



MODELLING CRUDE OIL PRICE VOLATILITY AND ITS IMPLICATIONS FOR THE NIGERIAN ECONOMY: A COMPARATIVE STUDY OF ASYMMETRIC GARCH AND LONG SHORT-TERM MEMORY (LSTM) MODELS

Mohammed Anono Zubair¹ and James Oyaonu Esther¹

¹Department of Statistics, Faculty of Physical Sciences, University of Abuja, Abuja,
Nigeria

Corresponding Author's Email: jamesesthero19@gmail.com

Abstract:

This study models the volatility of the global Brent crude oil price and examines its implications for the Nigerian economy, which depends heavily on crude oil exports for fiscal revenue and foreign exchange earnings. Using a monthly Brent crude oil spot price series (Europe) obtained from the U.S. Energy Information Administration via the Federal Reserve Economic Data (FRED) database, spanning May 1987 to July 2026 (471 price observations; 470 monthly log-returns), the study compares three GARCH-family conditional-variance models of increasing flexibility – symmetric GARCH(1,1), EGARCH(1,1) and GJR-GARCH(1,1), each with Student-t innovations – against a single-layer Long Short-Term Memory (LSTM) neural network built and validated from first principles. Descriptive statistics show pronounced excess kurtosis (9.06) and negative skewness (−0.49), with the Jarque-Bera statistic (719.84, $p < .001$) decisively rejecting normality. The Brent price level is non-stationary (ADF $t = -1.72$) while its return series is stationary (ADF $t = -9.02$), and both the Ljung-Box test on squared returns ($Q = 210.00$, $p < .001$) and the ARCH-LM test ($LM = 171.78$, $p < .001$) confirm strong volatility clustering. Formal structural-break tests at two independently documented oil-market events – the November 2014 OPEC non-cut decision and the March 2020 COVID-19 demand collapse – each find a highly significant post-break increase in return variance ($F = 2.034$ and $F = 2.552$, both $p < .001$) but no significant shift in the mean return ($p = .616$ and $p = .839$). Fitted on an 80% training window ($n = 376$), EGARCH(1,1) achieves the best in-sample fit by AIC (2653.32) while GARCH(1,1) achieves the best BIC (2675.44); neither asymmetric model's leverage coefficient is statistically significant (EGARCH $\gamma = -0.063$, $t = -1.16$; GJR-GARCH $\gamma = 0.111$, $t = 0.98$), although EGARCH's negative sign is stable across an independent post-2015 robustness re-fit. On a 94-month out-of-sample test window, EGARCH(1,1) achieves the lowest RMSE (427.77) and MAE (149.88), GJR-GARCH(1,1) the lowest Theil's U (0.529), and the LSTM the lowest MAPE (272.89%); a Diebold-Mariano test comparing EGARCH against the LSTM is not statistically significant ($DM = -1.067$, $p = .286$). The study concludes that Brent crude oil price volatility is highly persistent and asymmetric in sign, though not significantly so, and that a well-specified

asymmetric GARCH model remains competitive with, and is not clearly outperformed by, a deep-learning alternative, a finding of direct relevance to Nigerian fiscal and reserve-management planning.

Keywords: Crude oil price volatility; Brent; GARCH; EGARCH; GJR-GARCH; Long Short-Term Memory (LSTM); Structural break; Nigeria.

1.0 INTRODUCTION

Crude oil occupies a uniquely strategic position in the Nigerian economy. Petroleum exports have for decades accounted for the large majority of Nigeria's foreign exchange earnings and a substantial share of federal government revenue, and the Nigerian federal budget is set each year against an assumed benchmark crude oil price, so that a shortfall between the assumed and the realised price directly drives fiscal deficits, borrowing, and adjustments to the naira exchange rate. Nigeria's own export crude grades, principally Bonny Light, Qua Iboe and Forcados, are priced internationally at a differential to global benchmark crudes, of which Brent (sourced from the North Sea) is the most widely referenced for light, sweet crude oil globally. Movements in the Brent price accordingly feed through, with only a small and generally stable adjustment, into the price Nigeria actually receives for its own oil exports and into government revenue projections (Hamilton, 1983, 2003; Abdulkareem & Abdulkareem, 2016; Ozokede & Awogbemi, 2025).

Global crude oil prices are also, empirically, highly volatile, and that volatility is itself economically consequential quite apart from the price level: large, clustered swings in oil prices complicate Nigeria's budget planning, reserve management and exchange-rate policy in a way that a stable but merely high or low price would not. Abdulkareem and Abdulkareem (2016), applying GARCH-family models directly to oil price and Nigerian macroeconomic series, find significant volatility clustering in oil prices and evidence that this volatility spills over into domestic Nigerian aggregates. More recent evidence is mixed rather than uniformly confirmatory: Mgbomene et al. (2025), using quarterly GARCH estimates over 1990 to 2023, find that oil price volatility significantly reduces Nigerian GDP growth while simultaneously raising government revenue and foreign reserves, whereas Ozokede and Awogbemi (2025), using an annual ARDL model over the same broad period, find oil price volatility's effect on real GDP growth is positive but not statistically significant, and conclude that Nigeria “lacks effective transmission mechanisms” converting oil-price movements into growth outcomes. This literature supports treating oil price volatility as a variable of genuine Nigerian macroeconomic interest, which is the premise on which this study rests, without supporting a claim that its transmission into any specific Nigerian outcome is a settled, uncontested channel.

Volatility in energy prices is also widely documented to respond asymmetrically to shocks of different signs, a pattern often described as “rockets and feathers”: a price increase can raise volatility and downstream costs by more than an equivalent decrease reduces them. This asymmetry is precisely what the Generalised Autoregressive Conditional Heteroskedasticity (GARCH) family of models, and specifically its asymmetric variants EGARCH (Nelson, 1991) and GJR-GARCH (Glosten, Jagannathan & Runkle, 1993), was designed to capture, and that a symmetric GARCH(1,1) specification cannot. At the same time, Long Short-Term Memory (LSTM) neural networks (Hochreiter & Schmidhuber, 1997) have emerged as a data-driven alternative that makes no distributional assumption and can, in principle, learn non-linear volatility dependence directly from data. A growing literature compares GARCH-family and LSTM approaches, and hybrid GARCH-LSTM models, for oil and financial market volatility forecasting, including on Brent crude oil specifically (Fu, 2023; Ben Romdhane & Boubaker, 2026). Within Nigeria, Adams and Uchema (2024) compare an RNN-LSTM against EGARCH for naira exchange-rate volatility, finding EGARCH outperforms the LSTM, and Adams (2024) runs a closely comparable comparison on the Nigerian Stock Exchange All-Share Index, again finding the asymmetric GARCH specification ahead of the LSTM. What has not yet been done is to bring this same comparative design, together with formal structural-break testing, to a Brent series long enough to span the full documented history of major oil-market regime shifts since the late 1980s, framed explicitly around Nigeria's fiscal and macroeconomic exposure to that benchmark rather than around a Nigerian exchange-rate or equity series alone. This is the gap this study addresses.

