ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

Assessing Models Behaviors and the Forecasting Performance of Arima and Garch Models: An Empirical Study

Olusola S. Oluwabunmi[1] & Nasiru M. Olakorede[2]

[1,2]Department of Statistics, University of Abuja, Abuja, Nigeria

DOI: https://doi.org/10.51583/IJLTEMAS.2026.150400110

Received: 10 April 2026; Accepted: 16 April 2026; Published: 19 May 2026

ABSTRACT

This study investigates the statistical properties, volatility dynamics, and forecasting performance of crude oil returns using a comprehensive time series modeling approach. Descriptive analyses reveal that returns are centered around zero with small negative averages, exhibiting pronounced volatility clustering and episodes of extreme deviations driven by geopolitical and macroeconomic shocks. Diagnostic tests confirm the presence of nonlinear dependence and heteroskedasticity, making models such as GARCH suitable for capturing the persistent volatility patterns. Stationarity of the return series supports the application of ARIMA and GARCH-based models, which effectively accommodate the complex features observed in the data. Model selection across varying sample sizes consistently favors parsimonious ARIMA(0,0,q) structures with no autoregressive terms, while GARCH(1,1) captures the high volatility persistence evident in the market. Forecasting evaluations demonstrate that model accuracy improves with larger datasets, with combined ARIMA-GARCH models, especially those with higher-order ARIMA specifications, outperforming simpler models in larger samples. The findings underscore the importance of incorporating volatility modeling and selecting appropriate model complexity based on data availability to enhance forecast precision. Overall, the results provide valuable insights into the dynamics of crude oil prices and offer robust guidance for modeling and forecasting in volatile energy markets.

Keywortds: ARIMA, GARCH, Forecasting, Volatility, Parsimonious, Financial, Crude Oil.

INTRODUCTION

Forecasting in statistics is concerned with predicting future outcomes using historical observations arranged in time order. A time series differs from other datasets because its values are recorded at regular intervals, making it possible to detect systematic patterns such as upward trends, cyclical fluctuations, or random disturbances. Understanding these dynamics is essential in domains like agriculture, business, and finance where planning depends on reliable projections.

The Autoregressive Integrated Moving Average (ARIMA) framework is one of the most established approaches for such analysis. It combines two fundamental ideas: the autoregressive (AR) model, which relates current observations to previous values, and the moving average (MA) model, which expresses present outcomes in terms of past forecast errors. When data are non-stationary, the integrated (I) component applies differencing to stabilize the mean and variance before modeling. This makes ARIMA particularly useful for data that exhibit trends over time.

Since its development and popularization by Box and Jenkins (1976), ARIMA has become a versatile method for short-term prediction of linear time series. It can be implemented in two major forms: a non-seasonal ARIMA model, which captures trends and dependencies without seasonal adjustment, and a seasonal ARIMA (SARIMA) model, which incorporates recurring cycles such as monthly or yearly effects. These variations allow ARIMA to be adapted across contexts, from analyzing financial returns to modelling agricultural production levels.

Financial markets are often characterized by irregular fluctuations in asset prices, where periods of calm are interrupted by episodes of extreme turbulence. This feature, known as volatility clustering, implies that returns are not constant over time and that traditional models assuming fixed variance are inadequate. Understanding

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

and forecasting this variability is central to decision-making in risk management, asset pricing, and portfolio allocation.

The early attempt to capture changing variance was made by Engle (1982) through the Autoregressive Conditional Heteroskedasticity (ARCH) model, which demonstrated that the conditional variance of a return series could be linked to past error terms. Although successful, the ARCH model often required a large number of lagged terms to fit financial data adequately. To address this, Bollerslev (1986) proposed the Generalized ARCH (GARCH) model, which extended Engle's approach by allowing the conditional variance to depend both on lagged squared errors and on past variances. This development made the model more flexible and computationally efficient.

Since its introduction, the GARCH family of models has become a dominant tool in financial econometrics. They are widely applied to estimate Value-at-Risk (VaR), study exchange rate dynamics, analyze stock market returns, and support regulatory capital requirements. The core strength of GARCH lies in its ability to capture persistence in volatility and to reflect the clustering commonly observed in empirical financial data. In addition, several extensions such as Exponential GARCH (EGARCH), Threshold GARCH (TGARCH), and GJR-GARCH have been introduced to address asymmetry in volatility, where negative shocks tend to have a stronger effect than positive ones.

Recently, there have been increased attention in the literature towards analysis of volatility data using different approaches.

Mbonigaba et al. (2025) employed GARCH models for predicting market volatility in high-frequency trading. They studied data from 2020 to 2024. The research compared traditional GARCH models with newer versions like EGARCH and TGARCH. They found that while traditional GARCH models are a good starting point for forecasting volatility, they aren't accurate enough in high-frequency trading without improvements. The study showed that the newer GARCH versions work better at capturing how volatility changes in response to market events. To improve predictions in high-frequency trading, the authors suggest combining GARCH models with machine learning, using methods to reduce noise, and creating ways to adjust models in real time for better results.

Di-Giorgi et.al. (2025) proposed a hybrid approach to volatility prediction by combining GARCH models with deep neural networks (GRU, LSTM, BiLSTM). Their method also introduces an optimal sliding window mechanism. Through simulations and applications to real financial data (Chile's IPSA, S&P 500, ASX200), the approach showed improved accuracy in forecasting stock return volatility, particularly one week ahead.

Sirisha et al. (2024) Conducted a comparative analysis of time series forecasting methods using both statistical and deep learning approaches. Specifically, they applied the Autoregressive Integrated Moving Average (ARIMA), Seasonal ARIMA (SARIMA), and the deep learning-based Long Short-Term Memory (LSTM) neural network models to forecast profit data. To ensure the appropriateness of the statistical models, the dataset used was transformed to stationarity for ARIMA, while no such transformation was performed for SARIMA and LSTM. The models were trained and evaluated using test data to assess forecasting performance. Their findings revealed that the LSTM model outperformed both ARIMA and SARIMA, demonstrating superior predictive accuracy and highlighting the potential of deep learning methods in complex time series forecasting tasks.

Bunnag (2024) analyzed how gold acts as a safe-haven asset when the economy is uncertain. They used two models, ARIMA (2,1,3) and ARIMA (2,1,3)-GARCH (1,1), to predict gold prices. When they compared the models, the ARIMA (2,1,3)-GARCH (1,1) one performed better, with lower MAE (80.1371) and RMSE (96.8299), showing it's more accurate. The model predicts that gold will cost about 1942.094 USD per troy ounce by the end of 2024, which shows how well it can forecast trends in the gold market.

Hasanov et al. (2024) examined the empirical significance of structural changes in exchange rate data by replicating and extending the volatility forecasting study of Rapach and Strauss (2008). Using generalized autoregressive conditional heteroskedasticity (GARCH) models and out-of-sample tests, they incorporated recent U.S. dollar daily exchange rate data, employed alternative software, and applied updated forecast accuracy

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

tests and loss metrics. Their objective was to achieve scientific replication in a broad sense while assessing robustness. The findings broadly confirmed the original study, with additional evidence showing that models accounting for structural breaks consistently outperformed those that ignored instabilities across all forecast horizons and loss functions.

