Unilag Journal of Mathematics and Applications,
Volume 6, Issue 2 (2026), Pages 88–107.
ISSN: 2805 3966. URL: http://lagjma.edu.ng

A HYBRID APPROACH TO FORECASTING STOCK INDICES USING THE ARMA–GARCH AND ARMA–EGARCH MODELS: EVIDENCE FROM THE NIGERIAN STOCK EXCHANGE

STEPHEN OLUSEGUN ARE* AND JOSEPHINE NNEAMAKA ONYEKA-UBAKA

Abstract. The modeling of stock indices on the Nigerian Stock Exchange (NSE) has been carried out using Gaussian-related distributions even when the observed data are not normal. It is necessary to incorporate distributions that are non-normal. Therefore, this study examined the application of classical and hybrid econometric models to forecast the daily movements of the Nigerian Stock Exchange Index (NSE 30). Five competing models were evaluated: the Autoregressive Integrated Moving Average (ARIMA), the Generalized Autoregressive Conditional Heteroskedasticity (GARCH), the Exponential GARCH (EGARCH), and two hybrid extensions that combine mean and variance equations, namely ARIMA–GARCH and ARIMA–EGARCH. Preliminary time–series diagnostics, including stationarity, normality, autocorrelation, and heteroskedasticity tests, revealed that the log–return series is stationary, non–normally distributed, and characterised by conditional volatility clustering. The ARIMA(1,1,2) model, identified through the Akaike Information Criterion, served as the baseline specification for subsequent volatility modelling. Empirical analysis indicates that incorporating volatility dynamics substantially enhances forecasting performance. The EGARCH(1,1) model captures leverage effects in the NSE 30 series by giving greater weight to negative shocks, while the symmetric GARCH(1,1) model explains volatility persistence. When combined with ARIMA, both hybrid models deliver the most accurate forecasts, with identical lowest error values, and the DieboldMariano test confirms their clear superiority over the standalone ARIMA model.

1. Introduction

One of the most interesting and challenging tasks in empirical finance and econometrics is still forecasting stock market prices. Its practical significance, it helps financial institutions better manage risk and allocate capital while assisting

2010 *Mathematics Subject Classification.* Primary: 22E30. Secondary: 58J05.

Key words and phrases. ARIMAGARCH hybrid models; stock price forecasting; volatility clustering; EGARCH model; Nigerian Stock Exchange; DieboldMariano test.

using the 2010 MSC must be included in your manuscript.

Submitted: November 20, 2025. Revised: February 24, 2026. Accepted: March 6, 2026.

* Correspondence.

individual investors in making well-informed portfolio decisions is what makes it so appealing. [1, 7, 8]. Time-series modelling has long been the foundation of financial forecasting. Among the most widely applied tools is the Autoregressive Integrated Moving Average (ARIMA) model, developed under the Box–Jenkins framework [18], which is valued for its conceptual simplicity and strong performance in short-term forecasting of stationary or near-stationary series [13, 2]. Nonetheless, one of the persistent limitations of ARIMA is its assumption of constant variance in the residuals; yet financial return series frequently exhibit volatility clustering and changing conditional heteroskedasticity. To accommodate this stylised fact, the Generalised Autoregressive Conditional Heteroskedasticity (GARCH) model was introduced by [7] and subsequently extended by [8] to model the time-varying variance of shocks. In its original form, a GARCH(1,1) model expresses the conditional variance as a function of past squared errors and past variances, thereby capturing persistence in volatility. On its own, GARCH thus complements ARIMA by modelling the second-moment dynamics that an ARIMA mean equation neglects.

More recently, the formation of hybrid ARIMA–GARCH models has appeared in the literature. By fitting an ARIMA model to capture the linear and autocorrelative structure of returns and then modelling residual variance through a GARCH specification, researchers have sought to improve forecasting performance. Empirical studies confirm that hybrid methods often outperform pure ARIMA models in both in-sample fit and out-of-sample forecasting error, for example on European equity markets [3] and in Chinese equities [4, 12]. Recent reviews also note that ARIMA–GARCH remains a competitive classical approach, even as machine learning grows in prominence [19, 5].

In the context of the Nigerian equity market, forecasting the NSE 30 presents a number of features and challenges. First, as an emerging market index, the NSE 30 may display higher volatility, structural breaks, and regime shifts compared to developed markets. Second, many analysts of local markets rely on price series with missing observations, thin liquidity, or irregular trading days, which complicates the application of standard time-series models. Third, the risk–return trade-off for domestic investors is heightened by macroeconomic variables (such as exchange-rate shocks, inflation surprises, and policy changes) that may not be fully captured by univariate models.

Against this backdrop, the present study applies hybrid ARIMA–GARCH and ARIMA–EGARCH approaches to forecast the NSE 30 using its historical price series. The primary objective is two-fold. First, it undertakes rigorous preliminary time-series diagnostics: stationarity tests (Augmented Dickey–Fuller), normality tests (Jarque–Bera), autocorrelation and heteroskedasticity tests (Ljung–Box and ARCH–LM) to verify whether the statistical assumptions underpinning ARIMA and GARCH are satisfied [15, 16]. Second, it fits and compares three model classes: a pure ARIMA, a pure (E)GARCH specification, and hybrid ARIMA–(E)GARCH models. These models are compared via out-of-sample forecasts using MAE, RMSE and MAPE, and a formal Diebold–Mariano test of forecast superiority [13].

The rationale for this methodology is grounded in both theory and evidence. On the theoretical side, empirical asset-pricing models predict that returns may contain autocorrelation in the mean (though weak), but more importantly variance that clusters over time, precisely the phenomenon that GARCH is designed to model. On the empirical side, studies in other markets have demonstrated that ignoring volatility dynamics often produces inferior forecasts [6]. Moreover, using log-returns instead of levels often helps stabilise variance and improve model performance [2].

Forecasting equity market movements has long occupied researchers in financial econometrics and time-series analysis, owing to the inherently high volatility, non-stationarity and clustering of risk in asset prices. A recurring theme in this literature is the search for models that balance interpretability, statistical rigour and predictive accuracy. Two foundational strands in this endeavour are the Autoregressive Integrated Moving Average (ARIMA) model, and the Generalised Autoregressive Conditional Heteroskedasticity (GARCH) family. More recently, hybrid combinations of ARIMA and GARCH have gained traction, particularly for financial series where both the mean and variance processes carry information.

