

Volatility Clustering and Short-Term Forecasting of the Hang Seng Tech Index Based on the ARIMA-GARCH Model

Hongjie Lai*

Sydney Smart Technology
College, Northeastern University,
Northeastern University at
Qinhuangdao, 066004, China

*Corresponding author:
gbc1907@163.com

Abstract:

This study investigates the asymmetric risk characteristics of technology stocks by examining the Hang Seng Technology Index (HSTECH). Daily data from 2021 to 2026 are analyzed using a hybrid ARIMA-EGARCH model with a Student-t distribution, which combines linear mean modeling with non-linear volatility modeling to examine volatility clustering, leverage effects, and short-term forecasting performance. The results show that the HSTECH return series has clear heavy-tailed and leptokurtic characteristics, together with strong volatility clustering, indicating that the normality assumption is not suitable for this market. The analysis also finds a significant negative leverage effect, which shows that negative market shocks increase future volatility more than positive shocks of a similar size. Dynamic out-of-sample rolling forecasts further show that the hybrid ARIMA-EGARCH model produces lower prediction errors than the standard linear ARIMA model, demonstrating its ability to capture the volatility dynamics of the HSTECH more accurately. These findings provide useful evidence for dynamic hedging, Value-at-Risk (VaR) estimation, and downside risk management in technology-focused equity markets.

Keywords: Hang Seng TECH Index (HSTECH); ARIMA; EGARCH; volatility clustering; leverage effect

1. Introduction

The technology sector is an important driver of modern economic growth. In Hong Kong, it has become a fast-growing part of the offshore financial market and serves as an important channel for global investment in China's digital economy. The Hang Seng TECH Index (HSTECH) follows leading technology companies from different industries.

Between 2021 and 2026, the technology sector experienced large price swings. Rapid global interest rate increases, trade policy disputes, and the rapid growth of Large Language Models (LLMs) caused frequent market shocks. As the result, market risk changed by the time. Periods of high volatility lasted for a long time and often appeared together, a pattern known as volatility clustering. Traditional models that assume

constant variance often underestimate risk during market declines. So, models that allow volatility to change over time are needed to measure market risk more accurately.

Given that these market conditions, this study examines three questions. Does the HSTECH show persistent volatility clustering? Does the Exponential GARCH (EGARCH) model perform better than the standard GARCH model in describing the market's response to negative news? And can a hybrid ARIMA-EGARCH model provide accurate short-term forecasts in a highly volatile market?

The ARIMA model describes the average movement of prices, but it does not model changes in volatility. The standard GARCH model estimates changing volatility, but it assumes that positive and negative news have the same effect on future risk. This assumption does not match real financial markets, where negative news usually increases volatility more than positive news. Combining ARIMA with a GARCH-type model makes it possible to model both price movements and changes in market risk [1].

The EGARCH model is used because it can capture the leverage effect. This effect means that negative news has a bigger impact on future volatility than positive news of a similar size [2, 3]. Since the HSTECH shows large and changing risk, EGARCH provides a better way to describe its volatility. This approach also helps connect econometric modeling with risk management in the Hong Kong technology stock market.

Building on these methods, the key contributions of this paper are highlighted in three main areas: First, most past studies focus on broad indices like the Hang Seng Index. In contrast, this study focuses directly on the Hang Seng TECH Index (HSTECH) during a high-volatility period (2021–2026). This provides clear evidence of the unique heavy-tailed traits and leverage effects in the tech sector. Second, a hybrid ARIMA(2,0,3)-EGARCH(1,1) model with a Student-t distribution is built. This design fixes the flaws of the standard normal distribution assumption. It accurately measures both simple price moves and asymmetric responses to negative market news. Lastly, dynamic out-of-sample rolling forecasts show that this hybrid model lowers prediction errors (RMSE and MAE). These results give practical tools for short-term risk management and Value-at-Risk (VaR) estimation.

2. Literature Review

Research in finance has changed over time. Past studies mostly looked at price moves. But new studies focus more on tracking and guessing financial risk. Stock returns often show risk that changes over time, so showing these risk changes clearly is a key goal in finance.

