International Journal of Innovative Mathematics, Statistics & Energy Policies
14(3):201-223, July-Sept, 2026

© SEAHI PUBLICATIONS, 2026 www.seahipublications.org ISSN: 2467-852X

DOI:10.5281/zenodo.22877825

Modelling Non-Linear Volatility In The Nigerian Stock Market Using An Integrated EGARCH–MLP Approach

Loveth Bukola Kolawole[1]; Prof. Isaac D. Essi[2] & Dr. G. L. Tuaneh[3]

[1,2,3]Department of Mathematics, Faculty of Science, Rivers State University, Port Harcourt, Nigeria
[1]Kolawole.loveth@rsu.edu.ng

ABSTRACT
This study models the non-linear and time-varying volatility dynamics of the Nigerian stock market using an integrated Exponential Generalized Autoregressive Conditional Heteroskedasticity (EGARCH) and Multilayer Perceptron (MLP) framework. Daily Nigerian All-Share Index (ASI) observations from 30 January 2012 to 31 December 2024 were obtained from Investing.com, yielding 3,199 trading-day records. Log returns were computed from closing index values, while a range-based Parkinson volatility measure derived from daily high and low values was used as a common one-day-ahead volatility target for forecast evaluation. The return series was stationary but non-normal and heavy-tailed (kurtosis = 8.815; Jarque–Bera = 4557.796, $p < .001$), and the ARCH-LM test confirmed significant conditional heteroskedasticity (LM = 413.272, $p < .001$). An EGARCH(1,1) model with Student-t innovations provided a substantially better fit than the Gaussian specification (AIC = 7327.995 versus 7897.727). Its persistence parameter was high and significant ($\beta = 0.922$, $p < .001$), confirming clustered and persistent volatility, while the asymmetry coefficient was not statistically significant ($\gamma = 0.019$, $p = .343$), indicating that leverage effects were not robust in this sample. A standalone MLP and an integrated EGARCH–MLP were then trained chronologically using lagged returns, absolute returns, squared returns and range-based volatility, with the integrated model additionally receiving the one-step-ahead EGARCH volatility forecast. On the 2022–2024 out-of-sample period, the integrated EGARCH–MLP produced the lowest MSE (0.000017), RMSE (0.004114) and MAE (0.002920), and the only positive R^2 (0.0438), improving MLP MAE by 3.85%. The study concludes that Nigerian ASI volatility is dynamic and clustered rather than uniformly high: volatility peaked in stress periods such as 2015 and 2020, while 2024 annualized volatility was about 12.8%. The integrated framework therefore offers a statistically coherent way to combine EGARCH's volatility structure with MLP's non-linear learning for market-risk assessment and short-horizon volatility forecasting in Nigeria.
Keywords: Nigerian All-Share Index; volatility; EGARCH; Multilayer Perceptron; hybrid model; non-linear dynamics; volatility forecasting.

1. INTRODUCTION

Volatility is a central characteristic of equity markets because it represents the changing degree of uncertainty around returns. In an emerging market such as Nigeria, the All-Share Index (ASI) aggregates the performance of listed equities and therefore provides a broad indicator of market conditions. Recent Nigerian evidence continues to document volatility clustering, persistence, structural breaks and asymmetric responses in market returns, showing that market risk is not constant through time (Adegboyo & Sarwar, 2025).

The Nigerian stock market has passed through markedly different volatility regimes during the last decade. Macroeconomic adjustment, exchange-rate changes, oil-price shocks, political events and the COVID-19 disruption have all affected investor expectations and the dispersion of market returns. The statistical implication is that a single unconditional variance is inadequate for describing risk. Instead, volatility should be treated as a dynamic process that depends on past shocks and past volatility. (Adegboyo & Sarwar, 2025; Eme et al., 2024)

The Exponential Generalized Autoregressive Conditional Heteroskedasticity model is particularly relevant because it models the logarithm of conditional variance, accommodates volatility persistence and allows the sign of a shock to affect future variance differently. Recent Nigerian studies have found that EGARCH-type specifications remain useful for identifying persistence and asymmetry in local market data. However, econometric models impose a functional structure that may not capture all non-linear interactions in financial data. (Francq & Zakoian, 2019; Adegboyo & Sarwar, 2025)

The Multilayer Perceptron (MLP) provides a complementary approach. It is a feed-forward neural network that learns non-linear mappings between input features and a continuous output without requiring a pre-specified linear functional form. Recent stock-index forecasting studies show that MLP models can exploit non-linear patterns in lagged market information, although their performance depends on feature construction, tuning and out-of-sample validation. (James et al., 2021; Ecer et al., 2020)

A growing literature therefore combines volatility models with neural networks so that the econometric component acts as a structured feature extractor and the neural network learns remaining non-linear relationships. Hybrid GARCH–neural-network studies have reported gains over single models in several markets. This study applies that logic directly to the Nigerian ASI by integrating the one-step EGARCH volatility forecast into an MLP that predicts next-day range-based volatility. (Kim & Won, 2018; Koo & Kim, 2022)

Modelling volatility is the statistical representation, estimation and forecasting of changes in the dispersion of financial returns. Volatility does not indicate whether the market will rise or fall; it describes the expected magnitude and variability of price movements. Because the true variance of returns is not observed directly, it is inferred from return behaviour or approximated with a realized measure constructed from market prices. (Francq & Zakoian, 2019)

The principal dimensions of volatility examined in this study are time variation, clustering, persistence, asymmetry, heavy-tailed behaviour and forecast accuracy. Time variation means that risk changes rather than remaining constant. Clustering describes the tendency for turbulent observations to occur near other turbulent observations. Persistence measures how slowly the effect of a shock disappears, while asymmetry asks whether positive and negative shocks of equal magnitude affect future volatility differently. Heavy-tailed behaviour refers to extreme returns occurring more frequently than predicted by a normal distribution. Forecast accuracy concerns how closely a model's predicted volatility corresponds to the next observed volatility target. (Francq & Zakoian, 2019; Adegboyo & Sarwar, 2025)

The empirical target is one-day-ahead Parkinson volatility, which uses the daily high-low range. This range-based measure captures more within-day price movement than a measure based only on consecutive closing values. The target is kept identical for EGARCH, MLP and the integrated model so that differences in their reported errors reflect modelling performance rather than differences in what they were asked to predict. (Molnár, 2016)

Non-linear dynamics arise when the effect of market information is not adequately represented by a constant or proportional relationship. A large shock may affect risk differently from several small shocks; recent volatility may interact with return magnitude; and market behaviour may change between calm and stressed conditions. These patterns make the future volatility response dependent on both the size and the configuration of recent observations. (Lu et al., 2017; James et al., 2021)

The term non-linear in this dissertation has two related meanings. In EGARCH, conditional variance is represented on a logarithmic scale and responds to both the magnitude and sign of standardized innovations. In the MLP, non-linearity is introduced through hidden-layer activation functions that transform combinations of lagged market variables. The study therefore examines non-linearity through both an explicitly structured econometric equation and a data-driven learning function. (Francq & Zakoian, 2019; James et al., 2021)

The Nigerian stock market is the organized equity market through which listed companies raise capital and investors trade ownership claims. Its principal exchange is the Nigerian Exchange, while the Nigerian All-Share Index provides a broad measure of the aggregate price performance of eligible listed equities. The ASI is used in this study because it represents market-wide behaviour more directly than the price of a single company. (Adegboyo & Sarwar, 2025)

Emerging-market characteristics such as changing liquidity, market concentration and abrupt responses to information can produce non-normal returns and persistent volatility. The ASI is therefore an appropriate series for examining whether risk is time-varying, whether shocks decay slowly and whether a hybrid forecasting system can extract useful information from both econometric and nonlinear representations. (Adegboyo & Sarwar, 2025; Ali et al., 2022)

An integrated EGARCH-MLP approach combines the two models within one forecasting pipeline. The EGARCH model first converts the return history into a one-step conditional-volatility forecast. That forecast is appended to the

lagged market variables used by the MLP, allowing the network to learn whether the econometric variance signal improves prediction of the same next-day Parkinson volatility target. (Kim & Won, 2018; Koo & Kim, 2022)

This feature-level connection distinguishes an integrated model from a simple comparison of unrelated outputs. EGARCH contributes a compact and interpretable representation of persistence, shock magnitude and conditional variance. The MLP contributes the capacity to learn nonlinear interactions between that representation and recent return- and range-based features. Their roles are complementary rather than interchangeable. (Kartsonakis Mademlis & Dritsakis, 2021; Koo & Kim, 2022)

The integrated approach is assessed by comparing EGARCH, standalone MLP and EGARCH-MLP forecasts over identical out-of-sample dates using MSE, RMSE, MAE and R-squared. A lower error indicates incremental forecast value within the observed sample; it does not establish that every hybrid architecture will dominate in all markets or future regimes. (Petropoulos et al., 2022)

The central empirical concern of this study is to determine the extent to which volatility in the Nigerian All-Share Index exhibits persistence, temporal variation and non-linear dynamics. It further examines whether an integrated EGARCH-MLP model can forecast these evolving risk patterns more accurately than either model applied independently. EGARCH and MLP are therefore positioned as the principal analytical models, while alternative specifications are used only as limited benchmarks where necessary.

A second problem in volatility research is target inconsistency: an EGARCH model forecasts conditional variance, while many machine-learning studies evaluate return-level predictions. Comparing such outputs directly can produce misleading conclusions. To avoid this mismatch, this study evaluates EGARCH, MLP and the integrated EGARCH–MLP against the same one-day-ahead Parkinson range-based volatility proxy constructed from the ASI daily high and low values. This creates a common statistical basis for model comparison. (Petropoulos et al., 2022; Molnár, 2016)

1.6 Significance of the Study

The significance of this study lies in its integration of econometric modelling and non-linear machine learning within a common volatility-forecasting framework. For investors and portfolio managers, improved volatility forecasts can support more effective position sizing, risk-limit determination and portfolio allocation. For regulators, evidence of persistent volatility may identify periods that require closer market surveillance. Methodologically, the study contributes Nigerian evidence to the growing literature on hybrid volatility forecasting by employing EGARCH-generated estimates as inputs to an MLP rather than merely comparing the separate forecasts of unrelated models. (Koo & Kim, 2022; Araya et al., 2024)

1.7 Scope of the Study

The empirical analysis covers 3,199 daily ASI observations from 30 January 2012 to 31 December 2024. The period is determined by the complete historical observations available from the selected Investing.com download. The study focuses on index-level volatility and uses closing, high and low ASI values; it does not model individual listed equities or incorporate macroeconomic variables directly. The out-of-sample forecasting evaluation covers approximately the final 20% of usable observations, from June 2022 to December 2024.

2. Literature Review

2.1 EGARCH, MLP and Hybrid Volatility Forecasting

The conceptual framework is sequential. Historical ASI closes are transformed into log returns; high and low index values produce a range-based realized-volatility proxy; EGARCH estimates the persistent conditional-volatility process; the MLP learns non-linear relationships in lagged returns and volatility; and the integrated model adds the EGARCH one-step forecast as an additional feature. Forecast accuracy is then evaluated on a chronological hold-out sample.