Nigeria's fiscal position, exchange-rate stability and macroeconomic planning all depend, directly or indirectly, on the price of crude oil, yet the existing literature that could inform how that price's volatility should be modelled and forecast has three specific limitations. First, the Nigerian oil-macroeconomy literature establishes that oil price volatility is associated with Nigerian macroeconomic outcomes, with genuinely mixed findings on how strong and how reliable that association is, but none of these studies compares a modern asymmetric-GARCH framework against a machine-learning alternative on the oil price series itself. Second, where a rigorous GARCH-versus-LSTM comparison has been carried out on a Nigerian series, it has been applied to the exchange rate (Adams & Uchema, 2024) or the stock index (Adams, 2024); where it has been carried out on Brent crude oil specifically (Fu, 2023; Ben Romdhane & Boubaker, 2026), it has not been framed around Nigeria's fiscal exposure or tested for structural breaks across a multi-decade history. Third, a genuinely Nigeria-specific crude oil price series, tracked historically by both the Central Bank of Nigeria and the U.S. Energy Information Administration, could not be assembled as a sufficiently long, currently accessible bulk series for this study, which is part of why Brent, rather than a Nigeria-quoted price, is used here. This study therefore investigates the

volatility of the Brent crude oil price, the global benchmark most directly relevant to Nigeria's oil-export receipts and fiscal planning, and asks which of two modelling paradigms – asymmetric GARCH-family econometric models or an LSTM deep-learning model – better captures and forecasts that volatility.

1.1 Aim and Objectives

The broad objective of this study is to examine the volatility characteristics of the global Brent crude oil price and to compare the performance of asymmetric GARCH models and an LSTM model in capturing and forecasting that volatility. The specific objectives are to:

- examine the trend and volatility characteristics of the monthly Brent crude oil price over a long (1987 to 2026) sample, including its stationarity, distributional and autocorrelation properties;
- formally test for structural breaks or regime shifts in oil price volatility associated with major, independently documented oil-market events;
- model the asymmetric (leverage) effect of positive versus negative oil price shocks using EGARCH and GJR-GARCH specifications;
- design and train an LSTM neural network to forecast oil price volatility using the same data; and
- compare the forecasting accuracy of the asymmetric GARCH models and the LSTM model using standard out-of-sample evaluation criteria.

1.2 Scope and Limitations

The study's empirical focus is the monthly Brent crude oil spot price (Europe), sourced from the U.S. Energy Information Administration via FRED series MCOILBRETEU, spanning May 1987 to July 2026 (471 price observations, 470 monthly log-returns). Brent is a global benchmark price rather than a Nigeria-quoted domestic series; a genuinely Nigeria-specific alternative (Salisu, Dikko & Doguwa, 2022, via the Central Bank of Nigeria; the EIA's own Bonny Light landed-cost series) was investigated directly for this study and found not currently usable for a sample extending to 2026, because the CBN series requires interactive navigation this study's tools could not execute and the EIA Bonny Light series has been largely withheld since approximately 2011 for commercial-confidentiality reasons. Because a genuine, independently verifiable bulk monthly series for the naira exchange rate, the Nigerian stock index, or Nigerian consumer price inflation could not be assembled, a new Nigeria-specific transmission regression is not attempted here; the study's account of oil price volatility's Nigerian relevance instead rests on the existing verified literature, and this is disclosed as a scope limitation rather than presented as a completed analysis. The comparative modelling itself is limited to symmetric GARCH(1,1), EGARCH(1,1) and GJR-

GARCH(1,1), benchmarked against a single-layer LSTM architecture; other machine-learning methods and hybrid GARCH-LSTM ensembles are noted as extensions for future research.

2.0 LITERATURE REVIEW

2.1 Theoretical Framework

Hamilton's (1983, 2003) oil-shock hypothesis argues that energy price shocks propagate into the wider economy through cost-push and expectations channels, and, for oil-exporting economies, through a fiscal and export-revenue channel that runs in the opposite direction to the cost-push channel felt by oil importers. For Nigeria, this transmission logic has been applied directly: Abdulkareem and Abdulkareem (2016) model oil price and macroeconomic volatility linkages in Nigeria using GARCH-family methods and find that oil price volatility spills over into key domestic aggregates, providing the theoretical and empirical anchor for treating oil price volatility as a variable of direct Nigerian macroeconomic interest.

A long-standing asymmetric price-transmission literature documents that downstream prices, costs or volatility can respond differently to upstream increases than to decreases of the same size, the “rockets and feathers” pattern (Peltzman, 2000). Applied to crude oil, this motivates asking whether positive oil-price shocks move conditional volatility by a different amount than negative shocks of equal size, rather than assuming, as a symmetric GARCH(1,1) specification implicitly does, that the sign of a shock is irrelevant. This is the asymmetric, or leverage, effect that EGARCH and GJR-GARCH are designed to detect. Unlike the equity-market leverage effect, where a negative shock is well established to raise volatility by more, the direction of this asymmetry for oil and other commodities is not settled a priori and is accordingly tested, rather than assumed, in this study.

Engle's (1982) Autoregressive Conditional Heteroskedasticity (ARCH) model and Bollerslev's (1986) Generalised ARCH (GARCH) extension formalise volatility clustering: large price changes tend to be followed by further large changes of either sign, so the conditional variance of a series is itself time-varying and predictable from its own past. Nelson's (1991) EGARCH and Glosten, Jagannathan and Runkle's (1993) GJR-GARCH extend this framework to allow the asymmetric response that the price-transmission literature predicts. Together, symmetric GARCH as a baseline and EGARCH and GJR-GARCH as asymmetric alternatives form the econometric backbone of this study's third objective.