Old models like ARIMA assume that risk stays the same over time [1]. But this assumption is often wrong in real

markets. Stock returns usually have heavy tails, big price jumps, and risk that changes over time. So, the GARCH model was made to guess changing risk by looking at past shocks and past risk, making GARCH models are now a common tool to show how high-risk periods group together in stock returns [2, 3].

The basic GARCH model treats good and bad news as having the exact same effect on future risk. Bad news usually causes a much bigger jump in market risk than good news of the same size. This fact is called the leverage effect. The EGARCH model was made to track this unequal reaction, so it gives a more real picture of financial risk [4, 5]. Research on the Hong Kong stock market shows this trend. Most studies focus on the Hang Seng Index (HSI). They use GARCH models to show lasting risk changes and big groups of high risk across the whole market [6, 7]. These results show that ups and downs in the Hong Kong stock market are driven by world money flows and outside economic conditions.

But technology stocks' prices change more easily because of investor feelings, global rules, and changes in the tech industry. Few studies look at the Hang Seng TECH Index (HSTECH) as a separate group [6, 7]. Most research puts tech stocks only inside the whole market, and many studies still use models that treat all news the same way. This problem is very important during times of big economic and rule changes between 2021 and 2026, when tech stocks had big price shifts [4].

Many researches use machine learning and other complex computer tools to guess financial risk [8, 9]. These tools often make correct guesses, but their inner workings are hard to understand and give little economic proof for risk changes. And many studies using GARCH models still look only at general market indexes [6, 7]. This makes it hard to see the special risk habits of specific areas like technology. A combined ARIMA-EGARCH model is used on the Hang Seng TECH Index. ARIMA is used to track average price moves, and EGARCH is used to track changing risk and unequal reactions to market shocks [1, 2, 4, 10]. This tracks both simple and complex patterns in stock returns, and it shows risk changes better than using just one model [11].

This study shows unequal risk and the leverage effect in Hong Kong tech stocks more clearly by using a simple, clear model. At the same time, it helps people understand how risk spreads and how to guess short-term risk in a tech market. This adds new proof to past studies that mostly looked at the whole stock market.

3. Methodology

The basic data is collected from Investing.com and shows the daily closing prices of the HSTECH index from 2021 to 2026. Raw prices have moving trends and changing

noise. Because of this, the prices are turned into daily log returns () [6, 10]. This math subtraction method fixes scale problems and keeps a clear meaning:

$$R_t = \ln(P_t) - \ln(P_{t-1}) \quad (1)$$

where and are the closing prices on day and .

Before executing non-linear volatility frameworks, a diagnostic pre-testing stage is required to evaluate stationarity, normality, and autocorrelation, thereby preventing spurious regressions.

A test called the Augmented Dickey-Fuller (ADF) test checks if the data is stable [1, 10]. This test uses past changes to fix repeating errors. The test equation is:

$$\Delta R_t = \alpha + \beta t + \gamma R_{t-1} + \sum_{i=1}^m \delta_i \Delta R_{t-i} + \varepsilon_t \quad (2)$$

The test targets the null hypothesis (indicating a unit root). Rejecting establishes that the underlying series is integrated of order zero, or .

The Jarque-Bera (JB) test checks if the data fits a normal bell curve by measuring its skewness () and kurtosis () [6]:

$$JB = \frac{N}{6} \left[S^2 + \frac{(K-3)^2}{4} \right] \quad (3)$$

where denotes the sample size. Under the null hypothesis of normality (.), the statistic asymptotically follows a distribution.

The Ljung-Box Q-test identifies potential serial dependence up to a specified lag [10]:

$$Q(m) = N(N+2) \sum_{k=1}^m \frac{\rho_k^2}{N-K} \quad (4)$$

The resulting p-value determines whether the return series exhibits exploitable correlation or behaves as an unpredictable white noise process.

