Positive Science

Positive Science - Economics

Volatility Dynamics and Forecasting of the Nepal Stock Exchange Index: A Comparative Analysis of GARCH and EGARCH Models

Nischal Shrestha
Independent Researcher

Jeevan Pokhrel
Tribhuvan University

Publication Information

DOI: 10.65957/article.18

Positive Science - Economics, 2026, Article 18.

https://doi.org/10.65957/article.18

Citation: Article 18

Published: August 11, 2026

Publisher location: Strasbourg, France

Volatility Dynamics and Forecasting of the Nepal Stock Exchange Index: A Comparative Analysis of GARCH and EGARCH Models

ABSTRACT

This study examines the volatility dynamics and forecasting performance of the Nepal Stock Exchange (NEPSE) Index using 1,149 daily log-return observations derived from closing prices covering 30 May 2021 to 15 May 2026. Following tests for stationarity and conditional heteroskedasticity, symmetric GARCH (1,1) and asymmetric EGARCH (1,1) models were estimated with standardized Student-t innovations. A GJR-GARCH (1,1) model was additionally estimated as a robustness check for asymmetry. Model performance was evaluated using information criteria, diagnostic tests, fixed-window one-step-ahead forecasts, volatility half-life estimates, a News Impact Curve, and structural-break analysis. NEPSE returns exhibited non-normality, heavy tails, volatility clustering, and substantial volatility persistence. In-sample rankings varied by criterion: GJR-GARCH produced the highest log-likelihood and lowest AIC, whereas GARCH produced the lowest BIC. Out-of-sample results were also mixed: GARCH generated a slightly lower RMSE and higher Mincer–Zarnowitz R^2, whereas EGARCH produced a marginally lower MAE. The EGARCH signed-shock coefficient was negative but statistically insignificant, while the News Impact Curve was only directionally consistent with conventional leverage. Although the significant GJR-GARCH asymmetry coefficient provided supporting evidence, the leverage effect remained model-dependent rather than conclusive. A structural break was identified on 13 July 2022, after which both models exhibited lower persistence and shorter volatility half-lives. Overall, NEPSE volatility is persistent and regime-dependent, while the forecasting performances of GARCH and EGARCH are broadly comparable across different loss measures. The findings demonstrate the importance of evaluating model fit, forecasting accuracy, asymmetry, and parameter stability separately when assessing volatility in a frontier equity market.

Keywords: Nepal Stock Exchange (NEPSE); Volatility Modelling; GARCH; EGARCH; Volatility Forecasting; News Impact Curve; Structural Breaks.

INTRODUCTION

The financial market serves as an intermediary for mobilizing savings, allocating resources, and facilitating economic growth. The efficient functioning of these markets depends not only on the returns generated but also on the degree of uncertainty associated with those returns. In financial economics, this uncertainty is commonly measured by volatility, which reflects the magnitude and frequency of fluctuations in the price of the assets over time. Because volatility is closely associated with financial risk, understanding and forecasting volatility has become an important area of research for investors, portfolio managers, policy makers, and financial institutions.

Financial returns exhibit several features that depart from the classical assumptions of constant variance and normality. Among these characteristics, volatility clustering is one of the most prominent, whereby large price changes tend to be followed by larger price changes and small changes tend to be followed by small changes (Mandelbrot, 1963). Such characteristics imply that volatility is time-varying rather than constant. To address this, Engle (1982) introduced the Autoregressive Conditional Heteroskedasticity (ARCH) model, which allows current volatility to depend on past forecast errors. Bollerslev (1986) later extended this framework with the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model by including lagged conditional variances, providing a more effective approach to modeling persistent volatility. Since then, ARCH and GARCH models have become the standard econometric tools for examining volatility dynamics in financial time series.

In addition to volatility persistence, financial markets often display asymmetric volatility responses to information shocks. Negative and positive news of equal magnitude may generate different effects on future volatility, a phenomenon commonly referred to as the leverage effect (Black, 1976). To capture such asymmetries, Nelson (1991) proposed the Exponential GARCH (EGARCH) model, which allows volatility to respond differently to the sign and magnitude of market shocks. Understanding whether such asymmetries exist is important because they provide insight into investor behavior and market risk dynamics.

Although extensive evidence exists for developed financial markets, empirical findings from frontier and emerging markets remain relatively limited. These markets often exhibit unique characteristics, such as lower liquidity, higher information asymmetry, concentrated investor participation, and greater sensitivity to economic and political developments. Consequently, it remains unclear whether established volatility models adequately capture market behavior in such environments.

The Nepal Stock Exchange (NEPSE) is Nepal’s only stock exchange, as indicated by the market-infrastructure information published by the Securities Board of Nepal (SEBON, 2026). During the study period, the NEPSE Index experienced pronounced upward and downward movements, highlighting the importance of understanding its volatility dynamics. Although previous studies have examined volatility clustering and persistence in NEPSE returns, evidence concerning directional asymmetry and comparative out-of-sample forecasting performance remains limited and mixed (Thapa & Gautam, 2016; Gajurel, 2021; Dangal & Gajurel, 2021). Furthermore, relatively few studies combine recent market data, asymmetric-model robustness checks, rolling out-of-sample forecasting, and structural-break analysis within a single empirical framework.

Although previous studies have examined NEPSE forecasting using models such as ARIMA (Shrestha, 2026), relatively limited attention has been given to forecasting conditional volatility through ARCH-family models. Consequently, evidence regarding volatility persistence, asymmetric responses to market shocks, and out-of-sample volatility forecasting performance in the Nepalese stock market remains limited. To address these gaps, this study employs ARCH-family models to examine the volatility dynamics of the NEPSE index using daily observations from 30 May 2021 to 15 May 2026. Specifically, the study estimates symmetric and asymmetric volatility models, evaluates their forecasting performance using an out-of-sample rolling-window approach, and investigates volatility shock behavior in the Nepalese stock market. By doing so, the study contributes to the growing literature on volatility modeling in frontier markets and provides evidence relevant to investors, policymakers, and market participants.

The main objectives of this study are:

1. To examine the presence of volatility clustering and volatility persistence in NEPSE returns using ARCH-family models.
2. To investigate whether volatility responds asymmetrically to positive and negative market news through the EGARCH framework and News Impact Curve Analysis.
3. To evaluate and compare the out-of-sample forecasting performance of GARCH and EGARCH models using rolling-window volatility forecasts.

Literature Review

The theoretical foundation of this study draws on three complementary strands of volatility literature. First, the ARCH/GARCH framework of Engle (1982) and Bollerslev (1986) provides the basis for modeling volatility clustering, the central stylized fact motivating conditional heteroskedasticity models. Second, the leverage-effect literature originating with Black (1976) and formalized by Nelson (1991) through the EGARCH model addresses the possibility that volatility responds asymmetrically to the sign of return shocks; this is empirically tested in the present study via the News Impact Curve Framework of Engle and Ng (1993). Third, the structural break literature of Bai and Perron (2003) addresses the concern that volatility persistence estimated over a full sample period may mask regime-specific shifts in market behavior. Together, these three strands correspond directly to the three central empirical questions of this study: whether volatility clusters and persists, whether it responds asymmetrically to shocks, and whether its dynamics are stable across the sample period.

Asymmetric volatility and the leverage effect

The standard GARCH model assumes that positive and negative shocks affect volatility equally. However, the financial market exhibits asymmetric volatility, in which negative news tends to generate a larger increase in future volatility than positive news of the same magnitude, a pattern commonly referred to as the leverage effect (Black, 1976).

Nelson (1991) addressed this with the Exponential GARCH (EGARCH) model, which specifies the logarithm of conditional variance as a function of both the magnitude and the sign of past standardized shocks, allowing positive and negative shocks to exert different effects on volatility while also guaranteeing a positive conditional variance without parameter restrictions. Engle and Ng (1993) formalized the empirical evaluation of these asymmetric effects through the News Impact Curve, which plots the relationship between a standardized shock and the resulting conditional variance, holding lagged variance at its unconditional level, providing a direct visual and quantitative test of whether, and how strongly, bad news moves volatility more than good news of equal size.

Volatility forecasting and structural breaks

An important objective of volatility modeling is forecasting future market risk, besides explaining historical volatility. Engle and Patton (2001) argued that the practical value of a volatility model should be assessed through its out-of-sample forecasting performance rather than solely through in-sample goodness-of-fit measures. Another issue in volatility modeling is the presence of structural breaks. To resolve this issue, Bai and Perron (2003) developed a multiple structural break framework that allows researchers to identify changes in volatility regimes without imposing predetermined break dates.

