Science World Journal Vol. 21(No 2) 2026 https://dx.doi.org/10.4314/swj.v21i2.44
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University

COMPARATIVE EVALUATION OF GENERALIZED AUTOREGRESSIVE CONDITIONAL HETEROSKEDASTICITY-TYPE AND MACHINE-LEARNING MODELS FOR EXCHANGE-RATE VOLATILITY FORECASTING IN NIGERIA

*[1]Oyelami A.S., [2]Arowolo O.T., and [3]Akinbo R.Y.

[1]Department of Statistics, University of Lagos, Akoka, Nigeria
[2]Department of Mathematical Science, Lagos State University of Science and Technology
[3]Department of Mathematics and Statistics, Federal Polytechnic, Ilaro, Ogun State, Nigeria

*Corresponding Author Email Address: abbeytech4uall@gmail.com

ABSTRACT
This study investigates exchange-rate volatility forecasting in Nigeria by comparing conventional econometric models with machine-learning approaches using a common volatility target. Monthly USD/NGN exchange-rate data obtained from the Central Bank of Nigeria covering January 2004 to December 2025 were analysed. Monthly log returns were computed, while squared returns served as a proxy for realized volatility. The empirical framework combines an ARMA mean specification with GARCH(1,1) and EGARCH(1,1) variance models, alongside Support Vector Regression (SVR) and Artificial Neural Networks (ANN). Forecast performance was evaluated using rolling time-series cross-validation and standard error measures, including Mean Absolute Error (MAE), Mean Squared Error (MSE), Root Mean Squared Error (RMSE), and the Diebold–Mariano test. Results indicate that returns are stationary, leptokurtic, and exhibit volatility clustering with evidence of structural instability. ARMA–GARCH(1,1) outperforms EGARCH(1,1) in stability, while SVR achieves the best forecasting accuracy followed by ANN. However, the Diebold–Mariano test shows that differences in predictive accuracy are not statistically significant at the 5% level. The findings suggest that machine-learning methods, particularly SVR, complement traditional volatility models in exchange-rate risk management in Nigeria

Keywords: Exchange-rate volatility, GARCH, EGARCH, Support vector regression; Artificial neural network.

INTRODUCTION
Exchange-rate volatility remains a major source of macro-financial uncertainty in emerging economies because fluctuations in exchange rates influence inflation, external debt servicing, trade competitiveness, reserve adequacy, and financial-market expectations. In Nigeria, where foreign-exchange conditions are shaped by oil revenues, import dependence, administrative interventions, and recurrent policy reforms, exchange-rate movements are closely linked to monetary stability and real-sector performance. The official exchange-rate series published by the Central Bank of Nigeria (2026) reveals sustained depreciation interspersed with major adjustments and regime changes, underscoring the importance of modelling exchange-rate volatility alongside exchange-rate levels. For policymakers, financial institutions, and firms with foreign-currency exposures, volatility provides a direct measure of exchange-rate risk and uncertainty.

The econometric literature has traditionally relied on Autoregressive Conditional Heteroskedasticity, ARCH and Generalized Autoregressive Conditional Heteroskedasticity, GARCH, ARCH/GARCH-type models to capture conditional heteroskedasticity in financial time series. The ARCH model of (Engle, 1982) and the generalized GARCH framework of (Bollerslev, 1986) remain foundational because they effectively model volatility clustering, persistence, and mean reversion. Extensions such as the EGARCH model developed by (Nelson, 1991) and further discussed by (Pilbeam and Langeland, 2014) asymmetric volatility responses by allowing shocks of different signs and magnitudes to exert unequal effects on conditional variance.
Despite their strengths, GARCH-type models may be less effective when volatility dynamics are influenced by nonlinear relationships, structural breaks, and institutional changes. This limitation is particularly relevant in Nigeria, where exchange-rate behaviour has been shaped by oil-price shocks, exchange-control policies, market reforms, and episodes of currency devaluation. Consequently, machine-learning techniques have increasingly emerged as alternative forecasting tools capable of modelling complex nonlinear patterns without imposing restrictive functional assumptions.
Among these techniques, Support Vector Regression (SVR) and Artificial Neural Networks (ANNs) have attracted considerable attention. SVR, derived from the support-vector framework of (Cortes & Vapnik, 1995), employs kernel-based learning and regularization mechanisms that enhance predictive performance. Empirical evidence suggests that SVR can produce competitive volatility forecasts under nonlinear conditions (Chen et al., 2008; Peng & Albuquerque, 2018; Sun & Yu, 2020; Muchtadi-Alamsyah et al., 2024). Similarly, ANNs benefit from the universal approximation property established by (Hornik et al., 1989) and have demonstrated considerable success in modelling nonlinear financial relationships (Panda & Narasimhan, 2007; Adewole et al., 2011; Rammusi et al., 2025).
Although the international literature on volatility forecasting has expanded substantially, evidence from Nigeria remains relatively limited. Existing studies largely focus on conventional GARCH-family models and generally report persistent exchange-rate volatility with occasional evidence of asymmetry (Dahiru & Asemota, 2013). Comparative assessments of econometric and machine-learning approaches remain scarce, particularly within frameworks that employ a common volatility target and a