Where the Efficient Market Hypothesis (Fama, 1970) suggests predictable structure in a price series should not persist, Lo's (2004) Adaptive Markets Hypothesis argues predictability can persist, and vary over time, in markets subject to shifting participants, institutional friction and behavioural bias – a

reasonable characterisation of a global commodity market that has passed through OPEC-supply-policy regimes, the US shale boom, a pandemic-driven demand collapse and geopolitically driven supply-shock episodes. This motivates using flexible, non-parametric learners such as LSTM networks (Hochreiter & Schmidhuber, 1997) alongside parametric GARCH models. Finally, a long span of monthly commodity-price data of the kind used here (1987 to 2026) plausibly passes through more than one volatility regime; Chow's (1960) test for a break in estimated coefficients between two sub-periods, together with the wider structural-break literature (Box & Tiao, 1975), provides the formal apparatus for testing this directly, at two candidate dates drawn from the documented history of the oil market: the late-2014 OPEC decision not to cut output amid the US shale boom, and the March 2020 COVID-19 demand collapse.

2.2 Empirical Review

Abdulkareem and Abdulkareem's (2016) CBN Journal of Applied Statistics study is a key Nigerian precedent for this study, finding significant volatility clustering in oil prices and evidence that this volatility spills over into domestic Nigerian aggregates. Two more recent studies sharpen, rather than simply confirm, this picture. Mgbomene et al.'s (2025) study, using GARCH estimates on quarterly Nigerian data from 1990 to 2023, finds that oil price volatility significantly reduces GDP growth while simultaneously raising government revenue, foreign exchange reserves and per-capita income. Ozokede and Awogbemi's (2025) study, using an annual ARDL model over a comparable period, instead finds crude oil price volatility's effect on real GDP growth is positive but not statistically significant, and concludes that Nigeria “lacks effective transmission mechanisms” converting oil-price movements into growth. Read together, these studies establish that oil price volatility is associated with Nigerian macroeconomic outcomes more robustly in some dimensions (fiscal revenue, reserves) than in others (aggregate growth), and that the transmission mechanism is itself contested rather than settled; none pairs its findings with a modern asymmetric-GARCH-versus-LSTM comparison on the oil price series itself, which is the specific extension this study offers.

Salisu, Dikko and Doguwa's (2022) study compares six symmetric and asymmetric GARCH-family specifications across three error distributions on monthly Nigerian crude oil prices (January 1982 to March 2019, 448 observations sourced from the Central Bank of Nigeria), finding that asymmetric EGARCH under a Student-t distribution achieves the best fit; no machine-learning model is included. Adegboyoye and Sarwar (2025) similarly find that asymmetric specifications provide a statistically better fit than symmetric GARCH for Nigerian financial series generally. Beyond Nigeria, Sekati et al. (2020) apply symmetric and asymmetric GARCH-family models to South African crude oil price volatility and likewise find evidence of asymmetric response to shocks, while Ng et al.

(2025) and Mwanakatwe et al. (2023) provide further cross-country confirmation that asymmetric conditional-volatility modelling is well established for oil and energy-linked series, with genuinely mixed findings on the direction of the asymmetric effect.

Further Nigerian evidence from Zubair and Boyi (2025) demonstrates the applicability of GARCH-type models to commodity-price volatility, using cotton and rubber stock returns to examine volatility clustering, persistence and asymmetric behaviour through symmetric and asymmetric GARCH specifications. This provides additional methodological support for the present study's use of GARCH-family models to characterise time-varying volatility in a commodity market. Related evidence is provided by Micheal and Anono (2025), who model and forecast volatility in the US dollar–Nigerian naira exchange rate, reinforcing the usefulness of conditional-volatility models for Nigerian macro-financial series. In addition, Zubair, Adams and Aniagolu (2022) examine determinants of the Nigerian inflation rate and show the broader macroeconomic relevance of fluctuations in key economic variables. Taken together, these studies strengthen the methodological and Nigerian macroeconomic context of the present study, while the current research differs by focusing specifically on global Brent crude oil price volatility and by comparing symmetric GARCH, EGARCH and GJR-GARCH with an LSTM model over a long monthly sample.

On GARCH-versus-LSTM comparisons specifically, Kim and Won (2018) hybridise LSTM with multiple GARCH-type models for stock-index volatility and find the hybrid outperforms either component model alone. Fu (2023) builds a monthly GARCH-LSTM hybrid specifically for Brent crude oil (January 2000 to December 2022) and reports improved forecast accuracy over a GARCH-only benchmark. Ben Romdhane and Boubaker (2026) go further: using daily Brent data from January 2010 to December 2022 (over 3,000 observations), they explicitly benchmark GARCH(1,1), asymmetric GJR-GARCH(1,1), a stochastic-volatility model and a standalone LSTM against one another, essentially the same comparative exercise this study undertakes, before proposing a three-stage hybrid that outperforms every individual benchmark. This is the closest precedent for this study's comparative design; the contribution of the present study is therefore not that a GARCH-versus-LSTM comparison on Brent has never been run, but that the existing comparison uses a decade of daily data that does not span the 1998, 2008 or 2014–2016 oil-market regimes and is not framed around any national economy's exposure to the benchmark, whereas this study's nearly forty-year monthly series (1987 to 2026) is built specifically to span those regimes, to test formally for breaks between them, and to connect the exercise to Nigeria's own fiscal exposure to Brent. Within Nigeria specifically, Adams and Uchema (2024) directly compare an RNN-LSTM against EGARCH for naira exchange-rate volatility (261 monthly observations, December 2001 to August 2023), finding EGARCH(1,1) consistently outperforms the RNN-LSTM, and Adams (2024) runs a

closely related comparison on the Nigerian Stock Exchange All-Share Index (468 monthly observations), again finding the asymmetric GARCH specification ahead of the LSTM by AIC and BIC. Nsengiyumva, Mung'atu and Ruranga (2025), studying foreign exchange risk in Rwanda, add a complementary finding that a hybrid GARCH-LSTM model improves Value-at-Risk estimation accuracy over either standalone model.

2.3 Gap in the Literature

Stated precisely, the gap this study addresses is a specific combination of elements, no single one of which is new in isolation. Asymmetric GARCH-versus-LSTM comparisons already exist for Brent crude oil (Ben Romdhane & Boubaker, 2026), for the Nigerian exchange rate (Adams & Uchema, 2024), and for the Nigerian stock index (Adams, 2024); asymmetric GARCH-family modelling of a Nigeria-specific crude oil series already exists without an LSTM comparison (Salisu et al., 2022); and the Nigerian oil-macroeconomy literature already establishes that oil price volatility matters for domestic outcomes, with a genuinely open question about how strongly and through what mechanism. What has not been brought together is: first, a Brent crude oil series long enough in monthly, calendar-time span to include the major, independently documented oil-market regimes since the late 1980s; second, formal, dated structural-break testing of that history rather than an assumption that a single set of parameters describes it; third, a full three-specification asymmetric GARCH-family comparison – symmetric GARCH, EGARCH and GJR-GARCH together – rather than one asymmetric variant alone, benchmarked against a standalone LSTM; and fourth, explicit framing of that comparison around Nigeria's fiscal and macroeconomic exposure to the benchmark price. This is a real but modest gap, a matter of combination and scope rather than an unprecedented method.