Empirical Evidence

Evidence from other emerging and South Asian equity markets similarly indicates that volatility dynamics are persistent but not uniform across markets or model specifications. De Santis and Imrohoroglu (1997) documented volatility clustering, predictability, and persistence in emerging financial markets, together with a higher probability of large price movements than in developed markets. In India, Karmakar (2005) found strong time-varying volatility and persistence and reported that the GARCH (1,1) model produced reasonably informative out-of-sample forecasts, although significant leverage effects were detected for only a minority of the examined companies. In contrast, Goudarzi and Ramanarayanan (2011) found a clear asymmetric response in the BSE 500 Index, with negative news increasing volatility more strongly than positive news. These mixed findings demonstrate that volatility persistence, asymmetry, and forecasting performance may vary across emerging markets, supporting the present study's comparison of symmetric and asymmetric specifications in the Nepalese market.

One of the earlier studies on NEPSE Volatility was conducted by Thapa and Gautam (2016), who examined daily stock returns from July 1997 to December 2012 using GARCH-Family models. Their findings revealed significant volatility clustering and persistence in NEPSE returns, indicating that periods of high volatility tend to be followed by further periods of high volatility. This study also reported asymmetric responses to market shocks, suggesting that positive and negative news do not affect volatility equally.

Similarly, Gajurel (2021) investigated asymmetric volatility and international market influences on the Nepalese Stock Exchange. Using an asymmetric GARCH specification, the study confirms the presence of persistent volatility in NEPSE returns and provides evidence that the volatility behavior of Nepal is influenced by both domestic market conditions and external financial developments. The findings highlighted the importance of considering asymmetric volatility structures when modeling risk in the Nepalese stock market.

Dangal and Gajurel (2021) further examined volatility dynamics using several GARCH-family models, including GARCH-M, TGARCH, EGARCH, and PGARCH. Their results confirmed the existence of volatility clustering, the leverage effect, and highly persistent conditional variance in NEPSE returns. Interestingly, while the symmetric GARCH model performed better for the full sample period, the asymmetric specification produced superior results when the sample was divided into different market regimes. This finding suggested that the volatility behavior of NEPSE may change across market conditions and structural events.

More recently, Kadel and Tiwari (2025) analyzed NEPSE volatility using daily observations from 2020 to 2025 and compared several ARCH-family models. Their findings confirmed continued volatility clustering and persistence in the Nepalese stock market, with the Power ARCH (PARCH) specification providing the best overall fit among the competing models. The study emphasized that volatility remains a dominant feature of NEPSE returns and continues to have important implications for investors and regulators.

Collectively, previous studies consistently report volatility clustering and persistence in the Nepal Stock Exchange. However, findings regarding volatility asymmetry remain inconclusive, largely due to differences in sample periods, model specifications, and methods of interpretation. Moreover, relatively few studies have evaluated out-of-sample forecasting performance or examined structural changes in volatility dynamics. These inconsistencies highlight the need for a comprehensive assessment that combines model estimation, forecasting evaluation, asymmetry analysis, and structural break testing within a single analytical framework.

Although previous NEPSE studies provide substantial evidence of volatility clustering and persistence, comparatively limited evidence is available on out-of-sample volatility forecasting and break-adjusted persistence. This limits direct comparison of the present structural-break and forecasting results with earlier Nepalese studies and constitutes an area for further investigation.

Methodology

Research Design

This study adopts a quantitative research design to examine the volatility dynamics and forecasting performance of the Nepal Stock Exchange (NEPSE) index. Time-Series econometric techniques are employed to investigate the presence of volatility clustering, persistence, asymmetric volatility responses, and structural changes in market volatility. This study utilizes ARCH-family models, including the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) and Exponential GARCH (EGARCH) frameworks, which are widely used in financial econometrics for modeling time-varying volatility.

In addition to volatility estimation, the study evaluates forecasting performance through an out-of-sample rolling window framework. This approach allows the practical forecasting ability of competing models to be assessed under conditions that closely resemble real-world forecasting environments.

Data Source and Sample Period

The study is based on secondary data consisting of the daily closing value of the Nepal Stock Exchange (NEPSE) Index, obtained from the Nepsealpha website (https://nepsealpha.com/nepse-data). Nepsealpha was used because it provides a downloadable historical daily series; however, possible discrepancies relative to the exchange's official records are acknowledged as a data limitation. The data cover the period from 30 May 2021 to 15 May 2026, yielding 1150 daily price observations and 1149 daily log-return observations after first-differencing. The data were retrieved on 20 May 2026. Daily observations were selected because daily-frequency data are more suitable for capturing short-run fluctuations in volatility and provide sufficient observations for estimating ARCH-family models. The sample period was selected to provide a recent five-year window containing sufficient daily observations for reliable GARCH estimation and rolling out-of-sample forecasting. It also encompasses contrasting market conditions and volatility regimes, making it suitable for examining volatility persistence, asymmetry, forecasting performance, and structural change in the NEPSE Index.

Preliminary Data Analysis

Before estimating volatility models, several preliminary analyses were conducted to examine the statistical properties of the return series. Descriptive statistics, including the mean, standard deviation, skewness, kurtosis, minimum value, maximum value, and the Jarque-Bera test, were calculated to assess the distribution characteristics of returns (Jarque & Bera, 1980). These measures provide insights into the central tendency, dispersion, normality, and tail behavior of the data.

To determine whether the series satisfied the stationarity requirement necessary for time-series modeling, the Augmented Dickey-Fuller (ADF) test was performed on both the price series and the return series (Dickey & Fuller, 1979; Said & Dickey, 1984).

H_o: The series contains a unit root (non-stationary)

H_1: The series is stationary

Rejection of the null hypothesis at conventional significance levels indicates that the series is stationary and suitable for time-series modeling.

Mean Equations Specification

Proper specification of the conditional mean equation is essential because misspecification may affect volatility estimates. The Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) were initially examined to identify serial dependency in returns. Subsequently, the Auto Arima procedure was employed to determine the most appropriate mean equation based on information criteria and diagnostic performance.

The selected mean specification was incorporated into all subsequent ARCH-family models to ensure that volatility estimates reflected genuine conditional heteroskedasticity rather than unmodeled autocorrelation.

Testing ARCH Effects

Before estimating volatility models, the presence of conditional heteroskedasticity was examined using the ARCH-Lagrange Multiplier (ARCH-LM) test proposed by Engle (1982). The ARCH-LM test evaluates whether past squared residuals contain information about current variance.

Formally, the ARCH-LM test is specified as:

H_o: $\alpha_1=\alpha_2=\ldots=\alpha_q=0$ (no ARCH effect; constant conditional variance)

H_1: At least one $\alpha_i\neq0$ (ARCH effect present; time-varying conditional variance)

Rejection of the null hypothesis indicates the presence of volatility clustering and provides the statistical justification for estimating ARCH-family models.

Volatility Model

The ARCH Model introduced by Engle (1982) assumes that current conditional variance depends on past squared error terms. ARCH models capture volatility clustering by allowing periods of high volatility to follow previous large shocks. ARCH specifications with different lag structures were estimated and compared to identify the most appropriate model.

To capture longer-term volatility persistence more efficiently, the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model developed by Bollerslev (1986) was estimated. The GARCH (1,1) specification is expressed as:

$$h_t = \omega + \alpha\varepsilon_{t-1}^2 + \beta h_{t-1} \quad (1)$$

where,

h_t: Conditional Variance at time t,

ω: Constant term,

ε^2_{t-1}: Squared residual (shock) from previous period,

α: Impact of recent shocks on volatility (The ARCH effect)

β: Volatility persistence (The GARCH effect)

$\alpha+\beta$: Overall persistence of volatility shocks

The GARCH (1,1) model was selected as the benchmark symmetric specification because it provides a parsimonious representation of volatility clustering and persistence and has demonstrated strong empirical and forecasting performance in financial markets (Bollerslev et al., 1992; Hansen & Lunde, 2005). Its relevance to the present setting is also supported by applications to Indian and Nepalese equity returns, where GARCH (1,1) has effectively captured time-varying and persistent volatility (Karmakar, 2005; Dangal & Gajurel, 2021).

The AR (3) mean equation was incorporated into both volatility models. Both GARCH (1,1) and EGARCH (1,1) were estimated using standardized Student-t innovations to account for the heavy-tailed distribution of NEPSE returns and to ensure that their log-likelihood and information criteria were directly comparable.

To investigate asymmetric volatility behavior, the Exponential GARCH (EGARCH) model proposed by Nelson (1991) was estimated. Unlike the standard GARCH model, EGARCH allows positive and negative shocks to affect future volatility differently.

