH parameter estimates to enhance prediction accuracy and robustness.

3. Research Methodology

As illustrated in Figure 1, the proposed framework for stock index volatility prediction consists of two primary components: (1) parameter estimation from GARCH-family models to capture linear volatility patterns, and (2) volatility forecasting using an LSTM network that integrates low-frequency financial indicators and realized volatility derived from high-frequency data. Detailed descriptions of each component are provided in the following subsections.

3.1 Sample Selection and Data Preprocessing

3.1.1 Data collection

The study selects the CSI 300 Index as the research subject, as it covers most large-cap stocks in the Shanghai and Shenzhen markets, providing a representative and highly liquid benchmark that reflects mainstream investor returns in China. The dataset comprises both low-frequency daily trading data and high-frequency 5-minute data, sourced from the Wind database. The sample covers the period from September 24, 2022, to September 24, 2025, totaling 905 trading days. Eighty percent of the data is used for training and model tuning, while the remaining 20% is reserved for testing and performance evaluation.

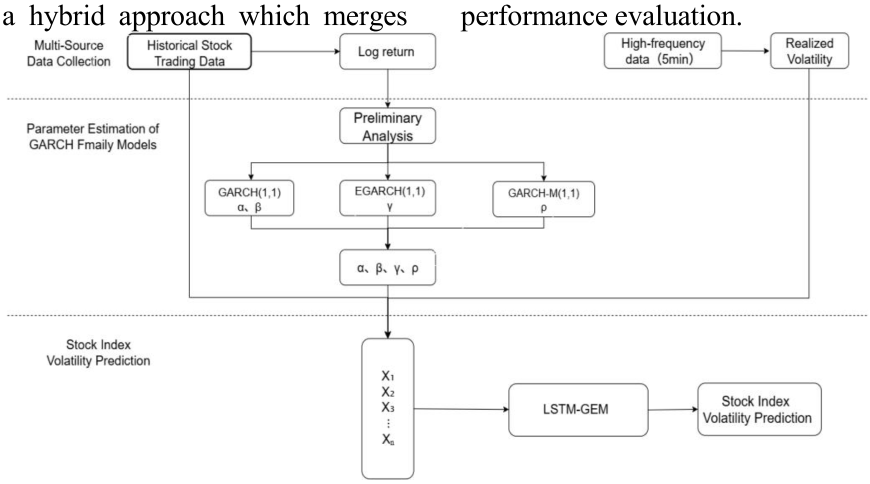

Figure 1. Stock Index Volatility Prediction Framework

3.1.2 Data processing

To eliminate scale differences among variables and accelerate model convergence, all features are normalized using the min–max method, represented by

$$\widetilde{x_i}=\frac{x_i-\min_{1\leq j\leq n} x_j}{\max_{1\leq j\leq n} x_j-\min_{1\leq j\leq n} x_j} \quad (1)$$

Where x_{min} and x_{max} denote the minimum and maximum of each variable. This transformation maps all features to the [0,1] range, improving model stability and predictive accuracy.

Furthermore, the stock index return is widely recognized as a fundamental measure for evaluating investment performance. Let p_t denote the daily closing price index; the return of the stock index can be calculated as:

$$r_t=100\ln\frac{p_t}{p_{t-1}} \quad (2)$$

where r_t represents the logarithmic return at time t.

What's more, Volatility, as a key indicator of stock price fluctuations, has received growing interest recently, particularly given the growing availability of high-frequency financial data at shorter intervals. Following the realized volatility framework proposed by Andersen et al.[20], the high-frequency data are processed through the following steps:

Use mean imputation to handle missing values and outliers

Remove data corresponding to non-trading periods, including market holidays, overnight gaps, and intraday breaks

Cleaned intraday prices are normalized using the min–max method to ensure consistent scaling and facilitate model convergence.