Recent Nigerian evidence confirms that stock-market volatility is time-varying and sensitive to market shocks. Ibrahim (2017) reported persistent volatility and found EGARCH useful for Nigerian stock-market forecasting. Kuhe (2018) showed that persistence and asymmetry estimates can be affected by structural breaks and found an EGARCH specification with Student-t innovations to be appropriate. More recently, Ucheoma (2024) demonstrated that explicitly treating structural breaks can reduce measured volatility persistence, reinforcing the need to interpret long-run persistence carefully.

Awogbemi et al. (2024) used EGARCH-X specifications with alternative innovation distributions and found that non-Gaussian innovations better accommodated skewness and kurtosis in Nigerian financial returns. Adegboyo and Sarwar (2025), using a period very close to the present dataset, documented volatility clustering and asymmetric effects in the Nigerian ASI and emphasized that heavy-tailed distributions improve volatility modelling. These

studies justify an EGARCH specification with Student-t errors as the primary econometric component of the present research.

2.1.4 EGARCH as the Core Econometric Model

EGARCH extends conventional conditional-variance modelling by expressing the logarithm of variance as a function of previous variance and standardized shocks. The log formulation guarantees positive conditional variance without requiring all variance parameters to be non-negative, while the sign term permits positive and negative shocks to have different volatility effects. This flexibility is useful for equity markets where large shocks and persistent variance are common, and recent studies continue to apply EGARCH to heavy-tailed and potentially asymmetric financial data (Francq & Zakoian, 2019; Adegboyo & Sarwar, 2025; Awogbemi et al., 2024; Eme et al., 2024).

2.1.4.3 Persistence, Mean Reversion and Shock Decay

Volatility persistence concerns the duration of a shock's influence. When the coefficient on lagged log variance is close to one in absolute value, the conditional-variance state changes slowly and shocks decay gradually. A lower coefficient implies faster reversion toward the model's typical variance level. Persistence is not equivalent to a unit root in returns; it is a property of the conditional second moment. The return series may remain stationary while volatility is highly persistent. (Francq & Zakoian, 2019; Adegboyo & Sarwar, 2025)

2.1.4.4 Asymmetry and the News-Impact Mechanism

The signed innovation coefficient tests whether shocks of equal absolute size have different effects according to direction. In a common EGARCH parameterization, the response to a positive standardized shock combines the magnitude and sign coefficients, while the response to a negative shock combines them differently. The exact interpretation depends on the equation and software convention used. Researchers must therefore identify the signed term explicitly rather than attaching a leverage label to a coefficient by position alone. (Francq & Zakoian, 2019) (Francq & Zakoian, 2019; Adegboyo & Sarwar, 2025)

2.1.4.5 Innovation Distributions and Heavy Tails

Maximum-likelihood estimation requires a conditional distribution for standardized innovations. The Gaussian assumption is convenient, but it can allocate too little probability to extreme observations. When a few large returns occur, the estimation may respond by altering variance parameters in order to accommodate tail events that are distributional rather than purely dynamic. A Student-t innovation provides a separate tail-thickness parameter and can distinguish heavy-tailed residual behaviour from volatility persistence more effectively. (Francq & Zakoian, 2019; Awogbemi et al., 2024)

2.1.5 Multilayer Perceptron as the Core Machine-Learning Model

MLP models are useful when the relationship between lagged market information and the forecast target is non-linear. Ecer et al. (2020) showed that MLP-based approaches can successfully model stock-index movement when network parameters are appropriately optimized. Gao et al. (2020) and Žmuk and Jošić (2020) similarly found that MLPs can provide useful stock-index forecasts, although forecast quality depends on the forecasting horizon, data window and network configuration. Nigerian machine-learning studies also show that non-linear methods can add predictive value in local equity applications (Oyewola et al., 2021).

The present study deliberately uses a relatively parsimonious MLP rather than a deep recurrent architecture. This choice aligns the model with the sample size and emphasizes the central methodological question: whether EGARCH-derived volatility information adds useful predictive content to a standard feed-forward non-linear learner. (James et al., 2021; Géron, 2022)

2.1.5.2 Feature Engineering for Volatility Forecasting

Machine-learning performance depends strongly on the information represented in the input vector. The baseline features in this study comprise five lags each of return, absolute return, squared return and Parkinson volatility. Raw returns retain direction, absolute returns represent movement magnitude, squared returns connect with conventional variance proxies, and lagged Parkinson volatility supplies recent information from the intraday range. Together they provide complementary descriptions of recent market behaviour. (James et al., 2021)

2.1.5.5 Time-Series Validation and Hyperparameter Selection

Randomly shuffling financial observations breaks the information ordering that defines a forecast. It can allow observations from a later regime to influence a model assessed on an earlier date and can distribute adjacent, highly related observations across training and validation sets. A time-series split preserves chronology by training on an earlier segment and validating on the next segment. Successive folds can use an expanding training window so that each decision resembles a real forecasting exercise. (Bergmeir et al., 2018)

2.1.6 Integrated EGARCH–MLP Modelling

The strongest justification for integration comes from recent hybrid-volatility research. Kartsonakis Mademlis and Dritsakis (2021) reported that a hybrid model using returns and EGARCH conditional-volatility estimates as neural-network inputs outperformed standalone alternatives for the Italian stock market. Amirshahi and Lahmiri (2023) likewise found that GARCH-family outputs can serve as informative features that improve feed-forward and recurrent neural-network volatility forecasts. Araya et al. (2024) reported substantial error reductions from hybrid GARCH–deep-learning architectures, while Bu et al. (2025) showed that GARCH outputs can be used as feature extractors for next-day stock-index volatility prediction.

These studies provide a direct methodological precedent for the present design: EGARCH is first estimated from ASI returns, its one-day-ahead volatility forecast is then passed to the MLP, and the hybrid forecast is evaluated against both standalone models on the same volatility target.

2.1.6.2 Why EGARCH Can Add Information to an MLP

Although the MLP receives lagged returns and volatility measures, it does not automatically reproduce the EGARCH recursion. EGARCH compresses the entire filtered history into a current conditional-variance state using parameters estimated from the likelihood. This state carries decaying information from shocks older than the five explicit MLP lags. It also standardizes shocks relative to predicted variance and incorporates the selected heavy-tailed distribution during estimation. (Kartsonakis Mademlis & Dritsakis, 2021; Koo & Kim, 2022)

2.1.6.3 Common-Target and Common-Horizon Requirements

A statistically meaningful comparison requires all models to forecast the same quantity over the same horizon. Conditional standard deviation from a return model, next-day return and range-based realized volatility are related but distinct targets. Ranking their numerical errors in one table without transformation would be misleading because the loss functions refer to different dependent variables. The common-target design is therefore a central methodological contribution rather than a minor reporting choice. (Petropoulos et al., 2022)

2.1.6.5 Interpretation of Hybrid Improvement

The hybrid achieves the lowest MSE, RMSE and MAE and the only positive out-of-sample R-squared. These results show that the EGARCH feature improves the selected MLP under the reported configuration. The magnitude of improvement is modest: this is consistent with complementarity at the margin rather than a radical change in forecastability. Small gains can still matter in repeated risk decisions, but they should not be described as universal dominance. (Petropoulos et al., 2022)

2.2 Empirical Evidence and Research Gap

The empirical literature is evaluated by distinguishing in-sample fit, parameter interpretation and genuine out-of-sample forecasting performance. These criteria answer different questions. Information criteria assess the balance between likelihood and complexity, EGARCH coefficients describe persistence and possible asymmetry, and forecast-loss measures evaluate proximity to a chosen volatility proxy. Consequently, the best-fitting specification need not produce the smallest error on a later period, and a neural network may reduce MAE without providing the parameter-level explanation available from EGARCH (Petropoulos et al., 2022; Francq & Zakoian, 2019).

Ibrahim (2017) examined Nigerian stock-market volatility over 1985–2014 using ARIMA and several ARCH-family models and reported that EGARCH was comparatively effective for volatility forecasting and the representation of asymmetric effects. That evidence agrees with the present study on volatility clustering and the usefulness of EGARCH. It differs, however, in sample period, forecast design and target construction: the current research uses daily ASI data through 2024, evaluates a common one-day-ahead range-based target and extends the econometric model with an MLP feature-level integration.

Usman et al. (2017) compared eleven GARCH-type specifications using monthly Nigerian market returns and found that model rankings differed between training and testing periods. Their result is relevant because it shows that superior in-sample fit is not a guarantee of superior forecasting performance. The current study addresses the same concern with a chronological hold-out sample, although it differs by using daily observations, a Parkinson volatility target and a restricted comparison centred on EGARCH, MLP and their integration.

Iwada et al. (2018) analysed daily prices of Access Bank and Fidelity Bank with symmetric and asymmetric GARCH specifications under Normal, Student-t and generalized-error innovations. Their preferred model varied with the security and evaluation measure, while EGARCH performed well under RMSE for part of the comparison. This supports the present study's attention to innovation distributions and multiple error metrics. The earlier work nevertheless used individual bank equities and did not investigate a neural-network hybrid, whereas the present analysis concerns the market-wide ASI and supplies the EGARCH forecast directly to an MLP.

Kuhe (2018) examined Nigerian stock-return persistence and asymmetry while accounting for exogenous structural breaks. The study showed that break treatment can materially influence estimated persistence and asymmetric

responses. This is consistent with the present conclusion that Nigerian volatility is time-varying, but it also introduces a difference: the current fixed-parameter EGARCH does not explicitly estimate break dates. The comparison therefore supports the recommendation for rolling or break-adjusted validation without invalidating the present controlled forecast experiment.

Emenyonu et al. (2023) modelled Fidelity Bank daily returns with ARMA-GARCH and ARMA-EGARCH specifications under alternative innovation distributions and selected an EGARCH model with Student-t innovations. This agrees with the current evidence that a heavy-tailed EGARCH fits Nigerian equity returns better than a Gaussian version. The variables and forecasting horizons are different, however, because their study used a single bank's price returns and MASE-based long-horizon projections, whereas the present study uses broad-market ASI returns and next-day range-based volatility.

Awogbemi et al. (2024) estimated EGARCH-X models for Nigerian stock prices under Normal, Student-t and skewed Student-t innovations. Their findings emphasize that non-Gaussian distributions and additional explanatory variables can affect fit, persistence and mean reversion. The present research is similar in testing heavy-tailed innovations, but it deliberately excludes macroeconomic covariates so that the incremental contribution of the EGARCH signal to the MLP can be isolated. The two designs therefore address related but distinct sources of forecasting information.

Ucheoma (2024) demonstrated that explicitly modelling structural breaks can reduce apparent long-run persistence in Nigerian stock-market volatility. Isenah (2024), using bagged GARCH-type models and ensemble averaging, reported that ensemble methods improved forecast errors relative to classical GARCH alternatives. Together, these studies indicate that both regime treatment and model combination can change forecast performance. The present study follows the combination logic but differs by using an interpretable EGARCH one-step forecast as the twenty-first MLP feature rather than averaging forecasts from multiple GARCH estimators.