After confirming that the data is stable and not normal, a two-part model is built to find the average trend with an ARIMA process and track changing risk with GARCH models [10]. The first part uses an ARIMA(p,0,q) model. The change setting is d=0 because the data is already stable:

$$\Phi(L)R_t = \Theta(L)\varepsilon_t \quad (5)$$

where and represent the autoregressive (AR) and moving average (MA) lag operator polynomials, respectively.

The best settings (p,q) are found by testing many options to find the lowest score on two tests called Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC). These tests help balance a good fit with a simple model [1, 11]:

$$AIC = -2\ln(L) + 2k \quad (6)$$

$$BIC = -2\ln(L) + k\ln(N) \quad (7)$$

where is the maximized likelihood, denotes the number of estimated parameters, and represents the in-sample

size. Minimizing these joint criteria penalizes overfitting and ensures that the extracted residuals () are insulated from systematic autocorrelation.

Then, these leftover errors () are checked for changing risk using the ARCH-LM test [3]:

$$\varepsilon_t^2 = \alpha_0 + \sum_{i=1}^q \alpha_i \varepsilon_{t-i}^2 + v_t \quad (8)$$

By regressing current squared residuals () on lagged values, significant coefficients () detect conditional heteroskedasticity. The resulting LM statistic () confirms the presence of volatility clustering, justifying an upgrade to conditional variance models.

To model this time-varying risk, the standard symmetric framework is established [2]:

$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2 \quad (9)$$

where parameter stability and positive variance require , and.

To accommodate the asymmetric shocks typical of offshore equity markets, Nelson's EGARCH (1,1) model is deployed [5]. By modeling the natural log of conditional variance, it bypasses non-negativity constraints while measuring the "leverage effect":

$$\ln(\sigma_t^2) = \omega + \gamma \left[\frac{\varepsilon_{t-1}}{\sigma_{t-1}} - E \left(\left| \frac{\varepsilon_{t-1}}{\sigma_{t-1}} \right| \right) \right] + \alpha \frac{\varepsilon_{t-1}}{\sigma_{t-1}} + \beta \ln(\sigma_{t-1}^2) \quad (10)$$

The setting measures this unequal reaction. A negative proves that bad news raises risk much more than good news.

The final segment defines the loss functions utilized to judge out-of-sample forecasting precision during the subsequent rolling-window simulation. The evaluation relies on RMSE and MAE:

$$RMSE = \sqrt{\frac{1}{T} \sum_{t=1}^T (\hat{y}_t - y_t)^2} \quad (11)$$

$$MAE = \frac{1}{T} \sum_{t=1}^T |\hat{y}_t - y_t| \quad (12)$$

where represents the actual observed baseline value, is the model's forward prediction, and is the total number of evaluation steps. Lower values mathematically establish superior forecast accuracy for dynamic risk management.

4. Empirical Results on the Volatility Characteristics of the HSTECH

4.1 Descriptive Statistics and Basic Stationarity Test

The investigation begins with a statistical evaluation of the processed HSTECH daily return series () during the in-sample window from 2021 to 2025, trading day ob-

servations. Table 1 systematizes the descriptive statistical properties of the baseline series.

Table 1. Descriptive Statistics of HSTECH (2021–2025)

Statistical Metric	Computed Value
Mean	-3.5172×10^{-4}
Standard Deviation	0.02523
Skewness	0.1617
Kurtosis	10.0095
JB p-value	<0.0001
JB Statistic	2,526.99

Table 1 shows that daily returns have a small negative mean of -3.5172×10^{-4} . And they have a standard deviation of 0.025. These numbers show the high price changes and risk in offshore markets.

The returns have a small positive skewness of 0.16. The kurtosis score is 10.0095, which is much higher than the normal score of 3.0. Showing that the HSTECH returns have a heavy-tailed shape. This means extreme market

shocks happen more often than in a normal setup [6]. The Jarque-Bera (JB) test rejects the idea of a normal distribution. This shows that the normal distribution does not fit the data. So, a fat-tailed model like the Student-t distribution should be used to calculate risk better.

And the ADF test is run to check for a steady trend. Table 2 shows these test results.

