



International Journal of Mathematics and Statistics

DOI: 10.5281/zenodo.23193897 | ISSN: 3141-6438 | Volume: 2 | Issue: 1 | 2026

Forecasting the Naira–pound sterling exchange rate: A comparative analysis of ARIMA, ARIMAX and ARIMA-GARCH models

Timinibife N. Charles¹, T. I. Ayakeme²

¹Department of Mathematics, Niger Delta University, Bayelsa State, Nigeria

²Department of Computer and Mathematics, University of Africa, Toru-Urua, Bayelsa State, Nigeria

Corresponding Author: charlesnoah615@gmail.com

Abstract

Accurate exchange-rate forecasting is important for financial planning, international trade, investment decisions and macroeconomic management. This study develops and compares three time-series forecasting approaches for the Naira–Pound Sterling exchange rate: autoregressive integrated moving average (ARIMA), autoregressive integrated moving average with exogenous variables (ARIMAX), and ARIMA combined with generalized autoregressive conditional heteroskedasticity (ARIMA-GARCH). A synthetic monthly dataset comprising 180 observations from January 2008 to December 2022 was generated specifically for methodological and forecasting-model evaluation. The synthetic observations are not presented as official historical observations. The analysis uses ARIMA(1,1,1) as the benchmark model, while ARIMAX incorporates the interest-rate differential, inflation differential, crude-oil price and Nigerian foreign-exchange reserves. ARIMA-GARCH combines an ARIMA conditional-mean specification with a GARCH(1,1) conditional-variance specification. The first 144 observations were used for model estimation and the final 36 observations were reserved for out-of-sample forecasting. Forecasting accuracy was evaluated using mean absolute error (MAE), root mean square error (RMSE) and mean absolute percentage error (MAPE). The simulated results show that ARIMAX achieved the lowest RMSE of 7.3931, while ARIMA recorded the lowest MAE and MAPE of 6.6043 and 3.4729%, respectively. ARIMA-GARCH produced results very close to the ARIMA benchmark.

Keywords: Naira–Pound Sterling exchange rate; ARIMA; ARIMAX; GARCH; synthetic data; forecasting; volatility; time-series modelling.

1. Introduction

Exchange-rate forecasting is an important component of financial and economic decision-making because exchange-rate movements affect international trade, investment, import costs, export competitiveness and financial planning. For economies such as Nigeria, exchange-rate fluctuations are particularly important because changes in the value of the Naira influence the domestic cost of imported goods and services and the profitability of international transactions.

The Naira–Pound Sterling exchange rate is therefore an important variable for understanding and forecasting currency movements between Nigeria and the United Kingdom. Conventional time-series methods provide useful tools for analysing exchange-rate behaviour because they can capture persistence, trends and stochastic movements in historical observations.

The objective of this study is to compare the forecasting performance of ARIMA, ARIMAX and ARIMA-GARCH under a common out-of-sample experimental design. An important feature of the study is that the dataset is synthetic. The data were generated specifically for methodological and forecasting-model evaluation and are not claimed to be official historical observations from the Central Bank of Nigeria, the Bank of England or another statistical institution.

1.1. Literature review

The literature on exchange-rate forecasting has evolved from classical linear time-series methods toward volatility models, macroeconomic specifications and nonlinear or machine-learning approaches. The Box–Jenkins framework established a systematic procedure for identifying, estimating, diagnosing and forecasting ARIMA models [1, 2]. Forecasting practice emphasizes that model selection should be judged using out-of-sample performance rather than relying only on in-sample fit [3]. For exchange rates, the empirical forecasting problem is particularly difficult: Meese and Rogoff [10] showed that structural exchange-rate models could perform poorly out of sample relative to simple benchmarks. More recent reviews document the expansion of exchange-rate forecasting research and the increasing use of nonlinear and computational approaches [16, 17].

For developing-country currencies, volatility modelling is especially relevant. ARCH was introduced by Engle [6] and generalized through GARCH by Bollerslev [7], while Nelson [8] proposed an asymmetric conditional-variance formulation. These models provide a framework for distinguishing conditional mean dynamics from time-varying uncertainty. Diagnostic testing is also essential: unit-root procedures such as the Dickey–Fuller test help determine whether differencing is required [4], while the Ljung–Box test can be used to examine residual serial correlation [5].

Recent Nigerian studies provide direct evidence for these approaches. Tanko, Musa, and Musa [13] modelled monthly official Naira/US-Dollar exchange rates using ARIMA and GARCH and reported that differencing was required before modelling; their preferred specification combined an ARIMA model with GARCH and a Student-t distribution. Magaji and Garba [14] similarly applied an ARIMA-GARCH framework to monthly Nigerian Naira/US-Dollar exchange rates. Adenomon and Apudi [15] compared ARIMA, GARCH and NNAR approaches using Nigerian exchange-rate data and showed that forecast performance can differ substantially across modelling families.