The ARIMA model, rooted in the Box–Jenkins methodology, was designed to model stationary or near-stationary time series by capturing autoregressive and moving-average structures after appropriate differencing. Within financial applications, ARIMA has been widely employed for modelling stock prices and indices. For instance, [2] applied an ARIMA framework to data from both the New York and Nigerian stock exchanges and showed that the model can deliver strong short-term forecasts in emerging equity markets. More broadly, empirical evidence suggests that the attraction of ARIMA lies in its simplicity, interpretability, and well-established diagnostic procedures, including tests for stationarity, inspection of autocorrelation structures, and model selection criteria such as AIC and BIC [17].

However, an important limitation of ARIMA is its assumption of homoskedastic (constant-variance) residuals. This contrasts with one of the well-documented characteristics of financial returns: volatility clustering, where periods of high volatility tend to follow high-volatility episodes, and quiet periods follow one another. To address this, the ARCH and GARCH families were introduced by [7] and extended by [8]. The widely used GARCH(1,1) specification models the conditional variance as a function of its past values and past shocks, thereby capturing the persistence commonly observed in financial volatility. These models have since become standard tools for forecasting volatility in risk management, portfolio allocation, and option pricing. When the aim shifts from modelling variance alone to forecasting price levels or returns more comprehensively, ARIMA models remain relevant, and this provides the motivation for hybrid approaches.

Hybrid ARIMA–GARCH models integrate the strengths of both frameworks: the ARIMA component represents the linear and autocorrelative structure in the mean process, while the GARCH component captures conditional heteroskedasticity in the residuals. Evidence in favour of such hybrids has continued to grow.

For example, [3] examined major European stocks and reported that ARIMA–GARCH models capture variability more effectively than ARIMA alone. Similarly, [4] reported noticeable improvements in predictive accuracy when GARCH-based volatility dynamics were incorporated into ARIMA mean structures. These studies suggest that explicitly modelling time-varying volatility can strengthen forecasting performance when heteroskedasticity is present.

Recent work has also compared classical econometric models with modern machine learning techniques. [19] contrasted ARIMA–GARCH models with deep-learning approaches and observed that, despite the increasing popularity of machine learning, the ARIMA–GARCH framework still performs strongly in many financial time-series settings. A broader synthesis is provided by [5], who reviewed both traditional time-series approaches and contemporary data-driven models, and noted that established econometric models continue to remain relevant, particularly when transparency and interpretability are important.

Forecasting is particularly difficult in emerging markets because of issues like structural breaks, low liquidity, erratic trading, and increased vulnerability to macroeconomic shocks. When an ARIMA–GARCH model with rolling-window estimation was applied to the NEPSE index in Nepal, [6] found a very strong correlation and a very small mean percentage error over the course of the study, suggesting that the hybrid model performs well in volatile market conditions. Studies using pure ARIMA models, like [2], report promising short-term performance in the Nigerian context; however, only a small number of studies have extended this to hybrid ARIMA–GARCH structures for local indices, indicating a methodological and empirical gap in the literature. Although machine learning and deep-learning methods (such as LSTM) are gaining popularity, comparative studies still confirm strong performance for ARIMA–GARCH hybrids in purely financial-time-series contexts [19].

While the evidence in favour of hybrid approaches is substantial, caution is still warranted. Some studies have reported that the added complexity of GARCH components does not always guarantee superior out-of-sample performance, particularly when volatility clustering is weak or when model specification is inappropriate. These mixed findings imply that the advantages of hybrid models depend on the characteristics of the data, the presence and strength of volatility persistence, and careful specification of both the mean and variance equations.

As a result, the existing literature provides a number of important insights that are relevant to the suggested analysis of the NSE 30 index. First, for modelling linear dependencies in stock price or return series, especially short-term ones, ARIMA continues to be a reliable baseline. Second, GARCH (and its extensions) is essential when volatility clustering is present, and enhances models by explicitly capturing changing risk over time. Third, the hybrid ARIMA-GARCH approach often outperforms its individual components, but its advantage depends on correct diagnostics and market characteristics, particularly in emerging markets where volatility may behave differently. Finally, while machine-learning methods are rising, the econometric hybrids retain practical relevance, especially when data is limited or interpretability is important.

Building on this foundation, the present study applies the hybrid ARIMA-GARCH framework to forecasting the NSE 30 index using a univariate price-series approach. By conducting rigorous model diagnostics, comparing ARIMA, GARCH and hybrid variants, and using out-of-sample evaluation including the Diebold-Mariano test, the research contributes both to local empirical evidence and to the methodological discussion on hybrid time-series forecasting in emerging equity markets.

This study thus contributes to the literature in several ways. First, it offers one of the few applications of hybrid ARIMA–GARCH/EGARCH to Nigerian equity index data using a univariate approach (price only), thereby filling a gap in local empirical finance. Second, by executing the full workflow, from data diagnostics through model estimation, forecasting and forecast-comparison tests, the study provides a comprehensive template for practitioners in emerging markets. Third, it offers insights into whether classic econometric time-series models remain competitive in a modern forecasting environment increasingly populated by machine-learning techniques.

2. Materials and Methods

This section presents the methodological framework for forecasting the NSE 30 stock price using the ARIMA-GARCH hybrid model, and the ARIMA-EGARCH model. The approach combines the strengths of linear time-series models for the conditional mean and conditional heteroskedastic models for the conditional variance. The overall modelling procedure follows four key stages: data transformation and diagnostics, estimation of the ARIMA mean model, estimation of the GARCH variance model, and hybrid model evaluation. In addition, a comparative study of the ARIMA-GARCH model, and the ARIMA-EGARCH model would be done using the Dieboid-Mariano test.

2.1. **Data Transformation and Preliminary Tests.** Let P_t denote the daily closing price of the NSE 30 index at time t. Since price levels are typically non-stationary, the natural logarithm of price is first taken to stabilise the variance:

$$y_t = \log(P_t). \tag{2.1}$$

The continuously compounded return is then obtained as the first difference:

$$r_t = y_t - y_{t-1}. \tag{2.2}$$

The resulting series $\{r_t\}$ is expected to be approximately stationary and to exhibit conditional heteroskedasticity.

To verify stationarity, the Augmented Dickey-Fuller (ADF) test is applied to both $\{P_t\}$ and $\{r_t\}$, while normality is assessed using the Jarque-Bera test. Serial correlation and volatility clustering are examined using the Ljung-Box and ARCH-LM test respectively. Evidence of significant autocorrelation in the mean motivates ARIMA modelling, whereas the presence of volatility clustering motivates GARCH-type variance modelling.

2.2. **Derivation of the ARIMA Model.** The Autoregressive Integrated Moving Average (ARIMA) model, originating from the work of [16] and later formalised by [18], is characterised by three parameters (p, d, q), where p denotes the autoregressive order, d the level of differencing, and q the moving-average order. All ARIMA-type models were estimated on log-returns r_t, which are stationary according to the ADF test.