Table 2. Augmented Dickey-Fuller (ADF) Unit Root Test Results

Test Statistic Component	Computed Value	1% Critical Value	5% Critical Value	10% Critical Value	Significance
(Unit Root Test)	-17.4622	-3.43	-2.86	-2.57	(p<0.01)
(Drift Intercept)	152.4664	6.43	4.59	3.78	(p<0.01)

The ADF test statistic is -17.46, which sits substantially below the strict 1% critical threshold of -3.43. Simultaneously, the joint test statistic confirms the statistical validity of the drift intercept. Therefore, the null hypothesis of a unit root is decisively rejected at the 1% significance level, establishing that the HSTECH return series is strictly stationary and integrated of order zero, or .

4.2 ARIMA Conditional Mean Filtering and Residual ARCH Testing

A grid search minimization of AIC and BIC identifies an configuration with a zero-mean restriction as the optimal architecture to filter conditional mean dependencies [1]. Table 3 summarizes the parametric coefficient estimations.

Table 3. ARIMA (2,0,3) Conditional Mean Estimation Results

Parameter Specification	Coefficient Estimate	Standard Error	t-statistic	Significance
ar1	-0.8233	0.0863	-9.5399	(p<0.01)
ar2	-0.8325	0.076	-10.9539	(p<0.01)
ma1	0.8267	0.0914	9.0449	(p<0.01)
ma2	0.8118	0.092	8.8239	(p<0.01)
ma3	-0.0645	0.0347	-1.8588	(p<0.10)
Model Fit Information	$\sigma^2=0.0006295$; Log-Likelihood = 2781.03; AIC = -5550.06; BIC = -5519.39			

The ar and ma coefficients exhibit high statistical significance, indicating that short-term price adjustments in the HSTECH index are driven by historical innovations. The

conditional mean filtering innovations are extracted and plotted in Fig. 1.




Fig. 1 ARIMA Residuals

Table 4. ARCH-LM heteroscedasticity

Lag Order Length	p-value	LM Test Statistic	Empirical Conclusion
Lag 1			reject (ARCH Effect Present)
Lag 3			reject (ARCH Effect Present)
Lag 6			reject (ARCH Effect Present)
Lag 12			reject (ARCH Effect Present)

The LM test results presented in Table 4 reveal consistently high statistics across all designated lag structures, with corresponding p-values falling well below standard significance thresholds. This robust empirical pattern firmly rejects the null hypothesis of residual homoskedasticity, validating the presence of strong volatility clustering within the HSTECH index. Consequently, these diagnostic outcomes provide a solid econometric justification for upgrading the empirical framework to incorporate time-varying, conditional variance models.

4.3 Parallel Comparative Evaluation of GARCH and EGARCH Formulations

After strong ARCH effects are confirmed, the constant variance assumption of traditional linear models is no longer valid, so GARCH-family models are used to describe time-varying volatility. Financial asset returns often show asymmetric shocks and heavy-tailed distributions, so different models are compared to avoid choosing a model based on subjective judgment. The GARCH (1,1) model and the EGARCH (1,1) model are tested under both the normal distribution and the Student-t distribution [2, 4].

Table 5. Fitness and Information Criteria of Alternative Volatility Specifications

Model Specifications	Log-Likelihood	AIC	BIC	Rank
ARIMA (2,0,3)-EGARCH (1,1)-Student-t	2896.257	-4.7068	-4.6609	1 (Final Selected Model — Economic Optimum)
ARIMA (2,0,3)-GARCH (1,1)-Student-t	2896.581	-4.7089	-4.6672	2 (Marginal Statistical Optimum)
ARIMA (2,0,3)-EGARCH (1,1)-Normal	2868.885	-4.6638	-4.6221	3
ARIMA (2,0,3)-GARCH (1,1)-Normal	2867.175	-4.6626	-4.6251	4

Table 5 shows that switching from a normal distribution to a Student-t distribution makes both GARCH and

EGARCH models fit much better. This change raises the Log-Likelihood score from about 2867 to over 2896, and it lowers both the AIC and BIC scores at the same time. These better scores prove that the HSTECH index has very heavy tails and many extreme price jumps. Normal distributions always miss rare "black swan" shocks when the market drops [12]. But the Student-t model tracks these extreme errors well, and this makes risk calculations much more reliable.