Jatinder et al. (2023) highlights the prominence of the ARIMA model as a widely accepted and robust time series forecasting method. They conducted an extensive analysis of ARIMA and hybrid models incorporating ARIMA, spanning fields such as environmental, health, atmospheric, and other disciplines. Their findings suggest that hybrid models, which integrate ARIMA with other approaches, outperform standalone ARIMA models by effectively capturing a broader range of patterns within time series data. Consequently, the adoption of hybrid modelling has become a standard practice, leveraging the combined strengths of multiple methods to enhance forecasting accuracy and robustness**.**

Quang et al. (2024) ground their study in financial time-series and volatility modeling theory, which explains Bitcoin price behavior through historical price movements and time-varying volatility. The framework assumes that Bitcoin prices exhibit non-stationarity, autocorrelation, and volatility clustering. Bitcoin price returns are treated as the dependent variable, with the ARIMA model capturing short-term return dynamics in the mean equation, while the GARCH model explains conditional volatility in the variance equation. The framework acknowledges the influence of macroeconomic and financial factors on Bitcoin prices but excludes them from the ARIMA–GARCH structure, indicating a limitation and direction for future research. Model specification and frequent recalibration are incorporated as moderating elements due to Bitcoin's high volatility. Overall, the framework highlights the effectiveness of the ARIMA–GARCH integration for short-term Bitcoin price forecasting while emphasizing the need for adaptive and extended modeling approaches for improved long-term accuracy.

The aim of this study is to evaluate and compare the forecasting accuracy of ARIMA, GARCH and ARIME-GARCH models through real-world data applications, thereby providing insights into their respective effectiveness in modeling and forecasting time series data exhibiting volatility dynamics.

The objectives were to;

i. Examine the trends and patterns of crude oil prices returns;
ii. Fit the crude oil prices returns to both ARIMA and GARCH models; and
iii. evaluate and compare the forecasting accuracy of ARIMA, GARCH and ARIMA-GARCH models.

METHODOLOGY

Autoregressive Integrated Moving Average (ARIMA) Model

The autoregressive integrated moving average (ARIMA) model is a sophisticated statistical tool that synthesizes autoregressive (AR) and moving average (MA) components with a differencing process to effectively model time series data. The differencing step is pivotal in transforming non-stationary data into a stationary form, thereby facilitating reliable statistical analysis. This integrated approach, combining AR, MA, and differencing, is termed ARIMA. This section elucidates the mathematical derivation of the autoregressive (AR) component, establishing the groundwork for its integration with MA and differencing processes. An autoregressive (AR) model is a representation of a type of random process. It is used to describe certain time-varying processes in time series data.

The ARIMA Model

ARIMA models are used for time series data that either already have a constant mean and variance over time or can be made to have these properties through a process called differencing. The model is written as ARIMA (p, d, q),

where:

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

- p: Order of the Autoregressive (AR) component.
- d: Degree of differencing required to make the time series stationary.
- q: Order of the Moving Average (MA) component.

Component of ARIMA

Autoregressive (AR): The AR part of the model looks at how an observation connects to past observations. It believes that the current value of the time series is made up of a mix of earlier values and some random error.

Mathematical Representation:

$$y_t = c + \varphi_1 y_{t-1} + \varphi_2 y_{t-2} + \cdots + \varphi_p y_{t-p} + \mathcal{E}_t \quad 2.1$$

Where:

y_t : Value of the time series at time (t)

(c): Constant term.

$\varphi_1, \varphi_2 \ldots \varphi_p$: Parameters (coefficients) of the lagged terms.

$\mathcal{E}_t$: Random error (white noise) at time (t).

(p): Number of lagged observations (order of the AR model).

Integrated (I): The "I" in ARIMA stands for integration, which refers to differencing the time series to make it stationary. A stationary time series has a constant mean, variance, and autocorrelation structure over time. Differencing removes trends and stabilizes the mean.

Differencing:

First-order differencing: $y'_t = y_t + y_{t-1}$

Second-order differencing (if needed): $y''_t = y'_t + y'_{t-1}$

(d): Number of times the data is differenced to achieve stationarity.

Moving Average (MA): The MA component shows how the current observation is connected to previous forecast errors. An MA(q) model shows the current value as a mix of past error terms.

Mathematical Representation:

$$y_t = c + \mathcal{E}_t + \theta_1 \mathcal{E}_{t-1} + \theta_2 \mathcal{E}_{t-2} \ldots + \theta_q \mathcal{E}_{t-q} \quad 2.2$$

Where:

y_t : The observed time series value at time (t).

c: The mean of the time series (often assumed to be 0 for simplicity in a stationary series)

$\mathcal{E}_t$: The white noise error term at time (t), assumed to be independent and identically distributed with mean 0 and variance δ^2

$\theta_1, \theta_2 \ldots \theta_q$: The moving average coefficients, which measure the impact of past errors on the current observation.

(q): The order of the MA model, indicating the number of lagged error terms included.

Hence, The ARIMA(p,d,q) model can be written as:

$$\varphi(B)(1 - B)^d y_t = \theta(B)\mathcal{E}_t \quad 2.3$$

Where:

y_t : The observed time series value at time (t).

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

$\varphi(B)$: The autoregressive (AR) operator, defined as $\varphi(B) = 1 - \varphi_1 B - \varphi_2 B^2 - \cdots - \varphi_P B^P$ where $\emptyset_1, \emptyset_2$ ... $\emptyset_p$ are the AR coefficients, and (B) is the backshift operator.

$(1 - B)^d$: The differencing operator is applied (d) times to make the series stationary by removing trends.

$\theta(B)$: The moving average (MA) operator, defined as $\theta(B) = 1 - \theta_1 B - \theta_2 B^2 \ldots \theta_P B^P$ Where $\theta_1, \theta_2 \ldots . \theta_p$ These are the MA coefficients.

ϵ_t : The error term (white noise) at time (t), assumed to be independent and identically distributed with mean 0 and variance σ^2

Forecasting in the ARIMA Model

In ARIMA modelling, forecasting is a crucial application for predicting future values in a time series. Moreover, we often want to predict the future value of a series, Y_t, based on past observations. Two common forecasting strategies are One-step-ahead forecasting and multi-step-ahead forecasting.

One-Step Ahead Forecasting in ARIMA

One-step ahead forecasting predicts the next value in the time series, which is the observation at time t+1, by using all the data available up to time t. It calculates the expected value of the time series based on the structure of the ARIMA model.

Mathematically,

$$\hat{Y}_{t+1|t} = \varphi_1 Y_t + \varphi_2 Y_{t-1} + \cdots + \varphi_p Y_{t-p+1} + \theta_1 \varepsilon_t + \theta_2 \varepsilon_{t-1} + \varphi_q \varepsilon_{t-q+1} \qquad 2.4$$

Where;

$\hat{Y}_{t+1|t}$: forecast of Y_{t+1} given data up to t.

φ_i: autoregressive coefficients.

ε_t :forecast error at time t.

Multistep Ahead Forecasting in ARIMA

Multistep ahead forecasting predicts values beyond one period (y_{t+h} where $h \geq 2$). Since future observations and errors are unknown, the forecast uses expected values, setting future errors to zero ($E[\epsilon_{t+h}] = 0$) and replacing future observations with their forecasts.