3.0 DATA AND METHODOLOGY

3.1 Research Design and Data Description

The study adopts a quantitative, ex-post-facto (non-experimental) time-series research design using secondary monthly data. Two modelling paradigms are estimated on the same univariate oil price return series and compared out-of-sample: three GARCH-family conditional-variance models of increasing flexibility – symmetric GARCH(1,1), EGARCH(1,1) and GJR-GARCH(1,1) – and a single-layer Long Short-Term Memory (LSTM) recurrent neural network. Crude oil price data were obtained from the U.S. Energy Information Administration (EIA) through the Federal Reserve Economic Data (FRED) database, series identifier MCOILBRETEU, covering the Europe Brent spot price, free on board (FOB), in U.S. dollars per barrel. The study covers May 1987 to July 2026 (471 monthly price

observations; 470 monthly log-return observations after the return transformation).

For the volatility analysis, the monthly Brent price series is transformed into continuously compounded returns:

$$R_t = 100[\ln(P_t) - \ln(P_{t-1})] \quad \dots (1)$$

where R_t denotes the monthly Brent crude oil return at time t , P_t represents the Brent crude oil price in month t , and P_{t-1} represents the price in the preceding month; multiplication by 100 expresses the return in percentage terms. The Brent series is an international benchmark rather than a Nigeria-specific price series; it is used because it is the best currently accessible proxy for Nigeria's oil-price exposure, since a genuinely Nigeria-specific series (Salisu et al., 2022, via the Central Bank of Nigeria; the EIA's Bonny Light series) could not be independently extended to 2026 using the tools available for this study.

3.2 Preliminary Diagnostics and Structural-Break Testing

Descriptive statistics (mean, standard deviation, minimum, maximum, skewness and kurtosis) and the Jarque-Bera normality test are computed for the return series. Stationarity is examined for both the price level and the return series using the Augmented Dickey-Fuller (ADF) test. The temporal dependence structure is examined using the autocorrelation function (ACF) and partial autocorrelation function (PACF), together with the Ljung-Box (Q) test for joint autocorrelation up to a specified number of lags, applied to both the returns and the squared returns. The presence of conditional heteroskedasticity is examined using the ARCH-LM test (Engle, 1982), which regresses the squared residuals on a constant and their own lagged squared values, with the corresponding Lagrange Multiplier statistic $LM = nR^2$ following a chi-square distribution with m degrees of freedom under the null of no ARCH effects.

Given the long, nearly forty-year span of the sample, formal structural-break tests are applied at two independently documented oil-market events – the November 2014 OPEC non-cut decision and the March 2020 COVID-19 demand collapse – using a variance-ratio F-test comparing pre- and post-break return variance, and a Chow (1960) mean-shift F-test comparing pre- and post-break mean returns. The resulting full-sample GARCH-family parameter estimates are additionally checked for interiority (i.e., not sitting at an estimation bound) as a joint check on whether a single-regime specification remains defensible despite any confirmed break.

3.3 Model Specification

All three GARCH-family models share the same mean equation:

$$R_t = \mu + \varepsilon_t, \quad \varepsilon_t = \sigma_t z_t \quad \dots (2)$$

where μ is a constant mean, ε_t is the mean-equation innovation, σ_t is its conditional standard deviation, and z_t is a standardised innovation assumed to follow a standardised Student-t distribution with v degrees of freedom, restricted to $v \in (4, 40)$; the Student-t distribution is chosen over the Normal because the descriptive statistics show pronounced excess kurtosis in the returns, consistent with the wider GARCH literature's preference for fat-tailed innovations in commodity and energy series.

The symmetric GARCH(1,1) conditional variance evolves as:

$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2, \quad \omega > 0, \quad \alpha, \beta \geq 0 \quad \dots (3)$$

where the persistence of a volatility shock is given by $\alpha + \beta$; a value close to 1 indicates highly persistent volatility. The EGARCH(1,1) log-conditional-variance evolves as:

$$\ln(\sigma_t^2) = \omega + \beta \ln(\sigma_{t-1}^2) + \alpha(|z_{t-1}| - E|z_{t-1}|) + \gamma z_{t-1} \quad \dots (4)$$

where $(E|z_{t-1}|)$ is the expected absolute standardised Student-t innovation, computed analytically from v . Because the recursion is in log-variance, no non-negativity constraint on ω , α or γ is required. The coefficient γ captures the asymmetric, or leverage, effect: a statistically significant γ indicates that positive and negative return shocks of equal magnitude move conditional volatility by different amounts. The GJR-GARCH(1,1) conditional variance evolves as:

$$\sigma_t^2 = \omega + \alpha \varepsilon_{t-1}^2 + \gamma \varepsilon_{t-1}^2 I_{t-1} + \beta \sigma_{t-1}^2 \quad \dots (5)$$

where the indicator

$$I_{t-1} = \begin{cases} 1, & \text{if } \varepsilon_{t-1} < 0, \\ 0, & \text{if } \varepsilon_{t-1} \geq 0. \end{cases}$$

. Here γ measures the additional effect of a negative shock: a positive γ means negative shocks raise volatility by more than positive shocks of the same size, while a negative γ means positive shocks raise volatility by more. For all three models, the parameter vector $\theta = (\mu, \omega, \alpha, \gamma, \beta, v)$, with γ omitted for symmetric GARCH(1,1), is estimated by maximising the conditional log-likelihood under the standardised Student-t assumption, with standard errors obtained from the numerical Hessian of the negative log-likelihood.

The LSTM is a single-layer recurrent network that maps a length-(L) lookback window of past returns to a one-step-ahead forecast of the same squared-residual volatility-proxy target used to evaluate the GARCH-family models in Section 3.6, giving a fair, like-for-like comparison.