With the Student-t distribution, the standard GARCH (1,1) model has slightly lower AIC and BIC scores. But the differences are very small, so they do not really matter. So, the ARIMA(2,0,3)-EGARCH (1,1)-Student-t model is chosen as the best tool for guessing future risk. This is because it tracks an unequal risk score (α) that the standard GARCH model misses completely.

The basic GARCH model forces a strict rule. It assumes that good and bad news change risk in the exact same way. But this assumption goes against how real investors act in the market. The technology area is very sensitive to unequal economic shocks. Testing the EGARCH model shows a strong unequal risk score (α). This proves that bad news causes much higher market risk than good news of the same size [5]. The standard GARCH model gets a slightly better score on some tests because it uses one

less math setting. But it completely misses the important downside risk of the tech market.

Ignoring this proven leverage effect (α) just for a tiny math gain would create big errors in future guesses and VaR risk calculations. This model choice balances a small math difference with a much better economic meaning. So, the EGARCH model rests on a solid math base. This careful choice gives a strong starting point for the moving forecasts in the next parts.

4.4 Post-Estimation Residual Diagnostics and Model Adequacy Verification

Post-estimation diagnostics on the standardized residuals are required to verify that the ARIMA (2,0,3)-EGARCH (1,1)-t specification fully captures both linear dependencies and time-varying variance. A well-specified model dictates that these residuals approximate a homoscedastic white noise process, safeguarding against spurious out-of-sample forecasting.

The Ljung-Box Q-test is applied to to evaluate remaining serial correlation, while the Ljung-Box and Engle ARCH-LM tests are executed on squared residuals (Z_t^2) to detect omitted conditional heteroscedasticity [3, 10].

Table 6. Quantitative Diagnostic Results for Standardized Residuals

Target Series	Test Type	Lag (k)	Test Statistic	p-value	Empirical Conclusion
Standardized Residuals ()	Ljung-Box Q (10)	10	5.5715	0.8499	No Serial Correlation
Standardized Residuals ()	Ljung-Box Q (20)	20	8.985	0.9831	No Serial Correlation
Standardized Residuals ()	ARCH-LM Test	5	4.5725	0.4702	Homoscedastic Residuals
Standardized Residuals ()	ARCH-LM Test	10	4.9481	0.8946	Homoscedastic Residuals
Squared Residuals ()	Ljung-Box (10)	10	5.2836	0.8714	ARCH Effects Eliminated
Squared Residuals ()	Ljung-Box (20)	20	6.9549	0.9968	ARCH Effects Eliminated

As indicated in Table 6, all asymptotic p-values ly exceed the 0.05 threshold. For the conditional mean equation, the Ljung-Box Q-statistics at lag 10 () and lag 20 () fail to reject the null hypothesis of no serial correlation, confirming that the ARIMA architecture exhausts linear predictability. Crucially, the Ljung-Box tests and ARCH-LM statistics at

lag 5 () and lag 10 () strongly support the null hypothesis of homoscedasticity. These metrics demonstrate that the specification successfully neutralizes volatility clustering and asymmetric leverage shocks.

Fig. 2 complements the quantitative findings through empirical Autocorrelation Function (ACF) plots.




Fig. 2 ACF

Visual diagnostics from the asset plots provide two primary insights regarding model specification. On one hand, the autocorrelation coefficients for the standardized residuals remain entirely within the 95% confidence intervals across all 20 evaluated lags, confirming that linear residual dependencies have been successfully eliminated. On the other hand, the ACF of the squared residuals demonstrates that the long-memory ARCH effects prominent in the raw index have completely disappeared, as all lag intervals now fall within stochastic boundaries. Taken together, these quantitative indicators and visual tests confirm that

the specified ARIMA (2,0,3)-EGARCH (1,1)-t framework achieves mathematical closure, making it a highly dependable tool for forward-looking risk projections.