The EGARCH specification models the logarithm of conditional variance and therefore guarantees a positive variance without imposing parameter restrictions. This framework is particularly useful for examining the leverage effects, where adverse market news generates a stronger volatility response than favorable news of equal magnitude.

$$\ln(h_t) = \omega + \beta \ln(h_{t-1}) + \alpha \frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} + \gamma \left(\frac{|\varepsilon_{t-1}|}{\sqrt{h_{t-1}}} - E\left[\frac{|\varepsilon_{t-1}|}{\sqrt{h_{t-1}}} \right] \right) \quad (2)$$

In the `rugarch` EGARCH parameterization, α is the signed-shock coefficient, γ captures shock magnitude, and β measures volatility persistence.

The study focuses on GARCH (1,1) and EGARCH (1,1) because they represent the benchmark symmetric and asymmetric volatility models widely used in empirical finance and permit a direct comparison between volatility persistence and asymmetric market responses while maintaining model parsimony.

Estimation and inference

All ARCH-family models were estimated in R version 4.6.0 using the 'rugarch' package (version 1.5-5) and its hybrid optimization solver. Parameters were estimated by maximum likelihood under standardized Student-t innovations. Statistical inference is based on the White (1982) robust sandwich covariance estimator reported by 'rugarch'; consequently, Tables 5 and 7 present robust standard errors and corresponding two-sided p-values.

Forecasting Framework

To evaluate forecasting performance, the dataset was divided into training and testing samples using an 80:20 split. The training samples were used for model estimation, while the remaining observations were reserved for out-of-sample evaluation.

A rolling-window forecasting procedure was implemented to replicate practical forecasting conditions. Under this approach, model parameters were continuously re-estimated as new observations became available, and one-step-ahead volatility forecasts were generated throughout the testing period. This procedure provides a more realistic assessment of forecasting performance than a single static estimation.

Forecast Evaluation Measures

Forecast accuracy was evaluated using three complementary measures. The Root Mean Square Error (RMSE) was employed to measure the average magnitude of forecast errors while assigning greater weight to larger errors. The Mean Absolute Error (MAE) was used to assess the average absolute deviation between forecasted and realized volatility.

The out-of-sample evaluation used a fixed-length rolling-window procedure. Of the 1,149 return observations, the first 919 observations (80%) formed the initial estimation window, while the remaining 230 observations (20%) constituted the out-of-sample evaluation period. Both models were re-estimated after every new observation, with the oldest observation removed so that the estimation window remained fixed at 919 observations. A one-step-ahead volatility forecast was generated at each iteration, producing 230 out-of-sample forecasts.

In addition, the Mincer-Zarnowitz regression was estimated to evaluate forecast efficiency and explanatory power. The regression is specified as:

$$RV_t = a + \beta FV_t + \varepsilon_t \quad (3)$$

Where RV_t is the realized volatility proxy (absolute return), and FV_t is the model-forecasted conditional volatility at time t. Under the joint null hypothesis of forecast unbiasedness (a=0, b=1), the R^2 indicates the proportion of variation in realized volatility explained by the forecast.

Volatility Half-Life

To provide an economically meaningful interpretation of volatility persistence, the volatility half-life was calculated. The half-life measure indicates the number of trading days required for half of a volatility shock to dissipate.

For the GARCH model, the half-life was calculated using:

$$HL = \frac{\ln(0.5)}{\ln(\alpha+\beta)} \quad (4)$$

For EGARCH, persistence is governed solely by the autoregressive coefficient on log conditional variance, β, rather than by $\alpha+\beta$ as in the symmetric GARCH case:

$$HL = \frac{\ln(0.5)}{\ln(\beta)} \quad (5)$$

A large Half-Life value indicates more persistent volatility and slower adjustment toward long-run equilibrium.

News Impact Curve

To examine the nature of volatility asymmetry, a News Impact Curve (NIC) was constructed based on the estimated EGARCH model. The News Impact Curve illustrates how positive and negative shocks of varying magnitudes affect future conditional volatility while holding other factors constant. This approach provides a direct visual assessment of whether unfavorable market news generates stronger volatility responses than favorable news of similar magnitude

Structural Break Analysis

Volatility dynamics may change over time due to shifts in market conditions, regulatory developments, or major economic events. To examine the stability of volatility behavior throughout the sample period, structural break analysis was performed using the Bai and Perron (2003) multiple breakpoint framework.

The procedure identifies statistically significant breakpoints in the variance process without requiring break dates to be specified in advance. Detecting structural breaks provides additional insight into whether volatility persistence remains stable across different market regimes and helps contextualize the interpretations of full-sample volatility estimates.

Results

This section presents the empirical findings of the study. The analysis begins by examining the statistical properties of the NEPSE return series before evaluating stationarity and conditional heteroskedasticity. Subsequently, ARCH-family models are estimated and compared, followed by model diagnostics, forecasting performance, volatility persistence, asymmetric volatility analysis, and structural break identification.

Characteristics of the NEPSE Return Series

Figure 1: Characteristics of the NEPSE Return Series

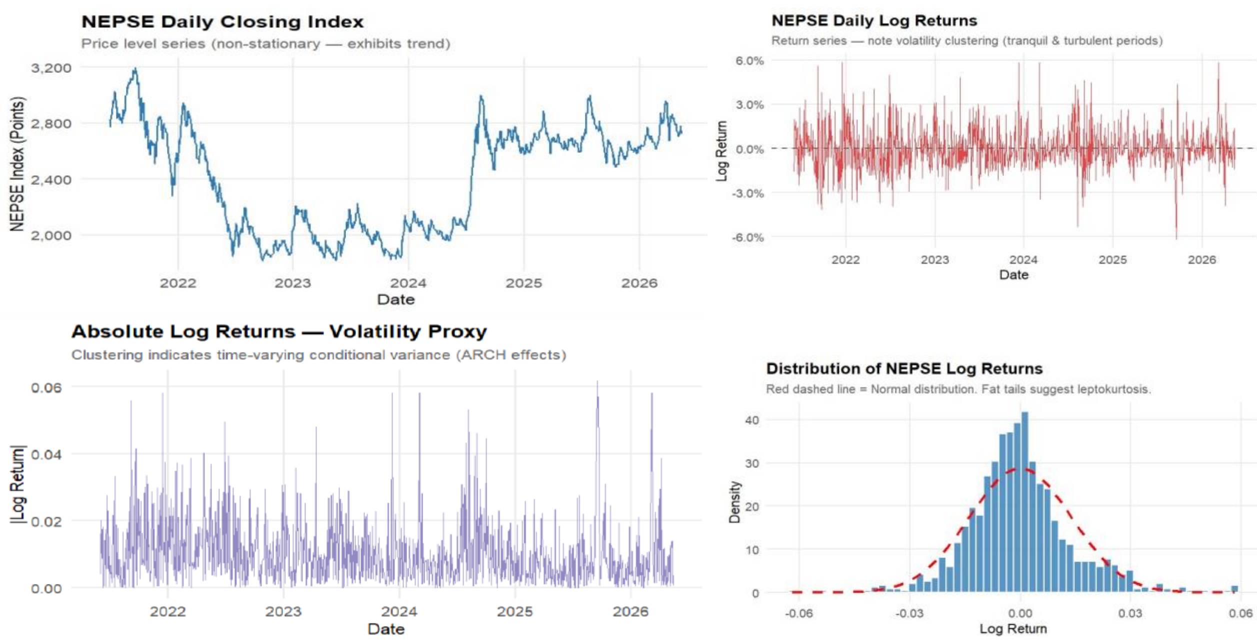

Source: Data Analysis Output, 2026.

Table 1: Descriptive Statistics of Daily NEPSE Log Returns

Statistic	Value
Observations	1149
Mean	-0.00003
Median	-0.00107

Std. Deviation	0.013882
Minimum	-0.06188
Maximum	0.058335
Skewness	0.4917
Excess Kurtosis	2.0621
Jarque-Bera Statistic	249.8829
JB p-value	P<0.001

Source: Data Analysis Output, 2026.

The statistical characteristics of the NEPSE return series were examined through graphical analysis and descriptive statistics before estimating the volatility models. Figure 1 (a) shows that the NEPSE Index exhibited pronounced upward and downward movements throughout the sample period rather than fluctuating around a constant level, indicating that the price series is non-stationary. After transformation into logarithmic returns, the series oscillates around a mean close to zero (Figure 1 (b)), suggesting that the underlying trend has been removed while preserving the day-to-day fluctuations required for volatility analysis.