For input x_t , previous hidden state h_{t-1} , and previous cell state c_{t-1} , the concatenated state-input vector is:

$$z_t = \begin{bmatrix} h_{t-1} \\ x_t \end{bmatrix}$$

The forget, input and output gates are:

$$f_t = \sigma(W_f z_t + b_f)$$

$$i_t = \sigma(W_i z_t + b_i)$$

$$o_t = \sigma(W_o z_t + b_o)$$

The candidate cell state is:

$$g_t = \tanh(W_g z_t + b_g)$$

The cell state is updated as:

$$c_t = f_t \odot c_{t-1} + i_t \odot g_t$$

The hidden state is:

$$h_t = o_t \odot \tanh(c_t)$$

where $\sigma(\cdot)$ is the logistic sigmoid function, $\tanh(c_t)$ is the hyperbolic tangent function, and $\odot$ is the elementwise Hadamard product.

The network output is a single affine read-out from the final hidden state:

$$\hat{y} = W_y h_L + b_y$$

The model is trained by minimising the mean squared error:

$$L = \frac{1}{2} (y - \hat{y})^2$$

via full backpropagation through time (BPTT), with parameters updated using the Adam optimiser. (Kingma & Ba, 2015).

3.4 Estimation Procedure

The full sample of 470 monthly log-returns is split chronologically 80/20 into a 376-observation training window (June 1987 to September 2018) and a 94-observation test window (October 2018 to July 2026), avoiding look-ahead bias. Symmetric GARCH(1,1) is fitted as a baseline, followed by EGARCH(1,1) and GJR-

GARCH(1,1), each by maximum likelihood on the training window, and compared in-sample using the Akaike and Bayesian Information Criteria (AIC and BIC). As a further robustness check on the structural-break findings, GARCH(1,1) and EGARCH(1,1) are re-estimated on a post-January-2015 sub-sample ($n = 139$, beginning immediately after the first candidate break date) and compared against the full-sample training-window estimates. For the LSTM, inputs are standardised using only training-sample mean and standard deviation; a lookback window of 6 months and a hidden layer of 8 units are used, with early stopping (patience of 25 epochs without validation-loss improvement) monitored on a chronologically held-out validation split.

3.5 Model Evaluation and Comparison Criteria

All models are evaluated out-of-sample against a common target, the squared, mean-adjusted return $\varepsilon_t^2 = (R_t - \hat{\mu})^2$, a standard realised-volatility proxy. Because a single metric cannot capture meaningful differences between models with such different functional forms, four complementary criteria are used together: Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE), which weight large forecast misses differently; Mean Absolute Percentage Error (MAPE), computed with a floor of 5% of the mean absolute target to prevent the denominator exploding in near-zero-volatility months; and Theil's Inequality Coefficient (U), which normalises RMSE against the scale of the actual and predicted series and is directly comparable across models fitted by entirely different numerical procedures. The Diebold-Mariano (1995) test, $DM = \frac{\bar{d}}{\sqrt{\widehat{\text{Var}}(\bar{d})}}$, where $d_t = |e_{1,t}|^2 - |e_{2,t}|^2$ is the loss differential between two competing forecasts and $\widehat{\text{Var}}(\bar{d})$ is a Newey-West heteroskedasticity-and-autocorrelation-consistent long-run variance estimator, is used to formally establish whether a difference in out-of-sample accuracy between the best-performing GARCH-family model and the LSTM is statistically significant.

4.0 EMPIRICAL RESULTS

4.1 Descriptive Statistics

Table 1 reports descriptive statistics for the full-sample monthly Brent log-return series ($n = 470$). The mean monthly return is 0.32% with a standard deviation of 9.90%; the distribution is mildly left-skewed (-0.49) and sharply leptokurtic (kurtosis = 9.06, against a Normal benchmark of 3), and the Jarque-Bera statistic (719.84, $p < .001$) decisively rejects normality. Notably, 0.0% of months in the series show an exactly-zero return, confirming Brent has been a continuously, genuinely market-determined price throughout the sample, which supports using

the full sample as the primary estimation window rather than a restricted sub-period.

Table 1: Descriptive Statistics of Monthly Brent Log>Returns (Full Sample, n = 470)

Statistic	Value
Mean (%)	0.320
Std. deviation (%)	9.904
Minimum (%)	−55.479
Maximum (%)	46.905
Skewness	−0.488
Kurtosis	9.061
Jarque-Bera statistic	719.84 (p < .001)
Share of exact-zero-return months	0.0%

Source: Author's computation, 2026, based on FRED series MCOILBRENTU.

4.2 Stationarity, Autocorrelation and ARCH Effects

Table 2 reports the stationarity and preliminary diagnostic tests. The price level is non-stationary (ADF $t = -1.723$, $p > .10$) while the log-return series is stationary (ADF $t = -9.020$, $p < .01$), confirming that modelling returns rather than levels is correct. Both the Ljung-Box test on squared returns ($Q = 210.00$, $p < .001$) and the ARCH-LM test ($LM = 171.78$, $p < .001$, $R^2 = .370$) reject the no-volatility-clustering null decisively, providing direct justification for GARCH-family modelling. The Ljung-Box test on the returns themselves is also significant ($Q = 52.06$, $p < .001$), indicating some linear autocorrelation remains in the mean of the series; the ACF and PACF (not tabulated here for brevity) show a significant lag-1 return autocorrelation of 0.268 and a much larger lag-1 squared-return autocorrelation of 0.595, decaying to within the confidence band by lag 3, a textbook GARCH(1,1)-consistent pattern.

Table 2: Stationarity and Preliminary Diagnostic Tests (n = 470)

Test	Statistic	p-value	Conclusion
ADF, price level (log)	$t = -1.723$ (2 lags)	$> .10$	Non-stationary, $I(1)$
ADF, log-return	$t = -9.020$ (6 lags)	$< .01$	Stationary, $I(0)$
Ljung-Box, returns (12 lags)	$Q = 52.06$	$< .001$	Significant autocorrelation
Ljung-Box, squared returns (12 lags)	$Q = 210.00$	$< .001$	Volatility clustering present
ARCH-LM (6 lags)	$LM = 171.78$ ($R^2 = .370$)	$< .001$	Conditional heteroskedasticity present

Source: Author's computation, 2026.

4.3 Structural Break Tests

Table 3 reports the variance-ratio and Chow-type mean-shift tests applied at the two documented oil-market events. At both dates, the post-break increase in return variance is highly significant ($p < .001$ in each case), while the shift in the mean return is not significant at conventional levels ($p = .616$ and $p = .839$, respectively). This confirms the Brent series, over its full nearly forty-year history, is not a single, textbook-stationary-variance process, motivating the robustness re-fit reported in Section 4.4.1.