Use Equation (1) to normalize the cleaned data

Construct five-minute intervals such that each trading day t is divided into 48 intervals, with d_0 denoting the opening time and $d_1, d_2, \cdots d_{48}$ representing subsequent 5-minute intervals. Let $P_{t,d}$ denote the price at time d on day t, where $P_{t,0}$ is the opening price and $P_{t,1}$ the closing price of the first interval. The intraday high-frequency return can be calculated as:

$$R_{t,d}=\ln P_{t,d}-\ln P_{t,d-1} \quad (3)$$

where $t=2,3\cdots 905$, $d=1,2\cdots 48$

The realized volatility (RV) of day t is then computed as the sum of squared intraday returns:

$$RV_t=\sum_{d=1}^{n_t} R_{t,d}^2 \quad (4)$$

Table 1. Descriptive Statistics of Daily Log Returns of the CSI 300 Index

	Mean	Standard Deviation	Skewness	Kurtosis	Jarque-Bera Statistic	p-value
r_t	0.0220	1.0362	0.3298	14.4509	6347.6373	0

where n_t=48 denotes the total number of intraday observations.

3.2 Parameter Estimation of the GARCH Family Models

3.2.1 Descriptive statistics

A preliminary descriptive analysis of the CSI 300 Index return series is conducted, including the computation of the mean, standard deviation, skewness, kurtosis, Jarque–Bera (J–B) statistic, and its associated p-value. The results are summarized in Table 1.

Table 1 and Figure 2 report the descriptive statistics of the dataset. The mean of daily log returns is 0.022, and the standard deviation is 1.0362, indicating that, for a given mean, larger standard deviation corresponds to higher stock price volatility. The skewness is 0.3298, suggesting a slight rightward skew, meaning that upward movements are marginally more probable than downward movements. The excess kurtosis is 14.4509, indicating a highly peaked distribution with fat tails, implying a higher likelihood of extreme returns. The Jarque-Bera test produces a value of 6347.6373 with a p-value below 0.01, which provides strong evidence against the null hypothesis of normality and suggesting that the return distribution significantly deviates from normality.

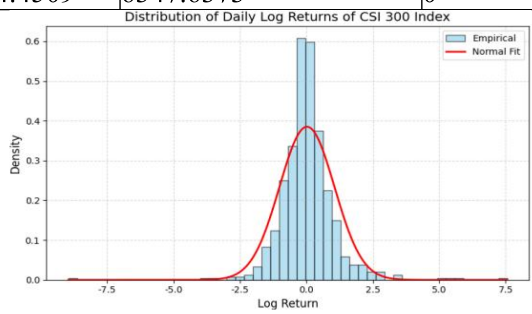

Figure 2. Log Difference in CSI 300 Index Histogram

3.2.2 Stationarity test

In this study, the stationarity of the CSI 300 Index daily logarithmic returns was examined using both the Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) tests. As reported in Table 2, the ADF test statistic is large in absolute value with a p-value below 0.01, providing sufficient evidence to reject the null hypothesis of non-stationarity at the 1% significance level. Likewise, the PP test yields a statistic of -23.1846 with a p-value of 0.0000, also rejecting the null hypothesis at the 1% level. These results collectively indicate that the daily log return series r_t is stationary.

Table 2. Stationarity Test of Daily Log Returns of the CSI 300 Index

	ADF Test		PP Test	
r_t	-11.4957	<0.01	-23.1846	<0.01

3.2.3 GARCH-family model diagnostics

In this study, the daily returns volatility of the CSI 300 Index was analyzed using GARCH-family models. Prior to model estimation, the suitability of the return series r_t for GARCH modeling was examined via the Ljung-Box and ARCH-LM tests. As shown in Table 3, the Ljung-Box test at lag 10 yields a statistic of 36.7432 with a p-value less than 0.01, constituting strong evidence against the null hypothesis of no autocorrelation at the 1% significance level. The ARCH-LM test reports an LM statistic of 73.9032 with a p-value under 0.01, indicating significant ARCH effects, indicating significant ARCH effects at the 99% level. These results confirm that the conditional heteroskedasticity exists within the daily returns of the CSI 300 Index, satisfying the prerequisites for GARCH-family modeling. Accordingly, standard GARCH(1,1), GARCH(1,1)-M, and EGARCH(1,1) specifications are employed for parameter estimation.