The absolute return series presented in Figure 1 (c) reveals that large and small return movements occur in clusters rather than being randomly dispersed over time. Periods of elevated volatility are followed by similarly volatile periods, whereas tranquil market conditions also persist for several consecutive trading days. This visual pattern provides preliminary evidence of volatility clustering, indicating that the volatility of returns changes over time.

The distribution of daily returns, illustrated in Figure 1(d), departs from the normal distribution by exhibiting a more pronounced central peak and heavier tails. This graphical evidence is consistent with the descriptive statistics reported in Table 1. The return series has a mean close to zero (-0.00003) and a standard deviation of 0.01388, indicating modest average returns but noticeable daily fluctuations. The positive skewness (0.4917) indicates that the return distribution has a relatively longer or heavier right tail, suggesting that extreme positive returns tend to be larger than extreme negative returns. Furthermore, the Jarque-Bera statistics (249.8829, $p < 0.001$) reject the null hypothesis of normality, confirming that the return series is non-normally distributed.

Taken together, the graphical evidence and descriptive statistics indicate that NEPSE returns exhibit non-normality, heavy tails, and periods of concentrated volatility. These characteristics provide initial empirical support for modelling the return series using ARCH-family models, which are designed to accommodate time-varying conditional variance.

Stationarity of the Return Series

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

Series	Test Statistics	P-value	Decision
NEPSE Closing Index	-2.3801	0.4174	Non-Stationary
Daily Log Returns	-9.2908	P<0.01	Stationary

Source: Data Analysis Output, 2026.

The stationarity of the NEPSE closing index and its corresponding logarithmic return series was examined using the Augmented Dickey-Fuller (ADF) unit root test. The results are presented in Table 2.

The ADF test indicates that the NEPSE closing index is non-stationary, as the null hypothesis of a unit root cannot be rejected ($p = 0.4174$). In contrast, the logarithmic return series rejects the null hypothesis at the 5% significance level ($p < 0.01$), confirming that the transformed return series is stationary.

Since the return series satisfies the stationarity requirement for ARCH-family modelling, subsequent analyses were conducted using the logarithmic returns.

Mean Equation Identification and ARCH Effects

Figure 2: ACF and PACF of Daily Log Returns

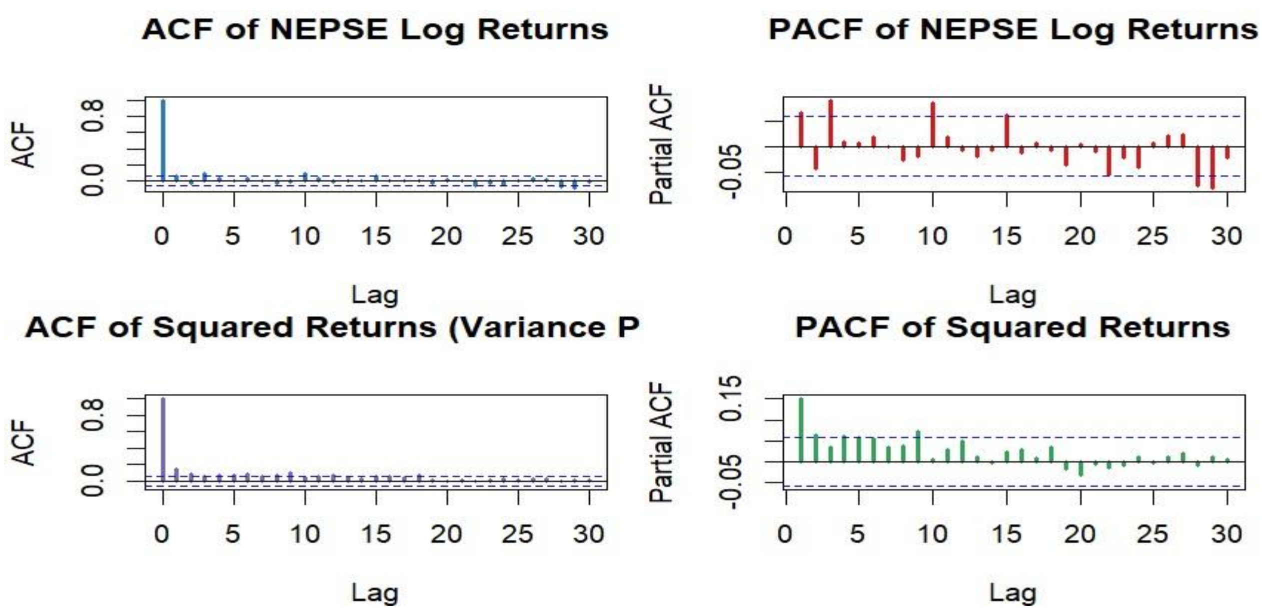

Source: Data Analysis Result, 2026.

Figure 2 presents the ACF and PACF patterns examined during the identification of the conditional mean specification and the assessment of dependence in squared returns.

The conditional mean of the NEPSE return series was specified before estimating the volatility models. Based on the model selection procedure, an AR (3) process was identified as the most appropriate specification for the conditional mean and was subsequently incorporated into all ARCH-family models estimated in this study.

To determine whether the return series exhibited conditional heteroskedasticity, the residuals from the selected mean equation were examined using the ARCH-LM test. The results are presented in Table 3.

Table 3: ARCH-LM Test Results

Lag	Chi-square	p-value	Decision
5	42.2621	<0.001	ARCH effects present
10	52.9475	<0.001	ARCH effects present
15	58.0326	<0.001	ARCH effects present

Source: Data Analysis Output, 2026.

The conditional mean of the NEPSE return series was identified before estimating the volatility models. Automatic model selection indicated that an AR (3) process provided the most appropriate specification for the conditional mean. Accordingly, the AR (3) mean equation was retained for all subsequent ARCH-family model estimations.

Following estimation of the mean equation, the ARCH-LM test was applied to examine the presence of conditional heteroskedasticity in the model residuals. As reported in Table 3, the null hypothesis of no ARCH effects was rejected at all lag lengths considered. The test statistics were significant at lags 5, 10, and 15 ($p < 0.001$), indicating that the variance of the residuals is not constant over time. Having confirmed the presence of conditional heteroskedasticity, the study proceeds to estimate and compare alternative ARCH and GARCH specifications.

Estimation and comparison of ARCH and GARCH Models

Table 4: Comparison of ARCH and GARCH Models

Model	Log-Likelihood	AIC	BIC
ARCH (1)	3361.39	-5.8388	-5.8080
ARCH (2)	3368.11	-5.8488	-5.8136
ARCH (3)	3373.26	-5.8560	-5.8164
GARCH (1,1)	3387.58	-5.8826	-5.8475

Source: Data Analysis Output, 2026.

Table 5: Estimated Parameters of the GARCH (1,1) Model

Parameter	Estimate	Robust p-value
μ	-0.000685	0.079
AR (1)	0.0985	<0.001
AR (2)	-0.0463	0.076
AR (3)	0.0726	0.018
ω	0.000009	<0.001
α_1	0.1287	<0.001
β_1	0.8336	<0.001
Shape (ν)	5.0473	<0.001

Source: Data Analysis Output, 2026.

Note: The model was estimated using standardized Student-t innovations. Robust p-values are calculated using the White (1982) sandwich covariance estimator.

Alternative ARCH-family models were estimated to identify the specification that best captures the conditional volatility of NEPSE returns. The estimated models include ARCH (1), ARCH (2), ARCH (3), and GARCH (1,1), each fitted using the selected AR (3) mean equations. Their comparative performance is presented in Table 4.

In the mean equation, μ ($p = 0.079$) and AR (2) ($p = 0.076$) were not statistically significant at the 5% level, whereas AR (1) and AR (3) were significant. Among the estimated models, the GARCH (1,1) specification achieved the highest loglikelihood (3387.58) together with the lowest AIC = -5.8826 and BIC = -5.8475. Based on these model selection criteria, the GARCH (1,1) model was selected as the preferred specification for modeling the conditional volatility of NEPSE return.

The estimated parameters of the selected GARCH (1,1) model are reported in Table 5. The variance equation parameters, ω, α_1, and β_1, are all positive and statistically significant at the 1%

significance level. The estimated model was therefore retained for further evaluation through diagnostic testing, forecasting assessment, and analysis of volatility dynamics.

