Explainable Artificial Intelligence for Inflation Forecasting in Nigeria: Interpreting A Multilayer Perceptron Model Using Shap Values.

SAIDAT FEHINTOLA OLANIRAN[1], ABDULRAHMAN FATIMOH BUKOLA[2]
[1,2]*Department of Mathematics and Statistics, Kwara State University, Malete, Kwara State, Nigeria.*

Abstract- ***Inflation forecasting is important for monetary policy, economic planning and decision-making, yet neural-network forecasts can be difficult to interpret because their internal relationships are not directly observable. This study applies SHapley Additive exPlanations (SHAP) to interpret a Multilayer Perceptron (MLP) model developed for forecasting inflation in Nigeria. Monthly macroeconomic data covering 2007–2024 were used, comprising inflation rate, interest rate, exchange rate, oil price and Consumer Price Index (CPI), together with lagged and temporal features used in the MLP. The SHAP analysis was used to explain an individual inflation prediction and identify the direction and relative magnitude of feature contributions. The model prediction examined increased from a baseline of 10.20 to 13.49. The largest positive contributions were associated with the third and first lags of oil price (OP_lag3 and OP_lag1), followed by seasonal and lagged inflation features. CPI variables also contributed positively, whereas exchange-rate and interest-rate variables showed smaller and mixed effects for the prediction examined. The result is consistent with Nigerian evidence that oil-price movements can transmit into domestic inflation, while the importance of lagged inflation supports inflation persistence. The study also demonstrates that SHAP can make an MLP more transparent without treating feature importance as causal evidence. Although the broader thesis found that the MLP did not outperform SARIMA and VAR in out-of-sample accuracy, the SHAP analysis shows that the model learned economically meaningful temporal and macroeconomic information. The study concludes that explainable machine learning can complement conventional inflation forecasting by providing interpretable information about the drivers of individual predictions.***

Keywords— inflation forecasting; multilayer perceptron; artificial neural network; SHAP; explainable artificial intelligence

I. INTRODUCTION

Inflation is a central macroeconomic concern because persistent increases in the general price level reduce purchasing power, affect investment and consumption decisions, and complicate monetary-policy implementation. Accurate inflation forecasts are consequently valuable to central banks, governments, firms, and households (Medeiros et al., 2021; Bank for International Settlements [BIS], 2021). In Nigeria, inflation forecasting is complicated by exchange-rate movements, oil-price shocks, monetary conditions, supply disturbances, and structural changes in the economy (Adenuga et al., 2012; Bawa et al., 2020). In particular, fluctuations in oil prices can transmit to domestic inflation through changes in production costs, government revenues, transportation costs, and the prices of energy-intensive goods and services (Adenuga et al., 2012; Bawa et al., 2020). Traditional approaches such as Autoregressive Integrated Moving Average (ARIMA), Seasonal Autoregressive Integrated Moving Average (SARIMA), and Vector Autoregression (VAR) have been extensively used because they provide established statistical frameworks for modelling temporal dependence and interrelationships among macroeconomic variables. However, inflation may also contain nonlinear patterns and complex interactions that are not adequately represented by conventional linear specifications. Machine-learning methods, including Artificial Neural Networks (ANNs), offer an alternative because of their ability to approximate complex nonlinear relationships among multiple predictors (Zhang, 2003; Haykin, 2009; Medeiros et al., 2021). The Multilayer Perceptron (MLP) is especially relevant because it can combine current and lagged economic information through interconnected neurons and nonlinear activation functions, enabling it to model complex relationships between explanatory variables and inflation (Haykin, 2009). Nevertheless, the flexibility of neural-network models creates an interpretability problem. A policymaker may obtain an accurate or informative forecast without knowing which economic variables caused the prediction to move upward or downward. This lack of transparency is commonly described as the "black-box" problem associated with many machine-learning models (Aras & Lisboa, 2022). SHapley Additive exPlanations (SHAP) provides a framework for addressing this interpretability problem by assigning contribution values to individual input features based on their influence on a model's prediction (Lundberg & Lee, 2017). SHAP has increasingly been applied to machine-learning forecasting models to identify the variables that contribute most strongly to predictions and to determine the direction and magnitude of their effects (Aras & Lisboa, 2022). This makes SHAP particularly useful for economic forecasting, where understanding the factors driving a prediction can be as important as the prediction itself. The present article focuses on the third objective of the underlying study: to use SHAP values to assess feature importance and interpret the MLP model in forecasting Nigerian inflation. The contribution of the article is therefore not another comparison of forecasting errors, but an examination of what the neural network learned from Nigerian macroeconomic data. Specifically, the SHAP analysis is used to identify the relative importance of the explanatory variables and to provide an interpretable account of their contributions to the MLP's inflation predictions. This approach provides additional economic insight into the determinants of inflation while improving the transparency and practical usefulness of nonlinear machine-learning forecasts.