transparent validation procedure.

A major challenge in volatility forecasting is ensuring methodological consistency across competing models. Forecast comparisons based on different volatility definitions can produce misleading conclusions regarding model superiority. To address this issue, the present study adopts squared returns as a common volatility proxy for all competing models. The empirical framework is further strengthened through ARMA mean modelling, rolling time-series cross-validation, structural-break diagnostics, and the Diebold–Mariano test of predictive accuracy. By adopting a unified and reproducible framework, the study provides a more rigorous assessment of exchange-rate volatility forecast in Nigeria.

Accordingly, this study compares the performance of GARCH and EGARCH models with SVR and ANN models in modelling and forecasting exchange-rate volatility in Nigeria. The findings are expected to contribute to exchange-rate risk management, forecasting practice, and evidence-based policy formulation.

Volatility modelling has been dominated by the GARCH family of models because of their ability to represent conditional variance as a function of past shocks and past volatility. The seminal works of (Engle, 1982) and (Bollerslev, 1986) established the theoretical foundations of ARCH and GARCH modelling, while empirical applications have demonstrated their effectiveness in capturing volatility clustering and persistence. In Nigeria, studies by (Dahiru and Asemota, 2013) reported substantial persistence in exchange-rate volatility, reflecting recurring macroeconomic imbalances, foreign-exchange shortages, and policy interventions.

The EGARCH model proposed by (Nelson, 1991) extends the conventional GARCH framework by modelling logarithmic variance and permitting asymmetric responses to shocks. (Pilbeam & Langeland, 2014) further emphasized the usefulness of asymmetric volatility models in financial markets where positive and negative innovations may carry different informational content. Although empirical evidence regarding exchange-rate asymmetry remains mixed, EGARCH continues to provide a valuable framework for capturing nonlinear volatility dynamics.

Support Vector Regression offers a flexible alternative to parametric volatility models. Building on the work of (Cortes & Vapnik, 1995), SVR employs kernel functions and structural risk minimization to model complex nonlinear relationships. Studies by (Chen et al., 2008; Peng and Albuquerque 2018; Sun and Yu 2020; and Muchtadi-Alamsyah et al., 2024) show that SVR can achieve strong forecasting performance, particularly when volatility dynamics deviate from conventional linear assumptions.

Artificial Neural Networks provide another widely used machine-learning approach. The theoretical foundation of ANNs derives from the universal approximation theorem of (Hornik et al., 1989), which establish their ability to approximate complex nonlinear functions. Empirical studies have shown that ANNs can effectively capture nonlinear dependencies and structural changes in exchange-rate dynamics (Panda & Narasimhan, 2007; Adewole et al., 2011). However, forecasting performance remains sensitive to network architecture, sample size, and tuning procedures. Integrating neural-network methods with econometric structures can further improve predictive accuracy (Rammusi et al., 2025).

Another important issue concerns the selection of volatility proxies. Because latent volatility is unobservable, researchers typically rely on squared returns, absolute returns, realized volatility, or range-based measures. The choice of proxy can significantly influence forecast evaluation, making a common target essential for valid model comparison. This consideration motivates the use of squared returns as the volatility measure adopted in the present study.

Model evaluation constitutes another important theme in the literature. (Nkemnole & Oyelami, 2025) compared volatility forecasts using Mean Absolute Error, (MAE), Mean Square Error, (MSE), and Root Mean Squared Error, (RMSE), demonstrating the usefulness of conventional forecast-loss measures. Nevertheless, descriptive loss functions alone cannot determine whether forecast differences are statistically significant. Consequently, the Diebold–Mariano test developed by Diebold and Mariano (1995) has become a standard procedure for assessing comparative predictive accuracy.