The forecast equation becomes recursive:

$$\hat{Y}_{t+1|t} = \varphi_1 \hat{Y}_{t+h-1|t} + \varphi_1 \hat{Y}_{t+h-2|t} + \cdots + \varphi_p \hat{Y}_{t+h-1|t} + \cdots + \theta_1 \breve{\varepsilon}_{t+h-1|t} + \theta_q \breve{\varepsilon}_{t+h-q|t} \qquad 2.5$$

Where:

$\hat{Y}_{t+1|t}$: forecast Y_{t+1} given data up to t.

For h>1, the past values and errors used are themselves **forecasts** rather than actual observations.

$\breve{\varepsilon}_{t+h-1|t}$: forecast error at step h, which is assumed to be zero for h>1

Generalized Autoregressive Conditional Heteroskedasticity (GARCH) Model

The GARCH model is commonly used in time series analysis to understand and predict volatility, especially in financial and economic data. Volatility often shows clustering, meaning big changes tend to follow big changes, and small changes follow small ones. Unlike ARIMA, which focuses on the average value of a time series, GARCH looks at how the variability changes over time. It builds on the ARCH model by adding lagged variance terms, making it more adaptable and efficient. The GARCH(p,q) model includes two key equations: one for the

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

average and one for the variance, where p represents the number of lagged variances and q represents the number of lagged squared errors.

Mathematically,

A time series y_t can be represented as:

$$y_t = \mu_t + \mathcal{E}_t,\ \mathcal{E}_t = \delta_t z_t \quad 2.6$$

Where:

μ_t: conditional mean (can be modelled as ARMA, ARIMA, etc.)

$\mathcal{E}_t$: error term,

$z_t \sim$ iid (0,1) Standardize White Noise

δ_t^2 : Conditional Variance at time t.

Hence, the GARCH (p,q) model for the conditional variance is:

$$\delta_t^2 = \omega + \sum_{i=1}^{q} \alpha_i \mathcal{E}_{t-1}^2 + \sum_{j-1}^{p} \beta_j \delta_{t-j}^2 \quad 2.7$$

δ_t^2: conditional variance (the forecasted volatility)

$\omega > 0$: constant term (baseline variance)

$\alpha_i \geq 0$: ARCH parameters

$\beta_j \geq 0$: GARCH parameters (long-run persistence of past variance).

$\alpha_1 + \beta_j > 0$: ensures stationarity (volatility does not explode)

Forecasting in the GARCH Model

Forecasting in GARCH focuses on predicting future conditional variances h_{t+h} for horizon h= 1,2... which is crucial for risk management, Value-at-Risk (VaR) estimation, option pricing, and portfolio optimisation. Unlike ARIMA models, which forecast the level of the series, GARCH captures volatility clustering, Value, and time-varying heteroskedasticity. ARCH forecasts are generated after estimating the model parameters using maximum likelihood estimation (MLE) on historical data. The process involves two components: the mean equation (often a simple constant or ARMA for returns) and the variance equation. Forecasting methods are recursive and rely on the autoregressive nature of the model. Below are the forecasting methods, along with their mathematical representations for the general GARCH(p,q) model and the commonly used GARCH(1,1) variant**.**

One-Step Ahead Forecasting

One-step ahead forecasting predicts the next period's volatility h_{t+1} using all information up to time (t). This is the most accurate short-term forecast since it incorporates the latest observed residual $\mathcal{E}_t$ and variance h_t

For h =1

$$E_t\left(\delta_{t+h}^2\right) = \omega + \alpha_1 \mathcal{E}_t^2 + \beta_1 \delta_t^2 \quad 2.8$$

ω : constant term (long-run average variance).

α_1: ARCH parameter: measures how much yesterday's shock ($\mathcal{E}_t^2$) affects today's volatility.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

β_1: GARCH parameter: measures how much yesterday's variance (δ_t^2) carries over into today's variance.

$\mathcal{E}_t^2$: the error/residual from the mean equation at time t.

δ_t^2 : the error/residual from the mean equation at time t.

Multi-Step Ahead Forecast

For horizons h≥2, forecasts are recursive because future residuals $\delta_t + K_{residual}$ (k>0) are unknown. The approach sets future shocks to their expectation ($E(\delta_{t+k}] = 0$ so ($\delta_{t+k}^2 \approx 0$) and uses previously forecasted variances. This leads to forecasts that decay toward the unconditional variance as (h) increases, reflecting mean-reversion in volatility.

For GARCH (1,1):

$$\mathcal{E}_t(\delta_t^2) = \omega + (\alpha_1 + \beta_1)E_t(\delta_{t+h-1}^2) \quad 2.9$$

This means:

As h→∞,

$E_t(\delta_{t+h}^2) \rightarrow \frac{\omega}{1-\alpha_1-\beta_1}$

Where:

ω: constant (baseline variance)

α: ARCH coefficient

β: GARCH coefficient

δ_t^2 : most recent conditional variance at time t.

h: forecast horizon (number of steps ahead).

$E_t(\delta_{t+h}^2) \rightarrow \frac{\omega}{1-\alpha_1-\beta_1}$: unconditional (long-run) variance. 2.10

ARMA- GARCH Model

The hybrid ARIMA-GARCH model combines an ARIMA(p, d, q) process for the conditional mean of the time series y_t with a GARCH(r, s) process for the conditional variance (volatility) of the residuals

ϵ_t. The Mean equation (ARIMA p,d,q Component), is:

$$\varphi(B)(1-B)^d(y_t - \mu) = \theta(B)\varepsilon_t \quad 2.11$$

Expanded form:

$$y_t = \mu + \sum_{i=1}^{p} \varphi_1(y_{t-i} - \mu) + \epsilon_t + \sum_{j=1}^{q} \theta_j \epsilon_{t-j} \quad 2.12$$

Where ϵ_t is the innovation (residual) term that is not white noise.

The Variance Equation GARCH component

$\varepsilon_t = \delta_t z_t$ z_t~i.i.d. (0,1)

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

$\sigma_t^2 = \omega + \sum_{i=1}^{r} a_i \varepsilon_t^2 + \sum_{j=1}^{s} \beta_j \sigma_{t-j}^2$ 2.13

The most widely used specification in practice is GARCH (1,1), i.e., r = 1, s = 1.

Where:

y_t = Observed value of the time series at time (t)

μ= Constant mean (drift) of the process

$\varphi(B) = 1 - \sum_{i=1}^{p} \vartheta_i B^i$ AR (autoregressive) polynomial of order (p)

$\theta(B) = 1 - \sum_{j=1}^{q} \theta_j B^j$ MA (moving average) polynomial of order (q)

φ_1 = AR coefficients

θ_j = MA coefficients

ε_t = Residual/innovation at time (t)

σ_t^2 = Conditional variance (volatility) at time (t)

(B) = Backshift (lag) operator, where $B^k y_t = y_{t-k}$

(d): Order of differencing (integration) to achieve stationarity

ω >Constant term in the variance equation

$a_i \geq 0$ ARCH coefficients (impact of past squared shocks)

$B_j \geq 0$ GARCH coefficients (impact of past conditional variances)

z_t Standardized white-noise innovation

(r): Order of ARCH terms in GARCH

(s): Order of GARCH terms

The onstraints for GARCH stationarity:

$\omega > 0, a_i \geq 0, B_j \geq 0, \sum_{i=1}^{r} \alpha_i + \sum_{j=1}^{s} \beta_j < 1$

Forecasting Model of ARIMA-GARCH

The hybrid model produces point forecasts for the mean and volatility forecasts for uncertainty.