II. RELATED LITERATURE.

Machine-learning methods have become increasingly relevant to inflation forecasting because they can model nonlinearities and interactions that may be difficult to specify in conventional econometric models. Medeiros et al. demonstrated that machine-learning methods can exploit a rich information set when forecasting inflation and can be useful in settings where relationships between predictors and inflation are nonlinear [8]. Their evidence supports the use of flexible predictive methods while also emphasizing the importance of appropriate model evaluation. Recent work has also considered neural networks for inflation forecasting. Paranhos examined recurrent neural networks for inflation prediction and reported that neural-network models can capture dynamic dependencies in inflation, although neural methods are not necessarily superior to conventional benchmarks in every forecasting environment [10]. This observation is relevant to the present study because the MLP in the broader analysis recorded higher forecasting errors than SARIMA and VAR. Thus, neural-network usefulness should not be equated automatically with superior predictive accuracy. The black-box nature of machine-learning models has stimulated research on Explainable Artificial Intelligence (XAI). Lundberg and Lee introduced SHapley Additive exPlanations as a unified framework for interpreting predictions from complex models [3]. SHAP is based on Shapley values from cooperative game theory and allocates the difference between a model's baseline output and a particular prediction among the input features. The basic additive representation can be written as:

$$f(x) = \varphi_0 + \Sigma(j=1 \text{ to } p)\ \varphi_j$$

where f(x) is the prediction for observation x, φ_0 is the baseline or expected prediction, and φ_j is the contribution attributed to feature j. A positive contribution moves the prediction upward relative to the baseline, whereas a negative contribution moves it downward. Importantly, the contribution is not the same thing as the numerical value of the feature. Aras and Lisboa applied explainable machine-learning techniques to inflation forecasting and showed the value of combining predictive models with Shapley-based explanations [4]. Their work is particularly relevant because it demonstrates that inflation forecasting can benefit from an explicit explanation of the information driving a machine-learning forecast. The economic interpretation of the present SHAP results is supported by Nigerian evidence. Adenuga et al. reported significant relationships between oil prices and inflation in Nigeria and also identified exchange-rate and interest-rate conditions as relevant inflation drivers [5]. Bawa et al. found asymmetric effects of oil-price movements on Nigerian inflation, providing further evidence that energy-market conditions can transmit into domestic prices [6]. These studies provide an economic basis for expecting oil-price variables to matter in a Nigerian inflation model. The present study extends that evidence by examining the contribution of lagged oil-price variables within a nonlinear MLP rather than interpreting a single econometric coefficient.

Research Gap

Despite the growing use of machine learning for economic forecasting, two gaps remain. First, much of the inflation literature focuses on forecast accuracy rather than explaining individual predictions. Second, explainable inflation forecasting using neural networks remains comparatively limited for Nigeria. The present study addresses both issues by applying SHAP to the MLP and relating the resulting feature contributions to established evidence on oil-price transmission, inflation persistence and macroeconomic conditions.