Structural instability also plays a crucial role in exchange-rate modelling. The Zivot–Andrews test (Zivot & Andrews, 1992) accommodates endogenous structural breaks in unit-root testing and is particularly relevant in economies characterized by frequent policy changes and market interventions. In Nigeria, where the foreign-exchange market has undergone repeated institutional and policy transformations, accounting for structural breaks enhances both model interpretation and forecast reliability.

Overall, the literature suggests that no single model consistently dominates across all volatility environments. GARCH-type models remain attractive because of their interpretability and theoretical grounding, whereas SVR and ANN models offer greater flexibility in capturing nonlinear dynamics. These considerations provide the motivation for a rigorous and methodologically consistent comparison of econometric and machine-learning approaches to exchange-rate volatility forecasting in Nigeria.

MATERIALS AND METHODS

The data used in this study were obtained from the monthly average exchange-rate series published on the website of the Central Bank of Nigeria (www.cbn.gov.ng). The series consists of monthly USD/NGN observations from January 2004 to December 2025, yielding 264 exchange-rate levels. After log-differencing, 263 monthly returns were obtained. To create a common modelling sample for all methods, three lags of returns and squared returns were used as predictors in the machine-learning models; this yielded 260 aligned observations, of which 208 were used for estimation and validation and 52 were reserved for out-of-sample forecast evaluation. The hold-out sample therefore covered September 2021 to December 2025.

Data Preprocessing

Let P_t denote the monthly USD/NGN exchange rate. Monthly continuously compounded returns were computed as: (Brooks, 2019)

$$r_t = 100[\ln(P_t) - \ln(P_{t-1})]. \qquad (1)$$

To ensure methodological comparability across all models, the volatility target was defined as squared return,

$$v_t = r_t^2. \qquad (2)$$

The two methods considered, GARCH-type and machine-learning models were evaluated against the same realized-volatility proxy. Before variance modelling, the conditional mean of returns was specified with an ARMA model selected on the training sample by the Akaike information criterion over a small (p, q) grid. The selected mean equation was $\mathrm{ARMA}(0,2)$, and the residuals from that mean process were used in the GARCH-type models.

GARCH-Type Models

The ARMA-GARCH(1,1) model is written as: (Bollerslev, 1986)

https://dx.doi.org/10.4314/swj.v21i2.44
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University

$$r_t = \mu_t + \varepsilon_t, \qquad \varepsilon_t = z_t\sqrt{h_t}, \tag{3}$$

where μ_t is the ARMA conditional mean, z_t is an i.i.d. standardized innovation, and the conditional variance h_t follows

$$h_t = \omega + \alpha\varepsilon_{t-1}^2 + \beta h_{t-1}, \omega > 0, \alpha \geq 0, \beta \geq 0. \tag{4}$$

Where ω is the constant term in the variance equation. It represents the long-run average volatility level, α and β are the ARCH and GARCH parameters

The ARMA-EGARCH (1,1) specification replaces the linear variance recursion with a log-variance equation:

$$\log\ h_t = \omega + \alpha(|z_{t-1}| - E|z_{t-1}|) + \gamma z_{t-1} + \beta \log\ h_{t-1}. \tag{5}$$

To reduce the influence of the heavy tails observed in the data, the GARCH-type models were estimated with a Student-*t* conditional likelihood. This choice responds to the referee's concern about purely Gaussian errors under non-normal returns.

Machine-Learning Models

Two machine-learning models were estimated: support vector regression and a feedforward artificial neural network. The predictor set consisted of three lags of returns and three lags of squared returns. For SVR, a radial basis function, RBF kernel was adopted and the hyperparameters (*C*,*γ*,*ϵ*) were tuned by grid search under rolling time-series cross-validation. For the ANN, a multilayer perceptron with early stopping was used, and the hidden-layer size, learning rate, and L_2 penalty were tuned by grid search. Standardization of predictors was carried out inside each training fold through a pipeline structure. The RBF kernel is defined as:

$$K(x_i, x_j) = \exp\left(-\gamma\|x_i - x_j\|^2\right) \tag{6}$$

where:

x_i and x_j are two observations,

$\|x_i - x_j\|^2$ is the squared Euclidean distance between them,

γ controls how far the influence of a training observation extends

Forecast Evaluation

Forecast accuracy was assessed with the usual loss functions

$$\text{MAE} = \frac{1}{n}\sum_{t=1}^{n}|v_t - \hat{v}_t|, \tag{7}$$

$$\text{MSE} = \frac{1}{n}\sum_{t=1}^{n}(v_t - \hat{v}_t)^2, \tag{8}$$

RM

$$\text{SE} = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(v_t - \hat{v}_t)^2}. \tag{9}$$