Diagnostic Evaluation of the Selected GARCH Model

Table 6: Diagnostic Tests for the Selected GARCH (1,1) Model

Diagnostic Test	Test Statistics	P-value	Decision
Ljung-Box (Residuals)	0.5484	0.459	No serial Correlation
Ljung-Box (Squared Residuals)	1.3590	0.244	No remaining Autocorrelation
ARCH-LM	0.7879	0.375	No remaining ARCH effects
Sign Bias Test	2.990	0.0029	Significant
Joint Sign Bias	9.319	0.0253	Significant

Source: Data Analysis Output, 2026.

The adequacy of the selected GARCH (1,1) model was evaluated using a series of postestimation diagnostic tests. The results are summarized in Table 6.

The weighted Ljung–Box test applied to the standardized residuals failed to reject the null hypothesis of serial independence, indicating that no significant autocorrelation remained in the residual series. Similarly, the Ljung–Box test on the standardized squared residuals was statistically insignificant, suggesting that the model adequately captured the dependence in conditional variance (Ljung & Box, 1978).

The weighted ARCH-LM test also failed to reject the null hypothesis of no remaining ARCH effects, indicating that the estimated GARCH (1,1) model successfully accounted for the conditional heteroskedasticity present in the return series.

The Sign Bias Test, however, was statistically significant, and the joint sign bias test likewise rejected the null hypothesis. These results indicate that the selected GARCH (1,1) model does not fully account for asymmetric responses of volatility to positive and negative shocks. Consequently, an asymmetric volatility specification was estimated in the subsequent analysis.

Asymmetric Volatility Modeling

Table 7: Estimated Parameters of the EGARCH (1,1)

Parameter	Estimate	Robust p-value
Mean equation		
μ	−0.0009	< 0.001
AR(1)	0.0934	< 0.001
AR(2)	−0.0477	< 0.001

Parameter	Estimate	Robust p-value
AR(3)	0.0772	0.026
Variance equation		
ω	−0.4849	< 0.001
α	−0.0379	0.146
β	0.9438	< 0.001
γ	0.2611	< 0.001
Shape (ν)	5.1415	< 0.001

Source: Data Analysis Output, 2026.

Note: The model was estimated using standardized Student-t innovations. Robust p-values are calculated using the White (1982) sandwich covariance estimator

The significant Sign Bias Test reported in the previous section suggested that the symmetric GARCH (1,1) model did not fully capture the asymmetric response of volatility to market shocks. To account for this behavior, an Exponential GARCH (EGARCH (1,1)) model was estimated.

Table 7 presents the variance-equation estimates of the EGARCH (1,1) model. The signed-shock coefficient is negative (α = -0.0379), which is directionally consistent with the conventional leverage effect. However, the coefficient is statistically insignificant under robust inference ($p=0.146$); therefore, the null hypothesis of no directional asymmetry cannot be rejected. In contrast, the shock-magnitude coefficient is positive and statistically significant (γ = 0.2611, $p<0.001$), indicating that larger innovations generate stronger volatility responses regardless of their sign. The persistence coefficient is also large and significant (β = 0.9438, $p<0.001$), demonstrating persistent conditional volatility. Overall, the results show significant magnitude dependence and strong persistence but do not establish a statistically significant leverage effect.

Robustness Check Using GJR-GARCH (1,1)

To examine the robustness of the asymmetry findings, a GJR-GARCH (1,1) model was estimated using the same AR (3) mean equation and standardized Student-t innovations (Glosten et al., 1993). The asymmetry coefficient was positive and statistically significant under robust inference (γ = 0.0986, p = 0.026), indicating that negative shocks generated a stronger volatility response than positive shocks of equal magnitude. However, because the EGARCH signed-shock coefficient was insignificant and the GJR-GARCH stability and joint sign-bias diagnostics indicated some remaining specification sensitivity, the evidence of conventional leverage is treated as suggestive and model-dependent rather than conclusive.

Table 8: GJR-GARCH(1,1) Robustness-Check Estimates

Parameter	Estimate	Robust p-value
μ	−0.000850	0.038

Parameter	Estimate	Robust p-value
AR(1)	0.093498	<0.001
AR(2)	−0.044023	0.086
AR(3)	0.069268	0.023
ω	0.000010	<0.001
α	0.092720	<0.001
β	0.826981	<0.001
γ	0.098614	0.026
Shape (ν)	5.121021	<0.001

Source: Data Analysis Output, 2026.

Note: The model was estimated using an AR(3) mean equation and standardized Student-t innovations. In the GJR-GARCH specification, a positive γ indicates a stronger conditional-variance response to negative shocks. Robust p-values are reported.

Comparison of Symmetric and Asymmetric Volatility Models

Table 9: In-Sample Comparison of GARCH (1,1), EGARCH (1,1), and GJR-GARCH (1,1)

Model	Log-likelihood	AIC	BIC
GARCH(1,1)	3387.580	−5.8826	−5.8475
EGARCH(1,1)	3389.730	−5.8847	−5.8451
GJR-GARCH(1,1)	3390.569	−5.8861	−5.8466

Source: Data Analysis Output, 2026.

Following the estimation of both symmetric and asymmetric volatility models, their in-sample performance was compared using the log-likelihood, Akaike Information Criterion (AIC), and Bayesian Information Criterion (BIC). The comparative results are presented in Table 9.

The EGARCH (1,1) model achieved a higher log-likelihood (3389.73) and a lower AIC (-5.8847) than the GARCH (1,1) model, indicating a marginal improvement in in-sample fit. In contrast, the GARCH (1,1) model recorded a slightly lower BIC (-5.8475) than the EGARCH (1,1) model (-5.8451), reflecting a stronger penalty for the additional parameter included in the asymmetric specification.

Both models provided an adequate representation of the conditional volatility of NEPSE returns. Their forecasting performance was subsequently evaluated using out-of-sample accuracy measures.

The GJR-GARCH (1,1) model achieved the highest log-likelihood and lowest AIC, while the symmetric GARCH (1,1) model retained the marginally lowest BIC. These small differences indicate that no model clearly dominated under every selection criterion. The significant GJR-GARCH asymmetry coefficient provides supporting evidence of conventional leverage, although this conclusion is not confirmed by the statistically insignificant EGARCH signed-shock coefficient.

News Impact Curve Analysis

Figure 3: News Impact Curve from the EGARCH (1,1) model

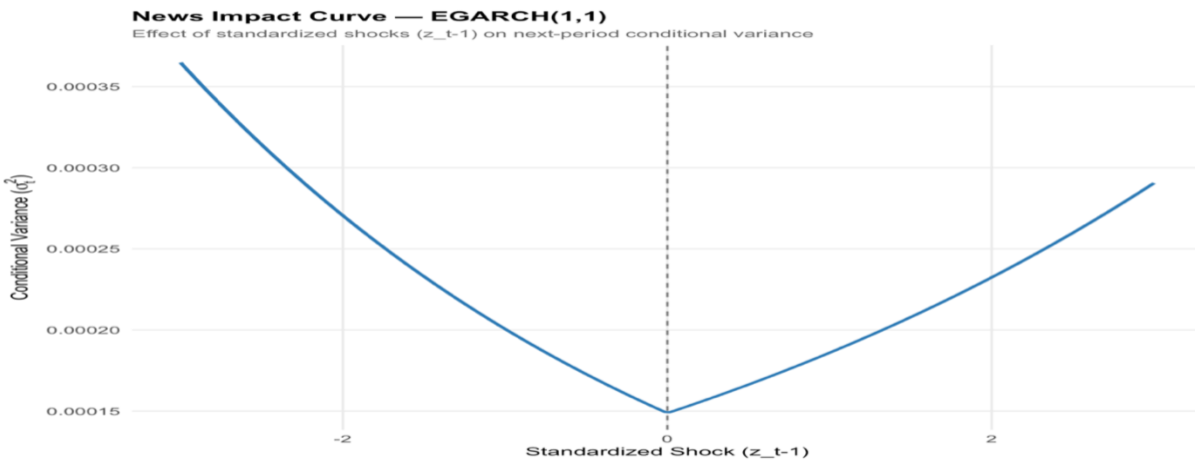

Source: Data Analysis Output, 2026.

Figure 3 presents the News Impact Curve derived from the EGARCH (1,1) estimates. The curve shows that larger shocks increase conditional variance regardless of their sign, while the higher left branch indicates a somewhat stronger estimated response to negative shocks. This pattern is directionally consistent with conventional leverage. However, because the signed-shock coefficient is statistically insignificant, the curve does not provide conclusive evidence of a leverage effect.