Table 3: Structural Break Tests at Two Documented Oil-Market Events

Statistic	Nov 2014 (OPEC non-cut / shale glut)	Mar 2020 (COVID-19 demand collapse)
n (pre / post)	329 / 141	393 / 77
Mean return, pre / post (%)	0.471 / -0.030	0.279 / 0.531
Std. deviation, pre / post (%)	8.662 / 12.353	8.860 / 14.154

Statistic	Nov 2014 (OPEC non-cut / shale glut)	Mar 2020 (COVID-19 demand collapse)
Variance-ratio F (post/pre)	2.034	2.552
Variance-ratio p-value	< .001	< .001
Chow mean-shift F	0.252	0.041
Chow mean-shift p-value	0.616	0.839

Source: Author's computation, 2026.

4.4 GARCH-Family Estimation Results

The full sample was split chronologically 80/20 into a 376-observation training window (June 1987 to September 2018) and a 94-observation test window (October 2018 to July 2026). Table 4 reports the maximum-likelihood parameter estimates for the three GARCH-family models fitted on the training window. All three models' parameters are interior to their admissible ranges. GARCH(1,1) finds persistence of 0.903 ($\alpha = 0.259$, $\beta = 0.644$), typical for monthly commodity return series. EGARCH(1,1) achieves the lowest AIC (2653.32), suggesting the best in-sample fit once model complexity is only lightly penalised, but GARCH(1,1) achieves the lowest BIC (2675.44), since BIC penalises the extra asymmetry parameter of EGARCH and GJR-GARCH more heavily; the two information criteria therefore disagree on a single best model. Neither asymmetric specification's γ coefficient is statistically significant at the 5% level: EGARCH's γ is negative (-0.063 , $t = -1.158$) and GJR-GARCH's is positive (0.111 , $t = 0.980$), each with a t -statistic below the conventional 1.96 threshold. A statistically significant leverage effect in Brent crude oil price volatility is therefore not confirmed on this evidence, though the negative sign of EGARCH's estimate is directionally consistent with the rockets-and-feathers theory (positive shocks raising volatility by somewhat more than equivalent negative shocks).

Table 4: GARCH-Family Estimation Results, Training Window (n = 376)

Parameter	GARCH(1,1)	EGARCH(1,1)	GJR-GARCH(1,1)
μ (mean, %)	0.527	0.407	0.460
Ω	8.896	0.585	9.616

Parameter	GARCH(1,1)	EGARCH(1,1)	GJR-GARCH(1,1)
A	0.259	0.449	0.182
γ (asymmetry)	n/a	-0.063	0.111
γ standard error	n/a	0.0542	0.1133
γ t-statistic ($H_0: \gamma = 0$)	n/a	-1.158	0.980
B	0.644	0.866	0.647
ν (Student-t d.f.)	13.28	15.56	13.27
Persistence ($\alpha + \beta$)	0.903	n/a (log-form)	0.885
Log-likelihood	-1322.89	-1320.66	-1322.41
AIC	2655.79	2653.32	2656.82
BIC	2675.44	2676.89	2680.40

Source: Author's computation, 2026.

4.4.1 Robustness Check: Post-2015 Sub-Sample Re-Fit

As a direct check on whether the full-sample specification is defensible given the confirmed structural breaks (Section 4.3), GARCH(1,1) and EGARCH(1,1) were independently re-estimated on the sub-sample beginning immediately after the November 2014 break date, January 2015 to July 2026 ($n = 139$). Table 5 compares the full-sample and post-2015 estimates. Symmetric GARCH(1,1), well-behaved on the full 376-observation training window, collapses to a boundary solution (α at its upper bound, β at its lower bound) on the shorter 139-observation sub-sample, reflecting the smaller number of observations available for estimation. EGARCH(1,1), by contrast, remains interior on the smaller sub-sample ($\alpha = 0.609$, $\beta = 0.601$), and its asymmetry coefficient is again negative and again not statistically significant ($\gamma = -0.130$, $t = -0.959$), closely consistent in sign and magnitude with the full-sample estimate ($\gamma = -0.063$, $t = -1.158$). This cross-sample consistency is the strongest evidence available that the EGARCH-based asymmetry finding is not an artefact of the specific training window chosen, and demonstrates why full-sample estimation is preferable to a shorter sub-sample for this series.

Table 5: Post-2015 Sub-Sample Re-Fit versus Full-Sample Estimates

Parameter	GARCH, full (n=376)	GARCH, post-2015 (n=139)	EGARCH, full (n=376)	EGARCH, post- 2015 (n=139)
α	0.259	0.400 (bound)	0.449	0.609
β	0.644	0.000 (bound)	0.866	0.601
γ	n/a	n/a	-0.063	-0.130
γ t-statistic	n/a	n/a	-1.158	-0.959
ν	13.28	6.36	15.56	6.70

Source: Author's computation, 2026.

4.5 Residual Diagnostics

Table 6 reports diagnostic tests on the standardised residuals of all three fitted models. All three models pass the squared-standardised-residual Ljung-Box test and the ARCH-LM test comfortably ($p > .57$ in every case): the conditional-variance equation in each model has fully captured the volatility clustering present in the raw returns. However, all three models fail the Ljung-Box test on the (non-squared) standardised residuals themselves ($p = .003$ to $.004$), confirming that the linear autocorrelation identified in the raw returns is not removed by any of the three variance specifications, because none of them models the conditional mean beyond a constant. This is disclosed as a genuine, specific limitation; a future extension incorporating an autoregressive mean term (e.g., AR(1) or AR(4)) would be a well-motivated next step. It does not undermine the volatility-focused findings, since every variance-equation diagnostic reported is clean.

Table 6: Diagnostic Tests on Standardised Residuals, Training Window (n = 376)

Test	GARCH(1,1)	EGARCH(1,1)	GJR-GARCH(1,1)
Ljung-Box, standardised residuals (6 lags)	Q=19.13, p=.004	Q=19.92, p=.003	Q=18.96, p=.004
Ljung-Box, squared standardised residuals (6 lags)	Q=4.53, p=.605	Q=4.73, p=.579	Q=4.15, p=.656
ARCH-LM, standardised residuals (4 lags)	LM=1.44, p=.837	LM=1.08, p=.898	LM=1.13, p=.890

Source: Author's computation, 2026.