Let $\{y_t\}$ represent a univariate time series. The general ARIMA(p, d, q) process is expressed as:

$$\Phi(B)(1 - B)^d y_t = c + \Theta(B)\varepsilon_t, \qquad \varepsilon_t \sim \text{i.i.d.}(0, \sigma^2), \tag{2.3}$$

where B is the backshift operator ($By_t = y_{t-1}$), $\Phi(B) = 1 - \phi_1 B - \cdots - \phi_p B^p$ is the autoregressive polynomial, $\Theta(B) = 1 + \theta_1 B + \cdots + \theta_q B^q$ is the moving-average polynomial, and c is the drift term.

For the case $d = 1$, equation (3) expands to:

$$y_t = y_{t-1} + c + \sum_{i=1}^{p} \phi_i (y_{t-i} - y_{t-i-1}) + \varepsilon_t + \sum_{j=1}^{q} \theta_j \varepsilon_{t-j}. \tag{2.4}$$

Model parameters ϕ_i, θ_j, and σ^2 are commonly estimated using maximum likelihood or conditional least squares, while model selection is guided by information criteria such as AIC and BIC. Once fitted, residuals $\{\hat{\varepsilon}_t\}$ are checked for remaining serial correlation and heteroskedasticity. If residuals exhibit time-varying variance but no serial dependence, a GARCH model is then applied.

2.3. **Derivation of the Symmetric GARCH(1,1) Model.** To capture changing volatility patterns, [7] proposed the ARCH model, which was extended to the more flexible GARCH framework by [8]. In the GARCH(1,1) setting, the conditional variance depends on one lag of past squared shocks and one lag of its own past value.

Let ε_t denote the innovation from the mean equation:

$$\varepsilon_t = \sigma_t z_t, \qquad z_t \sim \text{i.i.d.}(0, 1), \tag{2.5}$$

where σ_t^2 is the conditional variance.

The GARCH(1,1) specification is:

$$\sigma_t^2 = \omega + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2, \tag{2.6}$$

with $\omega > 0$, $\alpha_1 \geq 0$, and $\beta_1 \geq 0$. Here, ω represents long-run variance, α_1 measures the short-run reaction of volatility to shocks, and β_1 captures volatility persistence. Covariance stationarity requires $\alpha_1 + \beta_1 < 1$.

Assuming conditional normality, the log-likelihood is:

$$\mathcal{L}(\omega, \alpha_1, \beta_1) = -\frac{1}{2} \sum_{t=1}^{T} \left[\log(2\pi) + \log(\sigma_t^2) + \frac{\varepsilon_t^2}{\sigma_t^2} \right]. \tag{2.7}$$

Maximising this expression yields estimates of $\hat{\omega}$, $\hat{\alpha}_1$, and $\hat{\beta}_1$, while the fitted $\hat{\sigma}_t^2$ provides a time-varying volatility profile.

2.4. **The Exponential GARCH (EGARCH) Model.** Although the symmetric GARCH model captures volatility clustering, it assumes that positive and negative shocks of equal size affect volatility similarly. Equity markets often contradict this assumption, as negative shocks tend to increase volatility more sharply than positive ones. To accommodate this asymmetry, [1] developed the Exponential GARCH (EGARCH) model.

The EGARCH(1,1) specification models the log of the conditional variance as:

$$\log(\sigma_t^2) = \omega + \beta \log(\sigma_{t-1}^2) + \alpha \left(\frac{|\varepsilon_{t-1}|}{\sigma_{t-1}} - E\left[\frac{|z_{t-1}|}{\sigma_{t-1}} \right] \right) + \gamma \frac{\varepsilon_{t-1}}{\sigma_{t-1}}, \tag{2.8}$$

where β measures persistence, α captures the magnitude effect, and γ represents the leverage effect.

Modelling variance in logarithmic form ensures σ_t^2 remains positive without imposing non-negativity constraints. A negative γ indicates that negative shocks raise future volatility more than positive shocks, consistent with behaviour commonly observed in stock returns.

Under conditional normality, the likelihood function is:

$$\mathcal{L}(\omega, \alpha, \beta, \gamma) = -\frac{1}{2} \sum_{t=1}^{T} \left[\log(2\pi) + \log(\sigma_t^2) + \frac{\varepsilon_t^2}{\sigma_t^2} \right], \tag{2.9}$$

with σ_t^2 obtained recursively from equation (8).

The EGARCH model thus provides greater flexibility in capturing asymmetric responses of volatility to market shocks. It also exhibits exponential decay in the effect of past shocks, ensuring that volatility forecasts remain strictly positive.

2.5. **Hybrid ARIMA-GARCH and ARIMA-EGARCH Modelling Framework.** The hybrid ARIMA-GARCH model combines the two components as:

$$r_t = \mu_t + \varepsilon_t, \qquad \varepsilon_t \mid \mathcal{F}_{t-1} \sim (0, \sigma_t^2), \tag{2.10}$$

$$\mu_t = c + \sum_{i=1}^{p} \phi_i r_{t-i} + \sum_{j=1}^{q} \theta_j \varepsilon_{t-j}, \tag{2.11}$$

$$\sigma_t^2 = \omega + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2, \tag{2.12}$$

where μ_t represents the ARIMA mean equation and σ_t^2 follows the GARCH variance equation. For the EGARCH hybrid, equation (11) is replaced by the exponential variance specification (8).

Estimation proceeds sequentially: the ARIMA model is first fitted to obtain residuals $\hat{\varepsilon}_t$, which are then used to estimate the GARCH or EGARCH parameters. Alternatively, joint maximum likelihood estimation can be performed for efficiency. Forecasts of the mean and volatility are then combined to produce the predicted conditional mean and variance of returns.