III. MATERIALS AND METHODS

Data and Variables- The study used monthly secondary data covering January 2007 to December 2024, giving 216 observations. Inflation rate (INF) was the response variable. The principal macroeconomic predictors were interest rate (IR), exchange rate (ER), oil price (OP) and Consumer Price Index (CPI). The data sources reported in the underlying study include the Central Bank of Nigeria Statistical Bulletin, National Bureau of Statistics and World Bank databases. To represent temporal dependence, lagged variables were included in the MLP feature set. The SHAP plot examined in this article contains lagged oil price, lagged inflation, lagged exchange rate, lagged CPI and lagged interest-rate features. Temporal features were also incorporated through the time trend and sine/cosine representations of the month.

Multilayer Perceptron- The MLP is a feed-forward neural network consisting of an input layer, one or more hidden layers and an output layer. The input layer receives the selected macroeconomic and temporal variables, while the output layer produces the inflation prediction. The hidden layer permits nonlinear transformations of the inputs. For a generic hidden neuron, the network can be represented as: $h = g(\beta_0 + \beta_1 x_1 + \beta_2 x_2 + \ldots + \beta_p x_p)$ where g(.) is the activation function, β_0 is the intercept or bias term, β_j are connection weights, and x_j are input features. The final output is obtained by combining the transformed hidden-unit responses. Because the network contains multiple nonlinear transformations, individual weights do not provide a simple marginal-effect interpretation.

SHAP-Based Interpretation- SHAP was applied after fitting the MLP to explain the model output. The analysis considered a selected observation for which the model generated an inflation prediction of 13.49. The baseline prediction was 10.20. The difference of 3.29 percentage points represents the net contribution of the features for that observation. A local SHAP explanation is different from a global feature-importance analysis. The figure used here explains one prediction. Therefore, the ranking should be described as the relative contribution of features to the prediction examined rather than as a universal ranking for the entire dataset. A global mean absolute SHAP ranking would require the underlying SHAP matrix or a summary plot containing aggregated values.

Interpretation Rule- Positive SHAP contributions indicate that a feature pushed the model prediction above the baseline, while negative contributions indicate that it pushed the prediction below the baseline. The horizontal magnitude of each bar represents the strength of the contribution. The feature value displayed next to a variable is an input value and should not be confused with the SHAP contribution.

Forecasting Context-The underlying comparative analysis evaluated SARIMA, VAR and MLP using out-of-sample error measures. The reported MLP errors were RMSE = 8.2607, MAE = 7.6524 and MAPE = 42.1021%, compared with RMSE = 2.8461, MAE = 2.7532 and MAPE = 15.2748% for VAR.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1722756

These results are retained as context: the purpose of SHAP in this article is interpretation, not to claim that the MLP was the best forecasting model.

IV. RESULTS AND DISCUSSIONS

The SHAP analysis explained an MLP prediction of 13.49 from a baseline prediction of 10.20. Thus, the combined net contribution of the features was approximately +3.29. Figure 1 presents the local SHAP contribution plot used for the interpretation.

Fig. 1. Local SHAP contribution plot for the MLP inflation prediction.