Adegboyo and Sarwar (2025) provide the closest econometric comparison because they analysed daily Nigerian ASI observations beginning on 30 January 2012 and extending to October 2024. They reported clustering, heavy tails and significant leverage effects, with a Student-t GJR-GARCH performing strongly. The present findings agree on clustering, persistence and the value of a heavy-tailed distribution, but differ on asymmetry because the EGARCH signed-shock coefficient is not statistically significant in the current sample. Differences in sample endpoint, model family, parameterization and volatility target provide plausible explanations and show why leverage should not be treated as universal.

Research using MLP and related neural networks also provides mixed evidence. Ecer et al. (2020) showed that stock-index prediction depends materially on optimization and network configuration, while Gao et al. (2020) and Žmuk and Jošić (2020) found that performance varies by index, horizon and algorithm. Oyewola et al. (2021) supplied Nigerian evidence that nonlinear learning can extract useful information from oil-and-gas equity variables. These studies support the representational capacity of neural networks but mostly predict prices, returns or direction. The current study differs by training the MLP directly on next-day range-based volatility, thereby preventing a return-versus-volatility target mismatch.

International hybrid studies support integration but use different markets and variables. Kristjanpoller and Hernández (2017) combined ANN and GARCH forecasts with exchange rates, oil prices and international stock indices to forecast metal-price volatility. Kim and Won (2018) integrated LSTM with multiple GARCH-type estimates for a stock-price index, while Kartsonakis Mademlis and Dritsakis (2021) applied hybrid GARCH-neural-network models to the Italian market. Koo and Kim (2022) added a distribution-manipulation strategy to a GARCH-LSTM system. These studies generally reported benefits from integration, but their use of recurrent networks, exogenous variables or developed-market data differs from the parsimonious MLP and Nigerian ASI focus adopted here.

More recent hybrid evidence extends the comparison. Kakade et al. (2022) combined GARCH and LSTM components for commodity-return volatility and risk forecasting; Amirshahi and Lahmiri (2023) combined deep learning with GARCH-family models for cryptocurrency volatility; Araya et al. (2024) integrated GARCH-family models with convolutional and bidirectional recurrent components; and Bu et al. (2025) combined LSTM with multiple GARCH estimates for the CSI 300 Index. Their findings broadly suggest that structured conditional-variance information can improve nonlinear forecasting. Nevertheless, the assets, data frequencies, architectures and targets differ, so their results cannot establish in advance that an EGARCH–MLP must dominate for the Nigerian ASI.

Ekpete et al. (2026) applied a rolling ARIMA-GARCH hybrid to daily NGX ASI closing prices from 2021 to 2025 and reported lower forecast errors than standalone ARIMA and GARCH alternatives. This recent Nigerian evidence agrees with the present finding that integration can add predictive information and reinforces the importance of

evaluation on later observations. The designs remain distinct: their hybrid combines a conditional-mean model with GARCH and uses a rolling window, whereas the present model integrates an EGARCH volatility signal with an MLP and evaluates all models against a common range-based volatility target.

Across the reviewed studies, the most stable points of agreement are that Nigerian equity returns exhibit time-varying variance, volatility shocks can be persistent, and non-Gaussian innovations frequently improve conditional-variance modelling. The principal differences concern the statistical significance of asymmetry, the preferred GARCH variant, the use of structural breaks or exogenous variables, the neural-network architecture and the definition of the forecast target. These differences justify the present design: Student-t EGARCH supplies an interpretable persistence signal, the MLP learns nonlinear relations among lagged market features, and the integrated model is judged on identical dates and an identical one-day-ahead volatility measure.

The recent Nigerian literature is rich in GARCH-family analysis but comparatively limited in integrated EGARCH–MLP volatility modelling on the broad ASI. Existing machine-learning work often predicts prices or returns rather than volatility, while econometric studies often stop at ranking GARCH variants. The gap addressed here is therefore a common-target, out-of-sample test of whether an EGARCH volatility signal improves an MLP forecast of next-day ASI volatility. (Adegboyo & Sarwar, 2025; Koo & Kim, 2022)

The literature reviewed in this chapter establishes that stock-market volatility is persistent, clustered and frequently non-Gaussian, making constant-variance models inadequate for the Nigerian All-Share Index. The conceptual discussion identifies EGARCH as an appropriate econometric framework for modelling conditional variance and possible asymmetric responses, while the MLP provides a flexible means of learning non-linear relationships among lagged volatility indicators. The theoretical perspectives further explain why new information, changing expectations and market frictions may produce time-varying and non-linear volatility dynamics. (Francq & Zakoian, 2019; James et al., 2021; Koo & Kim, 2022)

The empirical evidence nevertheless shows that neither model class is uniformly superior in every market or evaluation period. EGARCH offers interpretable estimates of persistence and asymmetry, whereas an MLP can approximate relationships that are not fully captured by a fixed parametric specification. This comparison supports an integrated design in which the one-step-ahead EGARCH volatility estimate is supplied to the MLP as an additional feature and all competing models are assessed against the same realized-volatility target.

3. MATERIALS AND METHODS

A quantitative longitudinal time-series design was adopted. The dataset consists of 3,199 daily Nigerian All-Share Index records downloaded from the Investing.com historical-data page and covers 30 January 2012 to 31 December 2024. The variables used directly are Date, Close (reported as Price), High and Low. The source file also contains Open, Volume and percentage change, which are retained in the supplementary dataset but are not required by the primary model. (Petropoulos et al., 2022)

The statistical population comprises daily observations of the Nigerian All-Share Index over the period for which a complete and internally usable sequence of closing, high and low values was available from the selected source. The study does not draw a random sample of trading days. Instead, it uses the full sequence of 3,199 available records between 30 January 2012 and 31 December 2024. The unit of analysis is the trading-day observation, and the temporal order of those observations is an essential part of the research design.

Using the full available sequence reduces arbitrary sampling variation but does not remove temporal dependence. Adjacent records are related through returns, volatility and lagged features. Standard cross-sectional assumptions of independent sampling are therefore inappropriate. All transformations, partitions and validation procedures preserve chronology. The empirical conclusions apply to the ASI and the observed study period; they are not automatically generalizable to every listed security or to future market regimes. (Bergmeir et al., 2018)

The observations were sorted chronologically, numeric separators were removed, and the date field was parsed to a daily trading calendar. No duplicate trading dates were retained. Daily continuously compounded return was computed as: (Petropoulos et al., 2022)

$$r_t = \ln(P_t / P_{t-1}) \qquad (3.1)$$

where r_t is the log return on trading day t, P_t is the ASI closing value on day t, P_{t-1} is the previous trading-day closing value, and ln denotes the natural logarithm.

Data cleaning was designed to preserve genuine market movements while removing formatting problems that could obstruct numerical analysis. Dates were parsed into a consistent representation, thousands separators were removed from index values and the observations were arranged in ascending trading order. Duplicate

dates were checked so that no trading day received disproportionate weight. High, low and closing values were converted to numeric form and inspected for missing or impossible entries. (Petropoulos et al., 2022)

Logical checks required the daily high to be no lower than the daily low and all index levels to be positive, since logarithmic transformations are undefined for non-positive values. Non-trading calendar days were not imputed because the model operates on trading-time intervals. The previous observation in the return equation therefore means the previous recorded trading day, not necessarily the previous calendar day. This approach avoids manufacturing price paths over weekends and public holidays.

Extreme returns were retained because they may represent economically important stress rather than errors. Removing them solely because they are large would suppress the heavy-tailed behaviour the study is intended to model. Their influence is addressed through descriptive diagnostics and a Student-t conditional distribution. This treatment maintains a clear distinction between invalid data and valid but unusual market observations.

Because EGARCH forecasts volatility rather than the return level, the machine-learning models were evaluated on the same volatility target. The Parkinson range-based estimator was used because daily high and low ASI values are available and the measure is less noisy than a single squared close-to-close return. The daily range-based variance is: (Molnár, 2016; Petropoulos et al., 2022)

$$RV_t = [\ln(H_t / L_t)]^2 / [4 \ln(2)] \quad (3.2)$$

where RV_t is the Parkinson range-based variance, H_t is the ASI daily high, L_t is the ASI daily low, and ln(2) is the natural logarithm of 2. The volatility target is $\sigma^P_t = \sqrt{RV_t}$.

The Parkinson estimator uses the daily logarithmic range and is derived under assumptions that make the high-low interval informative about variance. Compared with a single squared close-to-close return, it incorporates the extent of movement within the day and is generally less noisy under its idealized setting. The available ASI dataset contains high and low observations, making the estimator practical without introducing a separate high-frequency source (Molnár, 2016).

The target is expressed as a non-negative daily volatility measure. It should be interpreted as a proxy for latent integrated volatility rather than a perfect observation. Opening jumps and market microstructure features are not represented explicitly, and the estimator's assumptions may not hold exactly in the Nigerian market. These limitations affect every competing model equally because all are evaluated against the same target. (Molnár, 2016)

Target alignment determines the timing of the modelling matrix. Features formed with information through trading day t are paired with Parkinson volatility on day t+1. The first observations are lost because five lags are required, and the final observation cannot serve as a feature origin without a subsequent target. After alignment and removal of transformation-induced missing values, 3,192 modelling observations remain. (Petropoulos et al., 2022)

The return distribution was summarized using mean, median, standard deviation, skewness and Pearson kurtosis. Normality was evaluated using the Jarque–Bera statistic: (Francq & Zakoian, 2019)

$$JB = n[S^2/6 + (K - 3)^2/24] \quad (3.3)$$

where n is the number of observations, S is sample skewness and K is Pearson kurtosis. The null hypothesis is that the returns are normally distributed.

The Augmented Dickey–Fuller (ADF) test and Kwiatkowski–Phillips–Schmidt–Shin (KPSS) test were used as complementary stationarity checks. For ADF, the null hypothesis is a unit root; for KPSS, the null hypothesis is stationarity. Joint rejection of the ADF unit-root null and non-rejection of KPSS stationarity supports modelling the return series in levels. (Petropoulos et al., 2022)

Conditional heteroskedasticity was tested by regressing squared residuals on q of their own lags: (Francq & Zakoian, 2019)

$$\varepsilon^{*2}_t = a_0 + \Sigma_{i=1}\text{^}q\ a_i\ \varepsilon^{*2}_{t-i} + u_t \quad (3.4)$$

where ε^{*2}_t is the squared mean-adjusted return residual, q is the number of lags, a_i are auxiliary-regression coefficients and u_t is an error term. The LM statistic is nR^2 and is compared with a chi-square distribution with q degrees of freedom. Ten lags were used.

The preliminary tests answer different questions and are interpreted jointly. Descriptive statistics indicate location, dispersion and tail behaviour. Jarque-Bera evaluates whether skewness and kurtosis conform to

normality. ADF and KPSS provide complementary evidence on stationarity because their null hypotheses are reversed. ARCH-LM tests whether past squared residuals explain current squared residuals. No single test is treated as a complete diagnosis of the series. (Francq & Zakoian, 2019)

The sequence matters. Modelling a non-stationary price level with a conditional-variance equation can confound trend and volatility, so returns are constructed first and tested for stationarity. Evidence of ARCH effects then establishes that constant variance is inadequate. Non-normality informs the comparison between Gaussian and Student-t innovations. The diagnostics therefore provide a chain of justification for the selected EGARCH specification rather than a disconnected collection of tests.