- h-step-ahead forecast for the mean ($\hat{Y}_{t+h|t)}$ Obtained directly from the fitted ARIMA(p, d, q) equation using the recursive ARMA forecasting formula on the differenced series, then integrated back (undo differencing).
- h-step-ahead forecast for conditional variance $\hat{\sigma}_{t+h|t}^2$Computed recursively from the GARCH equation
- h = 1

$\hat{\sigma}_{t+h|t}^2 = \omega + \sum_{i=1}^{r} \alpha_i \epsilon_{t+1-i}^2 + \sum_{j=i}^{s} \beta_j + \hat{\sigma}_{t+1-j}^2$

$\hat{\sigma}_{t+h|t}^2 = \omega + \left(\sum_{i=1}^{r} \alpha_i + \sum_{j=i}^{s} \beta_j\right) \hat{\sigma}_{t+h-1|t}^2$ 2.14

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

Since

$E(\epsilon_{t+h-i}^2|F_t) = \hat{\sigma}_{t+h|t}^2$ for future periods)

The full forecast interval for y_{t+h} is typically:

$\hat{Y}_{t+h|t)} \pm z_{\frac{a}{2}}.\hat{\sigma}_{t+h|t}^2$

Where:

$z_{\frac{a}{2}}$ = is the critical value from the distribution of z_t

$\hat{Y}_{t+h|t)}$ = h-step-ahead point forecast of the original time series

(h): Forecast horizon (number of steps ahead)

$\hat{\sigma}_{t+h|t}^2$ = h-step-ahead forecast of the conditional variance

ω =Constant (intercept) term in the GARCH variance equation.

α_i = ARCH coefficients

β_j = GARCH coefficients

ε_t = Residual (innovation) from the ARIMA mean equation at time t

σ_t^2 = Conditional variance at time t

(r): Order of the ARCH terms in the GARCH model

(s): Order of the GARCH terms in the GARCH model

$z_{\frac{a}{2}}$ = Critical value from the distribution of the standardized innovation

ANALYSIS AND RESULT

Crude Oil Prices Returns

Crude oil prices data use for this study consist of daily time series observations extracted from the Central Bank of Nigeria (CBN) website, spanning the period from 25th October 2021 to 11th November 2025. The dataset captures the day-to-day movement in crude oil prices over the study period, making it appropriate for time series analysis because it preserves the chronological ordering and short-run fluctuations inherent in oil price dynamics. As a daily frequency dataset, it is particularly useful for examining patterns such as volatility, trends, shocks, and possible structural changes in crude oil prices over time. The use of data sourced from CBN enhances the credibility and reliability of the study given the institution's role as an official repository of key macroeconomic and financial statistics in Nigeria. The analysis of the crude oil return data reveals several key insights relevant to real-world applications in financial econometrics and commodity market analysis.

Table 3.1: Descriptive Statistics and Preliminary Tests of the Crude Oil Price Returns

Sample Size	**50**	**100**	**250**	**500**	**1000**
Descriptive Statistics					
Mean	-0.0016355	-0.002128	-0.00075	-0.000740	-0.000276
Min.	-0.0430339	-0.094261	-0.09426	-0.09426	-1.086203

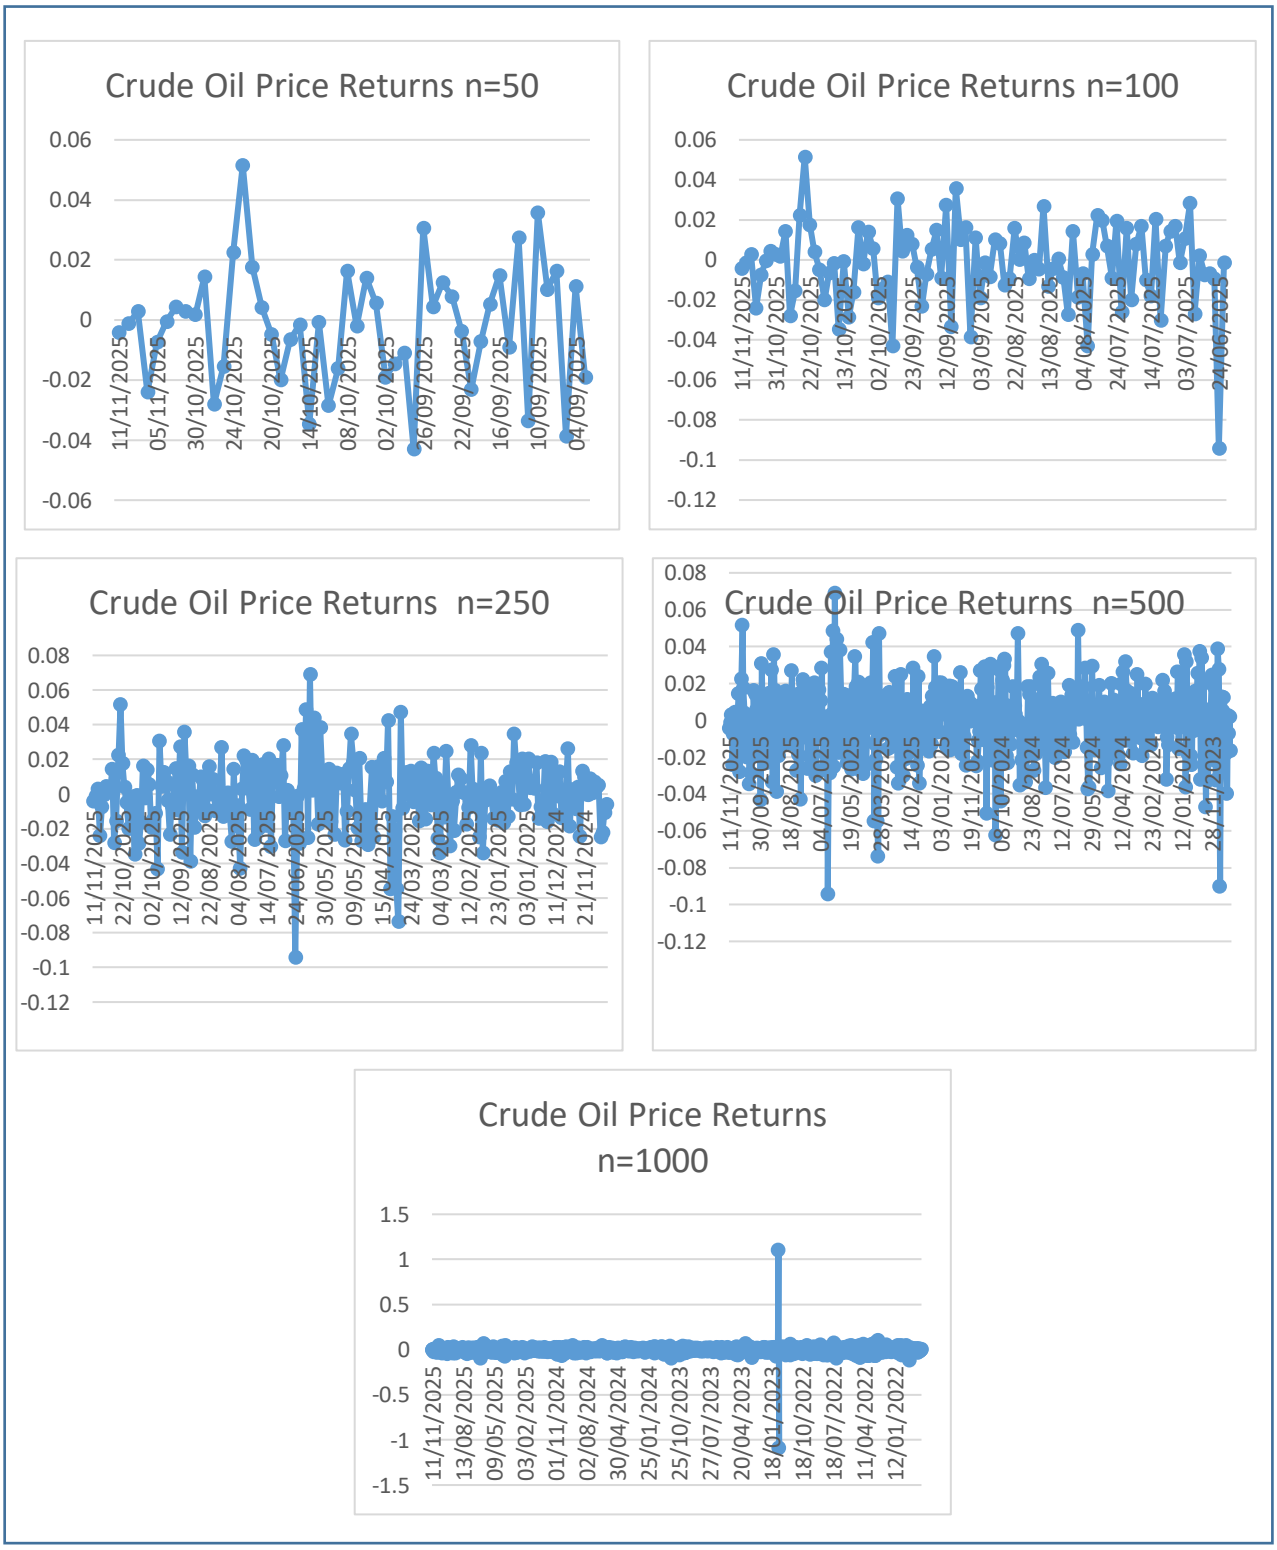
ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