For the GARCH-type models, in-sample information criteria were also reported:

$$\text{AIC} = 2k - 2\ln\ L, \qquad \text{BIC} = k\ln\ (n) - 2\ln\ L, \tag{10}$$

where k is the number of estimated parameters and L is the maximized likelihood.

To assess whether forecast differences were statistically significant, the Diebold–Mariano (DM) test of equal predictive accuracy was applied to the out-of-sample squared-error losses (Diebold & Mariano, 1995). The null hypothesis of the test is that two competing models possess equal predictive accuracy. Consistent with the volatility forecasting literature, the ARMA-GARCH model was adopted as the benchmark model because GARCH-type specifications are widely recognized as standard econometric benchmarks for modelling and forecasting financial market volatility (Bollerslev, 1986; Pilbeam & Langeland, 2014; Brooks, 2019).

Diagnostic and Robustness Tests

Stationarity was examined with the Augmented Dickey–Fuller test (Dickey and Fuller, 1979). Normality was examined with the Jarque–Bera test (Jarque and Bera, 1980). Serial dependence was assessed with the Ljung–Box statistic (Ljung and Box, 1978), the Zivot–Andrews test with an endogenous break was also applied (Zivot and Andrews, 1992). Finally, the ACF and PACF of returns and squared returns were inspected to verify serial dependence and volatility clustering.

RESULTS AND DISCUSSION

Descriptive Statistics and Preliminary Diagnostics

Table 1: Descriptive statistics of the exchange-rate data

Series	Mean	Std. Dev.	Minimum	Maximum	Skewness	Kurtosis
USD/NGN level	367.274	391.213	107.347	1669.470	2.284	7.135
Log return (%)	0.883	6.142	-34.164	56.573	4.300	44.745
Squared return	38.360	253.916	0.000	3200.509	9.736	108.125

Table 1 presents the descriptive statistics for the exchange-rate level, monthly return, and squared-return series. The exchange-rate level displayed a large spread and strong positive skewness, reflecting the prolonged depreciation of the naira and the concentration of observations at lower historical exchange-rate values. The return series was far more leptokurtic than the level series, indicating occasional extreme monthly changes. The squared-return series was even more right-skewed, which is expected when volatility is concentrated in a few turbulent episodes. These distributional features support the use of volatility

Science World Journal Vol. 21(No 2) 2026
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University
https://dx.doi.org/10.4314/swj.v21i2.44

models and justify caution in relying on Gaussian assumptions.

Figure 1: Time plot of the monthly USD/NGN exchange rate and log returns

The upper panel of Figure 1 shows a pronounced upward trend in the USD/NGN series, particularly in the later years, confirming the long-run depreciation of the naira. The lower panel indicates that returns fluctuate around a comparatively low mean but exhibit distinct bursts of turbulence. These bursts are consistent with volatility clustering and visually suggest regime changes rather than a single stable volatility process.

Table 2: Preliminary diagnostic tests

Test	Series	Statistic	p-value
ADF	USD/NGN level	0.6312	0.9884
ADF	Log return (%)	-4.0697	0.0011
Jarque–Bera	USD/NGN level	405.9284	<0.001
Jarque–Bera	Log return (%)	19161.9279	<0.001
Ljung–Box (lag 10)	USD/NGN level	2055.6177	<0.001
Ljung–Box (lag 10)	Log return (%)	47.6625	<0.001
Ljung–Box (lag 10)	Squared return	27.7469	0.0020
Zivot–Andrews	USD/NGN level	-4.3234	0.1820
Zivot–Andrews	Log return (%)	-4.5232	0.1127

Table 2 presents the results of the stationarity, normality, serial-correlation, and structural-break diagnostics. The ADF test fails to reject a unit root in the level series but rejects the null for returns, indicating that volatility modelling should be based on the stationary return process. The Jarque–Bera statistics strongly reject normality for both the level and return series, confirming fat tails and asymmetry. The Ljung–Box statistics reject the null hypothesis of no serial dependence for both returns and squared returns, supporting the presence of short-run dependence and volatility clustering. Although the Zivot–Andrews test identifies potential break dates, the null hypothesis of a unit root with structural break cannot be rejected at conventional significance levels, suggesting that the evidence for structural change should be interpreted cautiously.