Forecasting Performance Evaluation

Table 10: Out-of-Sample Forecast Performance of GARCH (1,1) and EGARCH (1,1) Models

Model	RMSE	MAE	Mincer–Zarnowitz R^2
GARCH (1,1)	0.009552	0.007521	0.058319
EGARCH (1,1)	0.009670	0.007406	0.045760

Source: Data Analysis Result, 2026.

Figure 4: Rolling One-Step Ahead Volatility Forecasts

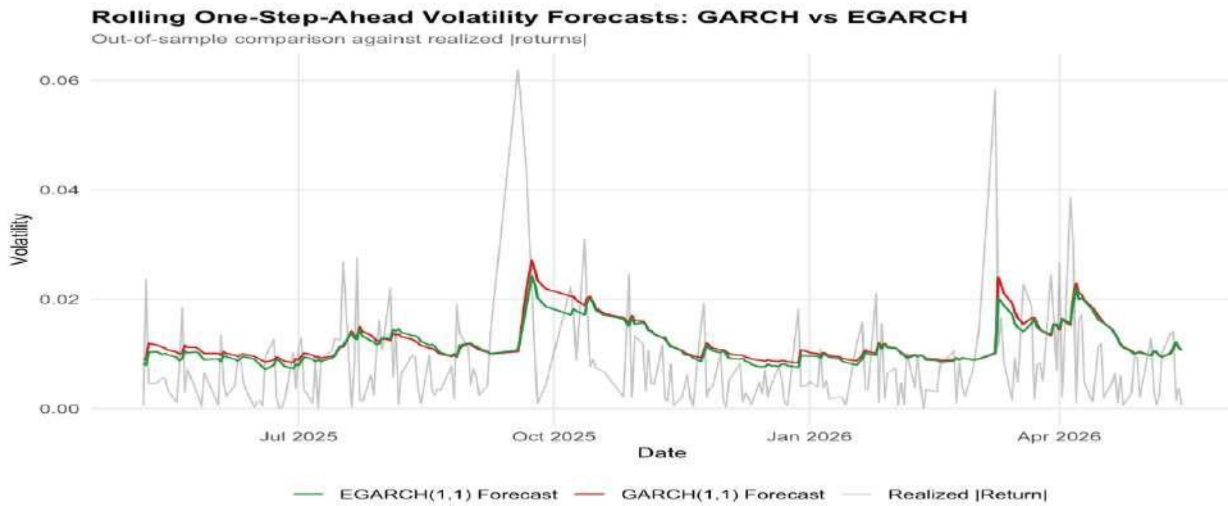

Source: Data Analysis Result, 2026.

The forecasting performance of the GARCH (1,1) and EGARCH (1,1) models was evaluated using a rolling one-step-ahead forecasting approach. Forecast accuracy was assessed using the Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and the coefficient of determination obtained from the Mincer–Zarnowitz regression (R^2). The comparative results are presented in Table 10, while the rolling volatility forecasts are illustrated in Figure 4.

The GARCH (1,1) model produced a lower RMSE (0.009552) than the EGARCH (1,1) model (0.009670), indicating slightly better overall forecasting accuracy. In contrast, the EGARCH (1,1) model recorded a marginally lower MAE (0.007406) compared with the GARCH (1,1) model (0.007521). The Mincer–Zarnowitz regression yielded an R^2 of 0.0583 for the GARCH (1,1) model and 0.0458 for the EGARCH (1,1) model, suggesting that the symmetric GARCH specification explained a slightly larger proportion of the variation in realized volatility during the forecasting period.

Overall, both models demonstrated comparable forecasting performance, with the GARCH (1,1) model showing a slight advantage according to RMSE and the Mincer–Zarnowitz R^2, while the EGARCH (1,1) model performed marginally better in terms of MAE.

Volatility Persistence Analysis

Table 11: Volatility Persistence and Half-Life Estimates

Model	Persistence Measure	Half-Life (Trading Days)
GARCH (1,1)	$\alpha + \beta = 0.9622$	18.0
EGARCH (1,1)	$\beta = 0.9438$	12.0

Source: Data Analysis Output, 2026.

The persistence of volatility shocks was assessed using the estimated persistence parameters and the corresponding volatility half-life for the GARCH (1,1) and EGARCH (1,1) models. The results are presented in Table 11.

The GARCH (1,1) model produced a persistence measure ($\alpha + \beta$) of 0.9622, indicating a high degree of volatility persistence in NEPSE returns. The corresponding volatility half-life was estimated at approximately 18 trading days. Similarly, the EGARCH (1,1) model yielded a persistence parameter (β) of 0.9438, with an estimated half-life of approximately 12 trading days.

The estimated half-life values indicate that the effects of volatility shocks remain in the market for several trading weeks before declining substantially. Compared with the EGARCH model, the GARCH model implies a longer persistence of volatility shocks.

Structural Break Test Results

Table 12: Structural Break Test Results

Number of Breaks	Break Date	Selection Criterion
1	13 July 2022	Bayesian Information Criterion (BIC)

Source: Data Analysis Output, 2026.

Figure 5: Estimated Structural Break in the NEPSE Return Series

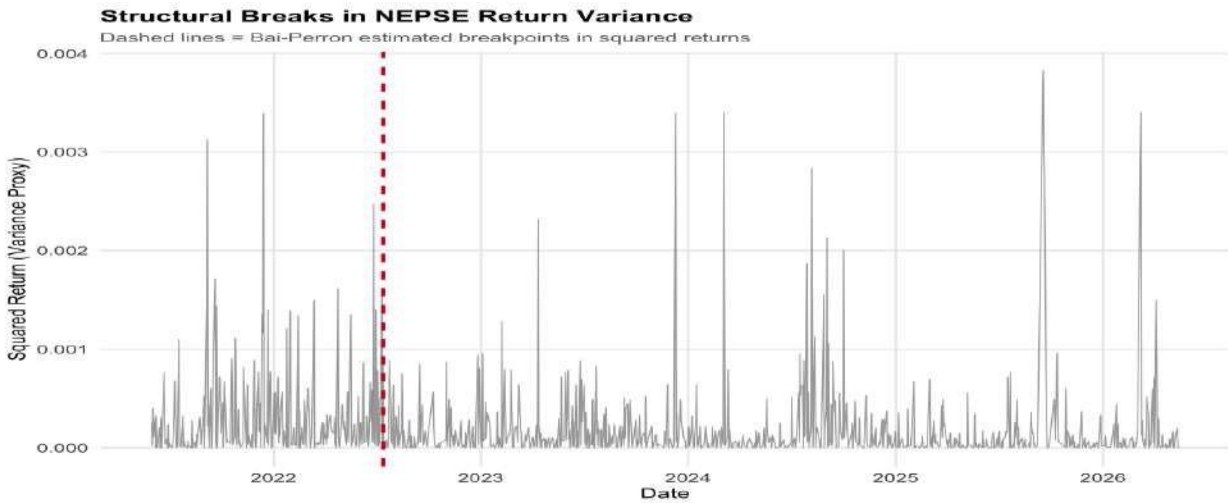

Source: Data Analysis Result, 2026.

The stability of the volatility process was examined using the Bai–Perron multiple structural breakpoint test. The results identified one statistically optimal structural break in the return series, occurring on 13 July 2022, as selected by the Bayesian Information Criterion. The estimated break point is presented in Table 12 and illustrated in Figure 5.

The identified break indicates that the volatility process was not constant throughout the sample period. Instead, the return series experienced a distinct shift in its volatility dynamics during mid-2022. This result suggests the presence of different volatility regimes within the study period and provides additional evidence that the behavior of NEPSE returns changed over time.

Pre- and Post-Break Volatility Persistence

Table 13: Pre- and Post-Break Volatility Persistence

Period	Observations	GARCH persistence	GARCH half-life	EGARCH persistence	EGARCH half-life
Full sample	1,149	0.9622	18.00	0.9438	12.00
Before 13 July 2022	269	0.9396	11.12	0.9500	13.52
From 13 July 2022	880	0.9211	8.43	0.9057	7.00

Source: Data Analysis Result, 2026.

The subperiod estimates confirm that volatility persistence changed after the structural break. GARCH persistence declined from 0.9396 before the break to 0.9211 afterward, while its half-life decreased from 11.12 to 8.43 trading days. EGARCH persistence declined from 0.9500 to 0.9057, with the corresponding half-life falling from 13.52 to 7.00 days. The full-sample GARCH persistence and half-life exceed both subperiod estimates, indicating that ignoring the break overstates persistence. The EGARCH full-sample estimate lies between the two regimes and therefore masks the substantial post-break decline. These findings demonstrate that volatility shocks dissipated more rapidly after July 2022.