Max.	0.0515610	0.051561	0.06921	0.06921	1.102399
SD	0.0195205	0.019979	0.01993	0.01942	0.05408
Preliminary Tests					
ADF Test	I(0) 0.0131*	I(0) 0.01*	I(0) 0.01*	I(0) 0.01*	I(0) 0.01*
BDS	0.0098*	0.01*	0.0326*	0.0097*	0.00002*
ARCH Test	0.7538	0.9095	0.0145*	0.0031*	2.20E-16*

Note: * denotes significant at 0.05 level

Source: Researchers' Compilations from R-Output

Figure 3.1. Time Series Plots of Crude Oil Price Returns



ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

Table 3.2: Estimation of Appropriate ARIMA Models across the Crude Oil Price Returns Sample Sizes

50	**100**	**250**	**500**	**1000**
ARIMA(0,0,0)	**ARIMA(0,0,0)**	**ARIMA(0,0,2)**	**ARIMA(0,0,3)**	**ARIMA(0,0,1)**
log-likelihood 97.47	log-likelihood 233.73	ma1 -0.0593*	mar1 -0.0244*	ma1 0.5436*
AIC -192.94	AIC -465.45	ma2 0.1575*	ma2 0.1501	log-likelihood 1607.16
-	-	log-likelihood 599.85	ma3 0.0322*	AIC 3210.31
-	-	AIC -1193.7	log-likelihood 1234.67	-
-	-	-	AIC 2467.33	-

Note: * denotes significant at 0.05 level.

Source: Researchers' Compilations from R-Output

Table 3.3: Estimation of GARCH(1,1) Model across the Crude Oil Price Returns Sample Sizes

Sample Size	**50**	**100**	**250**	**500**	**1000**
Mu	-0.001774	-0.001184	-0.000650	-0.00061	-0.000225
Omega	0.00001	0.00000	0.000094*	0.00004	0.000093
alpha1	0.00000	0.00000	0.123083*	0.10441*	0.164807
beta1	0.99899*	0.99900*	0.64626*	0.79678*	0.659608
logLikelihood	97.6143	233.9206	603.2151	1249.119	2285.302
AIC	-4.6807	-5.1093	-4.9935	-5.0821	-4.6087

Note: * denotes significant at 0.05 level

Source: Researchers' Compilations from R-Output

Table3.4: Estimation of ARIMAGARCH Models across the Crude Oil Price Returns Sample Sizes

50	**100**	**250**	**500**	**1000**
ARIMA(0,0,0) GARCH(1,1)	**ARIMA(0,0,0) GARCH(1,1)**	**ARIMA(0,0,2) GARCH(1,1)**	**ARIMA(0,0,3) GARCH(1,1)**	**ARIMA(0,0,1) GARCH(1,1)**
Mu -0.0018	Mu -0.001184	Mu -0.000588	Mu 0.000342	Mu -0.0002
Omega	Omega	ma1	mar1	ma1

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

0.000001	0.000000	-0.083115	-0.00243	-0.0236
alpha1 0.0000	alpha1 0.000000	ma2 0.125316	ma2 0.06254*	Omega 0.00009
beta1 0.99899*	beta1 0.999000*	Omega 0.000100	ma3 0.003292*	alpha1 0.16221
Loglikelihood 97.6143	LogLikelihood 233.9206	alpha1 0.121281	Omega 0.00002	beta1 0.66789
AIC -4.6807	AIC -5.1093	beta1 0.626822*	alpha1 0.08695	logLikelihood 2285.52
-	-	logLikelihood 605.44	beta1 0.84859*	AIC -4.582
-	-	AIC -4.9954	logLikelihood 1259.123	-
-	-	-	AIC -5.1811	-

Note: * denotes significant at 0.05 level. **Source:** Researchers' Compilations from R-Output

Table3.5: Comparison of Models' Forecasting Performance across the Crude Oil Price Returns At Different Sample Sizes

Sample Size	Models	Fitting Performance		Forecast Accuracy	
		Loglikelihood	AIC	RMSE	MAE
50	ARIMA(0,0,0)	97.47	-	0.000405	0.000390
		-	192.24	-	-
	GARCH(1,1)	97.6143	-	0.000404	0.000389
		-	4.6807	-	-
	ARIMA(0,1,0)+GARCH(1,1)	97.6143	-	0.000403	0.000389
		-	46807	-	-
100	ARIMA(0,1,1)	233.73	-	0.002724	0.001139
		-	465.45	-	-
	GARCH(1,1)	233.9206	-	0.002725	0.001138
		-	5.1093	-	-
	ARIMA(0,1,1)+GARCH(1,1)	233.9206	-	0.002725	0.001138
		-	5.1093	-	-
250	ARIMA(1,0,0)	599.85	-	0.000329	0.0003167
		-	1193.7	-	-
	GARCH(1,1)	603.2151	-	0.000300	0.000290