4.5 Out-of-Sample Forward-Looking Predictive Projections

Using the optimized specification selected for its structural non-normal and asymmetric capabilities, an out-of-sample multi-step forward simulation is executed to forecast market paths for the next 10 consecutive trading days (to).

Table 7. Forward 10-Day Out-of-Sample Return and Conditional Volatility Forecasts

Projection Time Horizon	Return Point Forecast ()	Conditional Volatility Forecast ()
T+1	-0.001522	0.01655
T+2	-0.000344	0.01679
T+3	-0.000216	0.01702
T+4	-0.002048	0.01724
T+5	-0.000641	0.01746
T+6	-0.000255	0.01766
T+7	-0.001765	0.01787
T+8	-0.000842	0.01806
T+9	-0.000328	0.01825
T+10	-0.001535	0.01844

The forecasting results presented in Table 7 indicate that the predicted conditional mean returns hover marginally below zero across the entire projection horizon. Meanwhile, the projected conditional volatility shows a steady upward trend, rising from about 0.017 at T+1 to 0.018 at T+10. This continuous rise implies that near-term market uncertainty is anticipated to intensify. While these multi-step simulations capture the model's theoretical forecast-

ing capacity based on historical parameters, assessing its real-world viability requires validation against actual out-of-sample data.

4.6 Dynamic Rolling Evaluation

To verify the operational reliability and predictive stability of the proposed specification independently of static

in-sample metrics, out-of-sample robustness is tested across the 2026 testing set comprising 97 independent trading days. A recursive dynamic rolling forecast simulation is implemented, expanding the estimation window by one observation at each increment to predict the immedi-

ate one-step-ahead log return ($\hat{r}_t$). The tracking performance of the hybrid framework is evaluated against a linear baseline model using RMSE and MAE [11]. Table 8 presents the performance comparison.

Table 8. Out-of-Sample Rolling Forecast Robustness Performance Matrix (2026)

Model	RMSE	MAE
ARIMA (2,0,3)	0.01789397	0.01415633
ARIMA-EGARCH-t	0.01777664	0.01407706

As summarized in Table 8, the proposed framework achieves lower error metrics across both dimensions compared to the standard linear ARIMA model, reducing the rolling forecasting RMSE to about 0.018 and compressing the MAE to about 0.014. This reduction demonstrates that capturing heteroscedasticity volatility clustering and non-normal thick tails allows the hybrid specification to effectively filter out asset noise, preventing the conditional mean equation over-reacting to extreme innovations and verifying its robustness.

5. Volatility Mechanisms and Forward-Looking Risk Discussion

The kurtosis score is higher than the normal score. This proves that HSTECH have very heavy tails and extreme price drops. This happens because fast US central bank rate hikes cut down global money, but hype around AI brought in speculative money. This clash between tight money and AI hype caused frequent, huge price changes. This shows that rare and extreme "black swan" events are a normal part of the market [9]. So, standard simple models underestimate these extreme risks, and this leaves tech portfolios in big danger when the market drops.

ARCH-LM tests reject the constant risk, and this proves that high-risk periods group together. Globally, risk does not disappear quickly, so high risk lasts a long time. And trade rule disputes, political tensions, and supply chain changes since 2021 have slowly affected the offshore market [10]. Also, Hong Kong's open financial market is filled with sensitive big investors. These investors often copy each other because they do not have clear information about cross-border rules. When market shocks happen, automatic selling tools and computer trades stop the market from reacting normally to news. This turns outside shocks into long periods of high risk [8].

The EGARCH model score shows a strong leverage effect, which means bad news raises risk much more than good news. This shows that investors are more sensitive to bad news. Short-selling can stop price bubbles when the market is growing, but it causes a bad loop of cash short-

ages and forced selling when the market crashes. Quick short-selling makes prices fall faster, and this leads to the huge risk jumps shown in the EGARCH model.