Figure 2 shows the ACF and PACF of monthly returns.

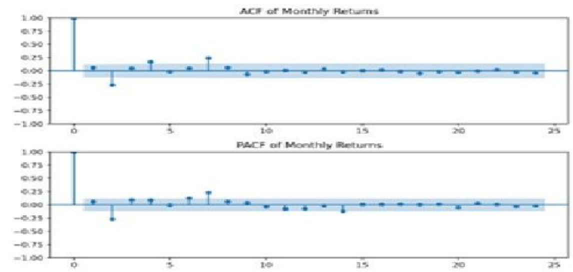

Figure 2: ACF and PACF of returns

The correlograms show modest serial dependence in returns and more persistent dependence in squared returns. This pattern is consistent with a series whose mean dynamics are relatively weak but whose second moment is time-varying.

Parameter Estimates
Table 3 summarizes the selected mean equation and the estimated parameters of the competing models.

Table 3: Parameter estimates and selected hyperparameters

Model	Parameter estimates / selected settings	Notes
ARMA mean equation	ARMA(0,2):constant = 0.873, MA(1) = 0.091, MA(2) = −0.159	Selected by AIC on training sample
ARMA-GARCH(1,1)	ω = 0.436, α = 0.625, β = 0.375, v = 3.998	Student-*t* likelihood
ARMA-EGARCH(1,1)	ω = 0.351, α = 0.781, γ = 0.287, β = 0.831, v = 3.764	Student-*t* likelihood
SVR	RBF kernel, C = 100, γ = 0.1, ϵ = 0.5	Selected by rolling CV
ANN	Hidden layers (10,5), learning rate = 0.001, L_2 penalty = 0.0001	Early stopping and rolling CV

The ARMA(0,2) mean specification indicates the presence of modest moving-average dynamics in the return series. Within the ARMA–GARCH framework, the estimated persistence parameter ($\alpha + \beta \approx 1.000$) suggests highly persistent volatility, implying that shocks to exchange-rate uncertainty decay slowly over time. The ARMA–EGARCH results also indicate strong persistence alongside a positive asymmetry coefficient. However, the forecasting results show that the EGARCH model becomes unstable in the post-2021 hold-out sample. For the machine-learning models, the SVR configuration exhibits strong nonlinearity while maintaining robustness through effective regularization via its margin-based formulation.

Model Comparison

Table 4 presents the in-sample fit statistics on the common volatility target $v_t = r_t^2$.

Table 4: In-sample performance on squared returns

Model	MAE	MSE	RMSE	AIC	BIC
ARMA-GARCH(1,1)	28.084	19082.618	138.140	837.272	850.622
ARMA-EGARCH(1,1)	228.127	7082914.246	2661.375	836.138	852.825
SVR	18.175	17941.779	133.947	——	
ANN	21.809	19801.693	140.718	——	

On the common volatility proxy, SVR achieved the lowest in-sample MAE, MSE, and RMSE. ANN ranked second on MAE but slightly underperformed ARMA-GARCH on MSE and RMSE. Among the GARCH-type models, ARMA-GARCH clearly dominated ARMA-EGARCH in terms of fit to the squared-return proxy, although the information criteria were close. This indicates that the EGARCH log-variance structure was not translating into superior approximation of the realized volatility proxy in this sample.

Table 5 presents the out-of-sample forecast results.