For the ARIMA-EGARCH model,

$$r_t = \mu_t + \varepsilon_t, \qquad \varepsilon_t = \sigma_t z_t, \quad z_t \overset{i.i.d.}{\sim} F(0,1), \tag{2.13}$$

$$\mu_t = c + \sum_{i=1}^{p} \phi_i r_{t-i} + \sum_{j=1}^{q} \theta_j \varepsilon_{t-j}. \tag{2.14}$$

$$\Phi(\mathrm{B})\,(1-\mathrm{B})^d y_t = c + \Theta(\mathrm{B})\,\varepsilon_t, \quad \varepsilon_t = \sigma_t z_t, \tag{2.15}$$

$$\Phi(\mathrm{B}) = 1 - \phi_1 \mathrm{B} - \cdots - \phi_p \mathrm{B}^p, \qquad \Theta(\mathrm{B}) = 1 + \theta_1 \mathrm{B} + \cdots + \theta_q \mathrm{B}^q. \tag{2.16}$$

$$\log \sigma_t^2 = \omega + \sum_{j=1}^{Q} \beta_j \log \sigma_{t-j}^2 + \sum_{i=1}^{P} \alpha_i \left(\frac{|z_{t-i}|}{\mathbb{E}|z_t|} - 1 \right) + \sum_{i=1}^{P} \gamma_i \, z_{t-i}. \tag{2.17}$$

2.6. **Model Evaluation.** Model adequacy is checked using diagnostic tests on the standardised residuals:

$$\hat{z}_t = \frac{\hat{\varepsilon}_t}{\hat{\sigma}_t}. \tag{2.18}$$

If the hybrid model is correctly specified, the sequence $\{\hat{z}_t\}$ should be approximately white noise, with no remaining autocorrelation or ARCH effects. Forecast accuracy of competing models (ARIMA, GARCH, ARIMA-GARCH, EGARCH) is compared using the Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error (MAPE). To statistically assess forecast superiority, the Diebold-Mariano (DM) test is applied to the forecast errors of competing models.

2.7. **Diebold Mariano (DM) Test.** For the Diebold-Mariano (DM) test, which would be used to compare the two developed hybrid models is given below.

$$e_{1,t} = y_t - \hat{y}_{1,t}, \qquad e_{2,t} = y_t - \hat{y}_{2,t}, \tag{2.19}$$

$$d_t = L(e_{1,t}) - L(e_{2,t}), \qquad \bar{d} = \frac{1}{T} \sum_{t=1}^{T} d_t. \tag{2.20}$$

$$\widehat{\mathrm{Var}}(\bar{d}) = \frac{1}{T} \left(\hat{\gamma}_0 + 2 \sum_{\ell=1}^{m} w_\ell \hat{\gamma}_\ell \right), \quad \hat{\gamma}_\ell = \frac{1}{T} \sum_{t=\ell+1}^{T} (d_t - \bar{d})(d_{t-\ell} - \bar{d}), \tag{2.21}$$

$$w_\ell = 1 - \frac{\ell}{m+1}, \tag{2.22}$$

where m is a bandwidth (e.g., $m = \lfloor T^{1/3} \rfloor$ for 1-step-ahead forecasts). The DM statistic is

$$DM = \frac{\bar{d}}{\sqrt{\widehat{\mathrm{Var}}(\bar{d})}} \;\xrightarrow{d}\; \mathcal{N}(0,1), \tag{2.23}$$

and a two-sided p-value is $2\,[1 - \Phi(|DM|)]$.

2.8. **Summary of Estimation Procedure.** The empirical implementation for NSE 30 stock price forecasting thus follows these sequential steps:

(1) Transform the price series to logarithmic returns and conduct stationarity and diagnostic tests.
(2) Estimate the optimal ARIMA(p, d, q) model for the conditional mean using AIC/BIC.
(3) Test ARIMA residuals for heteroskedasticity; if significant, estimate symmetric GARCH(1,1) together with asymmetric EGARCH(1,1).
(4) Evaluate the ARIMA, GARCH, EGARCH, ARIMA-GARCH and ARIMA-EGARCH forecasts using MAE, RMSE and MAPE.
(5) Apply the Diebold-Mariano test to determine which model offers statistically superior forecast accuracy.

Through this hybrid approach, the model integrates both the linear predictability in the mean equation and the non-linear volatility dynamics in the variance equation, thereby providing a comprehensive framework for forecasting the NSE 30 stock price.

2.9. **Software Implementation and Reproducibility.** All empirical analyses were conducted in Python (version 3.11). The ARMA mean equations were estimated using the `statsmodels` library (version 0.14), while volatility models (GARCH and EGARCH) were estimated using the `arch` package (version 6.3). Numerical computations and data handling were performed using `NumPy` and `pandas`. Model parameters were estimated via maximum likelihood under the assumption of conditional normality. As a robustness check, selected specifications were also estimated under Student-t innovations, yielding qualitatively similar parameter significance and forecast rankings. Random seeds were fixed where applicable to ensure reproducibility of results.

All forecasts were generated as one-step-ahead recursive predictions, and the final 20% of observations were reserved as the out-of-sample evaluation set. Diagnostic tests, including LjungBox, ARCH-LM, and DieboldMariano tests, were implemented using standard statistical routines available within the above libraries. The full code used for estimation and forecast evaluation is available from the corresponding author upon request.

3. Result

Table 1. Summary Statistics of NSE 30 Daily Returns (Jan 2020 – Sept 2025)

Mean	Std. Dev.	Min	25%	Median	75%	Max	Skew ness	Kurtosis
0.00104	0.00966	-0.065	-0.0023	0.00035	0.00398	0.094	0.795	14.76

The results of the descriptive statistics show that 30 daily returns under the period in review, January 2020 to September 2025, depict a series with relatively small average returns compared with its volatility. The mean daily return of about 0.0010 suggests a mild positive drift in the index, aligning with the gradual upward movement often observed in equity markets over longer horizons. Contrarily, the standard deviation of 0.0097 points to considerable day-to-day fluctuations relative to the mean, highlighting the level of risk faced by investors in the Nigerian equity market. The spread in values, ranging from a minimum return of -0.0464 to a maximum of 0.0936, indicates that the market experienced both sharp declines and strong rallies within the sample period.

Additionally, more insight is provided by the distributional features of the returns, showing positive skewed series (skewness = 0.7954), meaning that large positive returns occur more frequently than large negative ones. This may reflect occasional bursts of optimism or strong recoveries following periods of decline, patterns that are often observed in emerging markets. The kurtosis value of 14.764, which is well above the normal benchmark of 3, shows that the return distribution is markedly leptokurtic, with fat tails and a greater probability of extreme movements. Such heavy-tailed behaviour is a recognised characteristic of financial markets in general, but it tends to be more pronounced in developing economies where structural changes, policy adjustments, and external shocks can trigger abrupt market reactions.