Discussion

Volatility Behavior of NEPSE Returns

The empirical results show that NEPSE returns exhibit several characteristics commonly associated with financial time series. The return distribution is non-normal and heavy-tailed, while the absolute-return series displays volatility clustering. The significant ARCH-LM results further confirm that conditional variance changes over time. These findings support the use of ARCH-family models rather than models based on constant variance.

These results are consistent with the volatility-clustering framework introduced by Mandelbrot (1963), the ARCH model of Engle (1982), and the stylized facts summarized by Cont (2001). Similar evidence of volatility clustering and persistence has been reported for NEPSE by Thapa and Gautam (2016), Gajurel (2021), and Dangal and Gajurel (2021). The present study therefore confirms that time-varying and persistent volatility remains an important feature of the Nepalese equity market during the more recent sample period.

The evidence concerning directional asymmetry is less conclusive. In the EGARCH model, the signed-shock coefficient is negative but statistically insignificant ($\alpha = -0.0379$, $p = 0.146$). Thus, although its sign is directionally consistent with conventional leverage, the null hypothesis of no directional asymmetry cannot be rejected. By contrast, the shock-magnitude coefficient is positive and significant ($\gamma = 0.2611$, $p < 0.001$), indicating that larger shocks increase conditional volatility irrespective of their sign. The News Impact Curve similarly shows a higher estimated volatility response to negative shocks than to positive shocks of equal magnitude. However, because the curve is derived from the estimated EGARCH coefficients, it should be interpreted as a visual representation of the estimated response rather than as an independent statistical confirmation of leverage.

The GJR-GARCH robustness check provides additional but model-dependent evidence. Its asymmetry coefficient is positive and statistically significant ($\gamma = 0.0986$, $p = 0.026$), indicating a stronger conditional-variance response to negative shocks. Nevertheless, the insignificant EGARCH signed-shock coefficient and the remaining sensitivity indicated by the GJR-GARCH stability and joint sign-bias diagnostics prevent a definitive conclusion. Accordingly, the results provide suggestive evidence of conventional leverage, but the strength and statistical significance of this effect depend on the asymmetric specification employed.

This qualified conclusion is consistent with the mixed evidence reported in previous NEPSE studies. Thapa and Gautam (2016) and Dangal and Gajurel (2021) reported asymmetric volatility effects, whereas Gajurel (2021) documented a different pattern in which positive shocks produced a stronger volatility response. Differences in sample periods, model parameterizations, and market regimes may contribute to these contrasting findings. Overall, the present results establish volatility clustering, heavy-tailed returns, substantial persistence, and sensitivity to shock magnitude, while evidence of directional asymmetry remains suggestive rather than conclusive.

Forecasting Performance of GARCH and EGARCH Models

The in-sample comparison does not identify a single model that dominates under every selection criterion. GJR-GARCH(1,1) recorded the highest log-likelihood (3390.569) and lowest AIC (−5.8861), followed closely by EGARCH(1,1), with a log-likelihood of 3389.730 and an AIC of −5.8847. In contrast, GARCH(1,1) produced the lowest BIC (−5.8475), reflecting the stronger complexity penalty imposed by this criterion. These differences are small and indicate that the relative ranking depends on the criterion used. Because GJR-GARCH was introduced as an asymmetry robustness check and was not included in the rolling forecasting exercise, the out-of-sample comparison is restricted to GARCH and EGARCH.

The out-of-sample results are mixed rather than indicating an unequivocal winner. GARCH (1,1) produced a slightly lower RMSE than EGARCH (1,1) (0.009552 versus 0.009670) and a somewhat higher Mincer–Zarnowitz R^2 (0.0583 versus 0.0458). EGARCH (1,1), however, recorded a marginally lower MAE (0.007406 versus 0.007521). Therefore, GARCH performed slightly better when larger forecast errors received greater weight, whereas EGARCH performed marginally better according to average absolute forecast error. The small differences suggest that the two models have broadly comparable short-term forecasting performance.

The relatively low Mincer–Zarnowitz R^2 values indicate that the forecasts explain only a limited proportion of the variation in the realized-volatility proxy. This result should be interpreted cautiously because daily absolute returns are a noisy proxy for the latent volatility process and are affected by unpredictable market innovations. Moreover, the reported R^2 values alone do not establish forecast unbiasedness or statistical significance. Accordingly, the forecasting conclusions should be based primarily on the relative RMSE and MAE values and should not be interpreted as evidence of strong absolute predictive power.

The mixed ranking is consistent with the broader volatility-forecasting literature. Hansen and Lunde (2005) found that GARCH (1,1) can be difficult to outperform in some applications, although its relative performance depends on the asset, competing specifications, volatility proxy, and loss function. Similarly, Poon and Granger (2003) emphasized that no volatility model consistently performs best across all markets and evaluation settings. The present findings support this conditional interpretation: the preferred model changes according to the forecast-error measure employed.

GARCH and EGARCH therefore provide complementary information. GARCH offers a relatively parsimonious specification and shows a small advantage under RMSE and Mincer–Zarnowitz R^2. EGARCH allows volatility responses to shock sign and magnitude to be represented explicitly and achieves a marginally lower MAE. However, because its signed-shock coefficient was statistically insignificant, its forecasting performance should not be attributed to a statistically confirmed leverage effect. The rolling-window exercise extends the NEPSE volatility literature reviewed in this study by evaluating model performance outside the estimation sample rather than relying exclusively on in-sample information criteria.

More broadly, the study combines conditional-volatility estimation, diagnostic testing, fixed-window one-step-ahead forecasting, News Impact Curve analysis, volatility half-life estimation, and structural-break testing within a common empirical framework. This integration permits model fit, predictive performance, persistence, asymmetry, and parameter stability to be evaluated separately, thereby avoiding the assumption that superiority under one criterion necessarily implies superiority under the others.

Volatility, Persistence, and Market Dynamics

The full-sample estimates indicate substantial volatility persistence in NEPSE returns. GARCH(1,1) produced a persistence measure of 0.9622 and a half-life of approximately 18 trading days, whereas EGARCH(1,1) produced a persistence coefficient of 0.9438 and a half-life of approximately 12 trading days. Thus, volatility shocks are not immediately absorbed but continue to affect conditional volatility over subsequent trading sessions. The longer GARCH half-life indicates slower estimated shock dissipation under the symmetric specification.

The evidence concerning the direction of market shocks is more qualified. The News Impact Curve in Figure 3 shows a somewhat higher estimated volatility response to negative shocks than to positive shocks of equal magnitude. This pattern is directionally consistent with the conventional leverage effect described by Black (1976). However, the EGARCH signed-shock coefficient is statistically insignificant ($\alpha = -0.0379$, $p = 0.146$), meaning that the curve does not establish statistically significant directional asymmetry. The significant GJR-GARCH asymmetry coefficient provides supporting evidence that negative shocks may produce a stronger volatility response, but the overall evidence remains specification-dependent rather than conclusive.

The Bai–Perron analysis identifies a structural break on 13 July 2022, indicating that the volatility process was not stable across the full sample. The subperiod estimates support this conclusion. GARCH persistence declined from 0.9396 before the break to 0.9211 after the break, while its half-life decreased from 11.12 to 8.43 trading days. EGARCH persistence similarly declined from 0.9500 to 0.9057, with its half-life falling from 13.52 to 7.00 trading days. These estimates indicate that volatility shocks dissipated more rapidly in the post-break period.

The full-sample GARCH persistence estimate exceeds both subperiod estimates, suggesting that combining distinct volatility regimes can overstate persistence under the symmetric specification. The full-sample EGARCH estimate lies between the pre- and post-break estimates and therefore obscures the substantial decline in persistence after July 2022. This finding is consistent with Lamoureux and Lastrapes (1990), who showed that unmodelled structural changes can be mistaken for persistent volatility. Directly comparable evidence on break-adjusted volatility persistence in NEPSE remains limited; therefore, the subperiod findings extend the existing Nepalese literature. Nevertheless, they should be interpreted cautiously because the pre-break sample contains considerably fewer observations than the post-break sample.

The breakpoint test identifies a change in volatility dynamics but does not, by itself, establish its economic or institutional cause. Accordingly, the break should not be attributed to a particular policy, event, or investor response without additional evidence. Overall, the findings establish persistent but regime-dependent volatility, while evidence concerning directional asymmetry remains suggestive and model-dependent. Future research could examine the sources of the identified break using event-based analysis, regime-switching models, or models incorporating relevant macroeconomic and institutional variables.