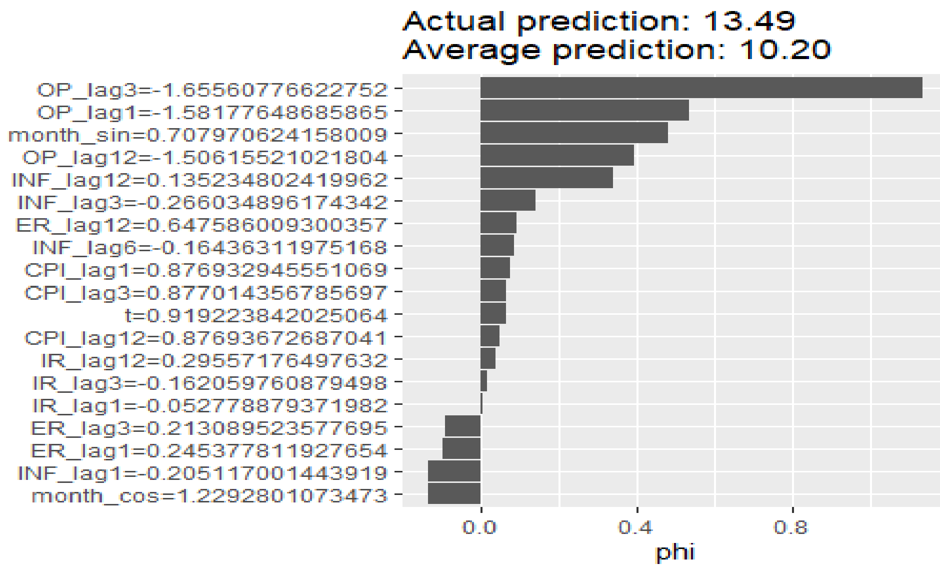

The strongest positive contributions in the displayed prediction were associated with OP_lag3 and OP_lag1. OP_lag12 and INF_lag12 also made positive contributions, while month_sin contributed positively and month_cos contributed negatively. The inflation lags therefore showed different directions, illustrating that the MLP did not impose a uniform effect across all lags.

Feature Contribution Table

Feature	Direction	Relative strength	Interpretation
OP_lag3	Positive	Very high	Largest displayed positive contribution
OP_lag1	Positive	Very high	Second major positive contribution
month_sin	Positive	High	Seasonal contribution to this prediction
OP_lag12	Positive	High	Longer lagged oil-price contribution
INF_lag12	Positive	High	Historical inflation contribution
CPI_lag1 / CPI_lag3	Positive	Moderate	Past consumer-price information
ER_lag1 / ER_lag3	Negative	Low–moderate	Reduced the prediction for this case
INF_lag1	Negative	Moderate	Short-run inflation lag reduced prediction
month_cos	Negative	Moderate	Seasonal component reduced prediction

Table 1 summarizes the relative contributions visible in the local explanation. Because the thesis figure does not provide the numerical SHAP φ-values for each bar, numerical SHAP contributions are deliberately not fabricated here. The relative-strength labels are based on the visual magnitude and ordering in the supplied figure.

The dominance of OP_lag3 and OP_lag1 is the central empirical finding of the local SHAP analysis. Both variables pushed the prediction upward, and their contributions were visibly larger than those of most other features. OP_lag12 also contributed positively, indicating that information from different historical horizons was used by the network. This pattern suggests that the MLP captured delayed information from the oil market rather than relying solely on contemporaneous oil-price conditions. The result is economically plausible for Nigeria because oil-price movements can affect foreign-exchange conditions, fiscal revenues, production costs and domestic prices through several transmission channels. The finding is consistent with Adenuga et al. [1], who documented oil-price pass-through to Nigerian inflation, and Bawa et al. [4], who found asymmetric effects of oil-price movements on inflation. Historical inflation was also important. INF_lag12 and INF_lag3 pushed the prediction upward, while INF_lag1 pushed it downward. This mixed pattern is important because it demonstrates the nonlinear character of the MLP. A conventional linear interpretation might expect a single coefficient sign for each lag, but SHAP shows how the observed feature value contributed to this particular prediction within the fitted nonlinear model. The importance of lagged inflation is consistent with the concept of inflation persistence, in which previous inflation conditions contain information about subsequent inflation. The result also complements the machine-learning evidence reported by Medeiros et al. [8], who emphasized the value of nonlinear information contained in lagged macroeconomic variables. CPI_lag1 and CPI_lag3 made positive contributions in the displayed prediction. This indicates that historical consumer-price information was useful to the MLP when constructing the forecast. The result is intuitive because CPI summarizes movements in consumer prices and therefore provides information closely related to the inflation process. Exchange-rate effects were mixed. ER_lag12 contributed positively, whereas ER_lag1 and ER_lag3 contributed negatively. This does not mean that exchange rate is economically unimportant. Rather, the local SHAP plot shows that its contribution varied by lag for the prediction examined. Nigerian evidence has established the importance of exchange-rate movements for inflation transmission [1], [4]. Interest-rate features made smaller contributions than the leading oil-price and inflation features. This result should be interpreted cautiously. A small local SHAP contribution does not establish that monetary policy has no effect on inflation; it only means that the interest-rate values for the selected observation contributed less to this model prediction than the dominant features. The seasonal variables also had opposing effects. month_sin contributed positively, while month_cos contributed negatively. Their contributions were smaller than those of the dominant oil-price lags. This suggests that the MLP used some seasonal information, but the displayed prediction was driven more strongly by macroeconomic and historical variables.