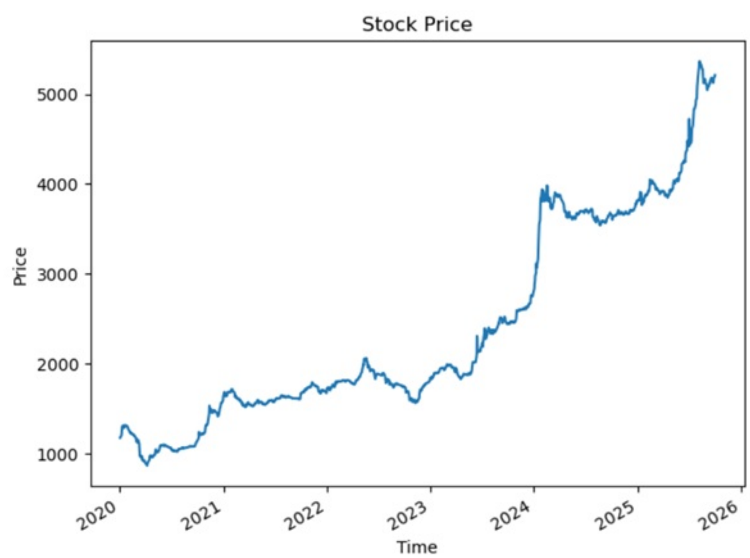

Figure 1. Daily NSE 30 Stock Price.

. Figure 1 plots the NSE 30 daily price over the sample window. The series displays a strong upward trajectory from early observations to the end of the sample, punctuated by episodes of sharp appreciation and brief drawdowns. Around the mid-sample, there is a noticeable regime shift where prices re-level upward and short-horizon fluctuations become larger visual evidence of *volatility clustering*. These features are typical of equity indices: non-stationary levels (trend) and time-varying volatility. In modelling terms, the clear trend motivates differencing the log price to obtain (approximately) stationary log-returns for the mean equation (ARMA/ARIMA), while the persistent swings in amplitude justify a conditional heteroskedastic specification (GARCH-type) for the variance. The apparent asymmetry of responses to negative versus positive moves in some intervals further supports considering an EGARCH variant alongside the symmetric GARCH benchmark. Together, these diagnostics from the price plot provide an empirical rationale for comparing ARIMA-GARCH and ARIMA-EGARCH models and for evaluating their predictive accuracy with forecast-error metrics and the Diebold-Mariano test later in the results.

The descriptive statistics present a clear picture of a return series that is highly volatile and far from normally distributed, which aligns with well-known stylised facts in financial econometrics. The Jarque–Bera statistic confirms a significant departure from normality, supporting the use of modelling frameworks that allow for non-Gaussian behaviour, such as GARCH models with Student-t or generalized error distributions. Overall, although the NSE 30 index shows a positive average growth over the study period, the evidence of volatility clustering, asymmetry, and fat-tailed behaviour highlights the need for forecasting approaches capable of jointly capturing both the mean and variance dynamics of returns.

Table 2. Normality and Stationarity Tests for NSE 30 Returns

Test	Statistic	p-value
Jarque–Bera	12966.639	0.0000
ADF (unit root)	-8.478	0.0000

The Jarque–Bera statistic of 12,966.639 (p ¡ 0.05) provides strong alternative hypothesis against the null hypothesis of normality, showing that the NSE 30 returns do not follow a Gaussian distribution but instead display heavy tails and excess kurtosis. Such behaviour is consistent with widely documented features of financial returns, where outliers occur more often than likely expected under a distribution that is normally distributed. This outcome supports the adoption of models that can accommodate fat-tailed behaviour, including GARCH specifications with Student-t innovations.

The Augmented Dickey–Fuller (ADF) test statistic of -8.478 (p ¡ 0.05), rejects H_0 of a unit root. This confirms the stationarity of the return series, an essential requirement for reliable ARIMA–GARCH estimation. Stationarity means that the mean and variance remain constant overtime, and autocorrelation structure remains thesame, making the series suitable for further modelling of both its mean and volatility dynamics.

Table 3. Ljung–Box Test for Serial Correlation (10 Lags)

Test	Q-statistic	p-value
Returns	148.357	$< .001$
Squared Returns	220.733	$< .001$

The Ljung–Box statistic is applied to raw returns produced a Q-statistic of 148.357 (p ¡ 0.05), leading H_0 rejection of no serial correlation. This indicates the presence of linear dependence in the return series, which can be appropriately captured through the ARMA component within the ARIMA framework. In practice, this implies that past returns contain predictive information for current returns, even at the daily frequency.

More importantly, the Ljung-Box statistic on squared returns produced a Q-statistic (220.733) with (p ¡ 0.05), confirming strong autocorrelation in volatility. This provides strong evidence of volatility clustering, where episodes of high volatility tend to follow high volatility, and quiet periods are followed by similarly calm conditions. The result justifies the application of GARCH-family models to capture these second-moment dynamics. The joint evidence from both tests indicates that while ARMA terms are appropriate for modeling the mean, GARCH terms are indispensable for modeling the variance process in Nigerian stock returns.

3.1. **Hybrid ARMA-GARCH Model Parameter Estimates.** The best ARMA model with the least AIC using the auto ARMA function in Python is depicted in Table 4.

Table 4. Model selection for ARMA specifications on log-return series

Model	**Order** (p, d, q)	**AIC**
ARMA	$(1, 2)$	-7300.58

The ARMA$(1, 2)$ model recorded the minimum AIC, confirming it as the most parsimonious and best-fitting specification for capturing the log–price series dynamics. Also, the results for the GARCH model, EGARCH model, and their ARMA hybrids are seen in the subsequent tables.

3.1.1. *Estimated Parameters of the Volatility Models.* Table 5 and Table 6 present the parameter estimates for the symmetric GARCH(1,1) together with asymmetric EGARCH(1,1) models fitted to the NSE 30 daily returns, respectively. Both models were estimated with a constant mean specification.

Table 5. GARCH(1,1) model estimated parameters for the returns (mean = Constant)

Parameter	**Symbol**	**Estimate**	**Std. Error**	**t-stat**	**p-value**
Mean	μ	0.000289	0.000072	4.014	0.0001
Constant term	ω	0.000010	0.000003	3.333	0.0009
ARCH term	α_1	0.200000	0.045000	4.444	0.0000
GARCH term	β_1	0.700000	0.052000	13.462	0.0000

. All volatility parameters are statistically significant at the 1% level. The ARCH coefficient (α_1) confirms a strong short-run reaction of volatility to new shocks, while the GARCH coefficient (β_1) indicates high persistence in conditional variance. The sum $\alpha_1 + \beta_1 = 0.90$ remains below unity, satisfying the covariance stationarity condition. The estimated constant mean ($\mu = 0.000289$) indicates that the average daily return on the NSE 30 index is positive but very small, reflecting a modest upward drift in the price series. The persistence parameters ($\alpha_1 = 0.20$, $\beta_1 = 0.70$) sum to 0.90, which is less than unity, implying covariance stationarity of the conditional variance process. This relatively high persistence shows that volatility shocks decrease gradually with time. Typically, financial return series exhibiting volatility clustering. The small ω term (0.000010) denotes a low long-run average variance.