Table 3. GARCH-family Model Diagnostics of Daily Log Returns of the CSI 300 Index

	Ljung-Box Test		**ARCH-LM Test**	
r_t	36.7432	<0.01	73.9032	<0.01

3.2.4 Parameter estimation of GARCH-Family models

To account for the time-varying characteristics of volatility, the parameters of the GARCH-family models are estimated through a rolling-window procedure. Specifically, a sequence of 200 consecutive trading days is employed for parameter estimation, with the estimated values assigned to the last day of the window. The window then rolls forward one day at a time to generate dynamic parameter estimates. The key parameters include: α and β, representing the ARCH and GARCH effects, respectively, which indicate the immediate and persistent impacts on conditional variance; γ, reflecting the asymmetric impact of positive and negative shocks on volatility, with $\gamma>0$ indicating a greater impact of positive news on returns and vice versa; and ρ, capturing the influence of expected volatility on the mean return. Since the explanatory variables evolve gradually, the rolling estimates remain relatively stable, reflecting consistent volatility dynamics over time.

3.3 Construction of the LSTM Prediction Model

As illustrated in Fig. 3, the LSTM prediction model is designed to capture the nonlinear temporal dependencies of realized volatility (RV). The architecture consists of an input layer that incorporates RV along with low-frequency financial features, followed by the first LSTM hidden layer and a Dropout layer to mitigate overfitting. A second LSTM hidden layer and an additional Dropout layer are then applied for deeper feature extraction and further regularization. The network concludes with one fully connected layer and a final output layer that generates the predicted RV. As illustrated in Fig. 3, the LSTM prediction model is designed to capture the nonlinear temporal dependencies of realized volatility (RV). The architecture consists of an input layer that incorporates RV along with low-frequency financial features. The first LSTM hidden layer processes the sequential input data, succeeded by a Dropout layer to reduce overfitting. Subsequently, a second LSTM hidden layer extracts deeper temporal features, accompanied by an additional Dropout layer for further regularization. The network concludes with two fully connected layers: the first dense layer integrates the extracted features, while the second layer functions as the output layer to produce the predicted RV. This structure ensures that both short-term and long-term patterns in the financial time series are captured effectively while controlling overfitting.

For model training, the hyperparameters are determined through grid search to ensure optimal predictive performance. The grid search systematically explores multiple parameter combinations and selects the configuration that minimizes the validation error. As a result, the final settings are as follows: the time step is set to 22, the batch size to 32, the number of epochs to 50, the learning rate to 0.001, and the mean squared error (MSE) is adopted as the loss function. This configuration effectively balances model complexity and generalization, leading to robust and accurate volatility forecasting.

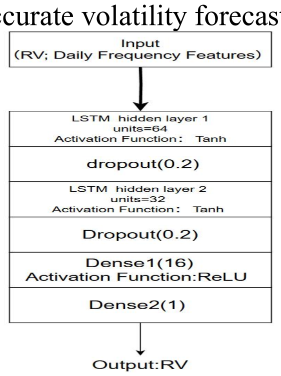

Figure 3. LSTM Model

3.4 Construction of the Hybrid Forecasting Model

This study integrates GARCH-family models with LSTM networks to construct a hybrid forecasting framework, leveraging the strengths of both approaches to enhance stock index volatility prediction. Specifically, the α coefficient representing volatility persistence and the β coefficient representing the magnitude of volatility shocks from the GARCH model are used as inputs to the LSTM, forming the LSTM-G model. Subsequently, the γ coefficient from the EGARCH model, which captures asymmetric information effects, is incorporated into the input layer to construct the LSTM-GE model. Finally, the ρ coefficient from the GARCH-M model, reflecting mean effects, is added to form the LSTM-GEM model. All LSTM and LSTM–GARCH hybrid models employ a rolling prediction scheme. Except for the differences in input features (GARCH-family parameters), the network architecture and hyperparameter settings remain consistent across models.