The strongest finding large contributions from lagged oil-price variables is broadly consistent with the Nigerian literature. Adenuga et al. [1] found that oil-price movements affected domestic inflation, while Bawa et al. [4] provided evidence of asymmetric oil-price effects. The present study adds an explainable machine-learning perspective: rather than estimating an average econometric coefficient, it shows that lagged oil-price information was among the strongest contributors to a particular MLP inflation prediction. The temporal pattern is also informative. The importance of OP_lag1, OP_lag3 and OP_lag12 indicates that the model used information from multiple historical horizons. This supports an interpretation in which oil-price shocks may transmit through the Nigerian economy with delays rather than instantaneously. The result is consistent with the broader machine-learning inflation literature. Medeiros et al. [8] showed that machine-learning methods can benefit from rich sets of lagged macroeconomic information and nonlinear relationships. The present MLP similarly relied strongly on lagged variables. Paranhos [2] found that neural-network approaches can capture inflation dynamics but are not necessarily superior to benchmark models. This is consistent with the broader study, where the MLP had higher out-of-sample errors than SARIMA and VAR. The implication is that model accuracy and model interpretability are distinct dimensions. A model can fail to minimize forecasting error while still providing useful information about the structure it learned. SHAP is therefore best viewed as a complement to forecast evaluation rather than as an alternative accuracy metric. The present result also aligns with the explainable inflation forecasting literature. Lundberg and Lee [7] established SHAP as a framework for assigning feature contributions to model predictions, while Aras and Lisboa [13] demonstrated the usefulness of Shapley-based explanations in inflation forecasting. The present study extends this line of work to an MLP using Nigerian macroeconomic data. An important contribution is the distinction between feature value and feature contribution. For example, OP_lag3 is displayed with a negative input value in the figure, yet its SHAP contribution is positive. This illustrates that SHAP does not simply reproduce the sign of a predictor. Instead, it explains how the fitted model transformed the observed feature value in combination with the other inputs. The prominence of lagged oil-price information suggests that policymakers should monitor international oil-market conditions when assessing future inflation risks. Historical oil-price movements may contain information about subsequent inflation through exchange-rate, fiscal and cost channels. The contribution of lagged inflation also supports the importance of persistence in inflation monitoring. CPI information reinforces the value of observing historical price developments, while the mixed exchange-rate contributions suggest that the timing of exchange-rate movements matters for individual predictions.

However, the SHAP results should not be interpreted causally. A large SHAP contribution means that the feature was important for the model's prediction; it does not prove that changing the feature would cause inflation to change by the same amount. Causal interpretation requires a separate identification strategy.

V. IMPLICATIONS, LIMITATIONS AND RECOMMENDATIONS

The study has several implications for inflation monitoring in Nigeria. First, oil-price movements and their lags deserve close attention in inflation surveillance. Second, historical inflation should be retained in forecasting systems because persistence contains predictive information. Third, CPI and monetary indicators remain useful complementary variables. Finally, explainable machine learning can help policymakers inspect why a particular forecast has moved above or below a model's baseline. The study contributes to knowledge in three principal respects. First, it applies SHAP to an MLP inflation forecasting model using Nigerian macroeconomic data. Second, it provides an interpretable account of the features contributing to an individual forecast rather than treating the neural network as an entirely opaque model. Third, it connects the machine-learning result with established Nigerian evidence on oil-price pass-through and inflation dynamics. The most important limitation is that the supplied SHAP output is a local explanation for one prediction. It should therefore not be described as a complete global ranking of feature importance. A full global analysis would require the SHAP values for all test observations and would normally report mean absolute SHAP values and, ideally, a beeswarm summary plot. A second limitation is that the MLP did not provide the best out-of-sample forecast accuracy in the underlying comparison. The MLP recorded RMSE of 8.2607, MAE of 7.6524 and MAPE of 42.1021%, whereas VAR produced lower errors. The interpretability result should therefore be presented as a complementary finding rather than evidence of forecasting superiority. A third limitation is that SHAP is not a causal method. Its contributions describe the behaviour of the fitted model for the observed data and should not be interpreted as structural economic effects