The primary econometric model is EGARCH(1,1) with standardized Student-t innovations. Returns were multiplied by 100 during numerical estimation to improve scale stability. The conditional mean is represented by $r_t = \mu + \varepsilon_t$, where $\varepsilon_t = \sigma_t z_t$. The conditional variance equation is: (Francq & Zakoian, 2019; Awogbemi et al., 2024)

$$\ln(\sigma^2_t) = \omega + \beta \ln(\sigma^2_{t-1}) + \alpha(|z_{t-1}| - E|z|) + \gamma z_{t-1} \qquad (3.5)$$

where σ^2_t is conditional variance at time t; ω is the variance constant; β measures volatility persistence; α measures the magnitude effect of standardized shocks; γ measures asymmetry or leverage; $z_t = \varepsilon_t/\sigma_t$ is the standardized innovation; and $E|z|$ centres the magnitude term. Under the Student-t specification, ν denotes the degrees-of-freedom parameter and $\nu > 2$ ensures finite variance.

The interpretation of γ is central to Objective 2. If γ is statistically different from zero, positive and negative shocks of the same magnitude have different effects on future volatility. The sign and size of γ determine the direction and strength of that asymmetry. A β estimate close to one indicates persistent volatility, while a smaller β implies quicker mean reversion.

The parameter vector $\theta = (\mu, \omega, \alpha, \gamma, \beta, \nu)$ was estimated by maximizing the standardized Student-t log-likelihood numerically. The optimizer minimized the negative log-likelihood subject to $\alpha > 0$, $|\beta| < 1$ and $\nu > 2$. Standard errors were obtained from the inverse numerical Hessian of the negative log-likelihood at the optimum. (Francq & Zakoian, 2019)

To verify the choice of heavy-tailed innovations, the EGARCH(1,1) model was also estimated under a normal distribution. The preferred innovation was selected by AIC and BIC: (Francq & Zakoian, 2019; Awogbemi et al., 2024)

$$AIC = 2k - 2\ln(\hat{L}) \qquad (3.6)$$

$$BIC = k\ln(n) - 2\ln(\hat{L}) \qquad (3.7)$$

where k is the number of estimated parameters, n is the number of observations and $\hat{L}$ is the maximized likelihood. Lower AIC and BIC values indicate a more parsimonious fit.

The EGARCH recursion requires an initial variance. A stable sample-based value is used to begin filtering, after which the influence of initialization diminishes as the recursion progresses. Returns are multiplied by 100 during likelihood estimation so that the optimizer works with percentage-point magnitudes rather than very small decimals. Conditional standard deviations are converted back to decimal scale before they are compared with Parkinson volatility or supplied to the MLP. (Francq & Zakoian, 2019)

Numerical optimization is conducted subject to the stated parameter conditions. Multiple components of the output are reviewed: convergence, likelihood, information criteria, coefficient estimates and standard errors. A successful convergence message alone is insufficient if the solution lies at an implausible boundary or produces inconsistent residual behaviour. The normal and Student-t models use the same return series and mean structure, allowing their information criteria to be compared directly.

The inverse numerical Hessian provides an approximation to the covariance matrix at the optimum. The resulting standard errors are used to calculate z-statistics and p-values. As with all asymptotic inference, these values depend on specification and sample properties. Parameter significance is therefore interpreted alongside model fit and economic meaning rather than as an isolated decision rule.

For each forecast origin, the conditional variance is updated from the preceding filtered variance and standardized innovation. The square root produces conditional standard deviation, which is expressed on the same decimal scale as the target. The one-step forecast uses only information available at the origin. This

timing convention is applied consistently in the standalone EGARCH evaluation and in the construction of the hybrid feature. (Francq & Zakoian, 2019)

EGARCH is fitted using the training observations so that test-period target values do not influence parameter estimation. During the test period, the fixed estimated parameters are used to filter the observed return sequence and produce one-step forecasts. This procedure evaluates parameter stability across the hold-out period. Rolling re-estimation could adapt to change more rapidly, but it would constitute a different experiment and is identified as a future extension.

The MLP was specified as a regression network for one-day-ahead Parkinson volatility. For each day t, the baseline feature vector contained five lags each of return r, absolute return |r|, squared return r^2 and Parkinson volatility σ^P, producing 20 input features. The prediction target was σ^P_{t+1}. (James et al., 2021; Géron, 2022)

$$h = \tanh(W_1 x + b_1) \quad (3.8)$$

$$\hat{y} = W_2 h + b_2 \quad (3.9)$$

where x is the input-feature vector, W_1 is the input-to-hidden weight matrix, b_1 is the hidden-layer bias vector, h is the hidden-layer activation vector, tanh is the hyperbolic-tangent activation function, W_2 is the hidden-to-output weight vector, b_2 is the output bias and $\hat{y}$ is the predicted next-day volatility.

A chronological 80:20 split was used. After lag construction, 2,553 observations formed the training sample and 639 observations formed the test sample. Hyperparameters were selected using time-series cross-validation on the training period. The selected standalone MLP used one hidden layer with 16 neurons, tanh activation, Adam optimization, L2 penalty $\alpha = 0.001$ and learning rate 0.001. (Bergmeir et al., 2018)

Four market-history transformations are created at each of five lags. Returns preserve signed movement; absolute returns and squared returns measure magnitude with different sensitivity to extremes; Parkinson volatility captures the recent daily range. Lag one refers to the most recent value available at the forecast origin, while lag five refers to the oldest observation in the selected window. The order of columns is fixed so that training and test matrices are identical in structure. (James et al., 2021; Géron, 2022)

The feature design intentionally excludes macroeconomic variables, sentiment indicators and firm-level data. This delimitation isolates the contribution of the ASI's own history and the EGARCH forecast. It also avoids combining series with different release schedules and revision histories. The resulting conclusions concern market-history-based volatility forecasting and should not be read as evidence that external variables have no predictive value.

3.6.2 Standardization and Training Protocol

Feature means and standard deviations are estimated on the training sample. Those parameters are stored and applied unchanged to validation and test data. The target transformation, where used within the numerical implementation, follows the same training-only principle and predictions are returned to the original volatility scale before evaluation. This procedure prevents the distribution of future observations from influencing the fitted network. (James et al., 2021)

The MLP uses one hidden layer of 16 tanh neurons, Adam optimization, a learning rate of 0.001 and L2 penalty of 0.001. Candidate settings are assessed through time-series cross-validation within the training period. The final model is then fitted on the designated training observations and evaluated once on the chronological test set. Reproducibility requires recording the software environment and random seed because neural-network optimization can vary with initialization. (Géron, 2022; Bergmeir et al., 2018)

The integrated model operationalizes the word "integrated" in the dissertation title. First, EGARCH parameters were estimated using the training period only. Second, the fitted EGARCH recursion generated a one-step-ahead conditional-volatility forecast σ^E_{t+1} using information available through day t. Third, this EGARCH forecast was appended to the 20 baseline MLP features: (Koo & Kim, 2022; Kartsonakis Mademlis & Dritsakis, 2021)

$$x^H_t = [x_t, \sigma^E_{t+1}] \quad (3.10)$$

where x^H_t is the 21-element hybrid feature vector, x_t contains the lagged market features and σ^E_{t+1} is the EGARCH one-step-ahead volatility forecast. The MLP then estimates $f(x^H_t) \approx \sigma^P_{t+1}$. This design allows the neural network to combine an econometrically structured volatility signal with non-linear information contained in recent returns and ranges.

The integrated procedure begins with the same cleaned and chronologically ordered data used by the standalone models. EGARCH parameters are estimated from the training return series. The fitted recursion produces a one-step conditional-volatility forecast for each eligible origin. This forecast is aligned with the same next-day Parkinson target used by the MLP and is appended to the baseline feature vector.

The hybrid MLP retains the standalone architecture and training procedure so that the comparison isolates the additional EGARCH information. Changing network size at the same time would make it difficult to determine whether improved performance resulted from integration or from greater model capacity. The 21st feature therefore represents the principal experimental treatment. (Petropoulos et al., 2022)

Forecasts from the three models are aligned by date before losses are calculated. Any observation unavailable to one specification is excluded from all three comparisons. This common-sample rule ensures that reported error differences are not caused by evaluating models on unequal sets of days.

All three models were evaluated on the same test dates and the same Parkinson volatility target. The following statistics were computed: (Petropoulos et al., 2022)

$$MSE = (1/n)\,\Sigma(y_t - \hat{y}_t)^2 \quad (3.11)$$

$$RMSE = \sqrt{MSE} \quad (3.12)$$

$$MAE = (1/n)\,\Sigma|y_t - \hat{y}_t| \quad (3.13)$$

$$R^2 = 1 - [\Sigma(y_t - \hat{y}_t)^2 / \Sigma(y_t - \bar{y})^2] \quad (3.14)$$

where y_t is realized Parkinson volatility, $\hat{y}_t$ is the corresponding model forecast, $\bar{y}$ is the test-sample mean target and n is the number of test observations. Lower MSE, RMSE and MAE indicate better forecasts, while higher R^2 indicates greater explanatory performance relative to the mean benchmark.

A Diebold–Mariano-style loss-difference statistic was also computed for the MLP and integrated model using squared and absolute forecast losses as a supplementary test of whether observed error reductions were systematic. Because the forecast horizon is one day, no overlapping-horizon correction was required.

Relative improvement is computed by comparing the hybrid error with the standalone MLP error. A reduction in MAE, for example, is expressed as the difference divided by standalone MAE. Relative measures help communicate the size of the gain, but they must be accompanied by the underlying errors because percentages can appear large when the denominator is small.

The loss-difference statistic examines whether the average difference between model losses is systematic relative to its sampling variability. Squared and absolute losses emphasize different aspects of performance. Because the forecast horizon is one day, the primary comparison does not require the overlapping-horizon correction associated with multi-step cumulative forecasts (Petropoulos et al., 2022).

Statistical significance and practical significance are distinguished. A small improvement may be operationally useful when forecasts are produced repeatedly, even if a formal test lacks power. Conversely, a statistically detectable difference may be economically minor. The dissertation therefore reports point metrics, relative improvement and supplementary loss-difference evidence together.

3.10.1 Benchmark Logic

A benchmark should be difficult enough to make improvement meaningful but simple enough to interpret. The standalone EGARCH represents structured conditional variance, and the standalone MLP represents nonlinear learning from lagged market variables. The hybrid is compared with both. The test-sample mean implicit in out-of-sample R-squared supplies an additional naive benchmark. Positive R-squared indicates that a model improves on this constant prediction under squared loss. (Petropoulos et al., 2022)

Benchmark consistency requires identical target scaling. EGARCH conditional standard deviation estimated from percentage returns is converted back to decimal units. MLP predictions are inverse-transformed where necessary. All metrics are then calculated in the natural scale of daily Parkinson volatility. This step prevents a numerical scale difference from being mistaken for predictive performance.