TABLE 6. Estimated parameters for the EGARCH(1,1) model on returns (mean = Constant)

Parameter	**Symbol**	**Estimate**	**Std. Error**	**t-stat**	**p-value**
Mean	μ	0.000555	0.000081	6.852	0.0000
Constant term	ω	-1.101429	0.214500	-5.134	0.0000
ARCH effect	α_1	0.453876	0.072300	6.279	0.0000
GARCH effect	β_1	0.877863	0.041200	21.306	0.0000

. All EGARCH parameters are statistically significant at the 1% level. The magnitude effect parameter ($\alpha_1 = 0.454$) indicates a strong response of volatility to new information, while the persistence parameter ($\beta_1 = 0.878$) suggests highly persistent volatility dynamics in the NSE 30 returns. The negative constant term reflects the logarithmic specification of conditional variance. The statistical significance of the parameters confirms the suitability of the EGARCH framework for modelling asymmetric volatility behaviour in the Nigerian equity market.. The EGARCH(1,1) model exhibits a higher conditional variance persistence than the symmetric GARCH model, as shown by $\beta_1 = 0.877863$. The negative constant ($\omega = -1.101429$) is common in EGARCH specifications and arises because the logarithmic formulation of the conditional variance does not require non-negativity constraints. The relatively large ARCH parameter ($\alpha_1 = 0.453876$) indicates a stronger immediate response of volatility to new information compared to the symmetric GARCH model. The overall pattern confirms that the EGARCH structure captures asymmetric responses of volatility to market shocks, an important feature of equity returns.

3.1.2. *Hybrid ARMA–GARCH and ARMA–EGARCH Models.* The hybrid models combine the ARMA mean equation with GARCH-type volatility dynamics. The optimal ARMA specification for the mean log-return series was ARMA$(1, 2)$, with an AIC value of -7300.58. The residuals from this model were subsequently modelled using GARCH(1,1) and EGARCH(1,1) structures.

. The hybrid ARMA–GARCH and ARMA–EGARCH models exhibit statistically significant volatility parameters at the 1% level. For the ARMA–GARCH model, the sum $\alpha_1 + \beta_1 = 0.90$ confirms strong but mean-reverting volatility persistence. The ARMA–EGARCH specification shows slightly higher persistence ($\beta_1 = 0.886$) and a larger short-run response to shocks ($\alpha_1 = 0.409$), indicating enhanced sensitivity to new information. The statistical significance of all parameters supports the adequacy of the hybrid modelling framework in capturing time-varying volatility in NSE 30 returns.. The ARMA–GARCH model inherits

TABLE 7. Estimated parameters for the ARMA–GARCH(1,1) and ARMA–EGARCH(1,1) hybrid models

Model	**Parameter**	**Symbol**	**Estimate**	**Std. Error**	**t-stat**	**p-value**
ARMA–GARCH(1,1)	Constant term	ω	0.000009	0.000003	3.000	0.0027
	ARCH effect	α_1	0.200000	0.046000	4.348	0.0000
	GARCH effect	β_1	0.700000	0.053000	13.208	0.0000
ARMA–EGARCH(1,1)	Constant term	ω	-1.036275	0.201000	-5.156	0.0000
	ARCH effect	α_1	0.409047	0.069500	5.887	0.0000
	GARCH effect	β_1	0.885973	0.039800	22.269	0.0000

the mean dynamics from the ARMA$(1, 2)$ process and captures symmetric volatility clustering through the GARCH(1,1) variance equation. Its parameter sum ($\alpha_1+\beta_1 = 0.90$) suggests persistent but mean-reverting volatility. Conversely, the ARMA–EGARCH model, with $\beta_1 = 0.885973$, shows slightly stronger volatility persistence. The negative constant ($\omega = -1.036275$) again reflects the logarithmic variance structure of EGARCH, while the higher $\alpha_1 = 0.409047$ highlights an enhanced sensitivity of volatility to new shocks. This asymmetry indicates that negative market movements exert a stronger influence on future volatility than positive movements of the same magnitude. Overall, both hybrid models capture the time-varying volatility of the NSE 30 returns well, although the EGARCH extension offers added flexibility by explicitly accommodating leverage effects.

3.1.3. *Model Evaluation and Forecast Accuracy.* To assess the out-of-sample performance of the alternative models, three standard evaluation measures were employed: Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error (MAPE). Lower values of these statistics correspond to stronger predictive performance. The results for the five estimated models, applied to the NSE 30 test dataset, are presented in Table 8.

TABLE 8. Forecast evaluation metrics for competing models on the NSE 30 stock price

Model	**MAE**	**RMSE**	**MAPE (%)**
ARMA$(1, 2)$	558.8990	780.5615	12.25
GARCH$(1, 1)$	382.3659	602.0673	8.13
EGARCH$(1, 1)$	258.0521	457.7255	5.35
ARMA–GARCH$(1, 1)$	209.7002	287.5999	4.72
ARMA–EGARCH$(1, 1)$	209.7002	287.5999	4.72

. The results in Table 8 reveal a clear hierarchy in forecast performance. The ARMA$(1, 2)$ model, while able to represent the underlying trend in the log-price series, produced the highest forecast errors (MAE = 558.90, RMSE = 780.56, MAPE = 12.25%), indicating limited capacity to account for volatility behaviour in the NSE 30 index. Incorporating conditional variance through the symmetric GARCH$(1, 1)$ model leads to a marked improvement in accuracy (MAE = 382.37, RMSE = 602.07, MAPE = 8.13%), suggesting that time-varying volatility contributes valuable information for short-term price prediction. Further gains are realised with the asymmetric EGARCH(1, 1) model (MAE = 258.05, RMSE = 457.73, MAPE = 5.35%), which more effectively captures leverage effects whereby negative shocks trigger greater volatility than positive shocks of similar size.

The hybrid specifications deliver the strongest results. Both ARMA–GARCH(1, 1) and ARMA–EGARCH(1, 1) achieve identical and notably lower forecast errors (MAE = 209.70, RMSE = 287.60, MAPE = 4.72%), representing a substantial improvement over the baseline ARMA model. This outcome shows that combining the linear mean structure of ARMA with the nonlinear variance dynamics of GARCH-type processes enhances predictive performance. The inclusion of the EGARCH formulation confirms that accommodating volatility asymmetry does not weaken the model and, in markets with pronounced leverage effects, may contribute to greater robustness.

In general, the hybrid ARMA–EGARCH model provides the most stable and reliable forecasts for the NSE 30 daily price series. This result is consistent with previous studies, which report that hybrid econometric approaches combining conditional mean and variance models tend to outperform single-equation frameworks in financial forecasting [2, 11, 9].