Future studies should compute mean absolute SHAP values across all test observations to obtain a global feature-importance ranking. Future research should investigate hybrid SARIMA–MLP and VAR–MLP models to combine strong forecasting performance with nonlinear learning. Additional macroeconomic variables, including money supply, unemployment, fiscal indicators and selected food-price measures, may be considered where data quality permits. Researchers should report both predictive accuracy and explainability so that machine-learning forecasts can be evaluated from statistical and policy perspectives. Overall, the study recommends that SHAP values should be incorporated into inflation forecasting frameworks as an interpretability and policy-support tool. The variables with the largest SHAP contributions should receive particular attention in inflation surveillance and early-warning systems. However, SHAP importance should not be interpreted as proof that a variable independently causes inflation; rather, it indicates how strongly that variable contributed to the MLP model's prediction within the data and model specification used in this study.

VI. CONCLUSION

This study used SHAP to interpret an MLP model for Nigerian inflation forecasting. The local explanation showed that the prediction increased from 10.20 to 13.49, with OP_lag3 and OP_lag1 providing the strongest positive contributions. Lagged inflation and CPI also contributed meaningfully, while exchange-rate and interest-rate variables showed smaller and mixed effects. The result is consistent with Nigerian evidence on oil-price transmission and with machine-learning research emphasizing lagged macroeconomic information. The central conclusion is that SHAP provides a practical way of opening the black box of neural-network inflation forecasting. While the MLP was not the most accurate model in the underlying comparison, the SHAP explanation demonstrated that the network learned economically meaningful temporal information. Explainable machine learning can therefore complement conventional econometric forecasting by showing not only what the model predicts, but also which information contributed to that prediction.

REFERENCES

[1] Adenuga, A. O., Hilili, M. J., & Evbuomwan, O. O. (2012). Oil price pass-through into inflation: Empirical evidence from Nigeria. *Economic and Financial Review, 50*(1). Crossref

[2] Aras, S., & Lisboa, P. J. G. (2022). Explainable inflation forecasts by machine learning models. *Expert Systems with Applications, 207*, 117982. Crossref

[3] Bank for International Settlements. (2021). What does machine learning say about the drivers of inflation? (BIS Working Papers No. 980).

[4] Bawa, S., Abdullahi, I. S., Tukur, D. S., Barda, S. I., & Adams, Y. J. (2020). Asymmetric impact of oil price on inflation in Nigeria. *CBN Journal of Applied Statistics, 11*(2), 85–113.

[5] Central Bank of Nigeria. (n.d.). Statistical bulletin (Various monthly and annual issues). Abuja, Nigeria: Author.

[6] Haykin, S. (2009). Neural networks and learning machines (3rd ed.). Pearson.

[7] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in neural information processing systems* (Vol. 30, pp. 4765–4774).

[8] Medeiros, M. C., Vasconcelos, G. F. R., Veiga, Á., & Zilberman, E. (2021). Forecasting inflation in a data-rich environment: The benefits of machine learning methods. *Journal of Business & Economic Statistics, 39*(1), 98–119. Crossref

[9] National Bureau of Statistics. (n.d.). Consumer Price Index and inflation reports (Various issues). Abuja, Nigeria: Author.

[10] Paranhos, L. (2025). Predicting inflation with recurrent neural networks. *International Journal of Forecasting.* Crossref

[11] Varian, H. R. (2019). Artificial intelligence, economics, and industrial organization. In *The economics of artificial intelligence.* University of Chicago Press.

[12] World Bank. (n.d.). World development indicators. World Bank DataBank.

[13] S. Aras and P. J. G. Lisboa, “Explainable inflation forecasts by machine learning models,” *Expert Systems with Applications*, vol. 207, 2022, Art. no. 117982, doi: 10.1016/j.eswa.2022.117982. Crossref