Table 4. Comparison of Prediction Results among Different Models

Model	MAE	MSE	MAPE	RMSE
GARCH	0.1834	0.0297	38.09	0.1722
LSTM	0.0692	0.0075	13.96	0.0865
LSTM-G	0.0597	0.0067	12.07	0.0816
LSTM-GE	0.0526	0.0058	10.63	0.0764
LSTM-GEM	**0.0457**	**0.0054**	**9.27**	**0.0737**
LSTM-GEM (without RV)	0.1309	0.0141	26.33	0.1186

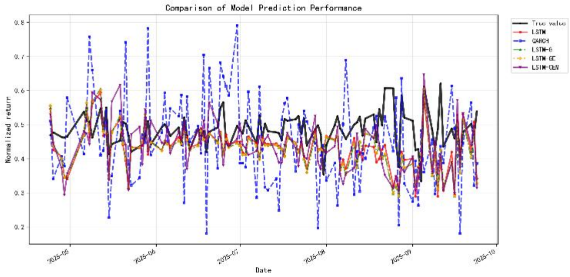

Figure 4. Comparison of Model Prediction Performance

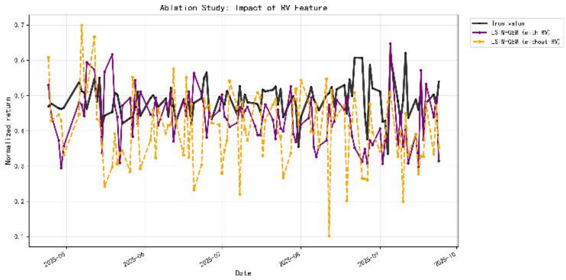

Figure 5. Ablation Study: Impact of High-frequency Data

4. Experiments and Results Analysis

4.1 Performance Evaluation Metrics

To comprehensively evaluate the predictive performance of different models, this study employs four error metrics: Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), Mean Squared Error (MSE), and Root Mean Squared Error (RMSE), defined as follows:

Mean Absolute Error (MAE):

$$MAE=\frac{1}{T}\sum_{t=1}^{T}\left|RV_t-\hat{\sigma}_t^2\right| \quad (5)$$

Mean Absolute Percentage Error (MAPE):

$$MAPE=\frac{1}{T}\sum_{t=1}^{T}\left|\frac{RV_t-\hat{\sigma}_t^2}{\hat{\sigma}_t^2}\right| \quad (6)$$

Mean Squared Error (MSE):

$$MSE=\frac{1}{T}\sum_{t=1}^{T}\left(RV_t-\hat{\sigma}_t^2\right)^2 \quad (7)$$

Root Mean Squared Error (RMSE):

$$RMSE=\sqrt{MSE} \quad (8)$$

4.2 Experimental Results and Performance Comparison

We select GARCH and LSTM as baseline models and conduct comparative experiments with the proposed LSTM-G, LSTM-GE, and LSTM-GEM models. Additionally, ablation experiments are conducted to evaluate the effect of incorporating high-frequency data on model performance.

Figure 4 and Table 1 present a summary of the models' forecasting accuracy. The hybrid approaches (LSTM-G, LSTM-GE, and LSTM-GEM) consistently outperform the baseline LSTM and GARCH models across all metrics. Notably, LSTM-GEM achieves the lowest errors, demonstrating the effectiveness of incorporating GARCH-family parameters, including asymmetry and mean-effect components, into the LSTM network. The error comparison plot clearly illustrates the progressive improvement from single-model to hybrid-model predictions.

Building on the comparative experiments, ablation studies were conducted to evaluate the contribution of high-frequency data (realized volatility, RV) to the hybrid model. Specifically, the LSTM-GEM model was re-evaluated without including RV as an input feature, and the results are summarized in Table 4 and visualized in Figure 5 The findings indicate that incorporating high-frequency data significantly enhances predictive performance, further demonstrating its importance in stock index volatility forecasting.

5. Conclusion and Future Work

This work introduces the LSTM-GEM hybrid, which incorporates GARCH-family model parameters, high-frequency realized volatility, and low-frequency financial features for forecasting the volatility of the CSI 300 Index. The experimental outcomes indicate that the proposed LSTM-GEM model considerably surpasses the baseline approaches, and the inclusion of high-frequency data further enhances predictive accuracy. These results highlight underscore the utility of integrating econometric models with deep learning approaches for stock index volatility forecasting, providing practical insights for risk management and derivative pricing.

However, this study has certain limitations. The three-year data span may restrict the model's capacity to reflect long-term trends and extreme market events. In addition, macroeconomic factors, such as investor sentiment, national policies, and market news, are not incorporated, although they may have a considerable impact on volatility dynamics.