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

			4.9935		
	ARIMA(1,0,0)+GARCH(1,1)	605.44	-	0.000299	0.000289
			4.9954	-	-
500	ARIMA(1,0,2)	-	-	0.0004689	0.000363
		1234.67	2467.3	-	-
	GARCH(1,1)	1249.119	-	0.0008069	0.000786
		-	5.0821	-	-
	ARIMA(1,0,2)+GARCH(1,1)	1259.123	-	0.0000867	0.000079
		-	5.1811	-	-
10000	ARIMA(4,0,4)	1607.16	-	0.0021442	0.002129
		-	3210.3	-	-
	GARCH(1,1)	2285.302	-	0.000399	0.000389
		-	4.6087	-	-
	ARIMA(4,0,4)+GARCH(1,1)	2285.52	-4.582	0.0003957	0.000387

Source: Researchers' Compilations from R-Output

DISCUSSION OF RESULTS

Table3.1 depicts the descriptive statistics and preliminary test of the crude oil price returns.

Firstly, the descriptive statistics indicate that crude oil returns are centered around a near-zero mean, with small negative average returns observed across different sample sizes. This aligns with the typical behavior of financial asset returns, which generally fluctuate around zero due to the efficient and unpredictable nature of market prices. The small drift observed may reflect gradual market trends or macroeconomic factors, but overall, the data exhibit a high degree of unpredictability.

Secondly, the volatility characteristics of crude oil returns demonstrate strong clustering behavior, a hallmark of financial time series affected by market shocks, geopolitical events, and macroeconomic fluctuations. The time series plots visually confirm periods of elevated volatility interspersed with calmer intervals, illustrating persistent volatility that can last over extended periods. These patterns are driven by exogenous shocks such as geopolitical conflicts, supply disruptions, OPEC decisions, and macroeconomic policy changes, which induce abrupt and significant price movements.

Thirdly, the diagnostic tests reinforce the presence of nonlinear dependence and heteroskedasticity in the data. The ARCH LM tests indicate significant conditional heteroskedasticity, especially in larger samples, confirming that volatility in crude oil markets is not constant but evolves over time. The BDS tests reveal strong nonlinear dependence, further suggesting that simple linear models like ARIMA are insufficient for capturing the complex dynamics of crude oil returns.

The stationarity tests consistently classify the return series as stationary (I(0)), which is an important property for modeling and forecasting. This stationarity implies that the statistical properties of returns do not change over time, making models like GARCH appropriate for capturing the persistent volatility patterns. It also supports the application of volatility models that focus on the conditional variance without requiring differencing or detrending.

The time series visualizations reinforce these statistical findings, with crude oil returns displaying irregular spikes corresponding to real-world shocks. Unlike simulated data, where shocks are generated within a controlled process, the spikes in actual market data are motivated by tangible geopolitical and macroeconomic

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

events. These irregular and often extreme deviations highlight the importance of models that can accommodate fat tails and asymmetric shocks.

Furthermore, the parameter estimates from GARCH(1,1) models consistently show high volatility persistence, with parameters close to unity. This indicates that shocks to volatility tend to persist over time, reflecting the structural nature of market uncertainty in crude oil trading. Such high persistence underscores the challenges in forecasting crude oil volatility and emphasizes the need for models that can adapt to prolonged periods of market turbulence.

Figure 3.1 presents the time series plots of the crude oil returns. The time series plot of crude oil returns, albeit with more pronounced spikes attributable to real market shocks. The plot displays clear periods of heightened volatility corresponding to global disruptions such as geopolitical conflicts, OPEC policy shifts, or macroeconomic crises.Volatility clustering is evident, reinforcing the presence of persistent volatility typical of commodity markets. The visual pattern demonstrates that the crude oil market behaves like other financial assets, with returns fluctuating around zero, interspersed with extreme deviations. These characteristics justify the use of volatility models such as GARCH and ARIMA–GARCH for modelling and forecasting such data.

Table3.2 examines the evolution of optimal ARIMA model structures as a function of sample size in the context of crude oil price return series. The model selection process, based on likelihood functions and the Akaike Information Criterion (AIC), reveals notable trends that inform understanding of the autocorrelation dynamics inherent in the data.

At all examined sample sizes, the selected models consistently feature an autoregressive (AR) parameter of zero, resulting in models of the form ARIMA(0,0,q). This consistent absence of autoregressive terms suggests that the return series exhibits minimal to no evidence of autocorrelation at prior lags, which aligns with the efficient market hypothesis where returns are often considered serially uncorrelated. The lack of AR terms indicates that past return values do not significantly predict future returns, and any autocorrelation present is better captured through moving average components rather than autoregressive structures.

At the smallest sample sizes (n=50 and n=100), the data favor a simple ARIMA(0,0,0) model, indicative of a white noise process with no significant autocorrelation detected. This suggests that limited data constrains the capacity to identify underlying temporal dependencies, resulting in a preference for minimal model complexity.

As the sample size increases to 250, the optimal model shifts to an ARIMA(0,0,2), incorporating MA(1) and MA(2) components. This transition indicates that larger datasets enhance the statistical power to detect short-term autocorrelation structures, which manifest as significant moving average terms. The inclusion of these additional parameters reflects a more nuanced understanding of the autocorrelation pattern, likely attributable to increased data-driven signal detection.

Further expansion of the sample size to 500 observations favors an ARIMA(0,0,3), suggesting the autocorrelation extends further into recent lags. This progression underscores the capacity of larger datasets to uncover more complex short-term dependencies within the return series.

Interestingly, at the largest sample size examined (n=1000), the optimal model reverts to the simpler ARIMA(0,0,1). This reversion may indicate a refinement of the autocorrelation structure, where the data support a more parsimonious model, or it may reflect overfitting concerns associated with higher-order models in smaller samples. The result underscores the importance of model parsimony and the potential for over-parameterization when datasets are overly complex relative to their true underlying structure.

Table3.3 investigates the stability and characteristics of volatility modeling in crude oil returns through the estimation of a GARCH(1,1) model across different sample sizes: 50, 100, 250, 500, and 1000 observations. The parameters estimated, mean return (Mu), constant variance (Omega), ARCH effect (alpha1), and GARCH effect (beta1), along with model fit indices such as log-likelihood and AIC, provide insights into the evolving dynamics of volatility as the dataset expands.

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

The estimated mean return remains negative across all sample sizes, decreasing in magnitude from approximately -0.00177 at n=50 to -0.000225 at n=1000. This consistent negative mean aligns with the typical characteristics of crude oil returns, which often exhibit slight negative drift over time. The diminishing magnitude suggests that as the sample size increases, the average return tends toward zero, reflecting a more stable estimate of the underlying return process.

Omega estimates are near zero for the smaller samples (n=50, 100), with a significant positive estimate at n=250 ($p<0.05$). For larger samples (n=500, 1000), Omega again hovers around 0.00009 but is not statistically significant. The initial significance at n=250 indicates a persistent baseline level of volatility that is more detectable with moderate sample sizes, while at larger sizes, the variance estimate stabilizes and becomes less distinguishable from zero, suggesting that the model attributes most of the volatility clustering to the ARCH and GARCH effects rather than a constant component.