3.10.3 Interpretation Rules for Hypothesis Decisions

The ARCH hypothesis is evaluated through the LM statistic and its p-value. Rejection indicates remaining dependence in squared mean-adjusted returns and justifies conditional-variance modelling. The asymmetry hypothesis is evaluated using the coefficient on the signed standardized innovation in the stated EGARCH

equation. A non-significant p-value leads to failure to reject symmetry within that specification, even when other model components are significant. (Petropoulos et al., 2022)

The hybrid-improvement hypothesis is assessed through out-of-sample loss and the supplementary loss-difference statistics. Lower point errors support directional improvement; statistical evidence depends on the loss-difference p-value. Because the two loss functions do not provide equally decisive evidence, the conclusion is framed as modest incremental improvement rather than universal superiority. This wording follows the complete pattern of results.

4. RESULTS

Figure 4.1. Nigerian All-Share Index level over the study period.

Figure 1. Nigerian All-Share Index level over the study period.

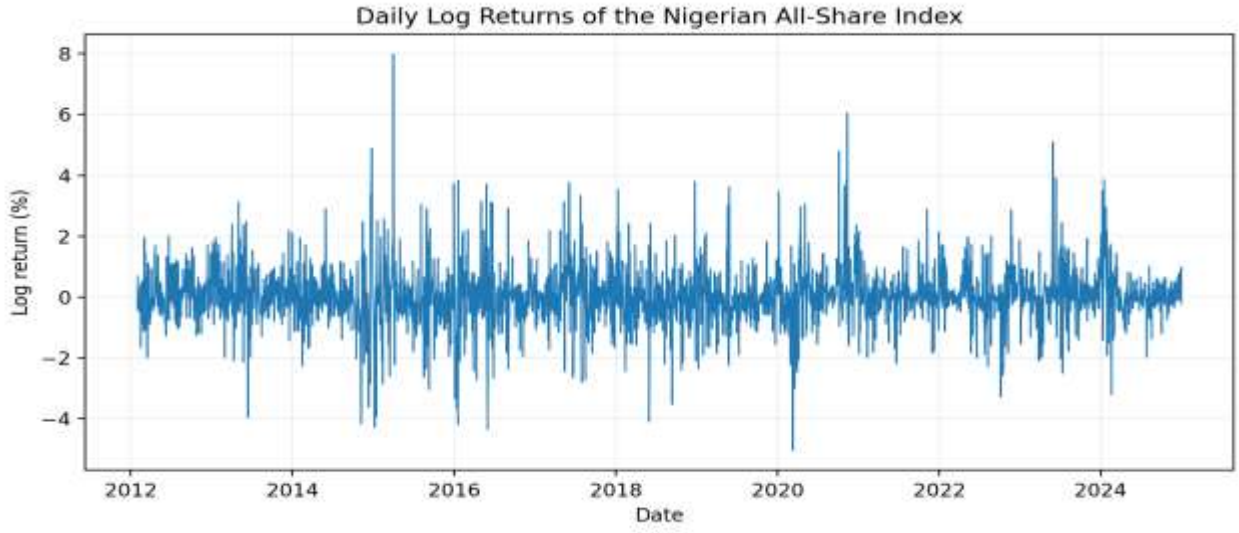

Figure 2. Daily log returns of the Nigerian All-Share Index.

Table 1. Descriptive statistics of ASI log returns

Statistic	Value
Observations	3,198 returns
Mean	0.000501
Median	0.000133
Standard deviation	0.009478
Minimum	-0.050329
Maximum	0.079848
Skewness	0.361728
Kurtosis	8.814578
Jarque–Bera	4557.796 (p < .001)

Table 2. Stationarity and conditional-heteroskedasticity tests

Test	Statistic	p-value	Decision
ADF	-20.4818	< .001	Reject unit root; returns are stationary
KPSS	0.2663	> .10	Fail to reject stationarity
ARCH-LM (10 lags)	413.272	< .001	Reject no-ARCH null

The ADF and KPSS tests give mutually consistent evidence that the return series is stationary. The ARCH-LM statistic is highly significant, establishing time-varying conditional variance and directly satisfying the statistical precondition for an EGARCH model.

Figure 4.3. Annualized standard deviation of ASI daily returns.

Table 4.3. Annualized ASI volatility by calendar year.

The annualized results answer the practical question "Is the ASI still volatile?" more precisely than a simple yes-or-no statement. Volatility is still present, but it is regime-dependent rather than uniformly high. The sample peak occurred in 2015 (20.02%), another high-volatility episode occurred in 2020 (18.70%), volatility fell sharply in 2021 (10.19%), and stood at 12.80% in 2024. Thus, the ASI continues to exhibit material volatility, but current risk should be interpreted relative to historical regimes.

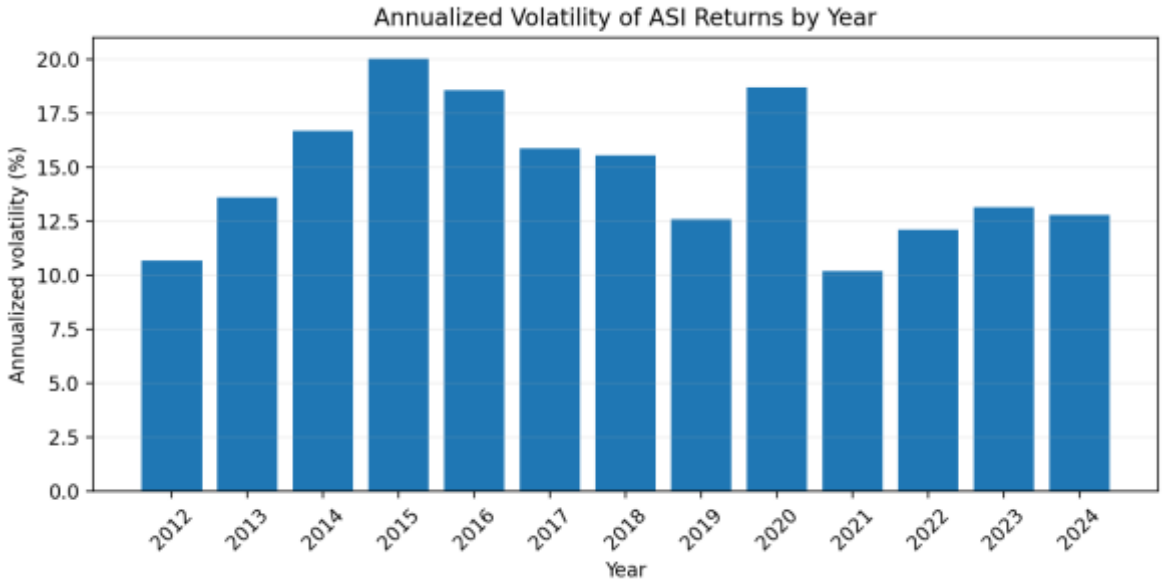

Figure 3. Annualized standard deviation of ASI daily returns.

Table 3. Annualized ASI volatility by calendar year

Year	Annualized volatility (%)
2012	10.69
2013	13.60
2014	16.70
2015	20.02
2016	18.57
2017	15.89
2018	15.57
2019	12.58
2020	18.70
2021	10.19
2022	12.09
2023	13.16
2024	12.80

4.1 EGARCH Results

4.2 Objective 2: EGARCH(1,1) Volatility Dynamics

Table 4.4. EGARCH innovation-distribution comparison.

The Student-t specification reduces AIC by approximately 569.7 points relative to the normal specification, providing strong evidence that the heavy-tailed innovation is more suitable for the ASI return distribution. The Student-t EGARCH(1,1) is therefore retained as the main econometric model.

Table 4.5. EGARCH(1,1)-t parameter estimates.

The magnitude parameter $\alpha = 0.4766$ is positive and highly significant, confirming that unusually large shocks increase subsequent volatility. The persistence coefficient $\beta = 0.9224$ is also highly significant and close to one, indicating that volatility shocks decay slowly; its approximate half-life is 8.6 trading days. By contrast, $\gamma = 0.0193$ is not statistically significant ($p = .343$). The null hypothesis of no asymmetric effect is therefore not rejected for this specification. This finding is important because it distinguishes strong volatility persistence from a statistically robust leverage effect.

Figure 4.4. EGARCH(1,1)-t estimated conditional volatility.

Figure 4.5. EGARCH news-impact curve implied by the estimated parameters.

The conditional-volatility series confirms clustering: elevated risk appears in episodes rather than as a constant state. The relatively mild asymmetry of the news-impact curve is consistent with the insignificant γ coefficient. Therefore, the evidence supports persistent non-linear volatility but does not support a strong direction-specific leverage effect in this sample.

Table 4. EGARCH innovation-distribution comparison

Specification	Log-likelihood	AIC	BIC
EGARCH(1,1), Normal	-3943.8636	7897.7272	7928.0771
EGARCH(1,1), Student-t	-3657.9975	7327.9950	7364.4148

Table 5. EGARCH(1,1)-t parameter estimates

Parameter	Estimate	Std. error	z-statistic	p-value
Mean (μ)	0.009738	0.010034	0.9704	0.3318
Constant (ω)	-0.006257	0.013429	-0.4659	0.6413
Magnitude effect (α)	0.476596	0.049652	9.5987	< .001
Asymmetry effect (γ)	0.019315	0.020355	0.9489	0.3427
Persistence (β)	0.922421	0.015911	57.9748	< .001
Student-t df (ν)	2.996760	0.194702	15.3915	< .001

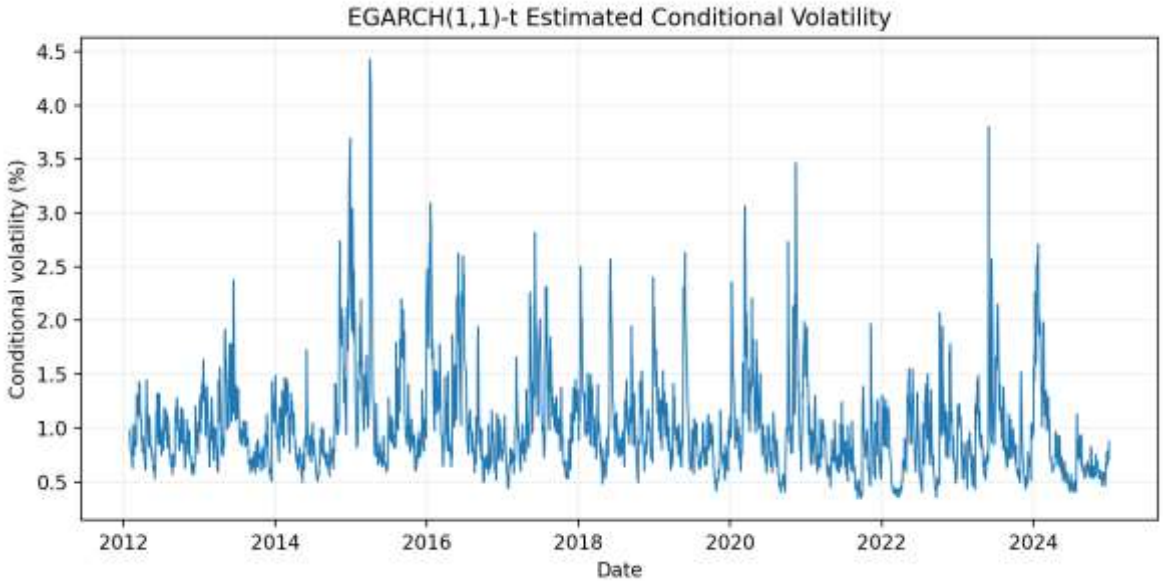

Figure 4. EGARCH(1,1)-t estimated conditional volatility.