Table 5: Out-of-sample forecast performance on squared returns

Model	MAE	MSE	RMSE
ARMA-GARCH(1,1)	194.988	337388.129	580.851
ARMA-EGARCH(1,1)	9.979×10^{19}	5.169×10^{41}	7.190×10^{20}
SVR	114.464	248407.300	498.405
ANN	115.162	248800.971	498.800

The out-of-sample comparison yielded three clear observations. First, both machine-learning models achieved materially lower forecast loss than ARMA-GARCH on the common squared-return target. Second, SVR was marginally better than ANN across all three loss measures, making it the strongest forecast model in the revised design. Third, ARMA-EGARCH became explosively unstable in the hold-out period. Because this instability persisted after the methodological revisions, it is reported transparently rather than suppressed. The result suggests that the EGARCH specification was highly sensitive to the large shocks and regime shifts that occurred in the later part of the sample. In that sense, the anomaly flagged by the referee was not merely a coding mistake; it reflected a genuine fragility of the fitted EGARCH recursion for this dataset.

Table 6 presents Diebold–Mariano tests using ARMA-GARCH as the benchmark.

Table 6: Diebold–Mariano tests of equal predictive accuracy (benchmark: ARMA-GARCH)

Comparison	DM statistic	p-value	Mean loss differential
ARMA-GARCH vs ARMA-EGARCH	-1.000	0.317	-5.169×10^{41}
ARMA-GARCH vs SVR	1.179	0.239	88980.829
ARMA-GARCH vs ANN	1.173	0.241	88587.158

The positive mean loss differentials for the SVR and ANN comparisons indicate that both machine learning models had lower average squared-error loss than ARMA-GARCH in the hold-out sample. However, the p-values exceeded 0.05, so equal predictive accuracy could not be rejected at conventional significance levels. The revised interpretation is therefore more cautious than in the earlier manuscript: SVR and ANN produced lower forecast loss, but the evidence is not strong enough to claim statistically significant dominance over ARMA-GARCH in this sample.

Discussion of Findings

The empirical results provide evidence on the relative performance of econometric and machine-learning models in capturing exchange-rate volatility in Nigeria. Both the ARMA–GARCH and ARMA–EGARCH specifications are able to model time-varying conditional variance; however, the ARMA–GARCH model exhibits greater stability and consistency across the sample. The estimated near-unit persistence indicates that shocks to exchange-rate volatility are highly persistent, with slow decay over time. Although the EGARCH model is designed to capture asymmetric responses, its performance weakens in the post-2021 period, suggesting that structural shifts in the foreign-exchange market reduce the reliability of purely asymmetry-based specifications.

The machine-learning models demonstrate strong capacity to capture nonlinear relationships between lagged variables and future volatility. The Support Vector Regression (SVR) model outperforms the Artificial Neural Network (ANN), consistent with evidence that SVR performs well in moderately sized samples due to its structural risk minimization and built-in regularization. The ANN model remains flexible but requires architectural constraints and early stopping to prevent overfitting, reflecting limitations associated with the relatively small monthly dataset. Despite this, its forecasting performance remains competitive, although consistently below that of SVR.

The common target evaluation framework shows that SVR delivers the best in-sample and out-of-sample forecasting performance among the stable models, while ANN ranks second and outperforms the ARMA–GARCH benchmark in terms of forecast loss. However, the Diebold–Mariano test indicates that these differences are not statistically significant at the 5% level. This suggests that although machine-learning models generate lower forecast errors, their superiority over traditional econometric approaches cannot be established conclusively. From a practical perspective, the findings highlight the importance of accounting for persistent volatility and structural changes, and the potential value of incorporating nonlinear machine-learning models as complementary tools for exchange-rate risk management.

Conclusion

This study establishes a unified empirical framework by aligning all competing models on a common volatility proxy, specifying an ARMA mean equation prior to variance estimation, incorporating systematic hyperparameter tuning for the machine-learning models using rolling cross-validation, and complementing the analysis with structural-break diagnostics and Diebold–Mariano tests of predictive accuracy. The empirical evidence indicates that the exchange-rate level series is non-stationary, while the return series is stationary. In addition, the return distribution exhibits pronounced non-normality and heteroskedasticity, consistent with well-known stylized facts of financial time series.

A comparison between the two conventional econometric models shows that the ARMA–GARCH(1,1) specification provides a more stable and reliable representation of naira exchange-rate volatility than the ARMA–EGARCH(1,1) model. Although the EGARCH model is theoretically designed to capture asymmetric volatility responses, its performance is less consistent under structural changes in the exchange-rate environment. Among the machine-learning approaches, the Support Vector Regression (SVR) model delivers the best forecasting performance, while the Artificial Neural Network (ANN) ranks closely behind it in out-of-sample evaluation.