New tests show that tracking changing risk gives lower error scores (RMSE and MAE) than a simple ARIMA model, and this makes guesses more stable. But the combined ARIMA-EGARCH model shows lasting risk and news reactions immediately [11, 12, 13].

6. Conclusion

The HSTECH index has returns that are not normal, with big price drops and times where high risk stays together. And bad news in this market causes much bigger risk jumps than good news. Simple models assume that risk stays the same, so they always overlook big market drops. Simple GARCH models fit the data well under a Student-t setup, but they treat all news the exact same way. But the EGARCH model is better because it tracks the real, unequal reactions to news. So, a combined ARIMA (2,0,3)-EGARCH (1,1)-t model is used to capture these risk changes. This combined model tracks risk much better than a basic ARIMA model, and it cuts down errors like RMSE and MAE. Adding this changing risk into the math setup makes guesses much more stable. And this gives a safer way to track risk and calculate VaR during highly risky times.

But this study has clear limits. It only uses daily closing prices, so it cannot see fast price changes during the day or sudden cash drops. And the model only looks at past prices without looking at outside economic facts. Future studies can build on this by using hour-by-hour data, looking at outside rules, and using smart computer tools to make risk guesses even better.

Future research can build on these findings in several ways. First, using high-frequency data can help capture sudden price jumps during trading hours. Second, including external facts, such as interest rate changes, policy updates, and market news sentiment, which can improve model accuracy. Finally, combining EGARCH models with smart machine learning tools can further increase prediction power and help build better risk management

tools.

References

- [1] Zhang L. The return rate prediction of China's CSI 300 index based on the ARIMA model. In SHS Web of Conferences, 2025, 225: 02011, EDP Sciences.
- [2] Bollerslev T. Generalized autoregressive conditional heteroscedasticity. *Journal of Econometrics*, 1986, 31(3): 307-327.
- [3] Bollerslev T., Engle R. F., Nelson D. B. ARCH models. *Handbook of Econometrics*, 1994, 4: 2959-3038.
- [4] Nelson D. B. Conditional heteroscedasticity in asset returns: A new approach. *Econometrica: Journal of the Econometric Society*, 1991, 347-370.
- [5] Brandt M. W., Jones C. S. Volatility forecasting with range-based EGARCH models. *Journal of Business & Economic Statistics*, 2006, 24(4): 470-486.
- [6] Lai Y. Volatility Analysis of Hong Kong Stock Hang Seng Index Based on GARCH Model. In 3rd International Conference on Internet Finance and Digital Economy, ICIFDE, 2023, 92-104, Atlantis Press.
- [7] Liu W., Morley B. Volatility forecasting in the hang seng index using the GARCH approach. *Asia-Pacific Financial Markets*, 2009, 16(1): 51-63.
- [8] Amirshahi B., Lahmiri S. Hybrid deep learning and GARCH-family models for forecasting volatility of cryptocurrencies. *Machine Learning with Applications*, 2023, 12: 100465.
- [9] Xu Z., Liechty J., Benthall S., et al. GARCH-informed neural networks for volatility prediction in financial markets. In *Proceedings of the 5th ACM International Conference on AI in Finance*, 2024, 600-607.
- [10] Phung Duy Q., Nguyen Thi O., Le Thi P. H., et al. Estimating and forecasting bitcoin daily prices using ARIMA-GARCH models. *Business Analyst Journal*, 2024, 45(1): 11-23.
- [11] Khang, V. H., Thuan, N. D. Forecast of Stock Prices with Arima, Rolling Forecast, and Garch: A Dynamic Approach in the Fluctuating Market. *Journal of Information Hiding and Multimedia Signal Processing*, 2025, 16: 3.
- [12] Nicholas N. The black swan: The impact of the highly improbable. *Journal of the Management Training Institut*, 2008, 36(3): 56.
- [13] Li W. J., Zhang D. Q. GARCH-FIS: A Hybrid Forecasting Model with Dynamic Volatility-Driven Parameter Adaptation, 2026, arXiv preprint arXiv:2603.14793.