4.2 MLP and Integrated EGARCH–MLP Forecasts

Table 4.6. Selected standalone MLP architecture and validation design.

The tuned standalone MLP generated an out-of-sample RMSE of 0.004228 and MAE of 0.003037. Its R^2 was -0.0099, indicating that although the model reduced errors substantially relative to standalone EGARCH under MAE and RMSE, it did not explain more test-period variation than a constant mean benchmark. This result demonstrates why error metrics and explanatory statistics should be interpreted jointly rather than using a single criterion.

4.4 Objective 4: Integrated EGARCH–MLP and Model Comparison

Table 4.7. Out-of-sample one-day-ahead volatility forecast performance.

The integrated EGARCH–MLP achieved the lowest error on all three point-forecast measures and the only positive R^2. Relative to the standalone MLP, the integrated model reduced MSE by 5.31%, RMSE by 2.69% and MAE by 3.85%. This indicates that the EGARCH volatility forecast contained incremental information that the lag-only MLP did not fully recover on its own.

A supplementary loss-difference test produced a squared-error statistic of 1.233 ($p = .218$) and an absolute-error statistic of 1.960 ($p = .050$) when comparing the standalone MLP with the hybrid. Thus, the direction of improvement favours the integrated model, with the absolute-error result at the conventional 5% boundary, but the evidence is not uniformly decisive across loss functions. This cautious interpretation is statistically preferable to claiming universal dominance.

Figure 4.6. Out-of-sample volatility forecasts, shown as 20-day moving averages for readability.

Figure 4.6 shows that all forecasts smooth the day-to-day realized-volatility target, but the hybrid tracks medium-term changes more closely than the standalone MLP in several periods. The EGARCH-only forecast is structurally informative but is less accurate on the range-based target because it is estimated from close-to-close returns rather than directly from intraday ranges.

4.5 Objective 5: What the Results Mean for the Nigerian Stock Market

The statistical evidence shows that the Nigerian ASI remains volatile, but "volatile" should not be interpreted as permanently high variance. The market alternates between low- and high-volatility regimes. Significant ARCH effects, high EGARCH persistence and the annualized volatility pattern all show that shocks can influence risk for several subsequent trading days. At the same time, 2024 volatility remained

below major stress years such as 2015 and 2020. For investors, this means that risk controls should respond to the current volatility regime rather than rely on a single long-run variance estimate.

Table 6. Selected standalone MLP architecture and validation design

MLP setting	Selected value
Input features	5 lags each of return, |return|, return² and Parkinson volatility (20 inputs)
Hidden layers	1
Hidden neurons	16
Activation	tanh
Optimizer	Adam
L2 penalty α	0.001
Learning rate	0.001
Validation	Time-series cross-validation
Training observations	2,553
Test observations	639

Table 7. Out-of-sample one-day-ahead volatility forecast performance

Model	MSE	RMSE	MAE	R^2
EGARCH(1,1)-t	0.000043	0.006522	0.005347	-1.4037
Standalone MLP	0.000018	0.004228	0.003037	-0.0099
Integrated EGARCH–MLP	0.000017	0.004114	0.002920	0.0438

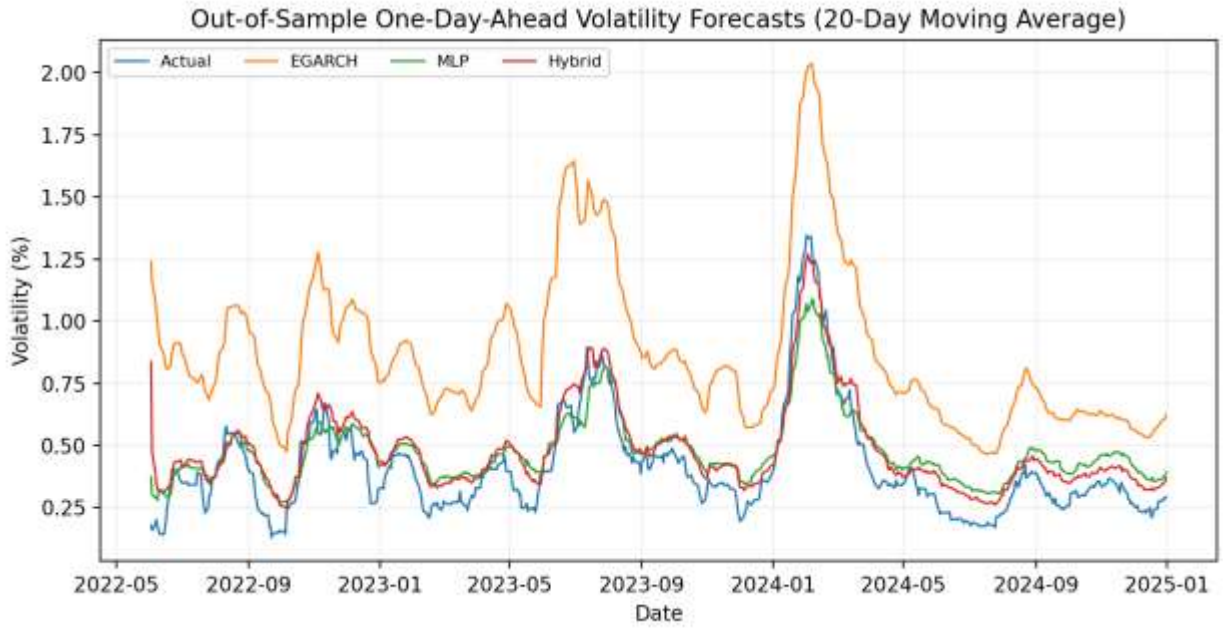

Figure 5. Out-of-sample one-day-ahead volatility forecasts.

DISCUSSION

Objective 1 was achieved. The return series was stationary, heavy-tailed and conditionally heteroskedastic. This combination is consistent with recent Nigerian studies that characterize equity-market risk as clustered and non-Gaussian (Adegboyo & Sarwar, 2025; Awogbemi et al., 2024). The present data also show that volatility is regime-dependent: annualized volatility was highest in 2015 and 2020 but materially lower in 2021 and 2024. This improves the substantive conclusion by showing not merely that the ASI is volatile, but how volatility evolves through time.

The strong rejection of normality indicates that ASI risk is not well summarized by a Gaussian distribution with a constant variance. Positive skewness in the full sample coexists with substantial extreme observations and high kurtosis. For practical risk analysis, kurtosis is particularly important because it indicates that large deviations occur more frequently than a normal benchmark would suggest. This supports the use of stress scenarios and tail-aware statistical distributions.

The stationarity results and ARCH evidence should be read together. Returns fluctuate around a stable mean representation, yet their conditional dispersion changes. This is precisely the setting in which volatility models are informative: transformation removes the non-stationary price level, while second-moment dependence remains available for estimation. The results therefore reject both the idea of a constant-risk market and the idea that volatility must be inferred from a trending index level.

Taken together, the evidence distinguishes volatility persistence from a permanently high level of risk. Persistence means that a shock continues to influence conditional variance over several subsequent trading days, whereas the annual estimates show that the absolute level of risk still changes substantially across market conditions. This distinction is important because a stationary return series can coexist with pronounced volatility clustering: returns may revert around a stable mean even though their conditional variance evolves slowly following a disturbance. The Nigerian ASI should therefore be treated as a market with changing risk intensity rather than one that can be adequately represented by a single unconditional variance.

The heavy tails and significant ARCH effect are consistent with the wider Nigerian evidence reported by Adegboyo and Sarwar (2025), which emphasizes volatility clustering and the need for conditional-variance models in the analysis of domestic equity returns. The present results add an important qualification by showing that clustering is not equivalent to uniformly extreme volatility: the peaks recorded in 2015 and 2020 were followed by lower-volatility periods, including the 2024 estimate. For risk assessment, the relevant issue is therefore both the magnitude of current volatility and the speed with which past shocks decay.

Objective 2 was achieved. EGARCH with Student-t innovations fitted the heavy-tailed returns much better than the Gaussian alternative, supporting recent Nigerian evidence that innovation distribution matters (Awogbemi et al., 2024; Adegboyo & Sarwar, 2025). The large and significant β coefficient indicates persistent volatility, while α confirms sensitivity to shock magnitude. However, γ was not significant, so this study does not find robust evidence that negative and positive shocks of equal size have different volatility effects. This differs from Adegboyo and Sarwar (2025), who report a leverage effect over a closely related period. The difference may reflect model specification, exact sample endpoint, data vendor construction or the choice of innovation distribution, and it demonstrates that asymmetry in the Nigerian market is specification-sensitive rather than a result that should be assumed in advance.

The persistence estimate of 0.922 implies that volatility shocks decay slowly within the fitted EGARCH process. For market participants, this means that a large movement should influence risk assessment beyond the day on which it occurs. The result supports dynamic position limits and repeated updating of risk measures. It does not imply that the market will remain permanently volatile, because the annualized estimates show substantial variation across years.

The difference between persistence and the average volatility level is central. A year may have a moderate annualized value while still containing clustered episodes. Conversely, a high annual figure summarizes the accumulation of turbulent days but does not mean every day was equally risky. EGARCH captures the conditional path, whereas annualized descriptive volatility provides a useful longer-period comparison. Together they show a market that is persistent after shocks but not uniformly high.

The non-significant signed coefficient does not weaken the evidence for EGARCH as a volatility model. The Student-t specification fits substantially better, and the magnitude and persistence components remain strongly informative. The result simply narrows the conclusion: shocks matter through their size and persistence, but the sample does not establish a stable difference between positive and negative shocks of equal standardized magnitude.

This finding differs from studies that report significant leverage in Nigerian or other equity data, but the difference is plausible. Asymmetry can vary with sample period, distribution, structural breaks and equation convention. A responsible discussion recognizes this heterogeneity rather than presenting leverage as a universal stylized fact. Future work may examine subperiods or rolling estimates to determine whether asymmetry emerges only in particular regimes.

The estimated persistence coefficient of approximately 0.922 gives economic content to the visual evidence of clustering. Its magnitude indicates that new disturbances are absorbed gradually, so a large forecast revision should not be expected to disappear immediately on the next trading day. This has practical implications for margin setting, exposure limits and short-horizon portfolio rebalancing, because risk controls based only on the latest return may understate the continuing influence of earlier shocks. At the same time, persistence below unity is compatible with eventual mean reversion in conditional volatility and does not imply that variance grows without bound.