Conclusion

This study examined the volatility dynamics and forecasting performance of the NEPSE Index using daily observations from 30 May 2021 to 15 May 2026. The empirical framework combined GARCH-family model estimation, diagnostic testing, in-sample model comparison, fixed-window one-step-ahead forecasting, News Impact Curve analysis, volatility half-life estimation, and structural-break testing. GJR-GARCH was additionally estimated as a robustness check on the asymmetry findings.

The results show that NEPSE returns are non-normal, heavy-tailed, and characterized by volatility clustering and conditional heteroskedasticity. The full-sample estimates also indicate substantial volatility persistence. GARCH (1,1) produced a persistence measure of 0.9622 and a half-life of approximately 18 trading days, whereas EGARCH (1,1) produced a persistence coefficient of 0.9438 and a half-life of approximately 12 trading days.

The evidence concerning directional asymmetry is mixed. The EGARCH signed-shock coefficient was negative but statistically insignificant, meaning that statistically significant leverage could not be established from the EGARCH model. Although the News Impact Curve was directionally consistent with a stronger response to negative shocks, it represents the estimated EGARCH relationship and does not provide an independent significance test. The significant GJR-GARCH asymmetry coefficient provided supporting evidence of conventional leverage, but the sensitivity of the findings across specifications indicates that the leverage effect should be regarded as suggestive and model-dependent rather than conclusive.

The in-sample model ranking also depended on the selection criterion. GJR-GARCH achieved the highest log-likelihood and the lowest AIC, while GARCH (1,1) produced the lowest BIC. The out-of-sample comparison between GARCH and EGARCH was similarly mixed. GARCH recorded a slightly lower RMSE and a higher Mincer–Zarnowitz R^2, whereas EGARCH produced a marginally lower MAE. These results show that a model's in-sample ranking does not necessarily determine its forecasting performance and that the preferred specification may vary with the evaluation criterion.

The structural-break analysis identified a change in volatility dynamics on 13 July 2022. Both GARCH and EGARCH persistence and half-life estimates declined in the post-break period, indicating faster dissipation of volatility shocks after the break. The full-sample GARCH estimate exceeded both subperiod estimates, while the full-sample EGARCH estimate obscured the substantial post-break decline. Persistence estimates derived from the complete sample should therefore be interpreted alongside the regime-specific results.

These findings have practical implications for risk assessment. The persistence estimates indicate that market shocks may affect conditional volatility for several trading sessions, while the structural-break results show that the rate of shock dissipation can change across market regimes. Investors and risk managers should therefore avoid treating volatility parameters as constant over time. Although the GJR-GARCH result suggests that downside shocks may warrant particular attention, the model-dependent nature of the asymmetry evidence requires caution when drawing risk-management conclusions.

The study is limited by its focus on the aggregate NEPSE Index, its use of daily absolute returns as a proxy for realized volatility, and the absence of macroeconomic or institutional explanatory variables. The rolling forecasting comparison was also limited to GARCH and EGARCH, while the pre-break subsample contained fewer observations than the post-break subsample. Moreover, the breakpoint test identifies the timing of a change but does not determine its economic cause. Future research could address these limitations through sector-level data, alternative volatility proxies, additional asymmetric and regime-switching specifications, macroeconomic variables, and explicit event-based analysis of structural changes.

Overall, the study establishes that NEPSE volatility is persistent, heavy-tailed, and sensitive to changes in market regime. Evidence of directional asymmetry is suggestive but not conclusive, and forecasting performance is broadly comparable across GARCH and EGARCH under different loss measures. The results therefore support selecting volatility models according to the specific purpose of the analysis rather than assuming that one specification is uniformly superior.

References

Bai, J., & Perron, P. (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, *18*(1), 1–22. https://doi.org/10.1002/jae.659

Black, F. (1976). Studies of stock price volatility changes. *In Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economics Statistics Section*, 177–181. American Statistical Association.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, *31*(3), 307–327. https://doi.org/10.1016/0304-4076(86)90063-1

Bollerslev, T., Chou, R. Y., & Kroner, K. F. (1992). ARCH modeling in finance. *Journal of Econometrics*, *52*(1–2), 5–59. https://doi.org/10.1016/0304-4076(92)90064-X

Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, *1*(2), 223–236. https://doi.org/10.1080/713665670

Dangal, D. N., & Gajurel, R. P. (2021). Volatility of daily Nepal stock exchange (NEPSE) Index return: A Garch family models. *Tribhuvan University Journal*, *36*(01), 31–44. https://doi.org/10.3126/tuj.v36i01.43514

De Santis, G., & Imrohoroglu, S. (1997). Stock returns and volatility in emerging financial markets. *Journal of International Money and Finance*, *16*(4), 561–579. https://doi.org/10.1016/S0261-5606(97)00020-X

Dickey, D. A., & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, *74*(366), 427–431. https://doi.org/10.2307/2286348

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, *50*(4), 987–1007. https://doi.org/10.2307/1912773

Engle, R. F., & Ng, V. K. (1993). Measuring and testing the impact of news on volatility. *The Journal of Finance*, *48*(5), 1749–1778. https://doi.org/10.1111/j.1540-6261.1993.tb05127.x

Engle, R. F., & Patton, A. J. (2001). What good is a volatility model? *Quantitative Finance*, *1*(2), 237–245. https://doi.org/10.1088/1469-7688/1/2/305

Gajurel, D. (2021). Asymmetric Volatility in the Nepalese Stock Market. *Journal of Comparative International Management*, *22*(1), 41–59. https://doi.org/10.7202/1075637ar

Glosten, L. R., Jagannathan, R., & Runkle, D. E. (1993). On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks. *The Journal of Finance*, *48*(5), 1779–1801. https://doi.org/10.1111/j.1540-6261.1993.tb05128.x

Goudarzi, H., & Ramanarayanan, C. S. (2011). Modeling Asymmetric Volatility in the Indian Stock Market. *International Journal of Business and Management*, *6*(3), 221–231. https://doi.org/10.5539/ijbm.v6n3p221

Hansen, P. R., & Lunde, A. (2005). A forecast comparison of volatility models: Does anything beat a GARCH(1,1)? *Journal of Applied Econometrics*, *20*(7), 873–889. https://doi.org/10.1002/jae.800

Jarque, C. M., & Bera, A. K. (1980). Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Economics Letters*, *6*(3), 255–259. https://doi.org/10.1016/0165-1765(80)90024-5

Kadel, D. R., & Tiwari, A. (2025). Modeling volatility dynamics of the Nepal stock exchange index: Evidence from GARCH-family models. *A Bi-Annual South Asian Journal of Research & Innovation*, *12*(2), 73–75. https://doi.org/10.3126/jori.v12i2.87974

Karmakar, M. (2005). Modeling Conditional Volatility of the Indian Stock Markets. *Vikalpa: The Journal for Decision Makers*, *30*(3), 21–38. https://doi.org/10.1177/0256090920050303

Lamoureux, C. G., & Lastrapes, W. D. (1990). Persistence in variance, structural change, and the GARCH model. *Journal of Business & Economic Statistics*, *8*(2), 225–234. https://doi.org/10.1080/07350015.1990.10509794

Ljung, G. M., & Box, G. E. P. (1978). On a measure of lack of fit in time series models. *Biometrika*, *65*(2), 297–303. https://doi.org/10.1093/biomet/65.2.297

Mandelbrot, B. (1963). The variation of certain speculative prices. *The Journal of Business*, *36*(4), 394–419. https://doi.org/10.1086/294632

Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, *59*(2), 347–370. https://doi.org/10.2307/2938260

Poon, S.-H., & Granger, C. W. J. (2003). Forecasting Volatility in Financial Markets: A Review. *Journal of Economic Literature*, *41*(2), 478–539. https://doi.org/10.1257/002205103765762743

Said, S. E., & Dickey, D. A. (1984). Testing for unit roots in autoregressive-moving average models of unknown order. *Biometrika*, *71*(3), 599–607. https://doi.org/10.1093/biomet/71.3.599

Securities Board of Nepal. (2026). Intermediaries and market infrastructure institutions. https://www.sebon.gov.np/

Shrestha, N. (2026). Time series modeling and forecasting of NEPSE Index: An ARIMA/SARIMA approach with market efficiency evidence. *Positive Science - Economics*, *1*(1), 1–20. https://doi.org/10.65957/article.16

Thapa, B. S., & Gautam, C. M. (2016). Stock return volatility and asymmetries in stock market of Nepal. *The Nepalese Management Review*, *16*(1), 117–131.

White, H. (1982). Maximum Likelihood Estimation of Misspecified Models. *Econometrica*, *50*(1), 1–25. https://doi.org/10.2307/1912526