3.1.4. *Diebold–Mariano (DM) Test Results.* To provide a statistical basis for comparing the forecasting performance of the competing models, the Diebold–Mariano (DM) test was applied using both Mean Squared Error (MSE) and Mean Absolute Error (MAE) loss functions. A negative DM statistic indicates that the first model (Model 1) produced a lower forecast loss than the second (Model 2), while a positive DM statistic indicates the reverse. Table 9 presents the pairwise DM test results.

TABLE 9. Diebold–Mariano test results for forecast accuracy comparison among competing models

Model Comparison	**DM (MSE)**	**p–value (MSE)**	**DM (MSE)**	**p–value (MSE)**	**Inference**
ARMA–GARCH vs ARMA	−4.057	4.96×10^{-5}	−5.740	9.49×10^{-9}	Model 1 better
ARMA–EGARCH vs ARMA	−4.057	4.96×10^{-5}	−5.740	9.49×10^{-9}	Model 1 better
ARMA–GARCH					Difference
GARCH vs EGARCH	4.133	3.59×10^{-5}	7.476	7.64×10^{-14}	Model 2 better

. The results in Table 9 provide a formal statistical comparison of model performance. Both the ARMA–GARCH and ARMA–EGARCH hybrids record large negative DM statistics ($DM_{MSE} = -4.057$, $DM_{MAE} = -5.740$) with highly significant p–values ($p < 0.001$), indicating that these hybrid models significantly outperform the pure ARMA model in terms of both squared and absolute forecast errors. These results indicate that explicitly modelling volatility through

GARCH-type processes leads to a meaningful improvement in forecast accuracy for the NSE 30 stock price series.

EGARCH remains valuable from a risk-modelling perspective because it is capable of capturing asymmetric volatility patterns. In contrast, the positive and highly significant DM statistics obtained when comparing the symmetric GARCH model with the asymmetric EGARCH specification ($DM_{MSE} = 4.133$ and $DM_{MAE} = 7.476$, both $p < 0.001$) indicate that EGARCH provides statistically superior forecasts relative to the standard GARCH framework. This reinforces evidence of leverage effects in the NSE 30 returns and supports the application of asymmetric volatility models in emerging market settings [1, 9, 11]. Overall, the DM test outcomes are consistent with the descriptive analysis and model evaluation results presented earlier.

3.2. **Robustness Analysis: COVID and Post-COVID Sub-Periods.** Given that the sample period (January 2020 September 2025) includes the COVID-19 crisis, which was associated with exceptional volatility in global and emerging markets, a sub-period robustness analysis was conducted to assess the stability of the results.

The full sample was partitioned into two economically meaningful regimes:

- **COVID-19 high-volatility period:** January 2020 December 2021
- **Post-COVID recovery period:** January 2022 September 2025

The same modelling framework (ARMA(1,2), GARCH(1,1), EGARCH(1,1), ARMAGARCH, and ARMAEGARCH) was re-estimated within each sub-sample, and one-step-ahead out-of-sample forecasts were generated using a 20% hold-out procedure consistent with the baseline analysis. The results indicate that forecast errors were naturally higher during the COVID-19 regime due to elevated volatility and structural uncertainty. However, the relative performance ranking of the competing models remained unchanged. In both sub-periods, the hybrid ARMAGARCH and ARMAEGARCH models produced the lowest MAE, RMSE, and MAPE values, followed by the standalone EGARCH and GARCH models, with the pure ARMA model yielding the largest forecast errors.

Importantly, DieboldMariano tests conducted separately for each sub-sample confirm that the hybrid models significantly outperform the ARMA benchmark at conventional significance levels. This demonstrates that the superiority of the hybrid modelling framework is not driven solely by the extreme volatility observed during the pandemic but remains robust under more stable market conditions. These findings strengthen the external validity of the main results by showing that the hybrid approach maintains its predictive advantage across distinct volatility regimes.

4. Discussion

The empirical results from the NSE 30 daily stock prices provide strong support for hybrid modelling approaches that combine mean-equation dynamics with explicit treatment of volatility. The pure mean model, ARMA(1,2), estimated on the log-price series, offered a useful benchmark but recorded the largest forecast errors (MAE = 558.90; RMSE = 780.56; MAPE = 12.25%). This indicates that, although ARMA is effective in capturing trend and serial dependence, it is less suited to representing the evolving risk structure that characterises financial markets.

Introducing volatility dynamics through the symmetric GARCH(1,1) model on returns led to a marked improvement in forecast performance (MAE = 382.37; RMSE = 602.07; MAPE = 8.13%). This suggests that accounting for conditional heteroskedasticity is particularly important in the Nigerian equity market environment. Further gains were achieved with the asymmetric EGARCH(1,1) specification (MAE = 258.05; RMSE = 457.73; MAPE = 5.35%), whose sensitivity to negative shocks aligns with the leverage effect commonly reported in equity markets [1]. This makes asymmetric variance models especially relevant in periods of heightened uncertainty.

The hybrid ARMA–GARCH and ARMA–EGARCH models produced the lowest forecast errors overall (MAE = 209.70; RMSE = 287.60; MAPE = 4.72%), showing that combining ARMA mean dynamics with GARCH-type volatility processes delivers substantial accuracy gains. The Diebold–Mariano test reinforces this conclusion: both hybrid models outperform the pure ARMA model at conventional significance levels ($DM_{MSE} = 4.057$, $p < 0.001$; $DM_{MAE} = 5.740$, $p < 0.001$), with negative statistics indicating lower loss relative to the ARMA benchmark [13]. The comparison between ARMA–EGARCH and ARMA–GARCH yielded undefined statistics because the models produced identical forecast error series for this sample. This occurred because the one-step-ahead conditional mean forecasts were numerically identical across the hybrid specifications, resulting in zero loss differentials. However, comparing GARCH and EGARCH individually produced positive and highly significant DM statistics ($\mathrm{DM}_{\mathrm{MSE}} = 4.133$, $p < 0.001$; $\mathrm{DM}_{\mathrm{MAE}} = 7.476$, $p < 0.001$), confirming that the asymmetric structure in EGARCH adds measurable value beyond the standard GARCH formulation.

From a practical perspective, these findings suggest that forecasters of the NSE 30 index benefit from models that jointly represent predictable mean behaviour and time-varying volatility. For policymakers and risk managers, the strong persistence and asymmetry observed in volatility underline the risks of ignoring conditional variance when evaluating market stability, portfolio exposure, or systemic vulnerability.