4.6 Out-of-Sample Forecast Comparison: GARCH-Family versus LSTM

Table 7 reports five forecast-accuracy criteria computed on the common squared-residual volatility-proxy target over the 94-month out-of-sample test window (October 2018 to July 2026, spanning the COVID-19 collapse and the 2026 geopolitically driven oil-price spike). No model dominates on every metric. EGARCH(1,1) achieves the lowest RMSE (427.77) and the lowest MAE (149.88), the two metrics most directly comparable to the AIC-based in-sample preference. GJR-GARCH(1,1) achieves the lowest (best) Theil's U (0.529), narrowly ahead of EGARCH (0.546) and GARCH (0.547). The LSTM achieves the lowest MAPE (272.89%), a full 100 percentage points below the best GARCH-family MAPE, but the highest (worst) RMSE and Theil's U of the four models, a pattern consistent with a model whose forecasts track typical, low-volatility months closely while under-reacting to the largest, sharpest volatility spikes in the test window.

Table 7: Forecast Accuracy Metrics, Out-of-Sample Test Window (n = 94)

Model	RMSE	MAE	MAPE (%)	Theil's U
GARCH(1,1)	461.27	172.35	454.54	0.547
EGARCH(1,1)	427.77	149.88	373.70	0.546
GJR-GARCH(1,1)	447.30	165.58	445.30	0.529
LSTM	492.70	160.94	272.89	0.764

Source: Author's computation, 2026.

A Diebold-Mariano test comparing EGARCH, the best-performing GARCH-family model by RMSE, against the LSTM gives $DM = -1.067$ ($p = .286$, $n = 94$): EGARCH's average forecast loss is lower, but the difference is not statistically significant at conventional levels. With 94 out-of-sample observations, this test is reasonably well powered, so this null result carries meaningful inferential weight; it is nonetheless a null result, and a numerically lower point estimate should not be read as a decisive win for EGARCH.

4.7 Discussion of Findings

Objective 1 is confirmed with strong statistical support: the Brent series shows no administratively-fixed sub-period, strong excess kurtosis and left skewness, a stationary return series with a non-stationary price level, and decisive evidence of volatility clustering, corroborated at the level of individual lags by the ACF and PACF. This is consistent with, and extends into a genuinely long sample, the transmission theory of Hamilton (1983, 2003) and the volatility-spillover

evidence Abdulkareem and Abdulkareem (2016) report for Nigerian oil and macroeconomic series specifically.

The finding also accords with the methodological evidence of Zubair and Boyi (2025), whose application of GARCH-type models to Nigerian commodity returns illustrates the usefulness of conditional-volatility modelling for detecting clustering, persistence and possible asymmetry in commodity markets. The relevance of volatility modelling beyond commodity prices is further supported by Micheal and Anono (2025) in the Nigerian exchange-rate context, while Zubair, Adams and Aniagolu (2022) demonstrate the wider macroeconomic importance of fluctuations in economic determinants through their analysis of Nigerian inflation. These studies therefore provide complementary Nigerian evidence for the present study's focus on volatility dynamics and their implications for macroeconomic stability.

Objective 2 is fully addressed: both the November 2014 and March 2020 candidate break dates show a highly significant increase in return variance but no significant shift in the mean return. Rather than treating this as disqualifying a single-regime GARCH-family specification, this study tested the consequence directly and found that EGARCH's full-sample and post-2015 sub-sample estimates are mutually consistent, while symmetric GARCH is the specification most vulnerable to sample-size-driven instability.

Objective 3, the asymmetric leverage effect, yields a clear, well-supported null: neither EGARCH's nor GJR-GARCH's asymmetry coefficient is statistically significant at the 5% level, and this null result is corroborated by the post-2015 robustness re-fit, where EGARCH's γ again fails to reach significance. On the evidence gathered here, a statistically confirmed leverage effect in Brent crude oil price volatility is not found, though the consistently negative sign of EGARCH's estimate across two independent samples is a directionally suggestive, if unconfirmed, hint that positive oil-price shocks may raise volatility somewhat more than negative shocks of equal size.

Objectives 4 and 5, LSTM design/training and comparative forecasting accuracy, are addressed together. On the 94-month out-of-sample test window, no single model dominates on every accuracy metric: EGARCH wins on RMSE and MAE, GJR-GARCH on Theil's U, and the LSTM on MAPE, and a Diebold-Mariano test comparing EGARCH against the LSTM finds the difference in average loss is not statistically significant ($p = .286$). This result sits between Adams and Uchema's (2024) finding that EGARCH outperformed an LSTM for Nigerian exchange-rate volatility and Nsengiyumva et al.'s (2025) finding that GARCH and LSTM combine to outperform either alone, motivating a hybrid GARCH-LSTM specification as a natural next step, following Kim and Won (2018) and Fu (2023).

5.0 CONCLUSION

This study examined the volatility of the global Brent crude oil price and compared asymmetric GARCH-family models against an LSTM neural network in capturing and forecasting that volatility, using a monthly Brent series of 471 price observations (470 log-returns, May 1987 to July 2026) sourced from the U.S. Energy Information Administration via FRED. Three broad conclusions follow. First, Brent crude oil price volatility is highly persistent (GARCH persistence ≈ 0.90) and characterised by strong excess kurtosis and left skewness, with decisive evidence of volatility clustering; formal structural-break tests confirm a significant post-break increase in return variance around both the November 2014 OPEC non-cut decision and the March 2020 COVID-19 demand collapse, though the full-sample GARCH-family estimates remain interior and well-behaved despite these confirmed regime shifts. Second, on this dataset, EGARCH(1,1) achieved the best in-sample fit by AIC and the lowest out-of-sample RMSE and MAE among the models compared, narrowly outperforming the LSTM on average forecast loss, but the margin over the LSTM was not statistically significant even with a well-powered 94-observation test window ($DM = -1.07$, $p = .29$); this sits alongside a genuinely mixed African and Nigerian literature – Adams and Uchemba (2024) similarly find EGARCH outperformed an LSTM for Nigerian exchange-rate volatility, while Nsengiyumva et al. (2025) find a hybrid GARCH-LSTM model improves on either component alone. Third, the asymmetric, leverage-type effect is directionally consistent (negative) across both the full sample and an independent post-2015 sub-sample, but is not statistically significant in either window, and should be treated as a plausible, sign-consistent but formally unconfirmed feature of the market over this period.

Because Nigeria's fiscal position, exchange-rate stability and macroeconomic planning all depend, directly or indirectly, on the Brent benchmark price to which its own export crude grades are priced at a differential, these findings carry direct implications for Nigerian fiscal and reserve-management planning, even though this study does not itself estimate a new Nigeria-specific transmission regression; that relevance rests on the existing, genuinely mixed verified literature (Abdulkareem & Abdulkareem, 2016; Mgbomene et al., 2025; Ozokede & Awogbemi, 2025).