The ARCH parameter (alpha1) is negligible and statistically insignificant in the smallest samples (n=50, 100), implying that short-term shocks do not significantly influence volatility in these datasets. As sample size increases to 250 and above, alpha1 becomes statistically significant ($p<0.05$), with estimates ranging from approximately 0.123 to 0.165. This indicates that recent shocks have a meaningful impact on current volatility, and the strength of this effect appears to grow with larger samples, reflecting the model's capacity to capture more pronounced volatility clustering in extensive datasets.

The GARCH parameter (beta1) is highly significant and close to unity across all sample sizes, ranging from approximately 0.646 to 0.999. Notably, at smaller sample sizes (n=50, 100), beta1 is extremely close to 1 (≈0.999), suggesting persistent volatility shocks that decay very slowly over time. As the sample size increases, beta1 decreases somewhat (to around 0.646 at n=250 and 0.660 at n=1000), indicating a somewhat faster mean reversion of volatility in larger datasets. The high significance of beta1 across all samples emphasizes the presence of strong volatility persistence in crude oil returns.

Log-likelihood values increase substantially with sample size, from 97.6143 at n=50 to 2285.302 at n=1000, reflecting improved model estimation with more data. Correspondingly, AIC scores improve (become more negative), indicating better model fit in larger samples. The AIC decreases from about -4.6807 at n=50 to approximately -5.1093 at n=100, then stabilizes around -4.6 to -5.08, suggesting that the model captures volatility dynamics more effectively as the sample size grows.

Both the ARIMA and GARCH(1,1) models demonstrate valuable insights into the dynamics of crude oil returns, yet each exhibits distinct strengths in capturing different aspects of the data. The ARIMA models consistently favor simpler structures with no autoregressive terms across all sample sizes, indicating that crude oil returns exhibit minimal autocorrelation and are predominantly driven by short-term shocks rather than persistent autoregressive effects. In contrast, the GARCH(1,1) models reveal strong volatility clustering and persistence, with high and significant GARCH parameters close to unity across all sample sizes, emphasizing the importance of modeling volatility dynamics explicitly. While ARIMA models effectively capture the mean process and short-term dependencies, the GARCH models excel in modeling the heteroskedastic nature of returns, which is critical for risk management and derivative pricing.

Table3.4 presents an extensive comparison of ARIMAGARCH models applied to crude oil price returns across varying sample sizes (50, 100, 250, 500, and 1000 observations), incorporating different ARIMA(0,0,0), (0,0,1), (0,0,2), and (0,0,3) specifications with GARCH(1,1) volatility modeling. The results highlight that models with minimal or no autoregressive components in the mean equation—particularly ARIMA(0,0,0)—yield high likelihoods and strong GARCH parameter estimates, with beta1 consistently close to unity (≈0.999), indicating persistent volatility. For smaller samples (n=50, 100), the models exhibit very high log-likelihoods and superior AIC scores, emphasizing their effectiveness in capturing the volatility clustering characteristic of crude oil returns. The inclusion of MA components (e.g., ma1, ma2, ma3) in the mean equation introduces additional dynamics, but their estimated coefficients are generally small or insignificant, suggesting limited impact on the mean process. Compared to the standalone ARIMA and GARCH models, the combined ARIMAGARCH models demonstrate enhanced flexibility by simultaneously modeling mean and volatility dynamics, with the best fit observed in models with no autoregressive terms and a simple GARCH(1,1) structure. Overall, the

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

ARIMAGARCH framework, particularly with minimal mean model complexity, outperforms the previous models in capturing the persistent volatility and subtle mean dynamics of crude oil returns, providing a more comprehensive and effective approach for modeling such financial time series. Consequently, for comprehensive modeling of crude oil, a combined approach that incorporates both ARIMA for mean dynamics and GARCH for volatility clustering provide the most robust framework, with the GARCH component clearly demonstrating superior performance in capturing the persistent volatility characteristic of crude oil markets.

Table3.5 provides a comprehensive evaluation of the forecasting performance of various time series models applied to crude oil price returns across differing sample sizes. The comparison encompasses both in-sample fitting metrics; log-likelihood and Akaike Information Criterion (AIC), and out-of-sample forecast accuracy measures; Root Mean Square Error (RMSE) and Mean Absolute Error (MAE). The results elucidate how model efficacy evolves with increasing data availability, offering valuable insights into model selection strategies for financial time series forecasting.

Across all sample sizes, the models demonstrate consistent in-sample fitting capabilities, with higher log-likelihood values indicating better fit. Notably, the ARIMA(1,0,2)+GARCH(1,1) model exhibits superior in-sample likelihood at larger sample sizes (e.g., 250 and 500 observations), suggesting its enhanced capacity to capture underlying data dynamics as the dataset grows. The AIC values, align with these findings, favoring models that balance fit and parsimony.

At the minimal sample size of 50, models achieve remarkably low RMSE (~0.0004) and MAE (~0.00039), indicating high forecast precision in short-term predictions with limited data. The ARIMA(0,1,0)+GARCH(1,1) model slightly outperforms others, highlighting its robustness in sparse data environments.

As the sample size increases, forecast errors decrease notably, with the ARIMA(1,0,2)+GARCH(1,1) model attaining the lowest RMSE (~0.0000867) and MAE (~0.000079) at 500 observations. This trend underscores the advantage of larger datasets in enhancing model accuracy, especially when volatility modeling via GARCH components is incorporated.

At the largest sample size examined, the ARIMA(4,0,4)+GARCH(1,1) model demonstrates superior forecast accuracy, evidenced by the lowest MAE (~0.000387). The results suggest that, with ample data, models capable of capturing complex dynamics, such as higher-order ARIMA components combined with GARCH, offer the most precise forecasts. Additionally, the substantial increase in in-sample likelihoods corroborates the improved fit and predictive power.

The findings highlight the importance of model complexity in relation to sample size. While simpler models like ARIMA(0,1,0)+GARCH(1,1) perform adequately with limited data, more sophisticated models with higher orders (e.g., ARIMA(4,0,4)+GARCH(1,1)) outperform in large-sample contexts, benefiting from their capacity to model intricate data structures. Incorporating volatility dynamics via GARCH consistently enhances forecast accuracy across different sample sizes, reaffirming the significance of volatility modeling in financial return series.

Overall, the analysis demonstrates that model performance in forecasting crude oil returns improves with increasing sample size, with complex models gaining prominence in larger datasets. The integration of GARCH components consistently contributes to reducing forecast errors, emphasizing the importance of capturing volatility in financial time series modeling. These insights provide a valuable framework for practitioners and researchers in selecting appropriate models based on data availability and the underlying data-generating process.

SUMMARY OF FINDINGS

This analysis presents a comprehensive examination of the statistical properties and modeling strategies for crude oil return series. Descriptive statistics reveal that returns are centered near zero with small negative averages, consistent with typical financial asset behavior characterized by high unpredictability. Visual inspections of the

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

time series indicate pronounced volatility clustering, with periods of elevated market turbulence often linked to geopolitical events, supply disruptions, and macroeconomic shocks. Diagnostic tests confirm the presence of nonlinear dependence and heteroskedasticity, notably through significant ARCH effects and nonlinear dynamics identified by the BDS test. Stationarity assessments verify that the return series are stationary (I(0)), supporting the suitability of volatility models such as GARCH.