The insignificant asymmetry coefficient must be interpreted separately from the strong persistence result. It indicates that, within this sample and specification, negative and positive return innovations of comparable magnitude did not produce statistically distinguishable effects on conditional volatility. This does not invalidate EGARCH, whose logarithmic variance formulation remains useful for modelling persistence and maintaining positive conditional variance. It instead cautions against treating the leverage effect as an automatic property of every emerging equity market. The superior fit of Student-t innovations is also consistent with Awogbemi et al. (2024), whose Nigerian evidence underscores the importance of accommodating non-Gaussian return behaviour in volatility modelling.

Objective 3 was achieved by constructing and validating a parsimonious MLP on a chronological training sample. The MLP improved MAE and RMSE relative to EGARCH when both were assessed against the same range-based target, confirming the value of non-linear learning. This is consistent with evidence that MLP models can exploit non-linear dependencies in stock-index data (Ecer et al., 2020; Gao et al., 2020; Žmuk & Jošić, 2020). The negative standalone R^2 nevertheless shows that an MLP is not automatically superior merely because it is non-linear; careful feature design and hybrid information can still matter.

The standalone MLP reduces MAE and RMSE relative to EGARCH on the Parkinson target, showing that the lagged range and return transformations contain useful nonlinear information. Its slightly negative R-squared nevertheless indicates that the improvement is not sufficient to beat the mean benchmark under total squared error. This combination is possible because metrics weight errors differently and because a small number of large misses can strongly affect R-squared.

The result illustrates a broader methodological lesson: model flexibility is a capacity, not a guarantee. Neural networks can approximate complex functions, but financial data contain noise, structural change and limited repeated examples of extreme regimes. A compact validated architecture is more credible than an unsupported claim that greater complexity will necessarily improve prediction.

The standalone MLP result illustrates both the value and the limitation of flexible non-linear prediction. Its lower point-forecast errors relative to EGARCH suggest that the lagged feature set contains interactions and threshold-type relationships that are not completely represented by the econometric variance recursion. However, the slightly negative out-of-sample R^2 shows that improved average error measures do not automatically imply strong explanatory power relative to a simple benchmark. This is a useful safeguard against interpreting computational flexibility as evidence of comprehensive predictability.

The result also reflects the difficulty of forecasting a noisy daily range-based target from end-of-day information. An MLP may learn useful local patterns while still being affected by shifts in market regime, limited sample size and the absence of contemporaneous macroeconomic or order-flow variables. Consequently, its performance should be viewed as evidence of incremental non-linear structure rather than as proof that daily ASI volatility can be predicted with high precision. The chronological test design strengthens this interpretation because the network was evaluated on observations that followed, rather than overlapped with, its training period.

Objective 4 was achieved through an actual model integration rather than a conceptual comparison. The one-step EGARCH volatility forecast was used as an explicit MLP feature. The resulting hybrid produced the best MSE, RMSE and MAE and moved R^2 from slightly negative under the standalone MLP to positive. The improvement is directionally consistent with hybrid studies in which GARCH-family forecasts act as structured features for neural networks (Amirshahi & Lahmiri, 2023; Araya et al., 2024; Kartsonakis Mademlis & Dritsakis, 2021; Bu et al., 2025). The modest size of the gain is also informative: EGARCH contributes useful but not overwhelming incremental predictive information beyond recent returns and ranges.

The hybrid's improvement across MSE, RMSE and MAE shows consistent directional advantage over the standalone MLP. Its positive R-squared further indicates that the EGARCH feature helps the network surpass the mean benchmark under squared loss. The gains are modest, which is unsurprising because the baseline already contains several volatility-related transformations. EGARCH contributes a compressed recursive state rather than entirely new market information.

The result supports complementarity. EGARCH supplies an interpretable representation of persistent conditional risk, and MLP learns a nonlinear mapping to the range-based target. The hybrid does not erase the limitations of either component, but it allows their information to be combined within one forecast. This is a stronger contribution than placing unrelated model outputs side by side.

The loss-difference evidence prevents overstatement. The absolute-error comparison is at the conventional boundary, while squared-error evidence is weaker. The appropriate conclusion is therefore that the hybrid produces the best observed point forecasts and adds incremental information in this sample, although further periods and rolling evaluation are needed before claiming stable superiority.

The integrated model improves performance because the EGARCH feature compresses information about recent shocks and persistent conditional variance into a single forward-looking signal, while the MLP combines that signal with the broader lagged feature set. The hybrid therefore does more than average two unrelated forecasts. It embeds an economically interpretable volatility state within a flexible non-linear mapping, allowing the network to adjust the weight attached to the EGARCH signal as surrounding market conditions change.

Nevertheless, the size and statistical strength of the improvement require a restrained conclusion. The hybrid records the lowest MSE, RMSE and MAE and the only positive R^2, which establishes the direction of the forecasting gain on this test sample. The supplementary loss-difference tests, however, do not reject equal predictive accuracy at conventional levels. The findings therefore support modest incremental value from integration, not universal dominance. Confirmation across alternative windows, volatility proxies and post-2024 market regimes is necessary before the model is adopted as an invariant forecasting rule.

Objective 5 was achieved. The results imply that risk in the Nigerian stock market changes through time and remains persistent after major shocks. The 2024 annualized volatility estimate of about 12.8% indicates that the market was not in its highest-volatility historical regime at the sample endpoint, but risk remained material and capable of changing quickly. For risk managers, the practical implication is to update volatility estimates dynamically. For investors, the integrated model can support short-horizon risk monitoring rather than direct return-direction trading. For policymakers, periods of rising conditional volatility may provide an additional statistical signal of market stress.

For investors, the central message is that market risk should be updated as conditions change. A persistent rise in forecast volatility can justify reduced position sizes, tighter concentration controls or additional stress testing. The forecast is not a buy-or-sell signal and should be combined with expected return, liquidity, investment horizon and transaction-cost considerations. Its contribution is to the risk side of the decision.

A hybrid forecast may be useful when risk systems already use statistical volatility estimates but require greater responsiveness to nonlinear market patterns. The modest error reduction should be evaluated against implementation cost and model-governance requirements. Because the MLP is less transparent than EGARCH, institutions should retain input definitions, validation records and monitoring thresholds.

For regulators and exchange authorities, clustered volatility means that surveillance should examine sequences of unusual movements rather than isolated daily changes alone. Persistent risk can affect margin adequacy, liquidity and investor confidence. A forecast framework may contribute to early-warning dashboards, although it cannot identify manipulation or causal events without additional transaction and market-quality data.

Heavy tails also indicate that normal-distribution thresholds may generate misleading assessments of rare events. Surveillance and stress-testing frameworks should account for the empirical frequency of extreme movements. The Student-t EGARCH result provides statistical support for tail-aware monitoring, while the hybrid result suggests that a nonlinear overlay may improve short-horizon estimates.

CONCLUSION AND IMPLICATIONS

The EGARCH(1,1) model with Student-t innovations provided a substantially better fit than the Gaussian specification. The persistence coefficient of 0.922 was high and statistically significant, indicating that volatility shocks exerted prolonged effects. The asymmetry coefficient was not statistically significant; consequently, the study does not establish leverage as a consistent feature of ASI volatility during the sample period.

The findings revealed that the ASI return series was stationary but strongly non-normal and heavy-tailed. The ARCH-LM test confirmed statistically significant time-varying variance. ASI volatility was clustered and persistent rather than uniformly high throughout the study period. Annualized volatility reached approximately 20.0 per cent in 2015 and 18.7 per cent in 2020, compared with approximately 12.8 per cent in 2024.

The aim of this study was to model and forecast the non-linear volatility dynamics of the Nigerian All-Share Index using an integrated EGARCH–MLP framework. That aim was achieved. Objective 1 established that ASI returns are stationary, heavy-tailed and conditionally heteroskedastic, confirming that market risk changes through time. Objective 2 showed that EGARCH(1,1) with Student-t innovations captures a highly persistent volatility process, although the direction-specific asymmetry coefficient was not statistically significant. Objective 3 demonstrated that a tuned MLP can learn useful non-linear relationships in lagged returns and range-based volatility. Objective 4 implemented a genuine hybrid by adding the one-step EGARCH volatility forecast to the MLP and showed that the integrated model produced the lowest out-of-sample MSE, RMSE and MAE. Objective 5 translated these findings into a market conclusion: the Nigerian ASI remains volatile, but its volatility is dynamic, clustered and regime-dependent rather than continuously high.

The most important contribution is therefore not a claim that machine learning replaces econometrics. Instead, the evidence shows that EGARCH and MLP perform complementary roles. EGARCH supplies an interpretable statistical representation of persistence and shock transmission, while MLP learns additional non-linear patterns in recent market information. Their integration yields a modest but measurable improvement in short-horizon volatility forecasting. At the end of the study period, annualized ASI volatility was lower than in the major stress years of 2015 and 2020, yet significant ARCH effects and high EGARCH persistence show that risk remains dynamic and capable of clustering after shocks.

Risk managers, investors and financial analysts should adopt time-varying volatility measures when monitoring the Nigerian equity market rather than relying solely on a single historical variance estimate. Such measures are better suited to capturing volatility clustering and changes in market risk over time.

When applying EGARCH models to Nigerian ASI data, researchers should consider heavy-tailed innovation distributions, particularly the Student-t distribution, instead of automatically assuming Gaussian errors. This is important because Nigerian stock-market returns exhibit non-normality, excess kurtosis and extreme observations that may not be represented adequately by a normal distribution.

Future volatility studies should compare competing models using a consistent realized-volatility target. Hybrid modelling should involve the statistical integration of econometric and machine-learning techniques rather than a comparison of models that predict different variables, such as returns and volatility.

Further research should examine rolling or walk-forward re-estimation to determine whether regularly updating model parameters improves forecasting accuracy. Structural-break indicators, macroeconomic variables and alternative realized-volatility measures should also be considered, particularly where reliable higher-frequency data are available.

By achieving the first objective, the study establishes with updated daily evidence that Nigerian ASI returns are stationary in level changes but non-normal, heavy-tailed and conditionally heteroskedastic. It further demonstrates that volatility is clustered and persistent without being uniformly high across the study period, thereby separating the persistence of shocks from the annual level of market risk.

Achievement of the second objective provides an empirically qualified account of EGARCH dynamics in the Nigerian market. The study shows that Student-t innovations fit the ASI return distribution materially better than Gaussian innovations and that the persistence coefficient is high and statistically significant. At the same time, the asymmetry coefficient is not significant, contributing evidence that leverage effects are sample- and specification-dependent rather than an automatic characteristic of Nigerian equities.

Through the third objective, the study develops and evaluates a standalone MLP that forecasts the same one-day-ahead Parkinson volatility target used for econometric comparison. The MLP reduces point-forecast error relative to EGARCH but retains a slightly negative out-of-sample R^2, documenting both the usefulness and the limitation of nonlinear learning from lagged market features.