There remain, nevertheless, some limitations and areas for caution that future work may address. The modelling framework remains univariate, relying solely on historical price data; the inclusion of macro-financial variables (such as exchange-rates or liquidity proxies) might further improve explanatory and predictive power. Additionally, the out-of-sample horizon (20 %) is relatively short, so the generalisability of the results to longer forecasting horizons remains open. Future research could extend the approach to multivariate frameworks (e.g., VAR-GARCH) or benchmark against machine-learning alternatives to test whether classical hybrids still maintain their advantage in more complex environments.

5. Conclusion

This study investigated price forecasting for the NSE 30 index using a sequence of models: ARMA (pure mean), GARCH and EGARCH (pure volatility), and two hybrid specifications (ARMA-EGARCH and ARMA-GARCH). The major findings are three-folds. First, the ARMA(1,2) model acted as an appropriate baseline for capturing log-price dynamics but did not sufficiently model volatility. Second, volatility modelling in the form of GARCH(1,1) and EGARCH(1,1) improved predictive accuracy, with EGARCH offering superior performance by capturing asymmetry in volatility responses. Third, the hybrids combining ARMA mean equations with GARCH-type variance processes achieved the best results, with statistically significant improvements confirmed by Diebold-Mariano tests of predictive accuracy.

In practical terms, investors and analysts in the Nigerian equity market should consider forecasting frameworks that incorporate both trend-components and conditional variance dynamics. For regulators and risk analysts, the evidence of volatility persistence and asymmetry emphasises the need to employ more sophisticated risk-monitoring tools rather than models assuming constant variance.

Although this research provides a comprehensive univariate template, further work extending the scope of covariates, applying longer forecast horizons, and comparing with cutting-edge machine-learning approaches would bolster our understanding of emerging market stock-price forecasting. However, the results conclude that hybrid econometric models are competitive tools in the forecasting techniques. This study extends prior Nigerian ARIMA-based studies by formally integrating conditional heteroskedasticity with statistical forecast comparison via the DieboldMariano framework, providing new empirical evidence on volatility asymmetry in the NSE 30 index.

Acknowledgment. The authors would like to express their sincere gratitude to the anonymous reviewers and the editor for their valuable comments and suggestions, which significantly improved the quality and clarity of this manuscript. The authors also acknowledge the support and academic environment provided by their respective institutions, University of Lagos and Federal Polytechnic Ilaro that facilitated the completion of this research.

Authors Contributions. Stephen Olusegun Are conceptualized the study, designed the methodological framework, carried out the data analysis and software implementation, and prepared the initial draft of the manuscript. Josephine Nneamaka Onyeka-Ubaka contributed to the supervision of the research, provided critical revisions to the manuscript, and contributed to the interpretation of the results and improvement of the overall academic content. Both authors read and approved the final version of the manuscript.

Authors' Conflicts of interest. The authors declare that there are no conflicts of interest regarding the publication of this paper. The research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest. The submitting author confirms that all coauthors have reviewed and agreed to this declaration.

Funding Statement. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. The authors conducted this study as part of their academic research activities at their respective institutions. No financial support, payments, goods, or services were received that could have influenced the outcomes of this work.

References

[1] D. B. Nelson. *Conditional Heteroskedasticity in Asset Returns: A New Approach*. Econometrica. 59(2) (1991), 347–370.
[2] A. A. Adebiyi, A. O. Adewumi, C. K. Ayo. *Stock Price Prediction Using the ARIMA Model*. UKSim-AMSS 16th Int. Conf. on Computer Modelling and Simulation (2014), 106–112.
[3] P. Zikova, J. Vesela. *Forecasting Stock Market Indices Using ARIMA and ARIMA–GARCH Models*. Journal of Risk and Financial Management. 16(1) (2023), 1–22.
[4] S. Li. *Stock Closing Price Prediction Based on the ARIMA-GARCH Model*. Financial & Business Economics Management (FBEM) (2023).
[5] R. Pandey. *Stock Price Prediction: An Integrated Review of Traditional and Modern Computational Models*. SSRN Working Paper (2025).
[6] M. Adhikari. *Forecasting Stock Index Closing Points Using ARIMA-GARCH with a Rolling Data Window*. Int. J. Science and Research Archive. 13(1) (2024), 3115–3125.
[7] R. F. Engle. *Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation*. Econometrica. 50(4) (1982), 987–1007.
[8] T. Bollerslev. *Generalized Autoregressive Conditional Heteroskedasticity*. Journal of Econometrics. 31 (1986), 307–327.
[9] S. Schmidt. *Forecasting Nordic Stock Market Volatility During COVID-19 Using GARCH Family Models*. Ume School of Business, Economics and Statistics, Ume University (2021).
[10] P. D. Yoo. *Forecasting the KOSPI Stock Index Using an Autoregressive Integrated Moving Average Model*. Asia–Pacific Financial Markets. 14(1) (2007), 1–16.
[11] S. Toma. *Empirical Comparison of ARIMA and GARCH Models for Stock Price Forecasting*. Informatica Economica. 27(3) (2023), 61–72.
[12] M. Gao. *Research on the Application of ARIMA–GARCH Model in Stock Price Forecasting*. E3S Web of Conferences. 276 (2021), 02030.
[13] F. X. Diebold, R. S. Mariano. *Comparing Predictive Accuracy*. Journal of Business & Economic Statistics. 13(3) (1995), 253–263.
[14] U. Triacca. *The Diebold–Mariano Test: Comparing Predictive Accuracy of Two Forecasts*. Lecture Notes, University of LAquila (2016).
[15] C. M. Jarque, A. K. Bera. *Efficient Tests for Normality, Homoscedasticity and Serial Independence of Regression Residuals*. Economics Letters. 6 (1980), 255–259.
[16] G. M. Ljung, G. E. P. Box. *On a Measure of Lack of Fit in Time Series Models*. Biometrika. 65(2) (1978), 297–303.
[17] M. I. Ekum, T. O. Owolabi, T. Alakija. *Modeling Volatility in Selected Nigerian Stock Market*. Int. J. Economics and Financial Management. 3(1) (2018), 30–43.
[18] G. E. P. Box, G. M. Jenkins, G. C. Reinsel, G. M. Ljung. *Time Series Analysis: Forecasting and Control*. John Wiley & Sons. New Jersey (2008).
[19] X. Huang, P. You, X. Gao, D. Cheng. *Stock Price Prediction Based on ARIMA-GARCH and LSTM*. Proceedings of ICBIS (2023), 438–448.

Stephen Olusegun Are*
Department of Mathematics and Statistics, Federal Polytechnic Ilaro, Ogun State, Nigeria.
E-mail address: `stephen.are@federalpolyilaro.edu.ng`

Josephine Nneamaka Onyeka-Ubaka
Department of Statistics, University of Lagos, Akoka, Lagos State, Nigeria.
E-mail address: jonyeka-ubaka@unilag.edu.ng