However, the Diebold–Mariano test results show that the observed differences in predictive accuracy between machine-learning models and the ARMA–GARCH benchmark are not statistically significant at the 5% level. This implies that, although machine-learning methods particularly SVR may yield lower forecast errors, their superiority over traditional econometric approaches cannot be established conclusively. Consequently, a complementary modelling strategy is recommended for exchange-rate volatility analysis in Nigeria. In this framework, GARCH-type models remain valuable for capturing persistence and providing interpretable structure, while machine-learning models enhance forecasting performance under nonlinear dynamics. Future research may extend this work by incorporating higher-frequency data for realized volatility estimation, exploring regime-switching or state-dependent volatility models, and applying broader forecast evaluation techniques such as model confidence sets.

REFERENCES

Adewole, A. P., Akinwale, Y. O., & Akintayo, O. A. (2011). Artificial neural network model for forecasting foreign exchange rates. *World Applied Programming, 1*(3), 110–118.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics, 31*(3), 307–327.

Brooks, C. (2019). Introductory Econometrics for Finance. Cambridge University Press, Cambridge. https://doi.org/10.1017/9781108524872

Central Bank of Nigeria. (2026). Monthly average exchange rates of the naira. Retrieved April 17, 2026, from https://www.cbn.gov.ng/rates/exrate.html

Central Bank of Nigeria. (2026). Exchange rate policy. Retrieved April 17, 2026, from https://www. cbn.gov.ng/intops/exchratepolicy.html

Chen, S., Jeong, K., & Ha¨rdle, W. K. (2008). Support vector regression based GARCH model with application to forecasting volatility of financial returns. *SFB 649 Discussion Papers, 2008 014*, 1–28.

Cortes, C. and Vapnik, V. (1995) Support-Vector Networks. Machine Learning, 20, 273 297. http://dx.doi.org/10.1007/BF00994018

Dahiru, B. A., & Asemota, J. O. (2013). Exchange-rates volatility in Nigeria: Application of GARCH models with exogenous break. *CBN Journal of Applied Statistics, 4*(1), 89–116.

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

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 heteroskedasticity with estimates of the variance of United Kingdom inflation. *Econometrica, 50*(4), 987–1007.

Hornik, K., Stinchcombe, M., & White, H. (1989). Multilayer feedforward networks are universal approximators. *Neural Networks, 2*(5), 359–366.

Jarque, C. M., & Bera, A. K. (1980). Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Economics Letters, 6*(3), 255–259.

Ljung, G. M., & Box, G. E. P. (1978). On a measure of lack of fit in time series models. *Biometrika, 65*(2), 297–303.

Muchtadi-Alamsyah et al., (2024). Support vector regression-based heteroscedastic models for volatility forecasting. *Applied Soft Computing, 161*, 111722.

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

Nkemnole, E. B., & Oyelami, A. S. (2025). Comparative analysis of stochastic models and machine learning algorithms for inflation rate prediction in Nigeria. Science World Journal, 20(2), 516–521. https://doi.org/10.4314/swj.v20i2.1

Panda, C., & Narasimhan, V. (2007). Forecasting exchange rate better with artificial neural network. *Expert Systems with Applications, 30*(3), 109–120.

Peng, Y., & Albuquerque, P. H. M. (2018). Best of two worlds: Forecasting high frequency volatility for cryptocurrencies and traditional currencies with support vector regression. *Expert Systems with Applications, 97*, 177–192.

Pilbeam, K. & Langeland, K. N. (2014). Forecasting exchange rate volatility: GARCH models versus implied volatility forecasts. International Economics and Economic Policy, 12(1), pp. 127-142. doi: 10.1007/s10368-0140289-4

Rammusi, N., Metsileng, D., Botlhoko, T., Tsoku, J. T., & Shogole, L. (2025). Enhancing forecast accuracy of exchange rate volatility using hybrid ANN-GARCH models: Evidence from South Africa, Brazil, and China. *International Journal of Economics and Financial Issues, 15*(6), 140–150.

Sun, H., & Yu, B. (2020). Forecasting financial returns volatility: A GARCH-SVR model. *Computational Economics, 55*(2), 451–471.

Zivot, E., & Andrews, D. W. K. (1992). Further evidence on the great crash, the oil-price shock, and the unit-root hypothesis. *Journal of Business & Economic Statistics, 10*(3), 251–270.