Recommendations

- Nigerian fiscal authorities, including the Budget Office of the Federation and the Debt Management Office, should treat crude oil price volatility, and not only its level, as a distinct planning risk when setting the annual budget benchmark price; the persistent, clustered volatility documented here (GARCH persistence ≈ 0.90) means a period of calm prices is not informative about the likelihood of a subsequent volatile period.

- Because this study finds a directionally consistent, though not statistically significant, negative asymmetry coefficient, fiscal and reserve-management contingency planning should not assume oil-price volatility risk is symmetric between upside and downside price movements, while stopping short of building policy exclusively around an unconfirmed asymmetry.
- The Central Bank of Nigeria and the National Bureau of Statistics should prioritise publishing genuinely bulk-downloadable, machine-readable monthly exchange-rate and consumer-price-index series, to enable future Nigeria-specific oil-to-macroeconomy transmission regressions that the present study's data constraints could not support.
- Given that EGARCH matched or narrowly outperformed a from-scratch LSTM even on a well-powered 470-observation sample, Nigerian public and industry forecasting units should not assume that adopting a deep-learning model automatically improves oil-price volatility forecasts; a well-specified asymmetric GARCH model remains a strong, cheaper-to-estimate baseline, with a hybrid GARCH-LSTM approach the better-evidenced next step for future work.
- Future research should extend the comparison to hybrid GARCH-LSTM ensemble models, re-run the core comparison using the daily Brent series to substantially expand the usable sample size, and extend the asymmetry testing to a Markov-switching or time-varying-parameter GARCH specification that jointly estimates the structural-break regimes documented here.

REFERENCES

- Abdulkareem, A., & Abdulkareem, K. A. (2016). Analysing oil price-macroeconomic volatility in Nigeria. *CBN Journal of Applied Statistics*, 7(1), 1–22.
- Adams, S. O. (2024). Asymmetric GARCH type models and LSTM for volatility characteristics analysis of Nigeria stock exchange returns. *American Journal of Mathematics and Statistics*, 14(2).
- Adams, S. O., & Uchema, J. I. (2024). Nigeria exchange rate volatility: A comparative study of RNN-LSTM and EGARCH models. *Journal of Artificial Intelligence and Big Data*, 4(2), 61–73.
- Adegboyo, O. S., & Sarwar, K. (2025). Modelling and forecasting of Nigeria stock market volatility. *Future Business Journal*, 11(1), Article 124.
- Ben Romdhane, W., & Boubaker, H. (2026). A hybrid HAR-LSTM-GARCH model for forecasting volatility in energy markets. *Journal of Risk and Financial Management*, 19(1), Article 77.

- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327.
- Box, G. E. P., & Tiao, G. C. (1975). Intervention analysis with applications to economic and environmental problems. *Journal of the American Statistical Association*, 70(349), 70–79.
- Chow, G. C. (1960). Tests of equality between sets of coefficients in two linear regressions. *Econometrica*, 28(3), 591–605.
- Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263.
- Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987–1007.
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2), 383–417.
- Federal Reserve Bank of St. Louis. (2026). Global price of Brent Crude Oil [MCOILBRENTU]. *FRED, Federal Reserve Economic Data*. <https://fred.stlouisfed.org/series/MCOILBRENTU>
- Fu, X. (2023). Oil price forecasting model based on GARCH-LSTM model. *Frontiers in Business, Economics and Management*, 10(3), 28–31.
- Glosten, L. R., Jagannathan, R., & Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance*, 48(5), 1779–1801.
- Hamilton, J. D. (1983). Oil and the macroeconomy since World War II. *Journal of Political Economy*, 91(2), 228–248.
- Hamilton, J. D. (2003). What is an oil shock? *Journal of Econometrics*, 113(2), 363–398.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
- 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.
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*.
- Lo, A. W. (2004). The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *The Journal of Portfolio Management*, 30(5), 15–29.

- Mgbomene, C., Ashakah, F. O., Metieh, F. C., Okoro, E. N., Dim, H. C., & Opara, F. N. (2025). Crude oil price volatility and the Nigerian economy. *International Journal of Energy Economics and Policy*, 15(3), 390–402.
- Micheal, N. O., & Anono, Z. M. (2025). Modelling and forecasting the volatility of the exchange rate between the US dollar and the Nigeria naira note. *International Journal of Statistics and Applied Mathematics*, 10(8), 276–286.
- Mwanakatwe, P. K., Daniel, J., & Nzungu, K. R. (2023). Forecasting crude oil price volatility by Heston model. *Journal of Mathematical Finance*, 13(3), 408–420.
- Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 59(2), 347–370.
- Ng, K. Y., Zainal, Z., Samsudin, S., & Wei, H. (2025). Taming the turbulence: Modelling and forecasting crude oil price volatility with GARCH and SARIMA models. *Journal of Science and Technology*, 17(2), 48–63.
- Nsengiyumva, E., Mung'atu, J. K., & Ruranga, C. (2025). Hybrid GARCH-LSTM forecasting for foreign exchange risk. *FinTech*, 4(2), Article 22.
- Ozokede, E. V., & Awogbemi, T. O. (2025). The impact of crude oil price volatility on economic growth in Nigeria. *International Journal of Research and Innovation in Social Science*, 9(9), 7385–7396.
- Peltzman, S. (2000). Prices rise faster than they fall. *Journal of Political Economy*, 108(3), 466–502.
- Salisu, Z.-A., Dikko, H. G., & Doguwa, S. I. (2022). Comparing different symmetric and asymmetric GARCH models in three error distributions using Nigerian crude oil export prices. *International Journal of Advanced Academic Research*, 8(10), 108–134.
- Sekati, B. N. Y., Tsoku, J. T., Metsileng, L. D., & Tokic, D. (2020). Modelling the oil price volatility and macroeconomic variables in South Africa using the symmetric and asymmetric GARCH models. *Cogent Economics & Finance*, 8(1), Article 1792153.
- U.S. Energy Information Administration. (2026). Europe Brent spot price FOB (compiled by the Federal Reserve Bank of St. Louis as FRED series MCOILBRETEU). *U.S. Energy Information Administration*. <https://fred.stlouisfed.org/series/MCOILBRETEU>
- Zubair, M. A., & Boyi, B. A. (2025). Modelling the volatility of cotton and rubber stock returns in Nigeria using GARCH type models. *FUDMA Journal of Sciences*, 9(5), 67–79.
- Zubair, M. A., Adams, S. O., & Aniagolu, K. S. (2022). Economic impact of some determinant factors of Nigerian inflation rate. *Journal of Economics and Financial Analysis*, 5(2), 23–41.