The present study extends this literature in two ways. First, it focuses specifically on the Naira–Pound Sterling exchange rate rather than the more commonly studied Naira–US-Dollar rate. Second, it compares a univariate ARIMA benchmark with an ARIMAX model that

incorporates external economic variables and an ARIMA-GARCH specification that allows conditional variance to evolve over time. This comparison is important because ARIMAX and GARCH address different sources of forecasting information: exogenous regressors can improve the conditional mean specification, whereas GARCH is designed primarily to capture changing conditional variance [6, 7]. The use of several accuracy measures is also consistent with the forecasting literature because alternative loss functions may rank models differently [3, 9].




Figure 1: Synthetic monthly Naira–Pound Sterling exchange rate.

Figure 1 presents the simulated monthly Naira–Pound Sterling exchange-rate series. The series exhibits substantial month-to-month variation and changes in its level over the sample. Because the observations are synthetic, the figure represents the simulated data-generating process rather than actual historical exchange-rate movements.

2. Model formulation

The study employs a synthetic monthly dataset containing 180 observations covering January 2008 to December 2022. The dependent variable is the Naira–Pound Sterling exchange rate, denoted by Y_t . The ARIMAX explanatory variables are the Nigerian–UK interest-rate differential, Nigerian–UK inflation differential, crude-oil price and Nigerian foreign-exchange reserves.

The first 144 observations were used for model estimation and the final 36 observations were reserved for out-of-sample evaluation.

The model-development strategy follows the Box–Jenkins approach of identification, estimation and diagnostic checking [1, 2]. Stationarity is assessed before estimation because differencing is central to ARIMA modelling [4]. The ACF and PACF provide complementary evidence for selecting autoregressive and moving-average orders [1, 2, 11]. The GARCH specification follows the conditional-variance framework introduced by Engle [6] and generalized by Bollerslev [7]; financial time-series modelling conventions are also discussed by Tsay [12].

$$\Delta Y_t = c + \phi_1 \Delta Y_{t-1} + \theta_1 \varepsilon_{t-1} + \varepsilon_t. \quad (1)$$

The benchmark estimates are $\phi_1 = 0.9237$, $\theta_1 = 0.9838$ and $c = -0.03725$.

$$\Delta Y_t = c + \phi_1 \Delta Y_{t-1} + \theta_1 \varepsilon_{t-1} + \beta_1 ID_t + \beta_2 INF_t + \beta_3 OIL_t + \beta_4 RES_t + \varepsilon_t. \quad (2)$$

$$\sigma_t^2 = \omega + \alpha_1 \varepsilon_{t-1}^2 + \gamma_1 \sigma_{t-1}^2. \quad (3)$$

2.1. Notation, assumptions and forecast evaluation

Here $\Delta Y_t = Y_t - Y_{t-1}$ is the monthly change in the simulated rate, measured as Naira per pound sterling; t denotes month, c is the conditional-mean intercept, ϕ_1 is the autoregressive coefficient, θ_1 is the moving-average coefficient, and ε_t is the innovation. The reported coefficient estimates are retained from the supplied manuscript. For a stationary and invertible ARMA(1,1) representation of the differenced process, the usual restrictions are $|\phi_1| < 1$ and $|\theta_1| < 1$; the reported estimates satisfy these restrictions, although the moving-average estimate is close to the boundary [2, 11].

In the ARIMAX equation, ID_t , INF_t , OIL_t and RES_t denote the interest-rate differential, inflation differential, crude-oil price and foreign-exchange reserves, respectively. The coefficients $\beta_1, \dots, \beta_4$ quantify their contributions to the conditional mean of ΔY_t . Their measurement units, transformations and fitted coefficients were not included in the supplied files. This equation specifies regressors in the differenced-response equation; software implementation must match that specification rather than silently fitting a regression in levels with ARIMA errors.

For the volatility component, write $\varepsilon_t = \sigma_t z_t$, where $E(z_t | \mathcal{F}_{t-1}) = 0$ and $E(z_t^2 | \mathcal{F}_{t-1}) = 1$, and $\mathcal{F}_{t-1}$ denotes the information available before month t . The parameters satisfy $\omega > 0$, $\alpha_1 \geq 0$ and $\gamma_1 \geq 0$; the restriction $\alpha_1 + \gamma_1 < 1$ ensures a finite unconditional innovation variance under the standard GARCH assumptions [7, 12]. The symbol γ_1 distinguishes the GARCH persistence coefficient from the ARIMAX regression coefficient β_1 .