Achievement of the fourth objective supplies the study's principal methodological contribution: a directly implemented EGARCH–MLP in which the one-step-ahead EGARCH volatility forecast becomes an additional MLP feature. The three specifications are evaluated on common test dates, forecast horizon and volatility target. The integrated model records the lowest MSE, RMSE and MAE and the only positive R^2, showing modest incremental information from econometric–machine-learning integration while avoiding a return-versus-volatility comparison.

By achieving the fifth objective, the study translates the statistical results into a cautious interpretation of Nigerian market risk. It shows that persistent conditional volatility supports time-varying risk limits and closer monitoring after large shocks, while the modest hybrid gain and statistically insignificant loss-difference tests prevent a claim of universal superiority. This balanced conclusion is itself a contribution because it connects forecasting performance with practical use while clearly identifying the boundaries of the evidence.

The dataset is limited to the historical coverage available from the selected data source and begins on 30 January 2012. The Parkinson target uses daily high and low values rather than intraday realized variance, so it remains an approximation to latent volatility. The EGARCH parameters are held fixed during out-of-sample filtering after being estimated on the training sample; rolling re-estimation may adapt more quickly to structural change. The MLP uses only market-history features and excludes macroeconomic, sentiment and liquidity variables. Finally, forecast improvements are modest and should be interpreted as evidence of incremental value rather than proof of universal hybrid-model superiority.

REFERENCES

Adegboyo, O. S., & Sarwar, K. (2025). Modelling and forecasting of Nigeria stock market volatility. Future Business Journal, 11, Article 124. https://doi.org/10.1186/s43093-025-00536-4

Ali, F., Suri, P., Kaur, T., & Bisht, D. (2022). Modelling time-varying volatility using GARCH models: Evidence from the Indian stock market. F1000Research, 11, 1098. https://doi.org/10.12688/f1000research.124998.2

Amirshahi, B., & Lahmiri, S. (2023). Hybrid deep learning and GARCH-family models for forecasting volatility of cryptocurrencies. Machine Learning with Applications, 12, 100465. https://doi.org/10.1016/j.mlwa.2023.100465

Araya, H. T., Aduda, J., & Berhane, T. (2024). A hybrid GARCH and deep learning method for volatility prediction. Journal of Applied Mathematics, 2024, 6305525. https://doi.org/10.1155/2024/6305525

Awogbemi, C. A., Adenomon, M. O., Nwikpe, B. J., Chajire, B. P., Ilori, A. K., Shitu, D. A., Dayo, V. K., Sani, Z. S., Tanimu, M., & Paul, V. B. (2024). Modeling Nigerian stock price volatility using EGARCH-X model with different innovations. Journal of the Royal Statistical Society Nigeria Group, 1(2).

Banerjee, S., & Mukherjee, D. (2021). Short term stock price prediction in Indian market: A neural network perspective. Vision, 26(1). https://doi.org/10.1177/2321022220980537

Bergmeir, C., Hyndman, R. J., & Koo, B. (2018). A note on the validity of cross-validation for evaluating autoregressive time series prediction. Computational Statistics & Data Analysis, 120, 70–83. https://doi.org/10.1016/j.csda.2017.11.003

Bodie, Z., Kane, A., & Marcus, A. J. (2021). Investments (12th ed.). McGraw-Hill Education.

Bu, T., Yan, T., & Yin, H. (2025). Forecasting the volatility of CSI 300 Index with a hybrid model of LSTM and multiple GARCH models. Computational Economics, 66(3), 1969–1999. https://doi.org/10.1007/s10614-024-10785-0

Ecer, F., Ardabili, S., Band, S. S., & Mosavi, A. (2020). Training multilayer perceptron with genetic algorithms and particle swarm optimization for modeling stock price index prediction. Entropy, 22(11), 1239. https://doi.org/10.3390/e22111239

Ekpete, M. S., Kenn-Ndubuisi, J. I., & Iyo, I. (2026). Forecasting the Nigerian Exchange Group's All Share Index: Application of ARIMA-GARCH hybrid models. FUDMA Journal of Accounting and Finance Research, 4(2), 21–39. https://doi.org/10.33003/fujafr-2026.v4i2.335.21-39

Eme, U. N., Enwere, G. M., & Williams, A. (2024). The nexus between exchange rate volatility and Nigerian stock prices: An EGARCH approach. African Journal of Economics and Sustainable Development, 7(1), 56–72. https://doi.org/10.52589/AJESD-ATAK411M

Emenyonu, S. C., Osu, B. O., & Olunkwa, C. (2023). Estimating volatility of daily price returns of Nigerian stock market. Journal of Mathematical Techniques and Computational Mathematics, 2(4), 163–169.

Fischer, T., & Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. European Journal of Operational Research, 270(2), 654–669. https://doi.org/10.1016/j.ejor.2017.11.054

Francq, C., & Zakoian, J.-M. (2019). GARCH models: Structure, statistical inference and financial applications (2nd ed.). Wiley. https://doi.org/10.1002/9781119313472

Gao, P., Zhang, R., & Yang, X. (2020). The application of stock index price prediction with neural network. Mathematical and Computational Applications, 25(3), 53. https://doi.org/10.3390/mca25030053

Géron, A. (2022). Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow (3rd ed.). O'Reilly Media.

Hu, Y., Ni, J., & Wen, L. (2020). A hybrid deep learning approach by integrating LSTM-ANN networks with GARCH model for copper price volatility prediction. Physica A: Statistical Mechanics and Its Applications, 557, 124907. https://doi.org/10.1016/j.physa.2020.124907

Ibrahim, S. O. (2017). Forecasting the volatilities of the Nigeria stock market prices. CBN Journal of Applied Statistics, 8(2), 23–45. https://doi.org/10.33429/2476-8472.1070

Isenah, G. M. (2024). Forecasting Nigerian stock market volatilities using bagging and ensemble averaging. CBN Journal of Applied Statistics, 15(2), 1–35. https://doi.org/10.33429/Cjas.15224.1/5

Iwada, E. M., Omoyeni, O. T., & Temitope, A. (2018). Modeling volatility in selected Nigerian stock market. International Journal of Economics and Financial Management, 3(1), 30–43.

James, G., Witten, D., Hastie, T., & Tibshirani, R. (2021). An introduction to statistical learning: With applications in R (2nd ed.). Springer. https://doi.org/10.1007/978-1-0716-1418-1

Kakade, K., Mishra, A. K., Ghate, K., & Gupta, S. (2022). Forecasting commodity market returns volatility: A hybrid ensemble learning GARCH-LSTM based approach. Intelligent Systems in Accounting, Finance and Management, 29(2), 103–117. https://doi.org/10.1002/isaf.1515

Kartsonakis Mademlis, D., & Dritsakis, N. (2021). Volatility forecasting using hybrid GARCH neural network models: The case of the Italian stock market. International Journal of Economics and Financial Issues, 11(1), 49–60. https://doi.org/10.32479/ijefi.10842

Kim, H. Y., & Won, C. H. (2018). Forecasting the volatility of stock price index: A hybrid model integrating LSTM with multiple GARCH-type models. Expert Systems with Applications, 103, 25–37. https://doi.org/10.1016/j.eswa.2018.03.002

Koo, E., & Kim, G. (2022). A hybrid prediction model integrating GARCH models with a distribution manipulation strategy based on LSTM networks for stock market volatility. IEEE Access, 10, 34743–34754. https://doi.org/10.1109/ACCESS.2022.3163723

Kristjanpoller, W., & Hernández, E. (2017). Volatility of main metals forecasted by a hybrid ANN-GARCH model with regressors. Expert Systems with Applications, 84, 290–300. https://doi.org/10.1016/j.eswa.2017.05.024

Kuhe, D. A. (2018). Modeling volatility persistence and asymmetry with exogenous breaks in the Nigerian stock returns. CBN Journal of Applied Statistics, 9(1), 167–196.

Lu, Z., Pu, H., Wang, F., Hu, Z., & Wang, L. (2017). The expressive power of neural networks: A view from the width. Advances in Neural Information Processing Systems, 30, 6231–6239.

Molnár, P. (2016). High-low range in GARCH models of stock return volatility. Applied Economics, 48(51), 4977–4991. https://doi.org/10.1080/00036846.2016.1170929

Oyewola, D. O., Ibrahim, A., Kwanamu, J. A., & Dada, E. G. (2021). A new auditory algorithm in stock market prediction on oil and gas sector in Nigerian stock exchange. Soft Computing Letters, 3, 100013. https://doi.org/10.1016/j.socl.2021.100013

Petropoulos, F., Apiletti, D., Assimakopoulos, V., Babai, M. Z., Barrow, D. K., Ben Taieb, S., Bergmeir, C., Bessa, R. J., Bijak, J., Boylan, J. E., Browell, J., Carnevale, C., Clements, M. P., Cordeiro, C., Cyrino Oliveira, F. L., De Baets, S., Dokumentov, A., Fiszeder, P., Frazier, D. T., … Ziel, F. (2022). Forecasting: Theory and practice. International Journal of Forecasting, 38(3), 705–871. https://doi.org/10.1016/j.ijforecast.2021.11.001

Ucheoma, E. C. (2024). Modeling volatility in the Nigeria Stock Exchange: The role of structural breaks. CBN Journal of Applied Statistics, 15(2), 71–92. https://doi.org/10.33429/2476-8472.1349

Usman, U., Auwal, H. M., & Abdulmuhyi, M. A. (2017). Fitting the Nigeria stock market return series using GARCH models. Theoretical Economics Letters, 7(7), 2159–2176. https://doi.org/10.4236/tel.2017.77147

Žmuk, B., & Jošić, H. (2020). Forecasting stock market indices using machine learning algorithms. Interdisciplinary Description of Complex Systems, 18(4), 471–489. https://doi.org/10.7906/indecs.18.4.6

Appendix A: Dataset Source and Reproducibility

The source dataset was downloaded from the Investing.com "NSE All Share Historical Data" page. The final study period is 30 January 2012 to 31 December 2024 and contains 3,199 daily trading records. The accompanying cleaned CSV supplied with this dissertation contains Date, Price, Open, High, Low, Volume, Change %, Log Return, Parkinson Volatility and Squared Return. The full compact source table used for the core calculations is reproduced below with Date, Close, High and Low.

Appendix B: Model Replication Summary

Compute $r_t = \ln(P_t/P_{t-1})$ from chronological ASI closes.

Compute Parkinson variance from daily High and Low and take its square root as the volatility target.

Test returns using descriptive statistics, Jarque–Bera, ADF, KPSS and ARCH-LM.

Estimate EGARCH(1,1) under both Normal and standardized Student-t innovations and retain the lower-AIC/BIC distribution.

Split the lagged modelling sample chronologically into 80% training and 20% testing observations.

Fit a baseline MLP using 5 lags each of r, |r|, r^2 and Parkinson volatility.

Filter the EGARCH recursion using parameters estimated on the training data and create the one-day-ahead EGARCH volatility forecast.

Append that forecast to the MLP features and fit the integrated EGARCH–MLP.

Evaluate all three models against the same next-day Parkinson volatility using MSE, RMSE, MAE and R^2.