Future work can be pursued along broader directions, including extending the data types and time horizon, incorporating additional market-relevant information, exploring alternative hybridizations of deep learning and econometric models, and applying the approach to different financial markets and asset classes to enhance predictive performance and generalization capabilities.

References

[1] A. García-Medina and E. Aguayo-Moreno, 'LSTM–GARCH hybrid model for the prediction of volatility in cryptocurrency portfolios', Computational Economics, vol. 63, no. 4, pp. 1511–1542, 2024.

[2] T. Bollerslev, 'Generalized autoregressive conditional heteroskedasticity', Journal of econometrics, vol. 31, no. 3, pp. 307–327, 1986.

[3] R. F. Engle, 'Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation', Econometrica: Journal of the econometric society, pp. 987–1007, 1982.

[4] B. Mandelbrot, 'The variation of some other speculative prices', The Journal of Business, vol. 40, no. 4, pp. 393–413, 1967.

[5] D. B. Nelson, 'Conditional heteroskedasticity in asset returns: A new approach', Econometrica: Journal of the econometric society, pp. 347–370, 1991.

[6] J.-M. Zakoian, 'Threshold heteroskedastic models', Journal of Economic Dynamics and control, vol. 18, no. 5, pp. 931–955, 1994.

[7] R. F. Engle, D. M. Lilien, and R. P. Robins, 'Estimating time varying risk premia in the term structure: The ARCH-M model', Econometrica: journal of the Econometric Society, pp. 391–407, 1987.

[8] R. F. Engle and K. F. Kroner, 'Multivariate simultaneous generalized ARCH', Econometric theory, vol. 11, no. 1, pp. 122–150, 1995.

[9] S. Hochreiter and J. Schmidhuber, 'Long short-term memory', Neural computation, vol. 9, no. 8, pp. 1735–1780, 1997.

[10] S. Siami-Namini and A. S. Namin, 'Forecasting economics and financial time series: ARIMA vs. LSTM', arXiv preprint arXiv:1803.06386, 2018.

[11] R. Xiong, E. P. Nichols, and Y. Shen, 'Deep learning stock volatility with google domestic trends', arXiv preprint arXiv:1512.04916, 2015.

[12] K. Chen, Y. Zhou, and F. Dai, 'A LSTM-based method for stock returns prediction: A case study of China stock market', in 2015 IEEE international conference on big data (big data), IEEE, 2015, pp. 2823–2824.

[13] T. Fischer and C. Krauss, 'Deep learning with long short-term memory networks for financial market predictions', European journal of operational research, vol. 270, no. 2, pp. 654–669, 2018.

[14] T. H. Roh, 'Forecasting the volatility of stock price index', Expert systems with applications, vol. 33, no. 4, pp. 916–922, 2007.

[15] H. Y. Kim and C. H. Won, 'Forecasting the volatility of stock price index: A hybrid model integrating LSTM with multiple GARCH-type models', Expert Systems with Applications, vol. 103, pp. 25–37, 2018.

[16] Y. Hu, J. Ni, and L. Wen, 'A hybrid deep learning approach by integrating LSTM-ANN networks with GARCH model for copper price volatility prediction', Physica A: Statistical Mechanics and its Applications, vol. 557, p. 124907, 2020.

[17] W. Cao and S. Ren, 'Research on RMB Exchange Rate Volatility Prediction Based on the Hybrid Model of LSTM and GARCH Family', Computer Applications and Software, vol. 37, no. S1, pp. 79–82, 2020.
[18] S. Verma, 'Forecasting volatility of crude oil futures using a GARCH–RNN hybrid approach', Intelligent Systems in Accounting, Finance and Management, vol. 28, no. 2, pp. 130–142, 2021.
[19] H. Pan, Y. Tang, and G. Wang, 'A Stock index futures price prediction approach based on the MULTI-GARCH-LSTM mixed model', Mathematics, vol. 12, no. 11, p. 1677, 2024.
[20] T. G. Andersen, T. Bollerslev, F. X. Diebold, and P. Labys, 'Modeling and forecasting realized volatility', Econometrica, vol. 71, no. 2, pp. 579–625, 2003.