Let $n = 36$, y_j be the observed synthetic test value, $\hat{y}_j$ its forecast, and $e_j = y_j - \hat{y}_j$. The evaluation criteria are

$$\text{MAE} = \frac{1}{n} \sum_{j=1}^n |e_j|, \quad (4)$$

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{j=1}^n e_j^2}, \quad (5)$$

$$\text{MAPE} = \frac{100}{n} \sum_{j=1}^n \left| \frac{e_j}{y_j} \right|, \quad y_j \neq 0. \quad (6)$$

MAE and RMSE have the units of the exchange rate, whereas MAPE is expressed as a percentage. A valid evaluation requires all estimation and preprocessing to use only the training observations. Future exogenous values must either be available at the forecast origin or forecast separately; using their subsequently observed values produces a conditional forecasting experiment [3].

2.2. Computational reproducibility

The supplied package contains the manuscript and six plotted figures, but no numerical dataset, data-generation program, fitting program or saved forecast vectors. Accordingly, the reported estimates, diagnostic p-value and forecast scores are preserved as supplied results and have not been independently reproduced. Full replication requires the random seed and data-generating

equations, all 180 observations and regressors, software versions, estimation settings, fitted coefficients, diagnostic lags, and the forecast-origin protocol. In particular, it must be stated whether the 36 test forecasts are fixed-origin multistep forecasts or sequential one-step forecasts. The accompanying validation script checks citation coverage and the internal ordering of the reported scores; it does not estimate the models or reproduce the experiment.




Figure 2: First-differenced synthetic exchange-rate series.

Figure 2 shows the first difference of the simulated exchange-rate series. Differencing reduces persistent level movements and produces a series fluctuating around a more stable level.




Figure 3: ACF of the first-differenced exchange-rate series.

Figure 3 presents the ACF of the differenced series. The ACF provides information about

remaining serial dependence and helps assess autoregressive and moving-average dynamics.




Figure 4: PACF of the first-differenced exchange-rate series.

Figure 4 presents the PACF of the differenced series. Together with the ACF, it supports assessment of a parsimonious low-order ARIMA structure.

3. Results and discussion

The three models were evaluated using the final 36 observations as the out-of-sample forecasting period. ARIMAX produced the lowest RMSE of 7.3931, compared with 7.4384 for ARIMA and 7.4385 for ARIMA-GARCH. However, ARIMA produced the lowest MAE of 6.6043 and MAPE of 3.4729%. ARIMA-GARCH produced results almost identical to ARIMA.

The manuscript reports an ARCH-LM diagnostic p -value of approximately 0.4773, but does not specify the residual series or lag order used. Since $p > 0.05$, there is insufficient evidence of an ARCH effect in the simulated series at the 5% significance level. For the tested residual series, this indicates no statistically detected conditional heteroskedasticity at that level. It does not establish an improvement or deterioration in point-forecast accuracy.

The reported RMSE advantage of ARIMAX relative to ARIMA is approximately 0.61%. No predictive-accuracy test is reported, so statistical significance of this difference is not established. When the ARIMA mean parameters are held fixed, adding GARCH changes the conditional variance rather than the conditional mean; close point forecasts therefore do not by themselves evaluate volatility forecasting performance. However, the improvement is modest, and ARIMA remains highly competitive according to MAE and MAPE. The results also demonstrate that greater model complexity does not automatically translate into improved forecasting accuracy.

The interpretation of the competing forecasts follows the principle that out-of-sample accuracy is more informative for forecasting selection than in-sample fit alone [3]. The use of MAE, RMSE and MAPE recognizes that models can have different rankings under different loss functions [3, 9]. The exchange-rate forecasting benchmark literature includes Meese and Rogoff [10], while the Nigerian evidence reported by Tanko et al. [13], Magaji and Garba [14], and Ade-

nomon and Apudi [15] also demonstrates the practical relevance of comparing ARIMA-family and volatility-based models for Nigerian exchange-rate data.

Table 1: Out-of-sample forecasting performance

Model	MAE	RMSE	MAPE (%)	RMSE Rank
ARIMA(1,1,1)	6.6043	7.4384	3.4729	2
ARIMAX(1,1,1)	6.7110	7.3931	3.5379	1
ARIMA-GARCH(1,1)	6.6043	7.4385	3.4729	3




Figure 5: Out-of-sample actual and forecast exchange rates.

Figure 5 compares the out-of-sample observations with forecasts from ARIMA, ARIMAX and ARIMA-GARCH. The supplied figure displays the simulated observations and the three forecast paths. ARIMAX gives the lowest RMSE, suggesting a modest advantage in controlling larger forecast errors.