Model selection procedures across varying sample sizes show a consistent preference for models with minimal autoregressive components, specifically ARIMA(0,0,q), indicating limited autocorrelation in returns' a finding aligned with the efficient market hypothesis. As sample size increases, models incorporating short-term moving average terms (e.g., ARIMA(0,0,2) and ARIMA(0,0,3)) become optimal, reflecting the improved detection of subtle autocorrelation structures with more data. Volatility modeling via GARCH(1,1) consistently captures persistent volatility clusters, with high GARCH parameters close to unity across all sample sizes, underscoring the structural and enduring nature of market uncertainty.

Furthermore, combined ARIMA-GARCH models demonstrate superior performance in both in-sample fit and out-of-sample forecasting accuracy, especially when model complexity is balanced to prevent overfitting. The results indicate that while simple models effectively capture the mean process, the integration of GARCH components is crucial for accurately modeling volatility dynamics and risk estimation. Notably, larger datasets enable the use of more complex models; such as ARIMA(4,0,4)+GARCH(1,1); which outperform simpler specifications in forecasting accuracy, as evidenced by lower RMSE and MAE values.

Overall, the findings suggest that for modeling crude oil returns, a parsimonious mean model combined with a GARCH(1,1) volatility component offers an effective framework. As data availability increases, employing higher-order ARIMA models alongside GARCH structures enhances predictive performance, capturing both the short-term dependencies and persistent volatility characteristic of crude oil markets. These insights provide valuable guidance for practitioners and researchers aiming to develop robust forecasting models in the context of volatile commodity markets.

CONCLUSION

Based on the comprehensive analysis conducted, the study provides valuable insights into the modeling and forecasting of crude oil returns, aligning with the stated objectives. The descriptive statistics and preliminary tests confirmed that crude oil returns are centered around zero, exhibit strong volatility clustering, and demonstrate nonlinear dependence and heteroskedasticity—characteristics typical of financial time series affected by exogenous shocks such as geopolitical events and macroeconomic fluctuations. These findings underscore the importance of employing models capable of capturing complex volatility dynamics.

In terms of model fitting, the results reveal that simple ARIMA models, particularly those with no autoregressive components, are sufficient to capture the mean dynamics of crude oil returns, consistent with the efficient market hypothesis. Conversely, volatility modeling through GARCH(1,1) consistently captures persistent volatility clustering and high volatility persistence, highlighting the necessity of incorporating heteroskedasticity in the modeling framework.

The comparison of forecasting performances indicates that model accuracy improves with increasing sample size. While simpler models perform adequately with limited data, more sophisticated models, particularly the ARIMA-GARCH combinations with higher-order ARIMA components offer superior forecasting accuracy in larger datasets. Notably, the ARIMA(4,0,4)+GARCH(1,1) model consistently outperforms others in terms of forecast precision, as evidenced by the lowest out-of-sample RMSE and MAE at larger sample sizes. The integration of volatility modeling via GARCH enhances predictive performance across all models, emphasizing its critical role in capturing the persistent volatility characteristic of crude oil markets.

In conclusion, the findings demonstrate that the effectiveness of time series models in forecasting crude oil returns depends significantly on data size and complexity. The combination of ARIMA for mean dynamics and GARCH for volatility provides a robust and flexible framework capable of accurately modeling and forecasting the complex behavior of crude oil prices. This study affirms that incorporating volatility models such as GARCH

ISSN 2278-2540 | DOI: 10.51583/IJLTEMAS | Volume XV, Issue IV, April 2026

into traditional ARIMA structures enhances forecast accuracy, thereby fulfilling the research objectives of evaluating and comparing the forecasting capabilities of these models in real-world data applications.

Implication and Future Research Directions

The findings have important policy and practical implications. The persistent volatility clustering and high volatility persistence highlight the need for market participants to adopt advanced volatility forecasting models, such as ARIMA-GARCH, for risk management and hedging. The minimal autocorrelation supports market efficiency, suggesting that policies promoting transparency can further stabilize markets. Recognizing the impact of geopolitical and macroeconomic shocks on volatility can inform policymakers in designing strategic interventions to mitigate market turbulence. For investors, incorporating volatility forecasts can enhance asset allocation and risk assessment.

Future research should explore models incorporating exogenous variables (e.g., macroeconomic indicators, geopolitical risk), asymmetric GARCH models to capture leverage effects, and regime-switching frameworks to account for structural breaks. Additionally, analyzing high-frequency data and comparing alternative volatility models, including machine learning approaches, could improve real-time forecasting accuracy and robustness. These avenues will deepen understanding of crude oil market dynamics and enhance predictive capabilities.

REFERENCE

1. Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, 31(3), 307–327. doi:10.1016/0304-4076(86)90063-1
2. Box, G.E.P. and Jenkins, G.M. (1976) Time Series Analysis: Forecasting and Control. Revised Edition, Holden Day, San Francisco.
3. Bunnag, T. (2024). The importance of gold's effect on investment and predicting the world gold price using the ARIMA and ARIMA-GARCH model. Ekonomikalia Journal of Economics, 2(1), 38–52. doi:10.60084/eje.v2i1.155
4. Di-Giorgi, G., Salas, R., Avaria, R., Ubal, C., Rosas, H., & Torres, R. (2025). Volatility forecasting using deep recurrent neural networks as GARCH models. Computational Statistics, 40(6), 3229–3255. doi:10.1007/s00180-023-01349-1
5. Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. Econometrica, 50(4), 987–1007. doi:10.2307/1912773
6. Hasanov, A. S., Brooks, R., Abrorov, S., & Burkhanov, A. U. (2024). Structural breaks and GARCH models of exchange rate volatility: Re-examination and extension. Journal of Applied Econometrics, 39(7), 1403–1407. doi:10.1002/jae.3091
7. Kaur, J., Parmar, K. S., & Singh, S. (2023). Autoregressive models in environmental forecasting time series: A theoretical and application review. Environmental Science and Pollution Research, 30, 19617–19641. doi:10.1007/s11356-023-25148-9
8. Mbonigaba, C., Vasuki, M., Kumar, A. D., & Asamoah, P. J. (2025). Applications of GARCH models for volatility forecasting in high-frequency trading environments. International Journal of Applied and Advanced Scientific Research, 10(1), 12–21. doi:10.5281/zenodo.14904200
9. Phung Duy, Q., Nguyen Thi, O., Le Thi, P. H., Pham Hoang, H. D., Luong, K. L., & Nguyen Thi, K. N. (2024). Estimating and forecasting Bitcoin daily prices using ARIMA-GARCH models. Business Analyst Journal, 45(1), 11–23. doi:10.1108/BAJ-05-2024-0027
10. Rapach, D. E., & Strauss, J. K. (2008). Structural breaks and GARCH models of exchange rate volatility. Journal of Applied Econometrics, 23(1), 65–90. doi:10.1002/jae.976
11. Sirisha, U. M., Belavagi, M. C., & Attigeri, G. (2024). Profit prediction using ARIMA, SARIMA and LSTM models in time series forecasting: A comparison. IEEE Access, 10, 124715–124727. doi:10.1109/ACCESS.2022.3224938