Figure 6: Forecasting error comparison.

Figure 6 compares MAE, RMSE and MAPE. ARIMAX has the lowest RMSE, while ARIMA records the lowest MAE and shares the lowest MAPE with ARIMA-GARCH. Thus, model ranking depends on the chosen loss criterion.

4. Conclusion

This study compared ARIMA, ARIMAX and ARIMA-GARCH models for forecasting the Naira–Pound Sterling exchange rate using a synthetic monthly dataset containing 180 observations from January 2008 to December 2022.

ARIMAX achieved the lowest RMSE of 7.3931, while ARIMA achieved the lowest MAE and MAPE of 6.6043 and 3.4729%, respectively. ARIMA-GARCH did not materially improve upon the ARIMA benchmark.

The principal conclusion is that ARIMAX is the preferred model when RMSE is used as the principal model-selection criterion within this synthetic experiment, while ARIMA remains competitive under MAE and MAPE.

The most important limitation is that the dataset is synthetic. The numerical results are simulation results and should not be interpreted as estimates of actual Naira–Pound Sterling exchange-rate behaviour. Future research should apply the framework to verified historical observations from authoritative Nigerian and UK sources.

Overall, the findings support the forecasting literature’s emphasis on empirical validation and parsimony [3, 9]. They also extend the Nigerian evidence by illustrating a synthetic Naira–Pound Sterling forecasting comparison; they do not establish new empirical evidence about historical Nigerian exchange rates.

Related regional research also examines exchange-rate comovements, as in the Colombian study by Sierra and Rodríguez [18].

Data and code availability

The supplied figures and manuscript are included in the publication package. The underlying synthetic observations and original model-fitting program were not supplied.

References

- [1] Box, G. E. P., & Jenkins, G. M. (1970). *Time series analysis: Forecasting and control*. Holden-Day.
- [2] Box, G. E. P., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time series analysis: Forecasting and control* (5th ed.). Wiley.
- [3] Hyndman, R. J., & Athanasopoulos, G. (2021). *Forecasting: Principles and practice* (3rd ed.). OTexts.
- [4] 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.
- [5] Ljung, G. M., & Box, G. E. P. (1978). On a measure of lack of fit in time series models. *Biometrika*, 65(2), 297–303.
- [6] Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4), 987–1007.
- [7] Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), 307–327.
- [8] Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 59(2), 347–370.
- [9] Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253–263.
- [10] Meese, R. A., & Rogoff, K. (1983). Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics*, 14(1–2), 3–24.
- [11] Hamilton, J. D. (1994). *Time series analysis*. Princeton University Press.
- [12] Tsay, R. S. (2010). *Analysis of financial time series* (3rd ed.). Wiley.
- [13] Tanko, N. A., Musa, G. K., & Musa, S. A. (2023). On the use of ARIMA and GARCH in modelling Nigeria's Naira–US Dollar monthly exchange rates. *Asian Journal of Probability and Statistics*, 22(2), 8–18. <https://doi.org/10.9734/ajpas/2023/v22i2479>
- [14] Magaji, B., & Garba, J. (2022). Forecasting the exchange rate of Nigerian Naira to United States Dollar using ARIMA-GARCH model. *Dutse Journal of Pure and Applied Sciences*, 8(3b), 87–96. <https://doi.org/10.4314/dujopas.v8i3b.9>
- [15] Adenomon, M. O., & Apudi, P. E. (2024). Comparison of ARIMA, GARCH and NNAR models for modelling exchange rate in Nigeria. *Journal of the Royal Statistical Society Nigeria Group*, 1(1).
- [16] Vasconcelos, C. de S., & Hadad Júnior, E. (2023). Forecasting exchange rate: A bibliometric and content analysis. *International Review of Economics & Finance*, 83, 607–628. <https://doi.org/10.1016/j.iref.2022.09.006>

-
- [17] Fang, S., & Wei, Y. (2024). 30 years of exchange rate analysis and forecasting: A bibliometric review. *Journal of Economic Surveys*. <https://doi.org/10.1111/joes.12572>
- [18] Sierra, G. M., & Rodríguez, N. J. M. (2019). Modelación y comovimientos de la tasa de cambio colombiana, 2011–2017. *Revista de Métodos Cuantitativos para la Economía y la Empresa*. <https://doi.org/10.46661/revmetodoscuanteconempresa.2966>

Copyright and License

Copyright © 2026 The Author(s).

This article is published by *International Journal of Mathematics and Statistics* under Ktrend Journals and is distributed under the terms of the Creative Commons Attribution License (CC BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Journal Website: <https://ktrendjournals.org/>